I M P E R I A L

Imperial College London

Department of Mechanical Engineering

In Vivo Imaging of Collagen Fibre Structures Using Dedicated Low-Field MRI

Author

Harry Lanz

Supervisor

Mihailo Ristic

August 2025

A thesis submitted for the degree of Doctor of Philosophy (PhD)

Abstract

Anisotropic collagen-rich tissues such as the anterior cruciate ligament (ACL) and meniscus are notoriously difficult to assess with conventional magnetic resonance imaging (MRI). These tissues often appear hypointense, yet, at certain orientations within an MRI scanner they appear bright due to the “magic angle” effect.

The magic angle effect arises from the anisotropic collagen fibre microstructure of connective tissues and their orientation relative to the main magnetic field, $\vec{B}_0$. Signal intensity is minimised when fibres align parallel to $\vec{B}_0$, and is maximised when they are oriented at the magic angle, 54.7°, relative to $\vec{B}_0$. Although conventionally regarded as an artefact, this phenomenon can, if systematically exploited, provide quantitative, non-invasive information on collagen fibre microstructure.

To this end, a rotatable low-field MRI system was developed, enabling reorientation of $\vec{B}_0$ relative to a static subject. This facilitated the first in vivo implementation of Magic Angle Directional Imaging (MADI). MADI utilises multiple acquisitions at different orientations to induce magic angle effect intensity variations, which are then analysed to estimate collagen fibre orientations in tissues such as the ACL and meniscus.

This thesis presents the development of a complete MADI workflow tailored to this prototype system. A bespoke multi-stage registration pipeline was developed to align multi-orientation images and correct geometric distortions intrinsic to permanent magnet systems. An extensible semi-automatic segmentation framework was implemented to reduce operator burden whilst retaining segmentation precision. A significantly accelerated collagen fibre orientation estimation algorithm was developed, enabling optimisation of scan orientation selection using a genetic algorithm. Novel visualisation approaches were also developed to interpret the unique MADI data.

The present work culminated in the first successful in vivo application of MADI in healthy volunteers and patients with ACL and meniscal pathology, with results validated against arthroscopic findings, demonstrating the feasibility and potential for clinical translation of MADI.

Statement of Originality

I declare that the work presented in this thesis is my own and has not been submitted for any other degree or professional qualification. Where work of others has been used or quoted, it is appropriately referenced.

I declare that the figures, diagrams, and images presented in this thesis are my own. Where figures, diagrams, or images produced by others are presented they have been appropriately referenced or attributed.

This thesis includes work carried out in collaboration with others. The contributions of collaborators have been clearly acknowledged.

This work was conducted in accordance with the research integrity and ethical guidelines of Imperial College London.

Copyright statement

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-NonCommercial 4.0 International Licence (CC BY-NC).

Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose.

When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes.

Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Papers and Presentations from this Thesis

The work from this thesis has resulted in the following publications:

H. Lanz et al. “Minimum number of scans for collagen fibre direction estimation using Magic Angle Directional Imaging (MADI) with a priori information”. In: *Array* 17 (Mar. 2023), p. 100273. DOI: 10.1016/J.ARRAY.2022.100273

M. Ristic et al. “First in-vivo magic angle directional imaging using dedicated low-field MRI”. In: *Magnetic Resonance in Medicine* 93.3 (2025), pp. 1077–1089. DOI: 10.1002/mrm.30332

The work from this thesis has resulted in the following presentations:

ISMRM Annual Meeting 2022, London (digital poster): H. Lanz et al. “Minimum number of scans for magic angle directional imaging based on a priori anatomical information”

ISMRM British Chapter Meeting 2023, Aberdeen (traditional poster): H. Lanz et al. “3D Distortion Compensation Using a 3D Printed Phantom and Soft Registration”

ISMRM Annual Meeting 2024, Singapore (power pitch presentation): H. Lanz et al. “Collagen Fibre Direction Estimation Using a Prototype Rotatable Low-Field Scanner”

ISMRM Annual Meeting 2025, Honolulu (traditional poster): H. Lanz et al. “Non-Invasive Visualisation of the Anteromedial and Posterolateral Bundles of the Anterior Cruciate Ligament In-Vivo using Virtual Arthroscopy”

Acknowledgements

I would first like to thank my supervisor, Dr Mike Ristic, for his unwavering support and guidance during this process, and his determination in driving this project forward. I am especially grateful to him for introducing me to the field of MRI and to our unique niche within it. Working on this project has been a fascinating opportunity to explore my academic interests, and I have learnt an enormous amount in the process. His passion and commitment have been instrumental in advancing this idea that our team truly believes in, and I hope that the research we have carried out will one day contribute to helping many people.

I am also grateful to the lab team, particularly Dr Karyn Chappell and Dr John McGinley, who worked tirelessly to develop our scanner and push it to the incredible level of performance it has reached today. Their technical expertise, persistence, and problem-solving were invaluable, and the achievements we have shared as a group go far beyond what I could have imagined when I first began this project.

I owe the deepest thanks to my mother, whose support, encouragement, perseverance, and sacrifices gave me and my brother the educational opportunities to pursue our passions. She is always willing to listen and provide a new perspective to any problem, and always shares in the joy of success with enthusiasm. I could not have reached this stage without her care and support.

I would also like to acknowledge my Grandad, Barry Lanz, who inspired and nurtured my love of science, engineering, and mathematics from a young age. He was an Imperial alumnus (Mining, 1958), and I was proud to follow in his footsteps when I was accepted to study at the college. I hope this thesis would have made him proud, and I regret that he is not here to see me complete this journey.

Finally, I am profoundly grateful to my partner, Kimberley, who has been by my side throughout the entirety of this PhD. She has shared in every high and low, helping me through my frustrations and celebrating my successes, and her support has meant more than I can express.

This experience would not have been as rich or complete without the people mentioned here, and I am enormously grateful for their involvement in such an important and rewarding chapter of my life.

Contents

List of Figures **14**

List of Tables **20**

Definitions **22**

Document Navigation **25**

1 Introduction **26**

1.1 MRI 27

1.2 MRI in this Thesis 28

1.2.1 Prototype Rotatable MR System 28

1.3 Research Motivation 30

1.4 Research Objectives and Contributions 33

1.4.1 Scan Protocol Design 33

1.4.2 Image Registration 33

1.4.3 Data Processing and Visualisation 34

1.4.4 Successful In Vivo Demonstration of MADI 35

1.5 Summary 35

1.6 Thesis outline 36

2 Theoretical Background **38**

2.1 Chapter Overview 38

2.2 MRI Principles 38

2.2.1 Quantum Mechanical Perspective 39

2.2.2 Source of MR Signal 40

2.2.3 Relaxation 41

2.2.4 Dipole-Dipole Interactions 43

2.2.5 T2 and T2* Relaxation 45

2.3 The Magic Angle Effect 46

2.3.1 Dipolar Interaction Strength 46

2.3.2 Collagen Structure 48

2.4 Orientation Dependent Intensity Model 52

2.5 Anatomy of the Knee 55

2.6 Summary 57

3 Literature Review **58**

3.1 Chapter Overview 58

3.2 Limitations of Conventional MRI 58

3.3 Diffusion Tensor Imaging for Examining Collagen Fibre Structure 63

3.4 Clinical Need for Dedicated MRI Systems 68

3.5 Low-Field MRI 71

3.6 Summary 76

4 Magic Angle Directional Imaging **78**

4.1 Chapter Overview 78

4.2 Magic Angle Directional Imaging Methodology 78

4.2.1 MADI Methodology Detailed Overview 80

4.3 Rotatable Low-Field System Overview 82

4.3.1 Installation 82

4.3.2 System Overview and Block Diagram 84

4.3.3 Motion Platform 87

4.3.4 Patient positioning 88

4.3.5 Magnet Design 89

4.3.6 Coordinate System 90

4.3.7 Motion Control Software 92

4.3.8 Patient Safety 94

4.3.9 Pulse Sequence 95

4.3.10 Examination Protocol 97

4.4 MADI Data Preprocessing 98

4.4.1 DICOM Intensity Rescaling 98

4.4.2 Yaw Intensity Rescaling 99

4.4.3 Denoising 99

4.5 Summary 102

5 Registration **104**

5.1 Chapter Overview 104

5.2 Registration Overview . . . 105
5.2.1 Registration Principles . . . 106
5.3 Registration Framework . . . 109
5.3.1 Mathematical Formulation of Image Registration . . . 109
5.3.2 Registration Framework Block Diagram . . . 110
5.3.3 Rigid Transform . . . 111
5.3.4 Affine Transform . . . 112
5.3.5 B-Spline Transform . . . 112
5.3.6 Registration Metrics . . . 113
5.3.7 Registration Penalties . . . 114
5.4 Registration in the Context of MADI . . . 115
5.4.1 Image Drift Compensation and Averaging . . . 117
5.4.2 Registration Pipeline . . . 119
5.4.3 Registration of Rotated Image Volumes . . . 122
5.4.4 Registration Masking . . . 124
5.4.5 Failure Detection . . . 124
5.5 Distortion Compensation . . . 125
5.5.1 3D Printed Calibration Phantom . . . 126
5.5.2 Distortion Correction using Non-rigid Registration . . . 128
5.6 Summary . . . 131

6 Segmentation **133**

6.1 Chapter Overview . . . 133

6.2 Metric-Assisted Segmentation . . . 134
6.2.1 Standard Deviation Metric . . . 135
6.3 Segmentation Graphical User Interface . . . 136
6.3.1 Manual Segmentation . . . 139
6.4 Summary . . . 139

7 Collagen Fibre Orientation Estimation 140

7.1 Chapter Overview . . . 140
7.2 Collagen Fibre Orientation Estimation Algorithm . . . 140
7.2.1 Previous Algorithms . . . 141
7.2.2 Updated Collagen Fibre Orientation Estimation Algorithm . . . 144
7.2.3 Collagen Fibre Orientation Estimation Output . . . 149
7.3 Scan Orientation Selection . . . 152
7.3.1 Previous Scan Orientation Selection Methods . . . 152
7.3.2 Collagen Fibre Orientation Estimation Accuracy Mapping . . . 156
7.3.3 Genetic Algorithm for Scan Orientation Selection . . . 157
7.4 Summary . . . 163

8 Results 164

8.1 Chapter Overview . . . 164
8.2 Tractography Visualisation . . . 164
8.3 Healthy ACL . . . 166
8.3.1 ACL Bundles . . . 167
8.4 Healthy Meniscus . . . 171

8.5 Partially Torn ACL . . . 175

8.6 Medial Meniscus Tear . . . 177

8.7 Summary . . . 180

9 Conclusions and Future Work **183**

9.1 Chapter Overview . . . 183

9.2 MADI Technique Summary . . . 183

9.2.1 Low-Field Rotatable MRI System . . . 184

9.2.2 Registration . . . 184

9.2.3 Segmentation . . . 184

9.2.4 Collagen Fibre Orientation Estimation . . . 184

9.3 In Vivo Results . . . 185

9.4 Limitations . . . 185

9.5 Future Work . . . 187

9.5.1 Hardware . . . 187

9.5.2 Acquisition and Pulse Sequence Optimisation . . . 189

9.5.3 Reconstruction and Image Processing . . . 190

9.5.4 Examination Protocol . . . 192

9.5.5 Broader System Utility . . . 193

9.5.6 Clinical Validation . . . 194

9.6 Final Conclusions . . . 195

References **197**

Appendices **213**

A Appendices Introduction **213**

A.1 Detailed List of Contributions 213

B MADI Appendices **215**

B.1 Effect of Denoising on Collagen Fibre Orientation Estimation 215

C Registration Appendices **218**

C.1 Registration: elastix Parameters 218

C.1.1 Rigid Registration Dummy Call Parameters 218

C.1.2 Rigid Registration Parameters 220

C.1.3 Affine Registration Parameters 222

C.1.4 B-Spline Registration Parameters 224

C.2 CadQuery Python Code for Calibration Phantom 226

D Segmentation Appendices **232**

D.1 Full Representative Dataset and Derived Metric Image Projections . . . 232

E Collagen Fibre Orientation Estimation Appendices **238**

E.1 Collagen Fibre Solver GUI 238

E.2 Golden Spiral Points on Sphere 238

E.3 Genetic Algorithm Binary Encoding 239

F Results Appendices **241**

G Volunteer Consent Form **242**

List of Figures

1.1 Conventional clinical Philips MRI scanner at the Sahlgrenska University Hospital, Sweden 27

1.2 Low-field rotatable scanner in the home configuration and a rotated configuration. 29

1.3 Bovine knee sample images acquired at two different $\vec{B}_0$ orientations. . 29

2.1 Dipolar magnetic field of a proton. 39

2.2 Net magnetisation generated by slight bias of proton spin orientation towards $\vec{B}_0$. 40

2.3 T2 relaxation caused by dephasing of initially coherent spins. 42

2.4 T2 signal decay example. 43

2.5 Dipole-dipole interaction between two neighbouring protons. 44

2.6 Contour plot of the z-component of the magnetic field around a proton. . 47

2.7 Hierarchical structure of tendon. 49

2.8 Diagram of orientationally restricted water bound to tropocollagen, illustrating the structural basis of the magic angle effect. 51

2.9 Intensity versus θ for $A = 1$ and $B = \{0.5, 1, 2\}$. 53

2.10 Diagram of the human knee joint connective tissue structures 56

4.1 Flowchart of MADI technique. 79

4.2 Installation of low-field rotatable system in RF shielded room. 83

4.3 Operator station and power supply for low-field system. 83

4.4 Close-up view of central section of low-field rotatable scanner. 84

4.5 Block diagram of low-field MR system. 86

4.6 Close up of roll and yaw motors of low-field system. 88

4.7 Patient positioning within the low-field rotatable scanner. 89

4.8 CAD model of low-field magnet design composed of an array of permanent magnet subunits. 90

4.9 Global coordinate system of the rotatable low-field scanner, and the $\vec{B}_0$ orientation when the scanner is in its home configuration. 91

4.10 Screenshot of motion control software used to set scanner roll and yaw angles. 93

4.11 Physical emergency stop button and scanner control software settings page with software emergency stop button. 94

4.12 Example of block matching denoising applied to a knee scan from the rotatable low-field scanner. 101

5.1 Image voids resulting from defining a transform from the moving to the fixed image. 106

5.2 Coordinate transform from fixed to moving image domain to avoid voids in the resampled registered image. 107

5.3 Transform-based registration demonstration. 108

5.4 Block diagram of the basic registration components of elastix. 110

5.5 Block diagram of affine registration for image drift compensation prior to averaging. 118

5.6 Difference images of excitations 1 and 2 before and after affine registration. 119

5.7 MADI registration pipeline. 121

5.8 Overview of inputs and output of the MADI registration pipeline. 123

5.9 Example of subtle registration failure. 125

5.10 CadQuery model of 3D distortion calibration phantom. 127

5.11 3D printed calibration phantom for registration based 3D distortion compensation. 128

5.12 Example of a subtractive boolean operation between a part of the calibration phantom CAD model and a simulated voxel. 129

5.13 Simulated and real scans of the 3D printed calibration phantom. 130

5.14 ACR phantom scan (axial slice 60/150) before and after non-rigid correction transform was applied. 131

6.1 Simple example of intensity variations across registered images due to the magic angle effect. 135

6.2 Standard deviation images from a MADI dataset showing a bright ACL and meniscus due to magic angle effect intensity variations. 136

6.3 Demonstration of metric-assisted segmentation software. 138

7.1 Low and high density sets of candidate fibre orientations. 146

7.2 Fibre plot of the ACL viewed from an oblique angle. 149

7.3 Fibre plots of the ACL from the representative dataset viewed directly against the y- and z-axes 150

7.4 Intensity versus θ for $A = 1$ and $B = 1$ with α marked to show the domain where two θ values map to a single intensity value. 153

7.5 Limited angular sampling range relative to an axially aligned collagen fibre due to the scanner's constrained yaw range. 155

7.6 Collagen fibre orientation estimation accuracy maps for sets of 4 and 9 scan orientations. 157

7.7 Convergence of total average deviation for genetic algorithm with $N = 6$. 160

7.8 Collagen fibre orientation estimation accuracy maps for $N = 4, 5, 6, 7, 8, 9$ genetic algorithm optimised scan orientations. 161

8.1 Intensity versus θ plot for the representative dataset. 166

8.2 Heatmap of estimated collagen fibre orientations from the representative dataset. 168

8.3 Fibre plot of anteromedial and posterolateral bundles of the ACL from the representative dataset. 169

8.4 Tractography results for the ACL and the anteromedial and posterolateral bundles of the ACL. 170

8.5 Fibre plot of the meniscus from a healthy volunteer. 171

8.6 Tractography plot of the healthy meniscus dataset. 173

8.7 Tractography plot of the healthy meniscus dataset, overlaid onto the mean image from the registered MADI images. 174

8.8 Fibre plot of the ACL from a volunteer with a partially torn ACL. 175

8.9 Tractography plot of a partially torn ACL, overlaid onto the original fixed image from the MADI dataset. 176

8.10 Fibre plot of the meniscus from a volunteer with a medial meniscus tear. 178

8.11 Tractography plot of a meniscus with a medial meniscus free edge tear. . 179

9.1 3D printed phased array receive coil whose total signal strength from the three channels combined is invariant with $\vec{B}_0$ orientation relative to the coil central axis. 188

9.2 Minimum, mean, and maximum intensity projections of a knee scan from the rotatable low-field scanner. 194

B.1 Histogram showing the deviation between corresponding fibres for the healthy ACL (representative) dataset with and without block matching denoising applied. 216

B.2 Histogram showing the deviation between corresponding fibres for the healthy meniscus dataset with and without block matching denoising applied. 216

B.3 Histogram showing the deviation between corresponding fibres for the pathological ACL dataset with and without block matching denoising applied.216

B.4 Histogram showing the deviation between corresponding fibres for the pathological meniscus dataset with and without block matching denoising applied. 217

D.1 Representative dataset $n = 1$. 232

D.2 Representative dataset $n = 2$. 232

D.3 Representative dataset $n = 3$. 233

D.4 Representative dataset $n = 4$. 233

D.5 Representative dataset $n = 5$. 234

D.6 Representative dataset $n = 6$. 234

D.7 Representative dataset $n = 7$. 235

D.8 Representative dataset minimum projection. 235

D.9 Representative dataset mean projection. 236

D.10 Representative dataset maximum projection. 236

D.11 Representative dataset standard deviation projection. 237

E.1 Solver GUI developed for fine tuning A and B parameters and also generating output figures for MADI process. 238

F.1 3D Slicer DTI tract advanced settings to control tract propagation. 241

List of Tables

2.1 Summary of each level of the hierarchical collagen structure. 50

3.1 Summary of field strength, isotropic resolution, scan time, and number of diffusion gradient encoding directions for three high resolution DTI studies of collagen fibre architecture. 67

3.2 Data obtained via Freedom of Information requests from various NHS Trusts across the UK detailing the total number and modality of MRI examinations performed. 70

3.3 Summary of typical MRI field strength categories, associated magnet technologies, and the key advantages and limitations of each. 75

4.1 Parameters for the 3D FLASH pulse sequence used for all acquisitions on the low-field rotatable scanner. 96

5.1 Summary of elastix inputs and outputs relevant to the MADI registration pipeline. 116

7.1 Reported speeds of previous collagen fibre orientation estimation algorithms and the updated collagen fibre orientation estimation algorithm. . 147

7.2 Approximate angles between scan orientation subsets and caprine patellar tendon from Lanz et al. 153

8.1 Roll and yaw angles for the representative dataset. 167

8.2 Roll and yaw angles for the healthy meniscus dataset. 172

8.3 Roll and yaw angles for the partially torn ACL dataset. 177

8.4 Roll and yaw angles for the medial meniscus tear dataset. 180

Definitions

ACL	Anterior cruciate ligament. A major stabilising ligament of the knee.
Arthroscopy	Invasive surgical procedure to visually inspect joints, used as the gold standard for joint assessment.
B_0	The main magnetic field of a magnetic resonance imaging system.
$\vec{B}_0$	The vector representation of the main magnetic field of a magnetic resonance imaging system.
$\hat{B}_0$	The unit vector in $\mathbb{R}^3$ representing the orientation of the main magnetic field within the imaging volume of a magnetic resonance imaging system. Conventionally the z-axis of the system is defined as parallel to the $\hat{B}_0$ direction.
CFOEA	Collagen Fibre Orientation Estimation Algorithm.
DOF	Degree(s) of freedom.
DSV	Diameter of spherical volume. The diameter of the homogeneous magnetic field region of a magnetic resonance imaging machine.
DTI	Diffusion tensor imaging.
DWI	Diffusion weighted imaging.
FOV	Field of view.
Genetic algorithm	An optimisation method inspired by natural selection, useful for large, non-differentiable search spaces.
GUI	Graphical user interface.
Isotropic voxel	A voxel with equal resolution in all three spatial dimensions.

MADI	Magic Angle Directional Imaging. The technique developed to utilise rotatable MRI to leverage the magic angle effect as a source of quantitative information on collagen-rich tissue microstructure.
Magic angle	54.7°. The angle between $\vec{B}_0$ and collagen fibres in anisotropic collagen-rich tissues that minimises dipolar interactions between protons in bound water within these tissues. At the magic angle these tissues will present maximally increased signal in MRI images.
Magic angle effect	The intensity variation of anisotropic collagen-rich tissues that is a function of the angle between $\vec{B}_0$ and the collagen fibres being imaged.
MATLAB	A high-level programming language and interactive environment used for scientific computing, data analysis, algorithm development, and visualization.
Meniscus	Anisotropic collagen-rich structure in the knee that distributes load and stabilises the joint.
MRI	Magnetic resonance imaging.
Pixel	Picture element. The smallest discrete element of a two-dimensional image.
Python	A high-level, widely used programming language, used for web development, data analysis, artificial intelligence, and scientific computing.
Registration	Process of aligning multiple images into a common coordinate system.
ROI	Region of interest.
Segmentation	Partitioning of an image into regions corresponding to anatomical structures.
SNR	Signal-to-noise ratio.

Tractography	Visualisation method that integrates local fibre orientations into continuous streamlines, commonly used in diffusion MRI.
Voxel	Volume element. The smallest discrete element of a three-dimensional volume image, representing an intensity value at a regular grid point in 3D space. Analogous to a 2D pixel but extended to 3D.
1D	One dimension or one-dimensional.
2D	Two dimensions or two-dimensional.
3D	Three dimensions or three-dimensional.

Document Navigation

When viewed electronically with suitable PDF reader software, all numerical references to chapters, sections, subsections, equations, figures, tables etc. are hyperlinked for direct navigation. In addition, the page numbers at the bottom of each page link back to the table of contents, from which all items are themselves clickable, facilitating rapid navigation across the document.

1 | Introduction

Magnetic Resonance Imaging (MRI) is a non-invasive imaging modality that exploits the properties of hydrogen nuclei to generate detailed images of biological tissues [1–3]. MRI scanners are employed as diagnostic tools across a broad range of clinical domains, including oncology, neurology, and musculoskeletal (MSK) imaging. MRI generally provides excellent soft tissue contrast; however, in MSK applications, MRI is not always reliable for imaging tissues composed of anisotropic fibrous collagen, such as tendons, ligaments, and menisci [4–6]. These tissues are challenging to image as they often appear hypointense in conventional MRI scans, whilst also exhibiting intensity variation dependent on their collagen fibre orientation relative to the magnetic field of the scanner. This phenomenon is known as the “magic angle effect” [7].

The magic angle effect is typically regarded as a source of unwanted contrast in clinical imaging, and is often described as an artefact, as it can obscure or even mimic true pathology [6, 8]. Strategies to suppress the effect include modification of imaging parameters and patient repositioning [8]. However, these approaches often present compromises in image quality or patient comfort, and for anatomical structures with inherently curved fibre architectures such as the meniscus, supraspinatus tendon, or articular cartilage, it is often geometrically infeasible to avoid the magic angle effect across the entire structure. Consequently, signal variation due to the magic angle effect is frequently dismissed as misleading or pathological [6], despite the fact that it encodes potentially useful quantitative information about collagen fibre microstructure.

The work presented in this thesis, *In Vivo Imaging of Collagen Fibre Structures Using Dedicated Low-Field MRI*, explores the use of the magic angle effect as a quantitative source of structural information for collagen-rich musculoskeletal tissues. Specifically, this work focuses on the in vivo application of a prototype rotatable low-field MRI scanner for analysing magic angle signal variation in the anterior cruciate ligament (ACL) and meniscus of the human knee joint.

1.1 MRI

MRI is a medical imaging technology that manipulates strong magnetic fields to generate two- or three-dimensional (2D or 3D) images of biological tissues. Unlike X-ray or CT imaging, which rely on tissue attenuation of ionising radiation to create images, MRI derives image contrast from differences in the distribution and molecular environment of hydrogen nuclei (protons) within the tissues. This makes MRI particularly well suited for safely imaging soft tissues such as the brain, muscles, ligaments, cartilage, and internal organs, due to their high water or fat content, of which hydrogen nuclei form a significant proportion.

Figure 1.1 shows a clinical whole-body MRI scanner manufactured by Philips. During a scan, the patient lies on a motorised table that slides into the scanner's central bore, which contains the main imaging components.



Figure 1.1: Conventional clinical Philips MRI scanner at the Sahlgrenska University Hospital, Sweden. The patient lies on the scan table (front centre of picture) and is then moved within the central bore of the scanner, where they are exposed to the homogeneous main magnetic field of the scanner and the gradient and radiofrequency fields.

Image attribution: Jan Ainali - Own work, CC BY 3.0, `https://commons.wikimedia.org/w/index.php?curid=3546051`

At the core of every MRI system is a powerful, homogeneous magnetic field known as the system's main field, denoted B_0 (or $\vec{B}_0$ as a vector quantity). When a patient is placed within an MRI scanner, $\vec{B}_0$ induces a weak net magnetisation in the body's tissues. By precisely applying additional magnetic gradient and radiofrequency fields, the net magnetisation is "excited" in such a way that its spatial variation can be measured, which is recorded and processed to reconstruct an image of the patient's tissues.

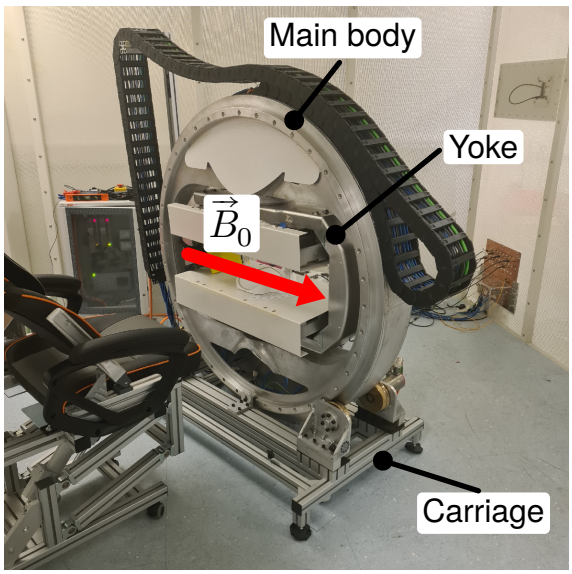
1.2 MRI in this Thesis

1.2.1 Prototype Rotatable MR System

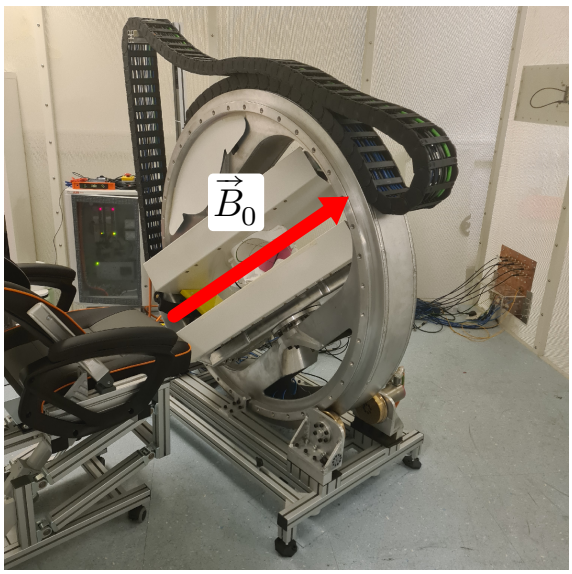
Conventional MRI struggles to image connective tissues composed of organised, anisotropic collagen fibres. These tissues typically exhibit low signal intensity, making them difficult to visualise and assess with MR imaging. Additionally, their fibrous microstructure gives rise to the magic angle effect, which causes MR signal variation that is dependent on fibre orientation relative to $\vec{B}_0$. Peak signal intensity is observed when the tissue fibres are oriented at the "magic angle" of 54.7° relative to $\vec{B}_0$. The magic angle effect is often regarded as an artefact that complicates image interpretation, rather than as a potential source of useful information.

However, with the right tools, the magic angle effect can become a useful tissue property rather than imaging artefact. This thesis concerns the application of a prototype, low-field, rotatable MRI scanner that exploits the magic angle effect as a source of quantitative information on the microstructure of collagen-rich tissues. The scanner was developed to enable the Magic Angle Directional Imaging (MADI) technique, a novel approach for assessing collagen fibre microstructure non-invasively.

Figure 1.2 shows the rotatable scanner in its home and a rotated configuration. All imaging components are housed within a central yoke which is capable of yawing within the scanner's main body, which itself is capable of rolling on its carriage. These two degrees of freedom, roll and yaw, enable precise reorientation of $\vec{B}_0$ relative to a static patient, who sits facing the scanner and places their leg through the transmit and receive coils at the centre of the scanner bore. Full details of this rotatable MRI system and imaging procedure are provided in Chapter 4.



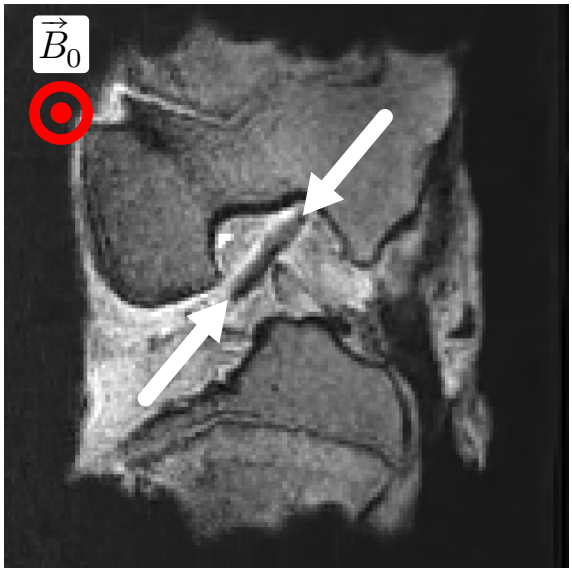
(a) Home configuration: no roll or yaw applied.



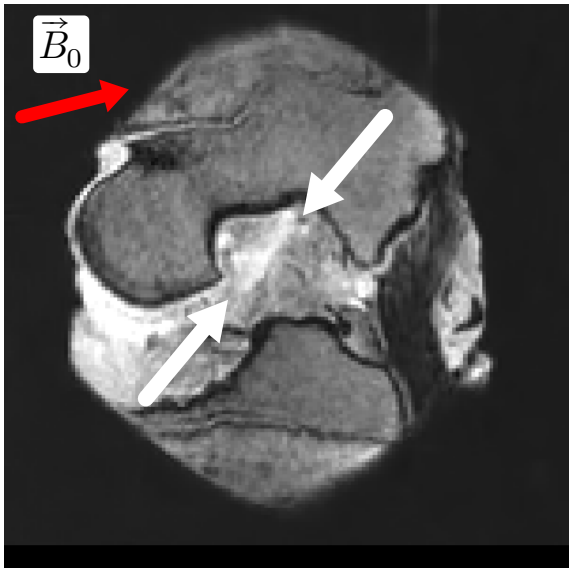
(b) Rotated configuration: 45° of roll and 20° of yaw applied.

Figure 1.2: Low-field rotatable scanner in the home configuration (Figure 1.2a) and a rotated configuration (Figure 1.2b). The $\vec{B}_0$ orientations for each configuration are indicated with a red arrow.

Figure 1.3 illustrates the magic angle effect in the posterior cruciate ligament (PCL) of an excised bovine knee sample, shown in images acquired using the rotatable scanner in its home and a rotated configuration.



(a) Bovine knee sample with dark PCL. Acquired with scanner in its home configuration.



(b) Bovine knee sample with bright PCL. Acquired with scanner in a rotated configuration.

Figure 1.3: Bovine knee sample images acquired at two different $\vec{B}_0$ orientations, indicated by red arrows. The visible section of the PCL, marked by white arrows, is shown to exhibit the magic angle effect: Figure 1.3a shows the PCL as dark with the scanner in its home orientation, and Figure 1.3b shows the PCL as bright with reoriented $\vec{B}_0$.

Figure 1.3a was acquired with the scanner in the home orientation and shows the visible PCL section as dark, with $\vec{B}_0$ pointing out of the page. Figure 1.3b was acquired with the scanner in a rotated orientation, and shows the PCL as bright due to the magic angle effect increasing its signal. $\vec{B}_0$ was oriented at approximately 90° relative to the predominant fibre orientation of the (dark) PCL section in Figure 1.3a, and in Figure 1.3b $\vec{B}_0$ was oriented at approximately 54.7°, the so-called "magic angle", relative to the same (now bright) PCL section.

The relationship between MR signal intensity and the orientation of collagen fibres relative to the main magnetic field is well understood [9–12]. A rotatable scanner enables acquisition of images at multiple $\vec{B}_0$ orientations relative to a static patient. The resulting intensity variations induced by the magic angle effect can then be compared to theoretical models to infer the underlying collagen fibre orientations. This methodology is the basis of the MADI technique, which transforms the magic angle effect from a diagnostic challenge into a powerful source of microstructural information.

1.3 Research Motivation

The research presented in this thesis was motivated by recent interest in the magic angle effect as a potentially useful tissue property. Chapter 3 also discusses other factors motivating the research such as: the limitations of conventional MRI for musculoskeletal (MSK) imaging, the use of diffusion tensor imaging (DTI) for MSK assessment, the clinical need for dedicated MRI systems, and the renewed interest in low-field MRI as a viable imaging technology. However, this section focuses specifically on the utility of the magic angle effect, which was the principal driver of this research.

Bydder et al. [8] noted that whilst the effect is typically considered an artefact and a source of diagnostic confusion, patient repositioning to orient the target tissues at the magic angle relative to the static B_0 field of a conventional scanner would boost the signal in those tissues. As such, the magic angle effect could become a source of image contrast. Rea et al. [13] built upon this work by developing devices to orient connective tissues in the hand, Achilles tendon, knee, and upper limb, at the magic angle relative to $\vec{B}_0$ in conventional closed bore scanners in order to boost the MR signal of the target tissues. Whilst these techniques were effective at achieving their goal of producing

increased signal from typically dark tissues, they faced several limitations: the target tissues must be straight and their orientation known a priori, and patient repositioning possibilities are restricted within the closed bore of a conventional MRI scanner.

Szeverenyi et al. [11] proposed that, rather than exploiting the magic angle effect to boost signal in a single acquisition, multiple scans at different $\vec{B}_0$ orientations could be used to analyse magic angle effect intensity variations and infer collagen fibre orientations. They demonstrated this approach using an excised goat meniscus, which was embedded in an epoxy sphere and imaged at 15 distinct $\vec{B}_0$ orientations by rotating the sphere within a conventional MRI system. All images were registered to the first acquisition to enable voxel-wise comparison. Fibre orientations were estimated by matching the observed intensity profiles at each voxel coordinate to an array of simulated intensity profiles generated for candidate fibre orientations under the same set of $\vec{B}_0$ orientations. The best fitting candidate orientation was assigned as the fibre orientation in each voxel. The resulting fibre architecture followed a circumferential path consistent with healthy meniscal anatomy. This study served as a successful proof of concept, demonstrating that the magic angle effect can reveal collagen fibre microstructure through orientation-dependent intensity analysis. However, the methodology used, requiring a large number of scans with orientations evenly distributed within three orthogonal planes, is not clinically practical due to long acquisition times and the infeasibility of such reorientations in vivo (even using a rotatable MRI system) due to ergonomic constraints.

Seidel et al. [10] conducted a similar study using an excised human iliotibial tract, held under tension in a custom jig within a conventional MRI scanner. Their approach to fibre orientation estimation differed from that of Szeverenyi et al. in two key ways. Firstly, they acquired 35 scans evenly distributed within only two orthogonal planes; and secondly, instead of comparing measured intensities with simulated candidate fibre orientations, they estimated fibre orientations by analysing the shape, maxima, and first derivatives of the measured intensity variation curves, and compared these to analytic models of magic angle intensity variation. Validation of the estimated fibre orientations was performed using polarised light microscopy (PLM), which showed strong agreement between the estimated fibre orientations and those directly observed. Both results showed fibre alignment with the macrostructure of the tissue, which is again consistent with the expected result in healthy anatomy. As with the study from Szeverenyi et al., this work served as a successful proof of concept for utilising the magic angle effect

to assess collagen microstructure. However, the large number of required acquisitions and the impracticality of the scan orientations again limit the clinical applicability of this technique.

These prior studies motivated further work by our research group to develop Magic Angle Directional Imaging (MADI) into a practical clinical reality. The approach proposed by Szeverenyi et al. was particularly influential, as its reduced reliance on an exhaustive sampling of $\vec{B}_0$ orientations made it more suited to clinical translation. Chappell et al. [14] adapted Szeverenyi et al.'s technique to investigate ligament maturity and injury in ex vivo caprine and canine knee specimens, designating this adapted technique as MADI. The specimens were mounted in custom rotatable spherical jigs and imaged in a conventional clinical scanner at nine distinct $\vec{B}_0$ orientations, distributed evenly over a hemisphere rather than restricted to two or three orthogonal planes. This study demonstrated that differences in ligament maturity could be detected through the degree of collagen fibre alignment, and that partial tears could be identified based on observed disruptions in the estimated fibre microstructure.

To assess the robustness of MADI under conditions of lower signal-to-noise ratio (SNR), Brujic et al. [15] added synthetic noise to the datasets from Chappell et al. and also performed Monte Carlo simulations on noisy synthetic data. Their findings confirmed that the MADI technique was sufficiently resilient to noise, producing accurate results at SNR values as low as 6, suggesting its suitability for use in vivo with the low-field rotatable MRI system developed by McGinley et al. [16].

These preceding studies laid the foundations for the research presented in this thesis, demonstrating that the magic angle effect can be exploited to estimate collagen fibre microstructure by analysing orientation-dependent MR signal variations. However, several critical challenges remained in translating this approach to successful in vivo application. Section 1.4 summarises the key research objectives, challenges, and contributions of the work presented in this thesis.

1.4 Research Objectives and Contributions

The overall research objectives of this thesis were to establish the theoretical and computational frameworks required for efficient application of MADI in vivo, and to demonstrate successful *In Vivo Imaging of Collagen Fibre Structures Using Dedicated Low-Field MRI*. To achieve these overarching goals, several specific research objectives were pursued, including the development of scan protocol design, image processing and registration pipelines, segmentation tools, and algorithms for estimating collagen fibre orientations. The key research objectives and novel contributions of this thesis are summarised below:

1.4.1 Scan Protocol Design

- **Research Objective:** To identify the minimum number and optimal distribution of $\vec{B}_0$ orientations required for accurate in vivo collagen fibre orientation estimation.
- **Challenges:** The rotatable scanner operates at low-field and has a constrained yaw range due to patient ergonomics and receive coil design limiting possible $\vec{B}_0$ orientations. Fewer acquisitions reduce scan time and improve patient comfort, but may compromise estimation accuracy. The protocol must therefore balance hardware constraints, acquisition time, and collagen fibre orientation estimation accuracy.
- **Contribution:** A simulation framework was developed to evaluate the suitability of $\vec{B}_0$ orientation sets. This included development of a custom fitness function to assess collagen fibre orientation estimation accuracy under noise. A genetic algorithm was implemented to optimise $\vec{B}_0$ orientation sets subject to constraints on number of acquisitions, SNR, and yaw range.

1.4.2 Image Registration

- **Research Objective:** To develop a fully automatic image registration pipeline capable of: aligning images acquired at different $\vec{B}_0$ orientations to a common reference, compensating for magnet heating image drift, correcting for soft-tissue deformations, and correcting for geometric distortions caused by B_0 inhomogeneity.

- **Challenges:** The registration pipeline needed to address both rigid transformations (e.g., global motion and image drift) and non-rigid deformations (e.g., soft-tissue movement and field-induced distortions). Additional complexity arose from the varying content at the image periphery between scans due to scanner rotation and cusps in B_0, as well as from extreme distortions near the image edges.

- **Contribution:** A fully automated multi-stage registration pipeline was developed, combining rigid and non-rigid registration methods. Rigid registration corrected for image drift and performed global rotation of the images, whilst subsequent non-rigid stages addressed soft-tissue deformation and spatial distortions due to B_0 inhomogeneity. The framework required no manual intervention and was robust to edge-of-field inconsistencies between scans.

1.4.3 Data Processing and Visualisation

- **Research Objective:** To develop methods for semi-automatic segmentation, efficient collagen fibre orientation estimation, and intuitive three-dimensional visualisation using glyphs and tractography.

- **Challenges:** Custom segmentation tools were required to leverage the unique intensity variation data produced by MADI. Existing fibre orientation estimation algorithms were too slow for practical use and lacked the spatial resolution necessary for clinical deployment. Furthermore, as MADI generates fundamentally different data from standard MSK MRI, new visualisation strategies were needed to convey this information in a clinically familiar and interpretable manner.

- **Contribution:** A semi-automatic segmentation tool was developed that uses voxel-wise signal variation metrics to highlight tissue exhibiting the magic angle effect. Improved fibre orientation estimation algorithms were implemented for increased computational efficiency and accuracy. Visualisation methods adapted from diffusion tensor imaging, such as glyph rendering and tractography, were integrated to present results in a format aligned with existing clinical imaging paradigms.

1.4.4 Successful In Vivo Demonstration of MADI

- **Research Objective:** To successfully apply the MADI technique in vivo in both healthy volunteers and patients with ACL or meniscal pathology.
- **Challenges:** In vivo imaging introduced several practical and technical constraints. Patient safety had to be ensured whilst in close proximity to the rotating system. Patient ergonomics also limited scanner positioning, in particular yaw rotation, which affected fibre estimation accuracy. Additionally, validation of results in vivo was inherently difficult, as MADI produces novel fibre orientation data not accessible through conventional, non-invasive methods.
- **Contribution:** MADI was successfully implemented in vivo, with results in healthy volunteers showing fibre orientations consistent with expected anatomy. In pathological cases, MADI findings were in agreement with arthroscopic findings, demonstrating the technique's potential for non-invasive microstructural assessment.

A detailed breakdown of individual and external contributions to the work in this thesis is provided in Section A.1.

1.5 Summary

MRI has long struggled to visualise connective tissues composed of organised, anisotropic collagen fibres such as tendons, ligaments, and menisci. These tissues typically exhibit low signal intensity and orientation-dependent signal variation due to the magic angle effect, which can obscure or mimic pathology. Recent interest in the magic angle effect as a potential source of microstructural information has motivated the development of a low-field, rotatable MRI scanner and associated multi-orientation imaging protocol, known as Magic Angle Directional Imaging (MADI), to enable non-invasive estimation of collagen fibre orientations in vivo. Realising this capability required overcoming several key technical and methodological challenges. The work undertaken to achieve successful MADI application in vivo is presented throughout this thesis, as outlined in Section 1.6.

1.6 Thesis outline

This thesis is structured to first introduce the key concepts underpinning MADI and the imaging challenges it is designed to address, before presenting a series of chapters that broadly mirror the structure of the MADI technique itself. The chapters are organised as follows:

Chapter 1: Introduction

Presents the fundamental principles of MRI, the limitations of conventional MRI for MSK imaging, and the significance of the magic angle effect. Introduces rotatable MRI as a means to address these challenges, outlines the motivation for this research, and defines the research objectives, challenges, and contributions of the thesis.

Chapter 2: Theoretical Background

Outlines the MRI principles most relevant to MADI, followed by a detailed explanation of the magic angle effect in the context of dipolar interactions and the anisotropic structure of collagen fibres. Concludes with an overview of the anatomy of interest in this thesis, focusing on the ACL and meniscus.

Chapter 3: Literature Review

Examines prior research to further motivate this work and identify the gap addressed by rotatable MRI and MADI. Reviews the limitations of conventional MRI for MSK imaging, assesses diffusion tensor imaging as another modality to analyse collagen fibre microstructure, outlines the clinical need for more diverse MRI systems, and surveys the renewed interest in low-field MRI technology.

Chapter 4: Magic Angle Directional Imaging

Provides an overview of the MADI methodology, followed by a detailed description of the low-field rotatable MRI system developed to enable it. Introduces the representative dataset, a MADI examination focused on the ACL, which is used throughout this and subsequent chapters to illustrate the technique and present intermediate results on real data.

Chapter 5: Registration

Describes the registration framework tailored for MADI using the low-field rotatable system, including a multi-stage pipeline for aligning multi-orientation acquisitions. The same registration techniques are further applied to correct geometric distortions, using a 3D-printed calibration phantom as ground truth.

Chapter 6: Segmentation

Presents the segmentation framework developed for MADI, which exploits standard deviation across multi-orientation images as a metric uniquely available through the MADI technique. This approach enables data-driven identification of orientation-sensitive tissues and reduces operator burden compared with conventional manual methods.

Chapter 7: Collagen Fibre Orientation Estimation

Describes the operation of the collagen fibre orientation estimation algorithm. Highlights the substantial efficiency gains of the updated implementation in this work, which enabled prospective evaluation of scan orientation sets via simulation studies and its use as a fitness function in a genetic algorithm for optimising scan orientations in MADI.

Chapter 8: Results

Presents MADI data acquired in healthy volunteers and patients with pathology, demonstrating the capabilities and clinical translational potential of the technique. Introduces new musculoskeletal MRI visualisation methods by adapting tractography from diffusion tensor imaging to depict collagen fibre structures. Correlation with arthroscopy is also provided in pathological cases to validate the in vivo MADI findings.

Chapter 9: Conclusions and Future Work

Summarises the successful in vivo application of MADI, demonstrating its promise as a non-invasive tool for assessing collagen fibre microstructure in connective tissues. While initial results with arthroscopic correlation are encouraging, the number of patients studied is not yet sufficient to establish the statistical power required for clinical replacement of arthroscopy, highlighting the need for larger trials. Beyond MADI, the development of the rotatable MRI system opens opportunities for a wide range of novel imaging applications, of which MADI represents only the first.

2 | Theoretical Background

2.1 Chapter Overview

This chapter outlines the theoretical background required to understand the concepts and methods presented in this thesis. It begins with Section 2.2, which outlines the fundamental MRI principles most relevant to the magic angle effect, including quantum spin, signal generation, and relaxation mechanisms. The physics underlying the magic angle effect is then covered in more detail in Section 2.3, including dipolar interactions between protons and the influence of the anisotropic collagen structure of tendons, ligaments, and menisci on T2 relaxation. The orientation-dependent intensity model that characterises the magic angle effect is then introduced in Section 2.4. Finally, the chapter concludes with an overview of knee anatomy in Section 2.5, with emphasis on the ACL and meniscus, providing anatomical context for the results presented in later chapters.

2.2 MRI Principles

This section outlines the key MRI principles necessary to support the subsequent discussion of the magic angle effect. It is intended for readers with a general understanding of MRI physics, including concepts such as T1 and T2 relaxation, image weighting, and pulse sequences. Rather than providing a comprehensive review, this section focuses specifically on the elements of MRI that are essential for understanding the mechanisms and implications of the magic angle effect. For a more detailed and rigorous discussion of MRI physics, the reader is directed to *MRI from Picture to Proton* by McRobbie et al. [17], and *Magnetic Resonance Imaging: Physical Principles and Sequence Design* by Haacke [3].

2.2.1 Quantum Mechanical Perspective

MRI exploits the quantum mechanical properties of hydrogen nuclei (protons) to generate images of biological tissues. Hydrogen is particularly well-suited for this purpose due to its abundance in the human body as a constituent of water and fat [18], making it an effective proxy for imaging biological tissues using MRI. The property central to MRI is the proton's intrinsic spin angular momentum, $\vec{I}$. Each proton produces a tiny dipolar magnetic field, arising from its intrinsic magnetic moment, $\vec{\mu}$, which is proportional to its spin angular momentum, following Equation 2.1 [19]:

$$\vec{\mu} = \gamma \vec{I} \tag{2.1}$$

Where γ is the gyromagnetic ratio of the proton, a particle-specific constant relating spin angular momentum to magnetic moment. Figure 2.1 illustrates the magnetic field around a proton, as well as its spin, charge, and magnetic moment[1].

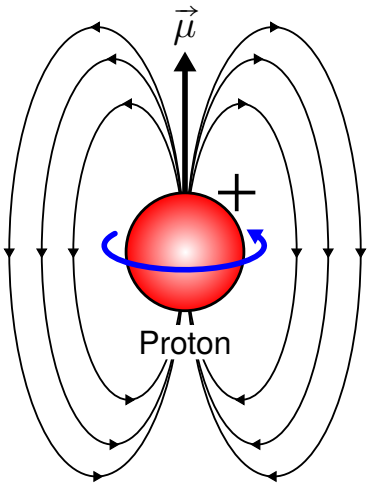


Figure 2.1: Dipolar magnetic field of a proton. The blue arrow represents the direction of the proton's spin as if it were a spinning ball in classical physics, the magnetic moment vector, $\vec{\mu}$ is then defined as upward based on the right-hand rule.

In a given population of protons, their spins are randomly oriented and their magnetic moments cancel out on average. However, when exposed to the main field of an MRI scanner, there is a slight induced bias of spin alignment towards $\vec{B}_0$ following the Boltzmann distribution, leading to a net magnetisation, $\vec{M}$, which is aligned with $\vec{B}_0$ at thermal equilibrium. The manipulation and detection of this net magnetisation forms the

[1]The spin angular momentum of a proton, $\vec{I}$, is parallel to the magnetic moment as $\vec{\mu} \propto \vec{I}$. $\vec{I}$ is a vector quantity whose direction corresponds to the proton spin axis.

basis of the MR signal. Figure 2.2 illustrates a vertical 2D slice through a spin system, showing spins with a common origin, and the resulting net magnetisation when subjected to $\vec{B}_0$. The z-axis is defined as parallel to $\vec{B}_0$, following MRI convention.

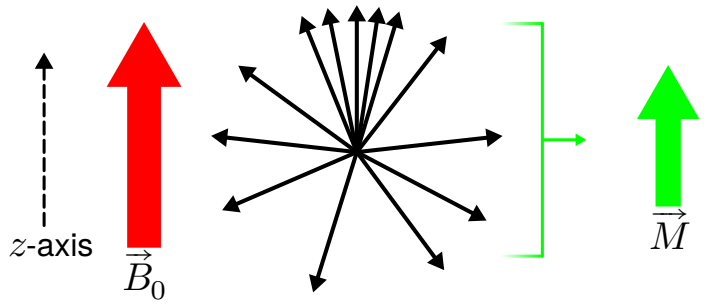


Figure 2.2: Net magnetisation, $\vec{M}$, generated by slight bias of proton spin orientation towards $\vec{B}_0$. $\vec{B}_0$ is illustrated by the red arrow, the proton spin distribution of the spin system is represented by the black arrows, showing a slight skew towards alignment with $\vec{B}_0$, and the resulting net magnetisation, $\vec{M}$, is illustrated by the green arrow.

2.2.2 Source of MR Signal

Protons exposed to $\vec{B}_0$ experience a torque that acts to align their spins with the direction of $\vec{B}_0$. However, due to conservation of angular momentum at the quantum scale, the spins do not align directly but instead precess around $\vec{B}_0$, analogous to the precession of a gyroscope under gravity. This precession occurs at the Larmor frequency [20], which is proportional to the strength of the magnetic field, as described in Equation 2.2.

$$f_0 = \gamma B_0 \tag{2.2}$$

Where f_0 is the Larmor frequency and γ is the gyromagnetic ratio; for protons, $\gamma = 42.58$ MHz/T. At thermal equilibrium, the net magnetisation $\vec{M}$ is aligned with $\vec{B}_0$ and does not precess, as the individual spins it represents have random relative phases that cancel in the transverse plane. Consequently, $\vec{M}$ has no transverse components. However, $\vec{M}$ can be tipped away from $\vec{B}_0$ by applying a secondary magnetic field, denoted B_1, oriented perpendicularly to $\vec{B}_0$ and rotated at the Larmor frequency.

B_1 is applied at the Larmor frequency to excite the spin system by delivering energy at its resonant frequency[2]. This coherent excitation tips $\vec{M}$ towards the transverse plane. The distribution of individual spins retains its slight initial skew, resulting in a

[2]This is the resonance of Magnetic "Resonance" Imaging.

net transverse component, $\overrightarrow{M}_{\perp}$, and, a net z-aligned component, $\overrightarrow{M}_{z}$. $\overrightarrow{M}_{\perp}$ precesses about $\vec{B}_0$ as it represents the precession of the skewed distribution about $\vec{B}_0$. This coherent, precessing transverse magnetisation produces a time-varying magnetic field, which is the detectable MR signal. The angle through which the spins are tipped is known as the "flip angle", which is a design parameter of pulse sequences as discussed in Section 4.3.9.

2.2.3 Relaxation

The tipped net magnetisation returns to its equilibrium alignment with $\vec{B}_0$ over time through a process known as "relaxation". Relaxation is characterised by two time constants, T1 and T2. T1 characterises the rate of spin-lattice relaxation, whereby energy in the spin system from excitation is lost to the surrounding molecular environment as the system returns to thermal equilibrium, and thus net magnetisation realigns with $\vec{B}_0$. T2 characterises the rate of transverse relaxation, where phase coherence between spins is lost and thus net transverse magnetisation decays.

Figure 2.3 shows a top-down view (looking against $\vec{B}_0$) of a horizontal 2D slice of spins in the transverse plane shown with a common origin. It illustrates the decay of the transverse component of net magnetisation, $\overrightarrow{M}_{\perp}$, due to phase decoherence over time in the spin system. $\vec{B}_0$ is represented by the red concentric circle and dot, and the spins are shown in a rotating reference frame such that $\overrightarrow{M}_{\perp}$ always points left to right.

Pulse sequences are designed to emphasise the contrast arising from different relaxation mechanisms. T1 weighted sequences are sensitive to tissues that have short T1 relaxation times. This is achieved with a short repetition time (TR, the time between successive B_1 excitations), and short echo time (TE, the time between excitation and signal acquisition). Under these conditions a spin that is more aligned with $\vec{B}_0$ prior to excitation has a greater signal contribution to $\overrightarrow{M}_{\perp}$ after excitation than a spin that is less aligned. The short TE minimises the effect of T2 decay on the acquired signal to emphasise T1 decay as a contrast mechanism. T1 weighted images thus show tissues with short T1 as bright.

T2 weighted sequences exploit differences in transverse relaxation times between tissues. This is achieved with longer TR and TE than T1 sequences. Tissues with longer T2

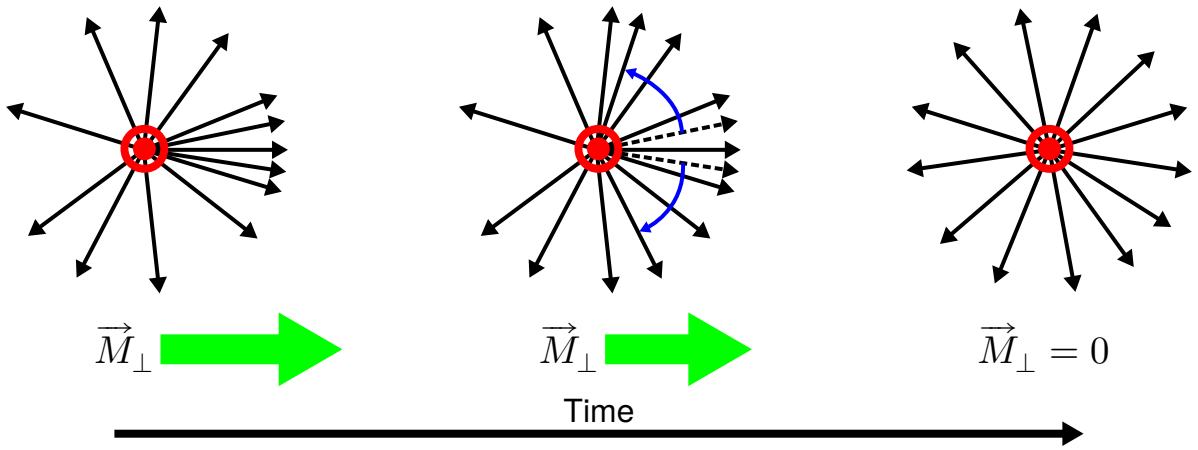


Figure 2.3: T2 relaxation caused by dephasing of initially coherent spins, viewed from above. $\vec{B}_0$ is represented by the concentric red circle and dot. Left: Spins that just have been rotated into the transverse plane. The initial skew in their alignment towards $\vec{B}_0$ before excitation now generates transverse magnetisation $\vec{M}_\perp$. Centre: Progressive loss of coherence reduces $\vec{M}_\perp$, the previous spin positions are shown as dotted arrows. The blue arrows represent spin dephasing causing spins to gain or lose phase relative to the total spin system. Right: Even distribution of spins yields no net transverse magnetisation. This diagram is shown in a rotating reference frame such that $\vec{M}_\perp$ always points left to right, in a static reference frame $\vec{M}_\perp$ would precess about $\vec{B}_0$ whilst decaying.

retain phase coherence for longer and thus yield more signal at acquisition time. T2 weighted images therefore show tissues with long T2 as bright. Variable T2 relaxation is the primary mechanism underlying the magic angle effect; the T2 of tissues displaying the magic angle effect is orientation-dependent, arising from the anisotropic structure of collagen fibres, which modulates the dephasing rate of coherent spins within the tissues. This is further discussed in Section 2.3.

The decay of $\vec{M}_\perp$, and consequently the measured MR signal, follows an exponential trend, described by Equation 2.3.

$$\left|\left|\vec{M}_\perp\right|\right|_2(t) = \left|\left|\vec{M}_\perp\right|\right|_2(0) \cdot \exp\left(-t/\mathrm{T2}\right) \qquad 2.3$$

Where $||\vec{M}_\perp||_2(t)$ is the magnitude of $\vec{M}_\perp$ at time t after excitation, and $||\vec{M}_\perp||_2(0)$ is the magnitude of $\vec{M}_\perp$ at $t = 0$. The T2 time constant characterises the decay rate, after T2 seconds $||\vec{M}_\perp||_2$ will be $\approx 37\%$ of its initial value. A larger T2 corresponds to a slower decay. Figure 2.4 illustrates the normalised signal decay for three hypothetical tissues with different T2 values.

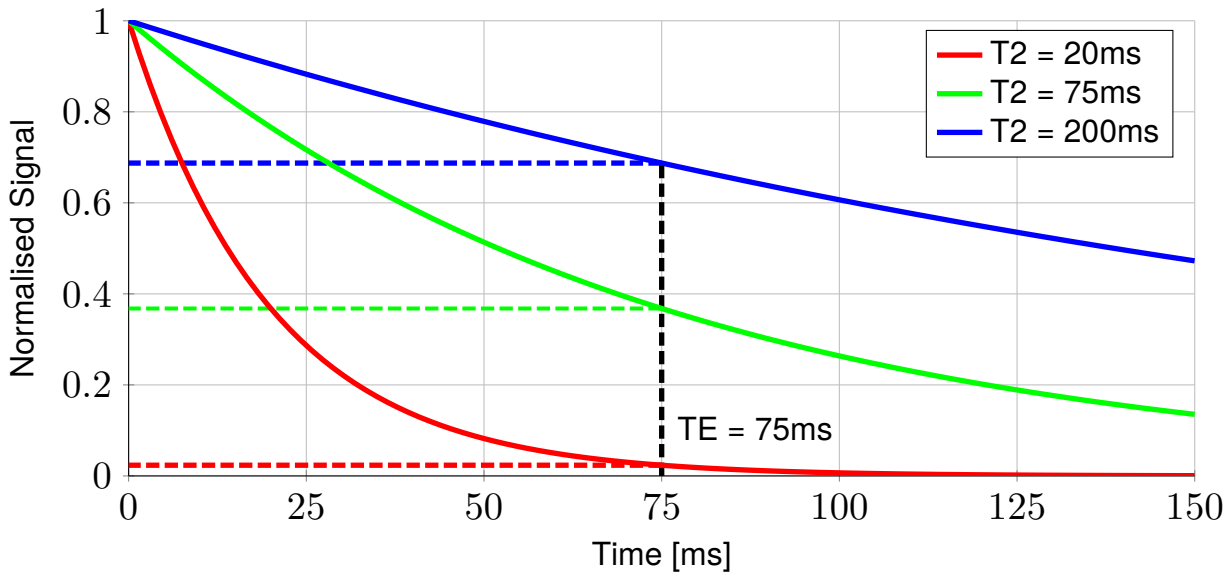


Figure 2.4: Example of T2 signal decay for three hypothetical tissues with T2 values of 20ms (red), 75ms (green), and 200ms (blue). At an echo time (TE) of 75 ms (illustrated with black dotted line), the red signal has largely decayed, the green signal has partially decayed, and the blue signal remains strong. In a corresponding MR image, the tissue with T2 $= 200$ms would appear brightest, followed by the tissue with T2 $= 75$ms, then the tissue with T2 $= 20$ms.

In Figure 2.4, the red trace corresponds to a tissue with T2 $= 20$ms, the green to T2 $= 75$ms, and the blue to T2 $= 200$ms. At TE of 75ms, the red signal has almost completely decayed, the green has decayed to ≈ 0.4, and the blue remains relatively strong at ≈ 0.7 on the normalised scale. In MR images, a stronger signal corresponds to increased brightness; therefore, tissues with longer T2 values appear brighter in T2 weighted acquisitions. Mechanisms that prolong T2, such as the magic angle effect, consequently increase tissue signal intensity when using appropriately designed pulse sequences.

2.2.4 Dipole-Dipole Interactions

T2 relaxation is primarily driven by magnetic dipole-dipole interactions between neighbouring protons. When two nearby protons are exposed to each other's magnetic fields, they affect each other's precession frequencies. The perturbed precession frequency of a proton, f_p, is given by Equation 2.4, which extends Equation 2.2 to account for dipolar interactions.

$$f_p = \gamma \left| \vec{B}_0 + B_{\mu z} \right| = f_0 + \delta f \qquad 2.4$$

Where $B_{\mu z}$ is the experienced z-component of a neighbouring proton's magnetic field. Note that only the z-component of a neighbouring proton's field is significant in this model, as $\vec{B}_0 + \vec{B}_\mu \simeq \vec{B}_0 + B_{\mu z}$ since $|\vec{B}_0| \gg |B_\mu|$. Depending on the relative spatial arrangement of nearby protons, δf may increase or decrease the precession frequency of neighbouring protons relative to the nominal Larmor frequency. In biological tissues, every proton experiences a unique local magnetic environment, and thus their precession rates diverge over time, leading to a loss of phase coherence across the spin system and a decay of detectable transverse magnetisation. Figure 2.5 illustrates how the dipolar fields of neighbouring protons perturb one another's local magnetic fields.

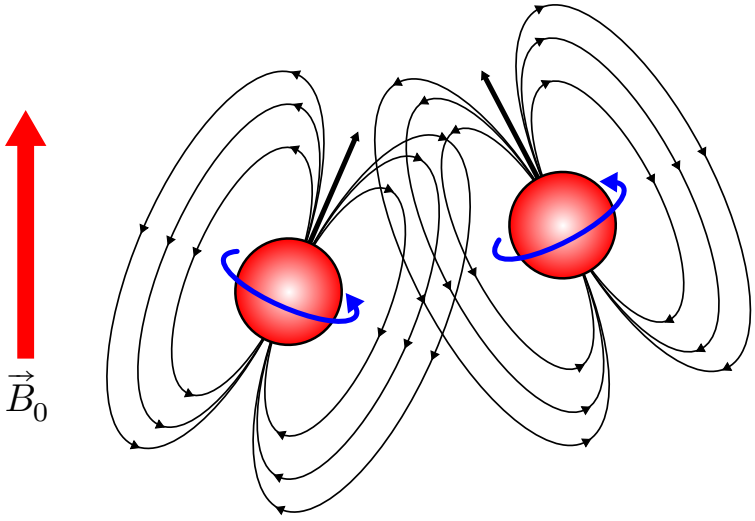


Figure 2.5: Dipole-dipole interaction between two neighbouring protons, each proton influences the local magnetic environment of its neighbour. In this example, the magnetic field of each proton serves to decrease the precession frequency of its neighbour due to the negative $B_{\mu z}$ component they each experience due to their relative displacement. This is evident as the vertical component of the field lines from each proton run against $\vec{B}_0$ when interacting with their neighbouring proton.

In biological tissues, unbound water is a principal source of the MR signal. Water molecules move and rotate freely, causing the dipolar interactions between their protons to occur with random relative orientations over short timescales. Precession frequency perturbations are cancelled out over time by these transient random interactions, thus preserving phase coherence and resulting in long T2 relaxation times and high signal intensity in MR images. In contrast, water molecules bound to tissue surfaces are spatially constrained, leading to persistent dipolar interactions that do not average out. This causes rapid dephasing, short T2 values, and consequently low MR signal from these protons. Tissues such as tendons, ligaments, and menisci contain predominantly bound water, and therefore typically appear dark in conventional MR images. However,

due to their anisotropic microstructure, dipolar interaction strength is modulated by collagen fibre orientation relative to $\vec{B}_0$, giving rise to the magic angle effect.

2.2.5 T2 and T2* Relaxation

The T2 time constant described in Sections 2.2.3 and 2.2.4 characterises the decay of transverse magnetisation arising from stochastic spin-spin interactions with the local molecular environment. In practice, the observed decay of transverse magnetisation is faster than that predicted by T2 alone. This accelerated decay is characterised by the time constant T2* which is always less than or equal to T2.

T2* relaxation is principally due to inhomogeneities in the B_0 field; either from imperfections in the B_0 field itself, or from susceptibility-induced distortions arising from the difference in magnetic susceptibility of adjacent tissues. Equation 2.5 models the T2* time constant following a linear relationship:

$$\frac{1}{\mathrm{T2^*}} = \frac{1}{\mathrm{T2}} + \gamma \Delta B_0 = \frac{1}{\mathrm{T2}} + \frac{1}{\mathrm{T2_i}} \tag{2.5}$$

Where γ is the gyromagnetic ratio, ΔB_0 is inhomogeneity in B_0, and $\frac{1}{\mathrm{T2_i}} = \gamma \Delta B_0$. It is apparent from Equation 2.5 that $\mathrm{T2_i} \geq 0$ will yield $\mathrm{T2^*} \leq \mathrm{T2}$. Since T2* arises from static B_0 inhomogeneities, it is possible to counteract this effect with an appropriate choice of pulse sequence. Spin echo sequences use a 180° refocusing pulse to invert the sign of accumulated phase for all spins simultaneously. Spins then continue to precess in the same direction as before the refocusing pulse, and therefore rephase at $\mathrm{TE} = 2\tau$, where τ is the time between the initial excitation and the 180° refocusing pulse. Thus, the dephasing caused by inhomogeneities is effectively cancelled out such that the residual signal decay reflects pure T2 relaxation. Gradient echo sequences, on the other hand, do not include this refocusing pulse and are thus sensitive to T2* decay.

At low field the difference between T2 and T2* is less than at high field. This is a direct result of the lower strength of the B_0 since susceptibility-induced fluctuations in B_0 are proportional to B_0 itself: $\Delta B_0 \propto B_0$ [21–23]. Therefore, the inhomogeneity contribution to transverse relaxation, $\gamma \Delta B_0$, is proportionally suppressed at low field. Consequently, T2* converges toward T2 at lower fields, and pulse sequences sensitive to

T2* decay yield contrast that more closely reflects the underlying T2 of the imaged tissues. The choice of pulse sequence used in the present work is discussed in Section 4.3.9.

2.3 The Magic Angle Effect

The magic angle effect is a well-documented phenomenon in MRI, first reported by Berendsen [24] in 1962, in which the signal intensity and relaxation properties of organised, anisotropic, collagen-rich tissues depend on their fibre orientation relative to $\vec{B}_0$ [6–9, 12, 25–29]. The effect arises from dipole-dipole interactions between protons in bound water molecules, whose orientation is restricted by the highly ordered, anisotropic structure of collagen fibres. This fibre arrangement prevents complete motional averaging of dipolar interactions, whilst also causing interaction strength to be a function of fibre orientation. This section outlines the physical basis of the magic angle effect, beginning with the angular dependence of dipolar interaction strength, followed by the role of collagen microstructure in mediating these interactions to produce the observed signal variation.

2.3.1 Dipolar Interaction Strength

The magic angle effect arises from the orientation-dependent strength of dipolar interactions between protons. As discussed in Section 2.2.4, it is the z-component, $B_{\mu z}$, of the interacting proton's magnetic fields that modifies their respective precession frequencies about $\vec{B}_0$, leading to short T2, and low signal intensity. The z-component of a proton's magnetic field in polar coordinates, $B_{\mu z}(r, \theta)$, is given by Equation 2.6 [30].

$$B_{\mu z}(r, \theta) = \frac{|\vec{\mu}|}{r^3} \cdot (3\cos^2\theta - 1) \tag{2.6}$$

Where r is the radial distance from the centre of the proton, and θ is the polar angle from the z-axis. If two protons are oriented such that they experience $B_{\mu z}(r, \theta) = 0$, the dipolar interaction strength between them is zero, minimising precession frequency perturbations and thus yielding an increase in T2 and signal intensity.

From Equation 2.6, $B_{\mu z}(r, \theta) = 0$ when $(3\cos^2\theta - 1) = 0$, or when $|\vec{\mu}|/r^3 = 0$ for the trivial case of $r \to \infty$. Solving $(3\cos^2\theta - 1) = 0$ for $\theta \in [-180°, 180°]$,

gives $\theta = \pm 54.7°, \pm 125.3°$. A contour plot of Equation 2.6 is shown in Figure 2.6, with $\theta = 54.7°, 125.3°$ emphasised with dotted lines.

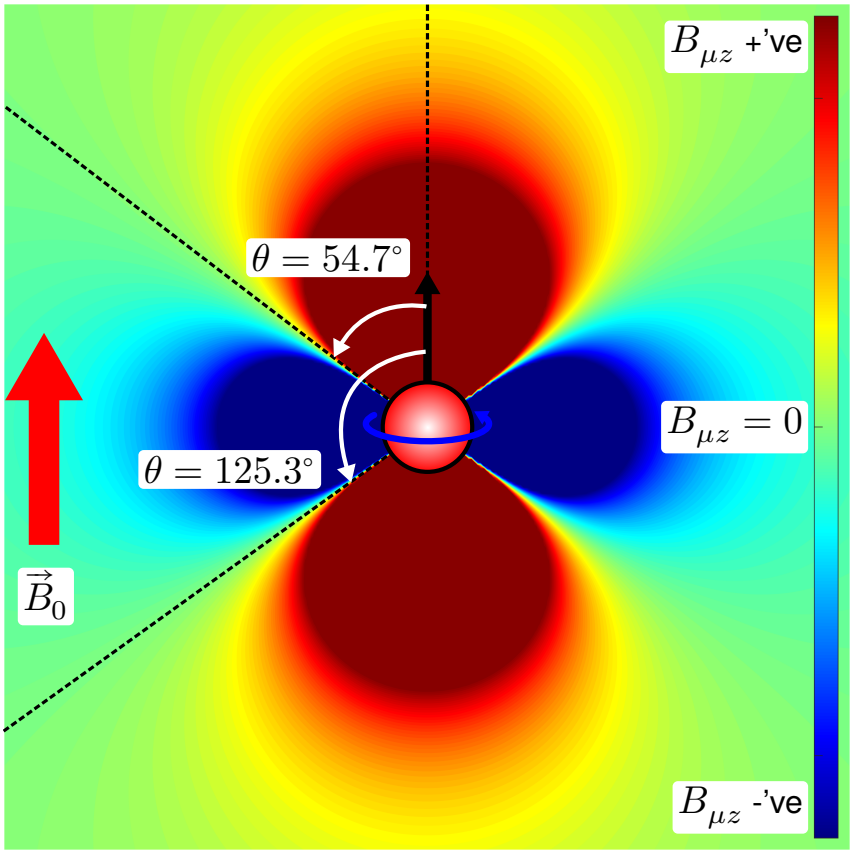


Figure 2.6: Contour plot of the z-component of the magnetic field around a proton with its spin aligned with $\vec{B}_0$. Red regions indicate positive $B_{\mu z}$, blue regions indicate negative $B_{\mu z}$, and green regions indicate $B_{\mu z} = 0$. Angles of $54.7°$ and $125.3°$ are shown with dotted lines, which align with $B_{\mu z} = 0$ spokes between $B_{\mu z} > 0$ and $B_{\mu z} < 0$ regions. Two similar $B_{\mu z} = 0$ regions are visible along the vertical mirror symmetry of the dipolar field.

The red lobes in Figure 2.6 indicate regions where $B_{\mu z} > 0$, the blue lobes where $B_{\mu z} < 0$. The narrow green zones between them represent angular orientations for which $B_{\mu z} = 0$, forming four "spokes" at $\theta = \pm 54.7°, \pm 125.3°$ from the vertical axis. From Equation 2.4, a proton experiencing a local field perturbation from the red lobes would precess at a higher frequency than the base Larmor frequency, f_0. Conversely, a proton perturbed by the blue lobes would precess at a frequency below f_0. Protons oriented along an interaction axis aligned with the thin zero intensity green zones would continue to precess (unperturbed) at f_0, as $B_{\mu z} = 0$ along this axis.

The angles $\theta = \pm 54.7°$ and $\pm 125.3°$ are the "magic angles", where dipole-dipole interactions between protons become null. When the dipoles of two protons are oriented at any of these angles relative to each other, their mutual interactions do not perturb

each other's precession frequencies. Although individual proton spins rarely align perfectly parallel (spin up) or antiparallel (spin down) relative to $\vec{B}_0$, they always precess around it, and their time-averaged orientation can be treated as purely spin up or spin down. Consequently, the contour plot shown in Figure 2.6 accurately depicts the time-averaged magnetic field experienced by surrounding protons, and it is this time-averaged interaction, not the instantaneous field, that governs the magic angle effect.

Figure 2.6 also illustrates the symmetry of the dipolar field surrounding a proton, which has both vertical and horizontal mirror symmetry. As a result, interacting protons can both be spin up, both be spin down, or have opposing spins, yet the magic angles for each condition remain the same. Moreover, due to these symmetries, the entire contour map can be reconstructed from a single quadrant via reflection, so only angles $\theta \in [0^\circ, 90^\circ]$ need be considered to fully characterise the dipolar field. Accordingly, the magic angle may simply be defined as 54.7°. This convention is particularly useful when considering the angle between $\vec{B}_0$ and a collagen fibre, as collagen fibres are bidirectional and lack a defined forwards and backwards direction (as a vector would have). The acute angle thus always characterises the relevant angular relationship between $\vec{B}_0$ and collagen fibre orientation.

2.3.2 Collagen Structure

Collagen is the most abundant protein in the human body and exists in multiple isoforms, each adapted to specific structural and functional roles. Type I collagen is the predominant fibrillar form, comprising the majority of load-bearing connective tissues including tendons, ligaments, menisci, and fibrocartilage [31]. Other types serve more specialised functions: type II collagen is found primarily in hyaline cartilage, type III contributes to the structure of reticular fibres in organs such as the liver and spleen, type IV forms part of basement membranes, and type V is involved in the regulation of fibril diameter and is often co-expressed with type I [31].

In structures such as the ACL and meniscus, type I collagen fibres align in an organised, anisotropic fashion to provide high tensile strength along their principal loading direction, analogous to how strands of a rope are oriented to resist tension. This alignment underpins both the mechanical function of these tissues, and their orientation-dependent MR intensity due to the magic angle effect. Figure 2.7 illustrates the hierarchical and

anisotropic structure of a tendon; ranging from tropocollagen to the macroscopic tendon structure.

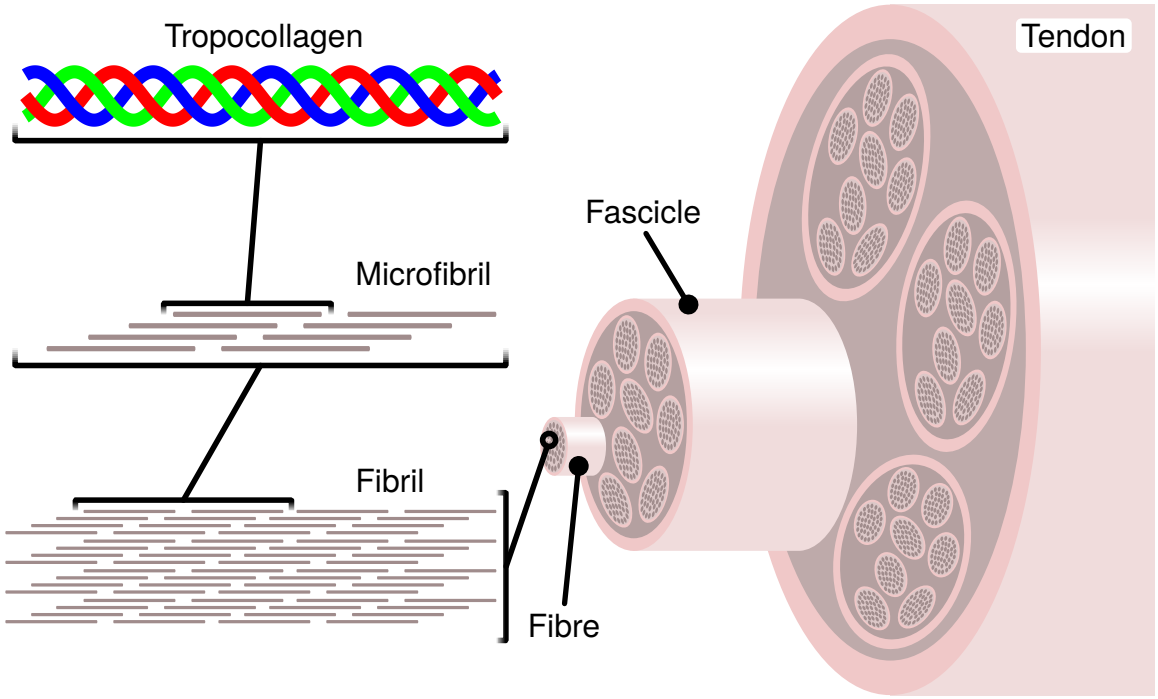


Figure 2.7: Hierarchical structure of tendon [31, 32]. Tropocollagen molecules assemble in a staggered arrangement to form microfibrils, which bundle into collagen fibrils. These fibrils group into fibres, which are further organised into fascicles. Multiple fascicles collectively form the macroscopic tendon. All levels of the hierarchy run parallel to each other, meaning the entire structure is anisotropic.

Collagen-rich tissues exhibit a hierarchical structure spanning multiple length scales: tropocollagen assembles into fibrils, which in turn form fibres, fascicles, and ultimately macroscopic tissue structures [31, 32]. Table 2.1 provides a summary of the key features at each hierarchical level, including approximate diameters, structural characteristics, and functional roles [31, 32].

In tissues composed of collagen, water is bound to tropocollagen molecules via hydrogen bonding [28]. These bound water molecules are orientationally restricted, held in the transverse plane relative to the collagen fibre axis by electrostatically charged sites on adjacent collagen chains. This restriction limits molecular motion, preventing dipolar interactions from averaging out, resulting in rapid dephasing and very short T2 relaxation times [28]. However, when the collagen fibres are oriented at the magic angle relative to $\vec{B}_0$, the $B_{\mu z}$ component of the local dipolar field becomes zero. This suppresses dipolar interactions, prolongs T2 relaxation, and yields markedly increased signal intensity in MR images. In experiments to determine the T2 variation between tendons at 0° and

Table 2.1: Summary of each level of the hierarchical collagen structure. The name, approximate diameter, structure, and function are given for each level of the hierarchy.

Name	Diameter	Structure	Function
Tropocollagen	1.5 nm	Triple-helical polypeptide composed of three α-chains	Fundamental molecular unit of collagen; forms the basis for higher-order assembly via staggered packing and cross-linking; binds water via hydrogen bonding
Microfibril	4-5 nm	Assembly of tropocollagen molecules with regular axial staggering	Intermediate substructure facilitating lateral and longitudinal stability in fibrils
Fibril	10-500 nm	Long cylindrical arrays of staggered tropocollagen with characteristic 67 nm D-banding periodicity	First structurally continuous level with significant mechanical integrity
Fibre	1-20 μm	Bundles of aligned fibrils surrounded by a sparse proteoglycan matrix	Functional subunits that transmit tensile forces
Fascicle	50-300 μm	Groups of fibres enclosed by the endotenon, a connective sheath containing blood vessels, nerves, and interstitial fluid	Mechanically distinct load-bearing units; allows tendon deformation and facilitates fascicular sliding under strain
Tendon	500 μm–1cm	Connective tissue composed of fascicles embedded within the epitenon	Transmits muscle force to bone; exhibits anisotropic mechanical properties and magic angle effect in MRI due to fundamental anisotropic structure

$54.7°$ relative to $\vec{B}_0$, Fullerton et al. [25] measured an increase in T2 from 0.6ms to 22ms in the Achilles tendon, and Henkelman et al. [12] measured a similar increase from 7ms to 23ms in a bovine Achilles tendon. Figure 2.8 shows a simplified model of the interactions that cause the magic angle effect at a molecular scale.

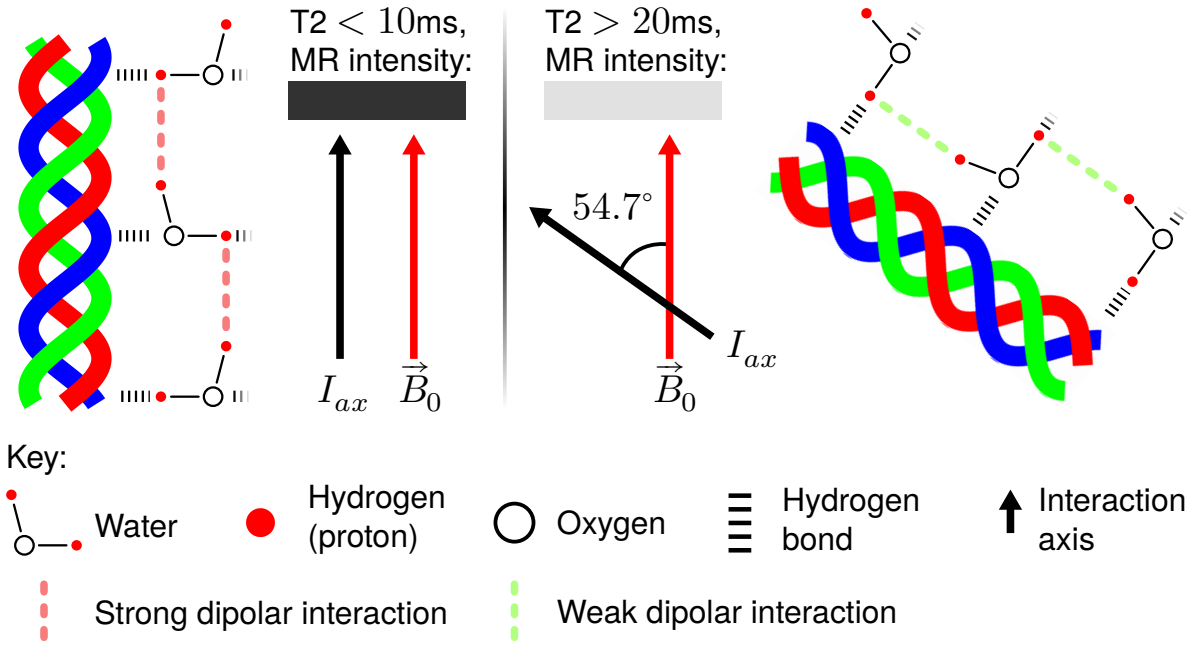


Figure 2.8: Orientationally restricted water bound to tropocollagen chains. Water molecules are shown bridging to (unseen) adjacent tropocollagen chains, and are thus restricted to orientations transverse to the collagen fibre axis, like rungs on a ladder. Dipolar interactions are shown by dotted lines, and act parallel to the interaction axis, I_{ax}. Left: the collagen fibre axis is aligned with $\vec{B}_0$. Dipolar interactions between bound water protons are strong and unaveraged, resulting in rapid T2 relaxation and low MR signal intensity. Right: the collagen fibre axis is oriented at the magic angle, $54.7°$, relative to $\vec{B}_0$. This minimises the dipolar coupling term $(3\cos^2\theta - 1)$, reducing spin dephasing, prolonging T2, and increasing signal intensity in the tissue in T2 weighted images.

Figure 2.8 shows water molecules hydrogen-bonded to tropocollagen chains, resulting in their orientational restriction to the transverse plane relative to the collagen fibre axis. For clarity, each water molecule is shown bound to a single collagen chain. On the left, the collagen fibre is aligned parallel to $\vec{B}_0$, meaning the dipolar interaction axis, I_{ax}, is also aligned with $\vec{B}_0$. The dipolar interactions between protons, shown by the dotted lines occur parallel to this interaction axis. In this configuration, dipolar interactions between water protons are strong and unaveraged, leading to rapid dephasing, short T2 relaxation times, and low MR signal intensity. On the right, the collagen fibre is oriented at the magic angle, $54.7°$, relative to $\vec{B}_0$. This orientation minimises the angular term

$(3\cos^2\theta - 1)$ in Equation 2.6, thereby reducing dipolar interaction strength, slowing dephasing, extending T2, and increasing MR signal intensity.

Recent studies by Momot [29] and Tadimalla et al. [33] suggest the relaxation mechanisms causing the magic angle effect are not completely understood at the molecular level, specifically regarding exactly which bound water molecules contribute to the phenomenon. However, orientation-dependent dipolar interactions between bound water are still understood to be the fundamental cause, and observed magic angle intensity variations are reliably and accurately captured by current models [9–11, 29].

2.4 Orientation Dependent Intensity Model

Equation 2.7 describes the widely accepted model [9–11, 29] for the relationship between MR signal intensity and the angle between collagen fibres and $\vec{B}_0$. The model incorporates both the isotropic signal contribution from free water within the tissue and the orientation-dependent signal contribution from bound water [11].

$$I(\theta) = A \cdot \exp\left(-B \cdot \left(3\cos^2\theta - 1\right)^2\right) \tag{2.7}$$

In Equation 2.7, I is the observed MR intensity from tissues exhibiting the magic angle effect, A and B are fitting coefficients, and θ is the angle between the collagen fibre axis and $\vec{B}_0$. A captures the peak intensity, when $\theta = 54.7^\circ$ and $(3\cos^2\theta - 1) = 0$, and B characterises the rate of intensity change. The $(3\cos^2\theta - 1)$ component encodes the effect of the dipolar coupling strength from Equation 2.6 on the observed intensity I. Unlike in Equation 2.6 the $(3\cos^2\theta - 1)$ component is squared to capture the fact that both positive and negative $B_{\mu z}$ perturbations will lead to dephasing and therefore shorter T2.

Figure 2.9 shows a plot of intensity versus θ, with $A = 1$ and $B = \{0.5, 1, 2\}$, for $\theta \in [0^\circ, 90^\circ]$. The peak intensity occurs at $\theta = 54.7^\circ$, with intensity decay as $\theta \to 0^\circ$ and $\theta \to 90^\circ$ away from the peak. This intensity variation illustrates that dipolar interaction strength is greater when protons interact along the z-axis ($\theta \approx 0^\circ$) than when protons interact along axes perpendicular to the z-axis ($\theta \approx 90^\circ$) as $I(0) < I(90)$. These observations correspond with the $B_{\mu z}$ contours in Figure 2.6, which indicate

stronger dipolar interactions along the z-axis than perpendicularly to it, and nullified dipolar interactions at the magic angle of 54.7° (and its symmetrical spokes).

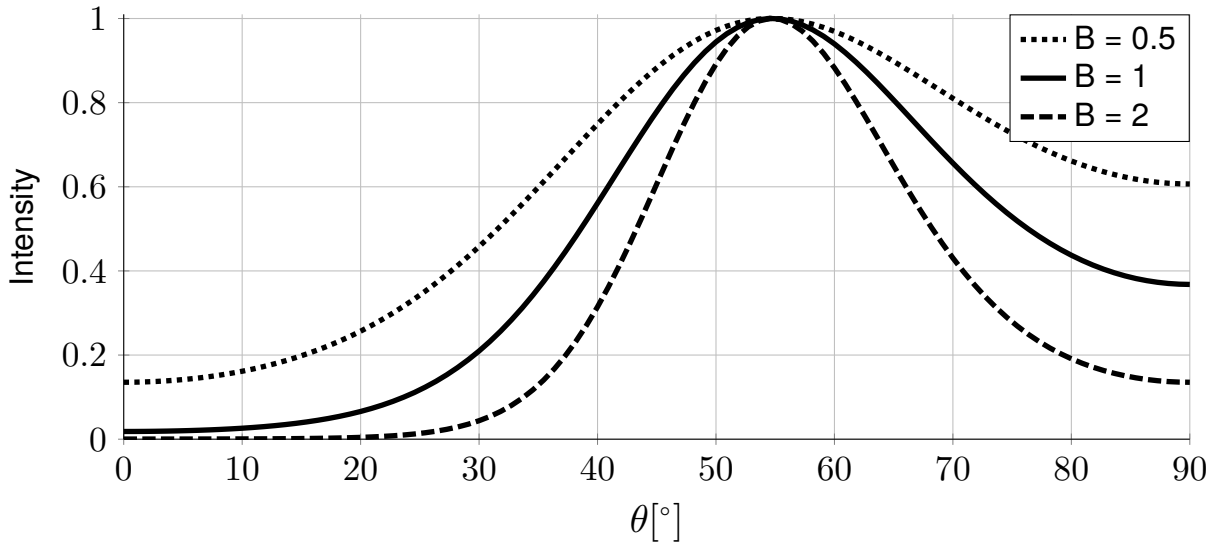


Figure 2.9: Intensity versus θ for $A = 1$ and $B = \{0.5, 1, 2\}$. The peak intensity is shown to occur at 54.7°, with intensity decaying away from the magic angle. Minimum intensity occurs at $\theta = 0^\circ$, with less intensity decay occurring at $\theta = 90^\circ$. Lower B values yield a smoother peak compared to higher B values. Higher B values cause a more rapid decay of intensity moving away from the magic angle.

In Figure 2.9, the θ domain is restricted to $\theta \in [0^\circ, 90^\circ]$ as this domain is sufficient to fully capture the orientation-dependent intensity response. The symmetries of the $B_{\mu z}$ field around a proton, combined with the bidirectionality of collagen fibres, mean that only the acute angle between $\vec{B}_0$ and the collagen fibres needs to be considered. This has practical implications in that the required $\vec{B}_0$ operating envelope[3] to achieve any θ relative to an arbitrary collagen fibre is a hemisphere rather than a full sphere. From the home position, the scanner therefore only needs to roll and yaw by $\pm 90^\circ$ (rather than $\pm 180^\circ$) to fully sample the intensity response of any arbitrary collagen fibre. In practice, however, the operating envelope is constrained by ergonomics and hardware limitations, as discussed in Chapter 4.

The fitting coefficients A and B have distinct roles in the model presented in Equation 2.7. A is simply a scaling coefficient to capture the peak intensity at the magic angle[4]. However, B captures several physical effects. Figure 2.9 shows a range of B values

[3]The operating envelope is the surface representing all attainable $\vec{B}_0$ orientations, with each orientation expressed as a point in 3D space. See Section 4.3.6 for more details.

[4]Raw MRI signal intensity is dimensionless and scanner-dependent, encoding relative signal strengths within an image rather than absolute tissue properties. A therefore serves as a scanner-dependent scaling coefficient rather than a direct measure of intrinsic tissue properties.

which illustrate its physical interpretation. B characterises the sensitivity of the observed intensity to θ, and it is determined jointly by tissue properties and pulse sequence parameters. In this orientation-dependent intensity model, non-orientation-sensitive tissues are characterised by $B = 0$; and collagen-rich tissues with organised collagen microstructure are characterised by $B > 0$, thus yielding intensity variation for $\theta \in [0^\circ, 90^\circ]$. B is also affected by the choice of TE. The longer the TE the more T2 / T2* signal decay occurs, thus reducing the dynamic range of $I(\theta)$. An optimal TE therefore maximises the contrast between the peak at $\theta = 54.7^\circ$ and the off-peak intensities, ideally yielding a high B. This in turn increases the efficacy of the MADI technique since it relies on intensity variation across scanner orientations, as discussed in Chapter 4.

In principle, B is a measurable and informative quantity for pulse sequence optimisation, however, quantitative estimation and optimisation of B in vivo remains a topic of future work (Section 9.5.2), as the feasibility of the MADI technique in vivo was the primary research objective. The present work utilised a pulse sequence based on the pulse sequence optimisation carried out by Chappell [34] as described in Section 4.3.9. Throughout this thesis $B = 0.5$ is used, as this value fit the measured intensity data well for the ACLs and menisci studied, as discussed in Section 7.2.3. Alternative pulse sequences would in general require B to be re-evaluated.

Equation 2.7 and Figure 2.9 represent an idealised, noise-free model of the magic angle effect. In practice, the low-field scanner used in the present work yielded noisy data. To assess and mitigate the effects of noise in the data, simulation studies with added synthetic noise and a genetic algorithm were utilised to select the $\vec{B}_0$ orientations employed in the MADI examination protocol. Each MADI dataset was chosen to comprise 7 images acquired across a diverse set of orientations, as described in Chapters 4 and 7, ensuring the technique was robust to noise by sufficiently sampling the domain of $I(\theta)$ during a complete MADI exam.

2.5 Anatomy of the Knee

The anatomical structures of interest for this thesis were the anterior cruciate ligament (ACL) and the meniscus. Both of these structures primarily consist of anisotropic fibrous type I collagen which exhibits the magic angle effect in MRI scans. These tissues were selected due to their clinical relevance, being among the most frequently injured structures in the knee, and also because they provide a useful spectrum of tissue morphology for imaging validation. Figure 2.10 illustrates the connective tissues present in the human knee joint. The ACL is shown in blue, connecting the femur and tibia, and the meniscus is shown in green, situated on the superior surface of the tibia.

The ACL is a predominantly linear, uniaxially aligned structure which resists anterior tibial translation and rotational forces [35]. In contrast, the meniscus exhibits a complex, curved figure-of-eight architecture, which functions to distribute and dampen loads between the femur and tibia by generating tensile hoop stress under joint compression [36, 37]. In a healthy joint, the collagen fibres of each structure are aligned parallel to the tensile forces that they experience: in the ACL the fibres run linearly along its length; and in the meniscus they flow circumferentially around the structure. Pathology of these tissues can involve complete or partial tears, or tissue degeneration. All of these pathologies would present at the microstructural level as disruption or disappearance of the tissue's collagen fibres.

Whilst this work focuses on the knee, the analysis of other extremity joints is equally feasible using the current prototype rotatable MR system. Future, larger-scale iterations of the system could extend the application of MADI to enable the study of a wide range of connective tissues throughout the body.

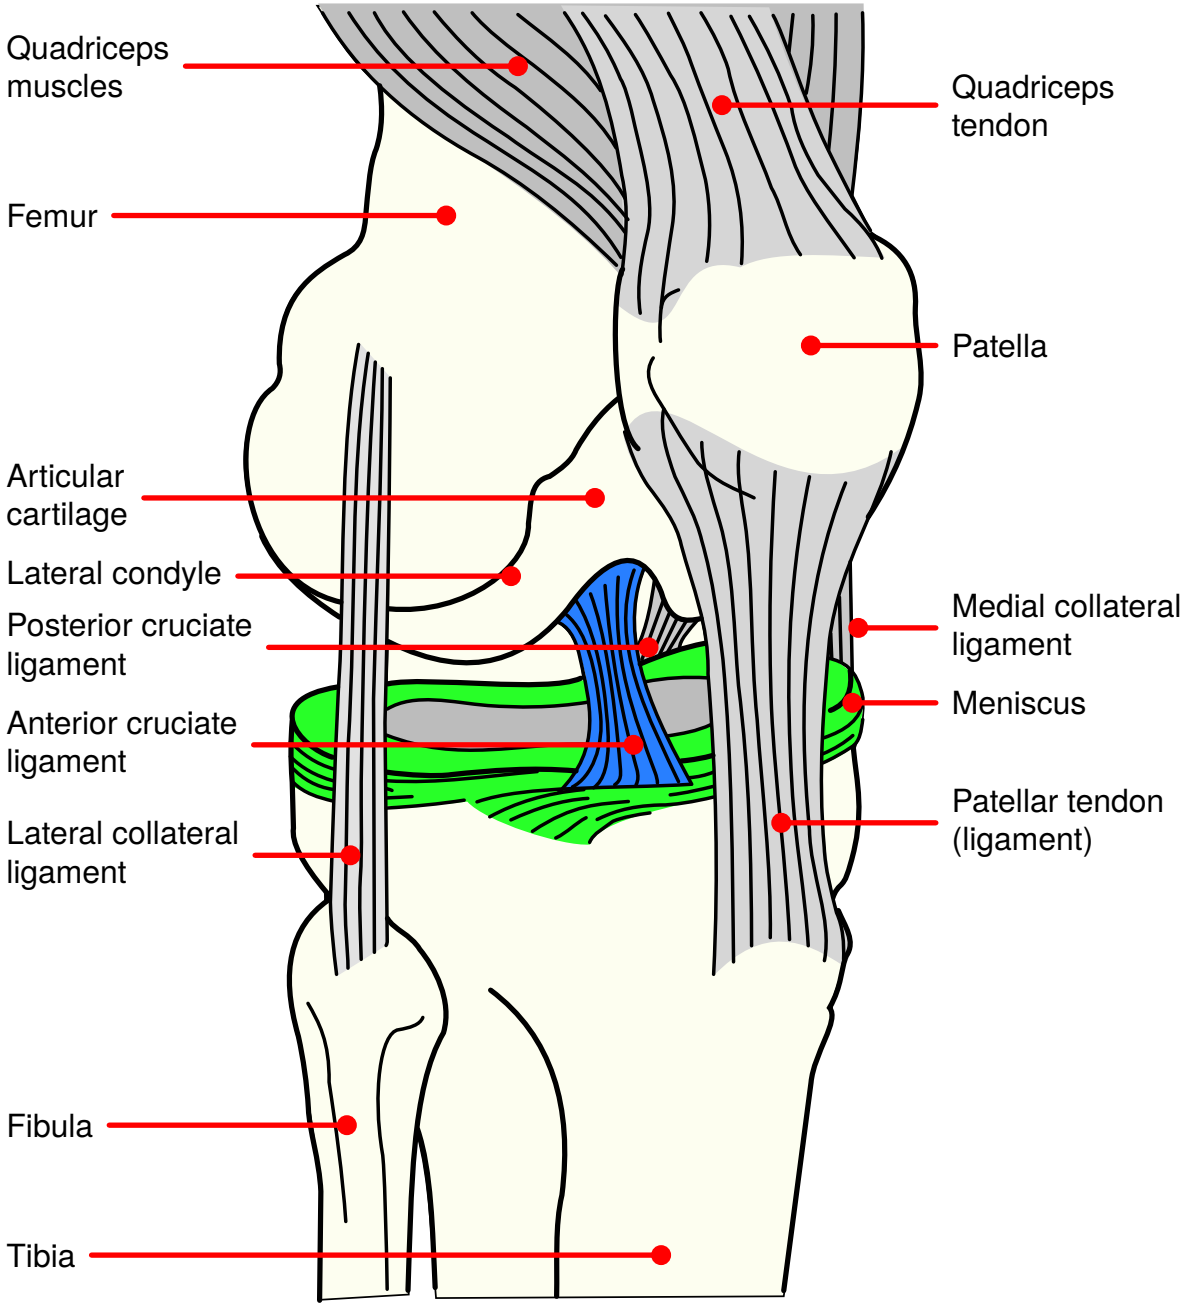


Figure 2.10: Diagram of knee joint connective tissue structures. The anterior cruciate ligament (ACL) is shown in blue originating from the posterior aspect of the femoral intercondylar notch and inserting anteromedially on the tibial plateau. The medial and lateral menisci are shown in green, capping the tibial plateau, interposed between the femoral condyles and the superior surface of the tibia. Lines on the soft tissue structures indicate the principal collagen fibre direction within the tissues, which is also the direction of tensile force the tissues typically experience. The gap between the femur and tibia has been exaggerated to better show the connective tissues within the joint.

Image attribution: Mysid - Vectorised and colourised in Inkscape, based on Image: Knee diagram.png, Public Domain, `https://commons.wikimedia.org/w/index.php?curid=4004042`

2.6 Summary

This chapter outlined the fundamental MRI principles underlying the magic angle effect, covering the quantum-mechanical origins of the MR signal, T2 / T2* decay, and the role of dipolar interactions in driving orientation-dependent T2 relaxation in anisotropic, collagen-rich tissues.

The strength of the z-component of the magnetic field around a proton gives rise to the "magic angle" of 54.7°, at which dipolar interactions between neighbouring protons are minimised. The reduced dephasing associated with minimised interaction strength leads to longer T2 values, and consequently higher signal intensity, if the axis of interaction can be constrained along these angles.

Tendons, ligaments, and menisci are primarily composed of type I collagen, the anisotropic, fibrous structure of which orientationally restricts the bound water within these tissues. Since the water molecules are constrained within this structure, the strength of dipolar interactions between their protons is mediated by the fibre orientation relative to $\vec{B}_0$. When oriented at the magic angle, this strength is minimised, T2 is extended, and signal intensity is maximised. This is the magic angle effect.

However, the magic angle effect is not simply an intensity spike when the angle between $\vec{B}_0$ and collagen fibres is 54.7°, it is a continuous orientation-dependent intensity variation which is a function of the angle, θ, between $\vec{B}_0$ and the collagen fibres being scanned. Equation 2.7 accurately models this intensity variation, characterising the entire intensity range with a θ domain of only $\theta \in [0^\circ, 90^\circ]$ due to inherent symmetries in a proton's magnetic field and the bidirectionality of collagen fibres.

The magic angle effect occurs in all tissues consisting of organised, anisotropic collagen fibres. Section 2.5 illustrates the anatomy of the human knee, with emphasis on the ACL and meniscus, which both exhibit the magic angle effect and are the focus of the work in this thesis.

3 | Literature Review

3.1 Chapter Overview

This chapter provides a focused literature review of the limitations of conventional MRI in musculoskeletal imaging of the knee, the use of diffusion tensor imaging (DTI) for examining collagen fibre microstructure, the clinical need for dedicated imaging systems, and the current state and future potential of low-field MRI in Sections 3.2 to 3.5.

The literature discussed here provides the rationale for the development of the prototype rotatable low-field scanner and the associated MADI technique presented in this thesis. Conventional MRI has well-recognised limitations in imaging connective tissues composed of organised, anisotropic collagen fibres, and its clinical availability is constrained by limited scanner capacity and radiologist resources. DTI has shown research promise for non-invasive collagen fibre assessment in MSK tissues. However, it faces fundamental technical barriers in MSK contexts including short T2, long acquisition times, and dependence on conventional MRI scanners, which limit its clinical translation. Low-field MRI offers potential solutions to several of these challenges. Dedicated, application-specific scanners can be developed cost-effectively, alleviating workload on conventional high-field systems, whilst providing a platform for techniques such as MADI that circumvent the short T2 and orientation sensitivity constraints that hinder both conventional MRI and DTI in this context.

3.2 Limitations of Conventional MRI

It is well known that connective tissues such as tendons, ligaments, and menisci, exhibit low signal and poor contrast in conventional MRI due to their short T2 relaxation times. In tendon and ligament, Bydder et al. [8] describe a baseline T2 of 2-8 ms when fibres are parallel with $\vec{B}_0$ ($\theta = 0^\circ$ in Equation 2.7). Even with echo times at the lower end of clinical sequences (TE $\approx$ 20 ms), this baseline T2 value is short enough such that tissues

consisting of organised collagen appear dark in conventional MR images. However, T2 relaxation time in these tissues is a function of their orientation relative to $\vec{B}_0$ due to the magic angle effect, and can thus be extended. As mentioned in Section 2.3.2, Fullerton et al. [25] observed an increase in T2 from 0.6 to 22ms in a human Achilles tendon, and similarly Henkelman et al. [12] observed an increase in T2 from 7 to 23ms in bovine Achilles tendon, when the tissues were oriented at $\theta = 0^\circ$ versus $\theta = 54.7^\circ$ relative to $\vec{B}_0$.

Whilst it may seem beneficial to increase MR signal, the variability of T2 is generally unwanted in conventional MRI exams as the magic angle effect is challenging to account for, and the increased intensity it yields is usually considered an imaging artefact as it can mimic or obscure pathology [6, 8]. For example, tears in connective tissues can cause oedema or haemorrhage, introducing fluids such as interstitial fluid or blood with longer T2 values into the tissues, thus increasing MR intensity which can be indicative of the pathology. A study of the meniscus by Peterfy et al. [6] found that for 31 of 42 patients assessed, increased signal in the upsloping portion of the posterior horn of the lateral meniscus was most probably due to the magic angle effect rather than tears or degeneration. However, due to the ambiguous source of the MR signal increase, pathology such as meniscal tears or degeneration in this region could not be ruled out either.

Many other studies have also reported difficulties in assessing connective tissues due to the complexity of the anatomy, poor tissue contrast, and the magic angle effect. Erickson et al. [26] and Hayes et al. [27] reported that due to the curvature of the supraspinatus tendon in the shoulder, collagen fibres near its insertion would tend to be oriented at the magic angle in conventional MRI, leading to diagnostic confusion as the increased intensity mimicked pathology. Van Dyck et al. [38] assessed the capability of 3T MRI to assess partial ACL tears, and found that differentiation between complete and partial tears of the ACL was challenging, and that MRI only had a maximum sensitivity of 77% for partial ACL tears due to the low signal intensity and complex morphology of partially torn ACLs. Temponi et al. [39] found that partial ACL tears were challenging to diagnose using physical clinical tests or MR imaging alone, and recommended a combination of clinical evaluation, MRI, and arthroscopy for diagnosis and management of partial ACL tears. Devitt et al. [40] found that MRI was not reliable for diagnosis of anterolateral ligament injuries in the setting of concomitant ACL injury “due to the inability to accurately

and consistently visualise the structure", and recommended physical examinations to assess suspicion of anterolateral ligament injury in this context. Li et al. [5] found that for multi-ligament knee injuries MRI had limited value in classification of pathology, and management should be based on intraoperative findings and physical examinations as well as imaging results.

Various approaches have been explored to mitigate the magic angle effect. Peh et al. [41] established a critical TE $= 37$ms, at which time the magic angle intensity increase in both human and bovine Achilles tendons was no longer observed. This approach is limited however, as not all sequences allow for modification of TE to this high a value whilst still retaining the desired image qualities for diagnostic imaging. Wright et al. [42] discuss the trade-off of using longer TE values, such as lower SNR, longer scan times, and a shift towards T2 weighting of the resultant image. Using TE > 37ms is therefore not a general solution to the magic angle artefact in MRI.

In their study of the patellar tendon, Karantanas et al. [43] suggest the magic angle effect can also be mitigated by reorientation of the target anatomy to ensure no fibres are at 54.7° relative to $\vec{B}_0$. However, this approach can be imprecise without a priori knowledge of the anatomy, and would equally not be possible in pathological scenarios where the tissues of interest may be displaced from their expected orientation, or for curved structures such as the meniscus and supraspinatus tendon which invariably lie at a range of angles relative to $\vec{B}_0$. Furthermore, the "magic angle" of 54.7° refers to the peak MR intensity expressed by collagen-rich tissues. Figure 2.9 demonstrates that there is a range of possible intensities between the baseline at $\theta = 0^\circ$ and $\theta = 54.7^\circ$, thus, simply avoiding the magic angle itself is no guarantee that increased tissue intensity will not be present in the images.

Bydder et al. [8] were among the first proponents for utilisation of the magic angle effect as "a determinant of image contrast, and a technique for imaging". They proposed "magic angle imaging", whereby tendons and ligaments would be deliberately placed at or near the magic angle to increase their MR intensity. However, they noted that this approach is limited to linear anatomical structures such as the Achilles tendon. For most collagen-rich tissues, patient reorientation to exploit the magic angle is not a practical or generalisable strategy, just as it is often equally unfeasible to completely avoid the magic angle effect in the same manner. In addition, techniques requiring patient reorientation

within a conventional static MRI scanner are constrained by the limited bore size and may compromise patient comfort if unnatural postures are needed. Consequently, whilst conceptually promising, magic angle imaging has remained confined to research settings and has not seen clinical adoption.

Numerous systematic reviews have evaluated the diagnostic accuracy of MRI for knee pathologies, particularly involving the ACL and meniscus. These studies consistently report that MRI exhibits good sensitivity and specificity for detecting lesions and degeneration; however, arthroscopy remains the diagnostic gold standard. Crawford et al. [44] performed a review of MRI diagnosis of knee pathology, with a focus on meniscal lesions and ACL tears, using arthroscopy as a reference standard. They found MRI had a diagnostic accuracy of 86%, 89%, and 93%, in the medial meniscus, lateral meniscus, and ACL respectively, with a total accuracy of 89% across the combination of these structures. Crawford et al. note that these relatively high metrics could be a function of publication bias, as studies with non-significant results are less likely to be published. They also note that the meniscus is challenging to visualise with MRI, with lower specificity in the medial meniscus potentially leading to unnecessary arthroscopies if MRI is used as the only form of pre-operative screening for medial meniscal lesions. Crawford et al. also discuss the fact that MRI studies have higher rates of false positives than false negatives, which could lead to unnecessary arthroscopies. They go on to discuss that MRI is not reliable for diagnosing recurrent meniscal tears, detecting only 66%, versus 88% using arthroscopy. Overall they conclude that MRI is a useful screening tool prior to therapeutic arthroscopy, and that a combination of clinical examination and MRI provides the most accurate non-invasive source of information about meniscal and ACL pathology.

Phelan et al. [4] performed a systematic review of the diagnostic accuracy of MRI for suspected ACL and meniscal tears, also with arthroscopy as the reference standard. They only reviewed prospective studies to mitigate verification bias, and found that MRI had a sensitivity and specificity of: 89% and 88% for medial meniscal tears, 78% and 95% for lateral meniscal tears, and 87% and 93% for ACL tears. Phelan et al. note that many studies had high or unclear risk of bias due to occurrences such as: MRI findings determining whether arthroscopy was performed (verification bias), surgeons that were unblinded to MRI findings (confirmation bias), and unclear processes for patient selection (selection bias), which could skew the MRI results towards appearing more accurate

overall. Whilst MRI is undeniably a useful screening and management tool for pathology of the knee, Phelan et al. warn that “surgeons should be aware of the level of evidence supporting its use when interpreting results and should question its applicability in the context of their clinical setting”.

Finally, Koo et al. [45] performed a systematic review that assessed the diagnostic performance of MRI for detecting ramp lesions in patients with ACL tears, with arthroscopy as the reference standard. They found that MRI demonstrated moderate sensitivity of 71% and an excellent specificity of 94%. The sensitivity was increased to 84% when the analysed studies were restricted to only studies that utilised 3T scanners, and with MR data interpretation by MSK radiologists rather than orthopaedic surgeons. This highlights the significant influence of interpreter expertise on the accuracy of knee pathology diagnosis via MRI, which was also reported by Krampla et al. [46]. From this review, Koo et al. concluded that “routine arthroscopic assessment is recommended for the presence of ramp lesion, regardless of whether it is suspected on MRI” and that “further clinical or radiological studies of the diagnostic algorithms for detecting ramp lesion, including high-resolution MRI with appropriate knee position, may be needed”.

This body of literature highlights the need for a non-invasive imaging technique capable of assessing collagen fibre microstructure in connective tissues that typically have low MR signal intensity, complex anatomical structure, and exhibit the magic angle effect. The proposed Magic Angle Directional Imaging technique developed in this thesis addresses this need by enabling controlled, repeatable manipulation of signal intensity through precise reorientation of $\vec{B}_0$ across multiple acquisitions. Combined with the analysis and visualisation framework presented later in this thesis, MADI offers functionality more akin to a “virtual arthroscopy” than conventional MRI. Rather than relying on indirect markers such as hyperintense signal from fluid ingress, MADI enables direct visualisation of the collagen fibre microstructure in both healthy and pathological tissue; whilst also simplifying data interpretation through automatic identification of tissues exhibiting the magic angle effect.

3.3 Diffusion Tensor Imaging for Examining Collagen Fibre Structure

Diffusion tensor imaging is an MRI technique that quantifies the directional dependence of water diffusion within biological tissues through acquisition and processing of a set of diffusion weighted images (DWI). Stejskal et al. [47] demonstrated that incorporating a pair of pulsed magnetic field gradient lobes into a spin-echo acquisition produces signal attenuation proportional to the displacement of water molecules along the gradient direction, with stationary spins fully refocused at the echo. This is the pulsed-gradient spin-echo (PGSE) sequence that underlies DWI. Signal attenuation is parameterised by the b-value: a scalar that encodes gradient amplitude, duration, and temporal separation[1]; and the apparent diffusion coefficient (ADC): which is proportional to the rate of diffusion along the gradient direction. The Stejskal-Tanner relation characterises the intensity of diffusion weighted images:

$$S = S_0 \exp\left(-b \cdot \text{ADC}\right) \tag{3.1}$$

Where S_0 is the signal intensity acquired with no diffusion weighting and S is the signal intensity acquired with weighting applied. More diffusion along the applied gradient yields a higher ADC and therefore a lower intensity, S. The PGSE sequence is spin-echo based, hence S carries inherent T2 weighting as well as the diffusion related attenuation.

Basser et al. [48] later developed DTI by acquiring diffusion weighted images along at least six non-collinear gradient directions, enabling estimation of a symmetric 3×3 diffusion tensor, **D**, per voxel. **D** can be represented as a diffusion ellipsoid, where the eigenvectors of **D** are the directions of the ellipsoid major axes, and the corresponding eigenvalues, $\lambda_1 \geq \lambda_2 \geq \lambda_3$, give their magnitudes. The eigenvectors describe the direction of diffusion within each voxel, and the corresponding eigenvalue quantifies the rate of diffusion along each eigenvector direction. Several scalar metrics can be derived from the components of **D**, with two key metrics being mean diffusivity (MD) and fractional anisotropy (FA):

[1] The magnitude of b can be increased by using stronger gradients, longer duration, or by waiting for longer between application of the first and second gradients

$$\text{MD} = \frac{\text{tr}\left(\mathbf{D}\right)}{3} \tag{3.2}$$

$$\text{FA} = \sqrt{\frac{1}{2}\left(\frac{\left(\lambda_1 - \lambda_2\right)^2 + \left(\lambda_2 - \lambda_3\right)^2 + \left(\lambda_3 - \lambda_1\right)^2}{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}\right)} \tag{3.3}$$

MD describes the average rate of diffusion within a voxel, with lower values indicating more restricted diffusion. FA characterises how anisotropic the diffusion is, and ranges from 0 (isotropic) to 1 (single-axis restriction). High FA is indicative of anisotropic tissue structure largely restricting diffusion to a single axis. Mori et al. [49] showed that three-dimensional fibre tractography could be carried out by propagating streamlines along the principal axis of each ellipsoid in each voxel of a DTI volume.

DTI is most established in neuroimaging, where myelin sheaths surrounding axonal bundles restrict transverse diffusion, thus yielding diffusion ellipsoids with high FA values very suitable for tractography [49, 50]. Clinical examples of DTI include pre-operative white matter mapping in neurosurgery [51, 52], stroke assessment via ADC reduction in infarcted tissue [53], and electrode targeting in deep brain stimulation [52].

The anisotropic collagen fibre architectures of tendons, ligaments, menisci, and articular cartilage, make them natural candidates for DTI [54–56]. However, recent comprehensive reviews from Raya et al. [54] and Saran et al. [55] identify a consistent set of technical barriers to application of DTI in these tissues: short T2 relaxation times in organised collagen-rich tissues limit available signal; sub-millimetric resolution requirements conflict with the thick-slice acquisitions needed for adequate SNR; susceptibility-induced distortion is problematic near bone; and the magic angle effect introduces orientation-dependent signal variation. As a consequence, DTI studies of ligaments have focused predominantly on ACL grafts rather than the native ACL, since grafts are thicker, present fewer susceptibility complications at the bone tunnels, and undergo clinically relevant remodelling [57, 58]. DTI has shown promise in larger, more accessible structures such as the Achilles and patellar tendons, where it is reported to detect microstructural abnormalities beyond the sensitivity of conventional MRI [55, 59, 60].

Yang et al. [61] were the first to demonstrate the feasibility of in vivo DTI of the ACL and ACL grafts. They utilised 3T MRI with $1.96 \times 1.96 \times 1.96$ mm^3 isotropic voxels with

scan times of 6 minutes and 36 seconds to 11 minutes and 45 seconds. Tractography of the ACL and ACL graft were successful, although the attachment points of the native ACL could not be visualised due to susceptibility effects near the bone cortex. Following this preliminary study, Yang et al. [62] performed a feasibility study of DTI of the ACL graft after reconstruction. This study utilised the same isotropic resolution as the previous, with a scan time of 11 minutes and 36 seconds. Yang et al. [62] were able to show some initial tractography and arthroscopic agreement for an ACL graft case study. They also presented some initial results suggesting that FA was higher for the portion of the graft within the bone tunnels than in the intra-articular space, whilst the opposite trend was observed for MD. These initial studies demonstrated the feasibility of DTI in ACL graft structural assessment as well as its ability to potentially provide other useful metrics such as FA and MD.

Building on the work of Yang et al., Van Dyck et al. [58] also performed a feasibility study of DTI of the ACL graft at 3T. They noted that typically DTI of the ACL was challenging due to the short T2 / T2* of the tissues and the small size of the target anatomy; thus, ACL grafts were a more suitable target as the bone tunnels reduce susceptibility effects at the graft apertures, and the grafts had larger cross-sectional area compared to the native ACL. Their study utilised $1.5 \times 1.5 \times 6$ mm^3 slices with a scan time of 7 minutes 25 seconds. Van Dyck et al. were able to perform tractography and also measure FA and MD on the intra-articular portion of the graft. They hypothesised that as the graft healed that FA would increase as the tissue became more anisotropic, and that MD would decrease as the collagen fibres became more densely packed.

Van Dyck et al. [57, 63] tested this hypothesis in subsequent longitudinal studies where DTI was performed on ACL grafts following reconstruction and also on the ACL following primary repair with internal bracing. These studies used the same DTI sequence as the original feasibility study [58]. Although DTI was successful in both instances, FA was found to actually decrease over time, which Van Dyck et al. suggested meant that healing was incomplete in both studies. Limitations of these studies included the thick slices, necessary for sufficient SNR, and the low in plane resolution. Van Dyck et al. concluded that DTI and its associated metrics could serve as a useful biomarker for ACL and ACL graft healing; however, more DTI sequence development is needed for MSK applications to increase the resolution (both in- and through-plane) and reduce the TE.

Liu et al. [64] was the only study found that attempted to link grades of ACL tear with DTI metrics. They utilised $0.625 \times 0.625 \times 2.5$ mm^3 voxels with a scan time of 9 minutes and 32 seconds. They found a strong negative correlation between FA and the grade (1–4) of ACL tear, suggesting that higher grade tears disrupted the anisotropic fibre structure of the ligament to a larger extent, offering a potential quantitative alternative to arthroscopy for tear grading.

The magic angle effect also plays a role in musculoskeletal DTI. The magic angle effect is considered to be artefactual across many DTI studies [54, 55], however, due to the extended T2 it can provide, some studies have utilised it to boost SNR. Wang et al. [65] demonstrated at 9.4T that SNR and DTI metrics in rat knee ligaments were strongly dependent on fibre orientation relative to $\vec{B}_0$, with FA and MD varying by approximately 32% and 11% respectively across $\vec{B}_0$ orientations, whilst no such dependence was observed in rat brain tissue. Wang et al. also noted that FA is overestimated at low SNR. This motivates deliberate orientation of collagen-rich tissue at the magic angle relative to $\vec{B}_0$ to prolong T2 and boost SNR, thereby increasing DTI accuracy and reducing FA overestimation. Wengler et al. [59] and Wengler et al. [60] adopted precisely this approach, intentionally positioning the Achilles and patellar tendons at the magic angle relative to $\vec{B}_0$ to obtain sufficient signal for in vivo DTI at 3T, however, $1.25 \times 1.25 \times 6$ mm^3 voxels were still necessary to ensure enough SNR whilst restricting scan times to less than 8 minutes. Using this technique Wengler et al. were able to assess the collagen microstructure of the Achilles and patellar tendons, revealing tendinopathy related collagen disruption in the Achilles that was not visible on conventional MRI.

Several recent studies have employed high-resolution DTI and tractography to characterise complex collagen fibre architectures in excised musculoskeletal tissues; key acquisition parameters are summarised in Table 3.1. Wang et al. [66] characterised the complex fibre architecture in the rat knee joint using high resolution diffusion imaging. Shen et al. [67] performed tractography of the porcine meniscus, identifying zonal variation in fibre orientation across the tissue volume. Zhao et al. [68] compared DTI with generalised q-sampling imaging (GQI), a reconstruction method capable of resolving multiple fibre orientations within a single voxel, unlike the standard DTI model which assumes a single dominant direction. Zhao et al. demonstrated that GQI resolved crossing fibre populations that DTI was unable to distinguish; both Shen et al. and Zhao et al. identified radial and circumferential fibre networks consistent with the known meniscal

architecture. Whilst these studies establish that high-resolution fibre mapping is feasible in anisotropic collagen-rich tissues, the acquisition parameters in Table 3.1 make clear that none are yet translatable to a clinical setting: all were performed ex vivo on small animal or tissue samples at field strengths of 7-9.4T, with scan durations ranging from approximately 9 to 20 hours.

Table 3.1: Summary of field strength, isotropic resolution, scan time, and number of diffusion gradient encoding directions for three high resolution DTI studies of collagen fibre architecture.

Study	Field strength	Resolution (isotropic)	Scan time	Number of diffusion gradients
Wang et al. [66]	9.4 T	45 μm	15.4 hours	147
Shen et al. [67]	7 T	125 μm	20.1 hours	81
Zhao et al. [68]	7 T	500 μm	9 hours	61

In summary, DTI shows genuine promise for characterising collagen microstructure in musculoskeletal connective tissues, but faces consistent barriers to clinical translation. High-resolution fibre mapping has been demonstrated ex vivo [66–68], but with scan durations that are clinically intractable. In vivo studies have achieved acceptable acquisition times in tendons, ACL grafts, and to a lesser extent the native ACL [57–61, 63], though at the cost of coarse anisotropic voxel geometries and limited SNR. Scalar metrics such as FA and ADC show sensitivity to injury grade [64] and graft remodelling [57, 63], though their precise biophysical interpretation remains uncertain in some cases, as reported trends do not always accord with histological expectations [57, 63]. In vivo DTI of the native meniscus remains largely unreported, reflecting the compounding challenges of short T2, small tissue dimensions, and complex fibre architecture. The magic angle effect has been reported as both an artefact and a useful tool to extend T2 and improve DTI performance [59, 60, 65]. Finally, all reviewed studies were conducted at 3T or above, which constrains clinical accessibility given the existing demand on high-field systems, as discussed in Section 3.4.

Relative to DTI, MADI is positioned as both an alternative and a complementary technology for musculoskeletal collagen assessment. DTI encodes fibre orientation through the diffusion ellipsoid, whereas MADI extracts equivalent orientation information from magic-angle-dependent T2 contrast. However, MADI is able to achieve this in vivo at an isotropic $1 \times 1 \times 1$ mm^3 voxel size (Section 4.3.9). As demonstrated in Chapter 8,

MADI has produced clinically promising results in both the native ACL and the meniscus in vivo, the latter remaining largely inaccessible to DTI. Operating at low field further reduces susceptibility-induced distortion near bone interfaces, as the magnitude of susceptibility-driven field perturbations scales linearly with B_0, potentially improving imaging access to regions that are problematic at 3T. The primary current limitation of MADI relative to in vivo DTI is total acquisition time, which remains longer, though reduction of this is the subject of ongoing work discussed in Section 9.5.2. DTI retains advantages in its established scalar metrics: FA, MD, ADC, which have demonstrated sensitivity to injury grade and graft remodelling. The low-field rotatable platform may also provide a basis for future low-field DTI with deliberate magic angle positioning to boost SNR, though this remains speculative given the limited literature on low-field DWI, currently only in neuroimaging contexts [53]. Ultimately, the two techniques probe collagen microstructure through independent contrast mechanisms, making them well suited to cross-validation rather than direct competition.

3.4 Clinical Need for Dedicated MRI Systems

In addition to the technological limitations of conventional MRI outlined in Section 3.2, there are significant unmet practical and clinical needs associated with current conventional MRI systems and their deployment. In the UK, MRI systems face long waiting times and a shortage of imaging staff, contributing to delays in diagnosis and treatment. A dedicated MSK imaging system, capable of producing data that are simpler to interpret than that from conventional MRI, could alleviate the burden on high-demand general-purpose scanners, thereby increasing conventional scanner availability for acute imaging needs such as for stroke and oncology patients. Furthermore, such a system could reduce reliance on highly specialised radiologists for musculoskeletal image interpretation, streamlining clinical workflows and allowing expert resources to be focused on more complex diagnostic tasks.

There is currently a shortage of radiologists and imaging systems in the UK. Data from the British Medical Association [69] show that the UK has 8.57 MRI scanners per million people, which is significantly lower than the EU average of 12.43. In 2020 an independent review of diagnostic services for NHS England [70] recommended increases of 2000 and 4000 to the workforces of radiologists and radiographers respectively, and in 2023

the government pledged to double the number of CT and MRI scanners by 2029 [71]. In 2024, The Royal College of Radiologists [72] published their yearly census which noted that whilst there had been 4.7% growth of the radiology workforce, there was still a 29% workforce shortfall. They also noted that demand for MRI and CT imaging grew by 8%, and if current trends continued there would be a 39% workforce shortfall by 2029. These figures highlight a need for more radiologists and MRI systems, whilst also emphasising the pressure placed on current scanners and radiologists to acquire and interpret medical images.

The total number of MRI scans performed in England between April 2022 and March 2023 was 4,084,055 [73], which was a 6.2% increase from the previous year. Additionally, Table 3.2 shows data gathered from freedom of information requests for the breakdown of the type of MRI scans performed across various NHS Trusts in 2023 (2021 for Rotherham data). Table 3.2 collates the data into knee scans, other extremity scans, and the total number of MRI scans. These data show that dedicated MSK MRI has the potential to alleviate 6% of knee scans from conventional scanners, and approaching 10% of total scans if all extremity MRI scans are offloaded to systems such as the prototype scanner developed during this thesis.

There is also a rising incidence of ACL and meniscal surgeries, which further underscores the clinical need for improved and more widely available musculoskeletal imaging. Abram et al. [74] investigated the rates of ACL reconstruction and meniscal repairs in England between 1997 and 2017. They found ACL reconstruction increased 12-fold from 2.0/100K population to 24.2/100K, and meniscal repair increased by over double from 3.0/100K to 7.3/100K from 1996-1997 to 2016-2017. These increased rates of ACL and meniscal surgical treatment suggest an increased clinical need for appropriate MSK imaging to plan surgery and monitor recovery. Alternatively, better imaging could also lead to fewer surgeries, with more confidence in conservative treatment options due to more informative scans, thus reducing surgical workload.

In addition to addressing unmet imaging needs, dedicated, low-field rotatable MRI and the MADI technique may also alleviate interpretation burden on radiologists. Conventional MRI requires expert interpretation, and radiologist training is both time-intensive and resource-demanding. With increasing imaging demand and a shortage of radiologists, technology that simplifies image interpretation is of substantial value. MADI produces

intuitive, three-dimensional visualisations of collagen fibre microstructure as well as the raw image data, serving as a form of "virtual arthroscopy" as previously discussed. These data are more readily interpretable than conventional MR images, potentially enabling MSK image review by a wider range of healthcare professionals with less specialised training. This could streamline diagnostic workflows, increase throughput, and reduce the reporting workload on radiologists, whilst also supporting deployment of MRI systems in a broader range of clinical settings beyond hospitals.

Table 3.2: Data obtained via Freedom of Information requests from various NHS Trusts across the UK detailing the total number and modality of MRI examinations performed [75–82]. "Knee MRI" refers to any MRI scan specifically targeting the knee joint, whilst "Other extremity MRI" excludes knee MRI counts and includes MRI of the lower limb distal to the knee (e.g. ankle, foot) and upper limb including the elbow and any anatomy distal to the elbow (e.g. forearm, wrist, hand). All data are from the year 2023, except for Rotherham NHS Trust, for which the most recent available data were from 2021. All MRI scans were performed on 1.5T or 3T scanners.

NHS trust location	**Knee MRI**	**Other extremity MRI**	**Total MRI**	**% Knee MRI**	**% Other extremity MRI**
Doncaster and Bassetlaw Teaching Hospitals	2600	1081	25781	10%	4%
Dudley Group	2417	1311	25458	9%	5%
Maidstone and Tunbridge Wells	1913	1708	39705	5%	4%
Rotherham	572	450	8540	7%	5%
Royal Cornwall Hospitals	1716	817	31417	5%	3%
Royal Devon University Healthcare	1826	1405	33282	5%	4%
Tayside	1436	474	22549	6%	2%
University Hospitals of North Midlands	2599	2191	73427	4%	3%
Totals	**15079**	**9437**	**260159**	**6%**	**4%**

3.5 Low-Field MRI

The prototype rotatable MR system central to the work in this thesis has a field strength of 0.15T, classifying it as a low-field scanner. Early MRI systems operated at magnetic field strengths that would also be classified as low-field by modern standards. The first clinically useful scanner, operating at 0.04T, required approximately 30 minutes to acquire a single 2D image with a resolution of 7-10 mm [83]. Efforts to improve image resolution, reduce acquisition time, and increase SNR subsequently focused on increasing the static magnetic field strength, as intrinsic SNR was understood to scale approximately linearly with B_0 [84]. As such, 1.5T and 3T MR systems quickly became the clinical standard. However, higher field strengths are associated with significant trade-offs, including increased manufacturing and maintenance costs, greater power consumption, more stringent safety requirements, and the need for specialised installation infrastructure. While high-field systems are undoubtedly essential in modern clinical practice, recent advances in hardware and software have renewed interest in low-field MRI as a more accessible, cost-effective, and versatile alternative. Table 3.3 summarises typical field strength categories, their associated magnet types, and their key advantages and disadvantages.

Review studies by Marques et al. [21], Arnold et al. [22], and Sarracanie et al. [23] comprehensively detail the recent history, current state, and future promise of low-field MRI. These works consistently highlight the enhanced accessibility and safety of low-field systems compared to their high-field counterparts, primarily due to reduced operational and setup costs, lower power consumption, diminished projectile risk, and lower specific absorption rate (SAR). Furthermore, the less stringent siting requirements of low-field MRI facilitate system installation in diverse environments, including the development of dedicated point-of-care devices. The studies also note that open magnet designs commonly found in low-field systems also offer benefits for claustrophobic patients, whilst reduced acoustic noise from gradient switching improves patient comfort, particularly in paediatric applications. Acknowledged limitations of low-field MRI across these reviews include lower SNR, reduced spatial resolution, and extended acquisition times compared to high-field MRI. However, the resurgence of interest in low-field MRI is primarily driven by advancements in image reconstruction and post-processing techniques. Specifically, the integration of deep learning for acquisition acceleration, denoising, compensation

for magnetic field inhomogeneity, and advanced artefact reduction strategies is actively addressing many low-field limitations, thereby enhancing its clinical appeal.

Marques et al. [21] focused on low-field systems operating between 0.25T and 1T. They highlighted that advancements in hardware, such as gradient systems, transmit and receive coil designs, and permanent magnet production, have collectively enhanced the quality of images obtainable from low-field MRI systems over time. They also noted that the inherent cost-effectiveness of low-field MRI makes it economically viable to develop dedicated systems, such as a knee-only MR system, in contrast to high-field systems which must perform many types of scan protocol to justify their cost. Furthermore, the more compact size of low-field scanners facilitates their potential deployment as point-of-care devices, akin to ultrasound machines.

Sarracanie et al. [23] focused their review on low-field MRI systems operating in the 1 to 199 millitesla range. They also suggest that low-field MRI could provide dedicated or point-of-care devices, that would alleviate the current bottlenecks in clinical workflows due to the limited availability of high-field systems. By deploying dedicated scanners tailored to specific departmental needs, the efficiency of imaging services could be significantly improved, reducing patient waiting times for MRI examinations. Similarly to Marques et al., Sarracanie et al. note that high-field MRI is an all-rounder imaging modality, and that dedicated low-field systems could fulfil similar roles to systems from other modalities such as X-ray, CT, and ultrasound, which all feature dedicated machines for various clinical applications.

Regarding image quality and resolution, Sarracanie et al. argue that these factors are not simply a matter of field strength. For example, they note that the image quality from 1980s 1.5T systems was considerably inferior to that achievable with modern 1.5T scanners, despite the exact same main field strength. This evolution underscores that performance is driven not merely by field strength, but by collective advancements in hardware (such as gradient systems, coil designs, permanent magnet technology), as well as software (such as image reconstruction, denoising, artefact reduction), and acquisition techniques (such as compressed sensing or novel k-space trajectories). Low-field MRI is seeing a similar increase in image quality due to translation of technology from high- to low-field, with recent developments in deep learning reconstruction and image processing improving MRI across all field strengths.

Arnold et al. [22] reviewed low-field MRI across various applications, including neuroimaging, MRI-guided procedures, paediatric imaging, and MSK imaging, at field strengths ≤ 1.5T. They highlighted that the cost of high-field MRI, at approximately 1 million USD/Tesla, contributes significantly to its inaccessibility, noting that this cost has also increased over time, contrary to typical technological trends. On the issue of lower SNR at low-field, Arnold et al. clarified that the primary concern is SNR per unit time rather than an absolute SNR limit, noting that increasing the number of signal averages can improve SNR at the cost of total acquisition time. They further emphasised that recent advancements in both hardware and software have boosted low-field SNR whilst reducing acquisition times. They also reiterate the findings of Marques et al. and Sarracanie et al., stating that reduced cost of low-field systems also facilitates the development of tailored scanner designs, such as those dedicated to specific body parts. They further argue that the cost-effectiveness and design flexibility of low-field MRI make it particularly well suited to future MSK applications, including the development of specialised scanners for weight-bearing examinations. They also note that low-field systems exhibit fewer metal or implant artefacts compared to high-field MRI, which is a significant advantage in MSK imaging where implants are common. However, Arnold et al. also highlight the practical limitations of low-field MRI in clinical settings, including inherent constraints on SNR and spatial resolution. They cite a meta-analysis of knee MRI studies demonstrating a positive correlation between field strength and diagnostic performance for ACL tears. Notably, these findings were based on conventional imaging protocols, the limitations of which MADI is specifically designed to overcome.

The prototype rotatable scanner developed for MADI operates at low-field strength, primarily due to the mechanical and engineering constraints that make the construction of a rotatable high-field superconducting system impractical. As highlighted in the reviewed literature, low-field MRI facilitates the development of dedicated imaging systems by reducing cost and technical complexity, thereby improving economic feasibility for targeted clinical applications. The compact footprint of the prototype scanner also supports its use as a point-of-care device, enabling deployment in non-hospital or mobile environments for the imaging of acute MSK pathologies. Importantly, the magic angle effect is independent of field strength and occurs across the entire field strength spectrum. This suggests that larger ultra-low-field systems could be developed for increased field-of-view (FOV) imaging, or, subject to future engineering advances,

even high-field rotatable implementations may become feasible. Overall, the resurgence of interest in low-field MRI, combined with progressing advanced reconstruction and acquisition acceleration techniques, provides a strong technical foundation for the future development of low-field imaging and the MADI technique introduced in this thesis.

Table 3.3: Summary of typical MRI field strength categories, associated magnet technologies, and the key advantages and limitations of each.

Field strength	**Magnet type**	**Advantages**	**Limitations**
Ultra low-field ($\leq$0.1T)	Permanent magnet array or resistive electromagnet	Extremely low cost, low power requirements, simple hardware, safer, reduced susceptibility artefacts and metallic implant artefacts	Lower signal-to-noise ratio (SNR), limited resolution, increased scan time, requires advanced reconstruction for clinically relevant data (e.g. AI)
Low-field (0.1-0.5T)	Permanent magnet array or resistive electromagnet	Low cost, open designs, portable, low power consumption, safer, reduced susceptibility artefacts and metallic implant artefacts	Lower signal-to-noise ratio (SNR), limited resolution, increased scan time
Mid-field (0.55-0.6T)	Low-helium superconducting electromagnet	Improved SNR over low-field, open-bore and reduced-footprint designs, lower cryogen consumption, reduced susceptibility artefacts relative to high-field, compatible with many metallic implants	Lower SNR than 1.5T, fewer available pulse sequences and hardware options than clinical standard systems, higher cost than permanent magnet systems
High-field 1.5 and 3T	Superconducting electromagnet	Clinical standard, high SNR, excellent soft tissue contrast, wide availability of pulse sequences and associated hardware	Expensive, requires cryogenics, more pronounced susceptibility artefacts and metallic implant artefacts
Ultra high-field $\geq$7T	Superconducting electromagnet	Very high SNR, superior resolution, advanced high field imaging techniques	Very high cost, significant safety and artefact concerns, regulatory limitations

3.6 Summary

This literature review has outlined the key motivations for the development of the prototype low-field rotatable MRI scanner and the associated MADI technique.

Conventional MRI faces fundamental limitations in imaging connective tissues such as tendons, ligaments, and menisci, due to their intrinsically low signal and susceptibility to the magic angle effect. Even when image quality is sufficient, interpretation remains highly specialised, often requiring experienced musculoskeletal radiologists. These challenges, combined with limited scanner availability and the prioritisation of more acute imaging needs, mean that MRI is often insufficient for confident assessment of MSK pathology, with arthroscopy remaining the diagnostic gold standard. There is thus a clear need for a non-invasive imaging approach capable of resolving collagen fibre microstructure whilst simplifying interpretation. A dedicated MSK system, such as the prototype rotatable scanner, could alleviate pressure on general-purpose MRI systems, reduce reliance on radiologists, and streamline clinical workflows.

MRI remains the modality of choice for MSK imaging, but faces both technical and systemic constraints that limit its effectiveness and accessibility. Low-field MRI offers a practical solution. Its lower cost, smaller footprint, and reduced siting requirements enable the development of dedicated systems that are more accessible and adaptable than conventional scanners. Recent advances in acquisition acceleration and deep learning reconstruction are also mitigating historical limitations such as low SNR, low resolution, and long scan times. Although these techniques were not the focus of this thesis, they represent promising avenues for future development of MADI.

DTI represents the most established alternative MRI approach for probing collagen fibre microstructure in musculoskeletal tissues, and is therefore the most direct point of comparison for MADI. Whilst DTI has demonstrated sensitivity to fibre organisation and pathology in tendons and ACL grafts in vivo, clinical translation remains constrained by the short T2 of collagen-rich tissues, resolution demands that conflict with SNR limitations, and the confounding influence of the magic angle effect. MADI addresses several of these limitations simultaneously: in utilising a dedicated low-field rotatable scanner, it reduces susceptibility-induced distortion, provides systematic control over the magic angle effect, and achieves isotropic $1 \times 1 \times 1$ mm^3 resolution in vivo in

tissues that remain largely unreported in MSK DTI literature, such as the meniscus. Overall the two techniques should be viewed as complementary rather than competing, probing collagen microstructure through independent contrast mechanisms and offering the prospect of mutual cross-validation.

MADI directly addresses the imaging challenges identified here by exploiting the magic angle effect as a source of quantitative information on the structure of organised, anisotropic collagen-rich tissues. Utilising the rotatable low-field system, MADI enables controlled reorientation of $\vec{B}_0$ without compromising patient comfort. MADI thus enables non-invasive visualisation of tissue microstructure, offering the potential to provide novel clinically relevant information and reduce the burden on conventional MRI systems.

4 | Magic Angle Directional Imaging

4.1 Chapter Overview

This chapter provides a high-level summary of the MADI methodology in Section 4.2, followed by an overview of the rotatable low-field MR system, including its supporting hardware and software components, in Section 4.3. Finally, Section 4.4 details the preprocessing steps required by the MADI technique as a consequence of the acquisition methodology and the rotatable scanner's hardware design.

The low-field rotatable scanner is a dedicated extremity system, capable of imaging anatomical regions with dimensions up to, and including, those of the knee. It is also currently the only device capable of performing the MADI technique, owing to its unique ability to reorient $\vec{B}_0$ relative to a static patient. The MADI methodology relies on this capability, and cannot be performed in vivo without such a system.

4.2 Magic Angle Directional Imaging Methodology

The MADI methodology encompasses both data acquisition using the rotatable low-field scanner and the subsequent data processing pipeline, developed using MATLAB (Natick, Massachusetts, USA), required to generate its unique outputs. An overview of the full process is presented in the flowchart in Figure 4.1.

Several steps of the MADI technique required substantial investigation and development to satisfy experimental design constraints and to increase the levels of accuracy and usability. Section 4.2.1 provides an ordered, high-level overview of each step in the MADI methodology, including key implementation details sufficient to understand the overall technique. Each entry includes a cross-reference to the relevant sections and chapters where the precise implementation and the associated development details are discussed in greater depth.

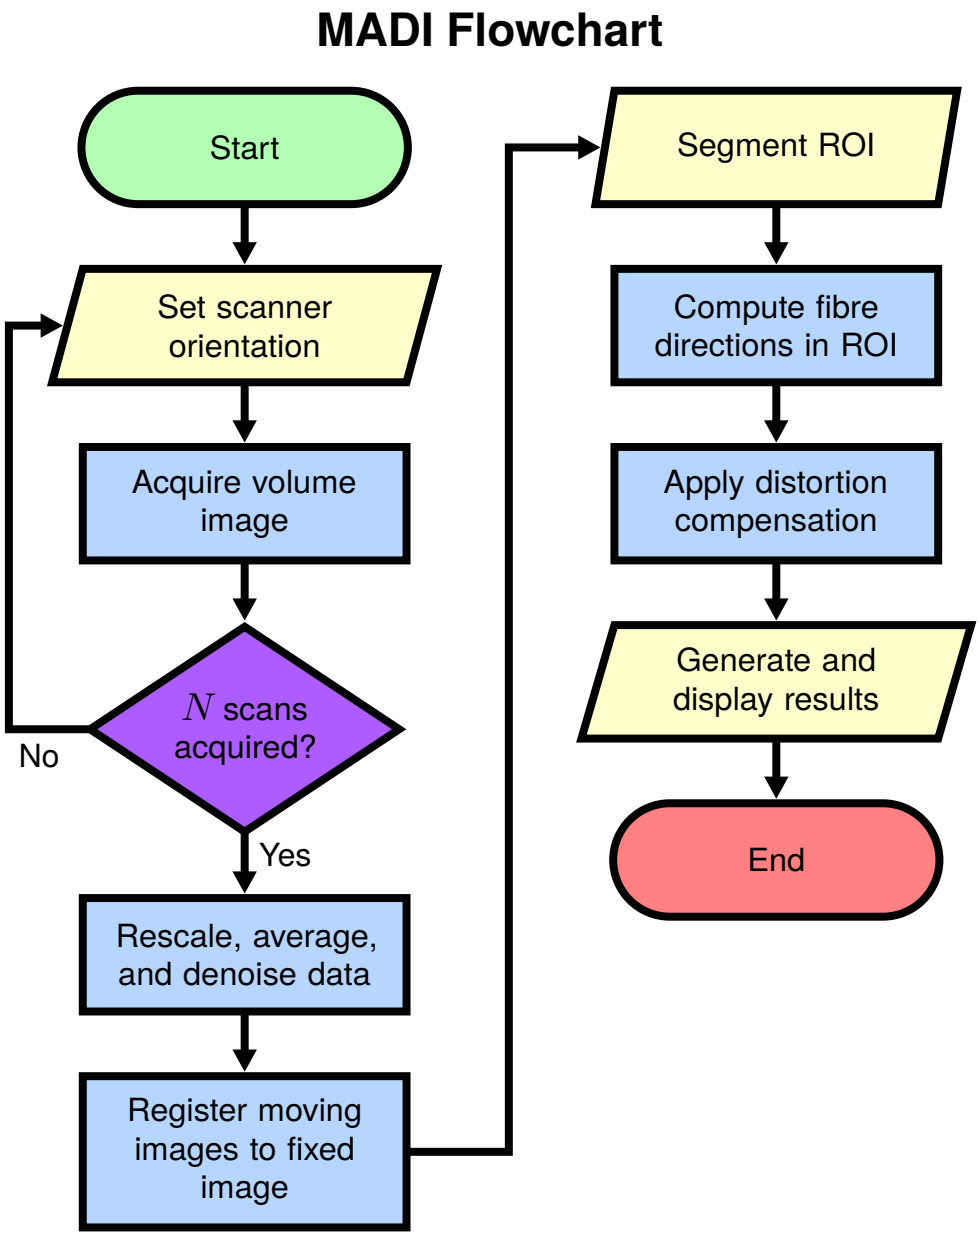


Figure 4.1: Flowchart of the MADI technique. The process begins with data acquisition, in which N scans are acquired at different scanner orientations; N denotes the number of orientations selected for a given examination. A typical MADI examination performed during this work had $N = 7$. Following acquisition, the data undergo preprocessing, and each moving image (acquired at rotated scanner orientations) is registered to the fixed image (the first scan, acquired with the scanner in its home orientation). The ROI is then segmented for analysis. Fibre orientations within the ROI are subsequently estimated, and the results are saved and visualised.

4.2.1 MADI Methodology Detailed Overview

1. **Set scanner orientation:** As described in Section 4.3.3, the orientation of $\vec{B}_0$ is uniquely characterised by the roll and yaw angles, θ_R and ϕ_Y, of the scanner body and yoke respectively. The scanner orientation was set by user input of $\theta_R \in [-90^\circ, 90^\circ]$ and $\phi_Y \in [-30^\circ, 30^\circ]$. The first scan was always acquired at $\theta_R = 0^\circ,\ \phi_Y = 0^\circ$ to provide a view of the anatomy in a conventional orientation and to verify appropriate patient positioning.

2. **Acquire data:** N volume images were acquired, one at each scanner orientation, using the parameters and pulse sequence as described in Section 4.3.9. Tissue structures exhibiting the magic angle effect displayed intensity variations across these N scans due to changes in the orientation of $\vec{B}_0$ relative to their collagen fibres. $N = 7$ for a typical MADI examination performed during this work. The selection of N and the corresponding scanner orientations are detailed in Chapter 7.

3. **Rescaling, averaging, and denoising:** These preprocessing steps ensured consistent and high-quality data input to the core processing stages of the MADI technique. Intensity rescaling was performed to recover raw signal values from the recorded data, signal averaging across multiple excitations was used to improve SNR, and denoising was applied to enhance the robustness of subsequent segmentation and analysis. A detailed description of these steps is provided in Section 4.4.

4. **Registration:** The first scan, acquired with the scanner in its home orientation ($\theta_R = 0^\circ$, $\phi_Y = 0^\circ$), was the fixed image. The remaining $N - 1$ scans, acquired at rotated orientations, were the moving images. Rotation of the scanner caused complementary rotation in the moving images. To enable segmentation and analysis of the acquired data, corresponding anatomical features in the fixed and moving images were spatially aligned. This was achieved using a multi-stage registration pipeline that employed both rigid and non-rigid (soft) registration to align features whilst correcting for inter-scan geometric distortions. Complete details of the registration process are provided in Chapter 5.

5. **Segmentation:** Following registration, semi-automatic segmentation was performed with a metric-assisted approach based on signal variation across the set of N registered images. For each voxel coordinate in the fixed image, a set of N intensity values was obtained; one from the fixed image, and $N-1$ from the corresponding coordinates in the registered images. The standard deviation of these intensities was computed to produce a standard deviation volume image, which was used to perform threshold-based segmentation. Tissues exhibiting the magic angle effect displayed high intensity variability across orientations, whilst tissues such as bone exhibited consistently low signal, resulting in low standard deviation. This metric was integrated into a graphical user interface (GUI), which enabled semi-automated segmentation of the ROI using a combination of standard deviation thresholding and user-defined bounding boxes. Further details are provided in Chapter 6.

6. **Collagen fibre orientation estimation:** For each voxel coordinate within the segmented ROI, the measured intensity profile (N intensities, one from each image) was compared to a dictionary of simulated intensity profiles from a dense ball of candidate fibre orientations. A pairwise correlation was computed between the measured and simulated intensity profiles, and the candidate orientation yielding the highest correlation was selected as the estimated fibre orientation at that voxel. A complete overview of this algorithm is provided in Chapter 7.

7. **Distortion compensation:** A calibration phantom with a regular internal geometry was designed, 3D printed, and filled with vegetable oil to provide uniform MR signal. The phantom was scanned on the low-field system and a simulated scan was performed to generate a distortion-free ground truth image. The image of the printed phantom was non-rigidly registered to the simulated scan of the CAD model. This enabled characterisation of geometric distortions arising from the B_0 inhomogeneity of the system. The non-rigid registration parameters were saved and subsequently applied to all subsequent images to correct for these systematic deformations. Further details of this correction method are provided in Chapter 5.

8. **Visualisation:** A principal output of the MADI technique was an estimation of collagen fibre orientations at each voxel coordinate in the ROI. This was represented as a vector field in which each fibre orientation was displayed at its correspond-

ing spatial coordinate. These orientations were also visualised using glyphs and tractograms, similarly to those used in DTI. To provide anatomical context, these visualisations were overlaid on the underlying MR images or displayed in context with bone segmentations. Macroscopic data visualisations were also produced, including an alignment index heatmap that visualises the dominant orientation patterns within the ROI. Additional outputs derived from the set of N registered acquisitions include volume images of standard deviation, and minimum, maximum, and mean intensity projections. These visualisations are described and presented across Chapters 6 and 8.

4.3 Rotatable Low-Field System Overview

This section provides an overview of the core components of the rotatable low-field system. A full technical breakdown of all components is beyond the scope of this thesis, however, sufficient detail is provided to give a high-level understanding of the system for readers familiar with MRI technology. Greater detail is provided for components and definitions unique to low-field rotatable MRI and the MADI technique, as these are essential for understanding the capabilities of the system and for interpreting the methodology and results presented in this thesis.

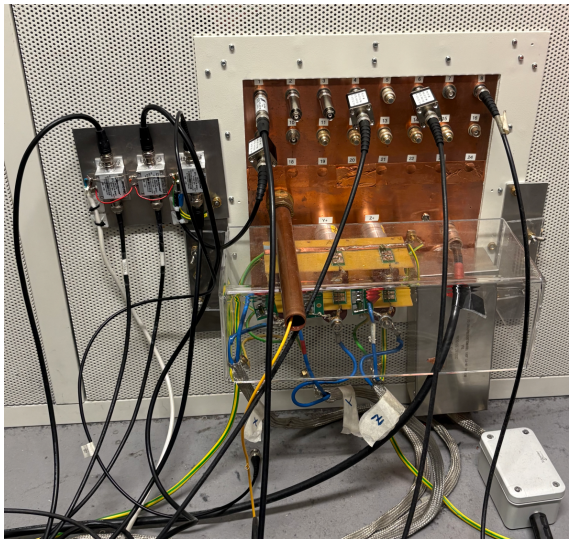
4.3.1 Installation

Figure 4.2 shows the installation of the system in a basement laboratory at the South Kensington campus of Imperial College London. Figure 4.2a shows the scanner and patient chair, located within a purpose-built RF shielded room designed to minimise electromagnetic interference from external sources. Figure 4.2b shows the patch panel that all filtered signal and power connections pass through.

Figure 4.3 shows the operator station and system power supply adjacent to the shielded room. Figure 4.3a shows the operator station where scan acquisition and motion control are managed via dedicated computers for each task respectively. Figure 4.3b shows the supercapacitor bank that powers the system during image acquisition. Further details of the operator station software and power supply hardware are provided in Section 4.3.2.

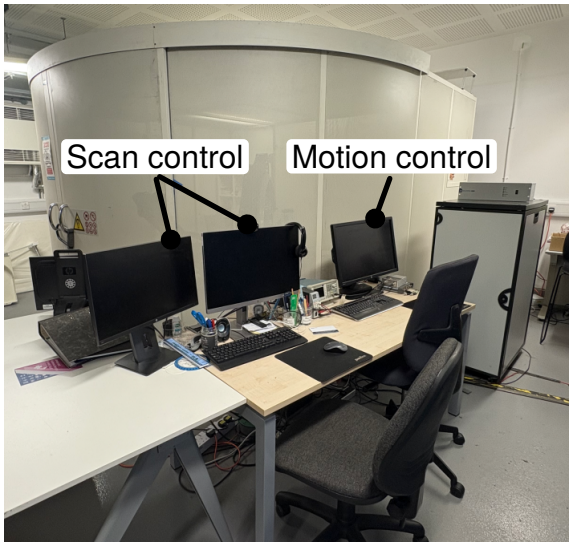


(a) Wide view of RF shielded room with sliding doors open. The patient chair is shown to the left of the low-field rotatable system in the centre of the shielded room.



(b) Patch panel with filtered connections between scanner and operating computers. Mains power and gradient cables are covered by acrylic protective box for safety.

Figure 4.2: Installation of low-field rotatable system in RF shielded room. Figure 4.2a shows the low-field rotatable scanner situated within the RF shielded room at Imperial College London South Kensington campus. Figure 4.2b shows the filtered signal and power cables passing through a patch panel in the rear right corner of the shielded room.



(a) Operator station adjacent to RF shielded room. Two dedicated computers control scan acquisition and scanner motion respectively.



(b) Supercapacitor bank that powers the gradient and transmit coils during image acquisition.

Figure 4.3: Operator station and power supply for low-field system. Figure 4.3a shows the operator station and Figure 4.3b shows the supercapacitor power supply adjacent to the shielded room containing the low-field rotatable MR system.

The components shown in Figures 4.2 and 4.3 are those that dictate the spatial requirements for installing the low-field rotatable MRI system. The system has a compact footprint, requiring only a single room for all components, and it requires no dedicated power supply infrastructure; Ristic et al. [85] designed the custom supercapacitor bank in Figure 4.3b to be charged from standard mains electrical outlets. The system can therefore be sited in any sufficiently large room with access to standard sockets, and without structural reinforcement of the floor as the scanner itself only weighs approximately 400 kg. Furthermore, the installation in Figures 4.2 and 4.3 is operated in a research environment which demands flexibility and rapid iteration; in future, a mature iteration of the system could easily be miniaturised to enable smaller-scale installations or portable deployment.

4.3.2 System Overview and Block Diagram

Figure 4.4 shows a close-up view of the central section of the low-field rotatable scanner. The receive (Rx) coil sits within the saddle transmit (Tx) coil at the centre of the scanner bore. The coils are centred between the two magnetic pole faces, which are enclosed by PVC covers. These poles are supported by the yoke assembly, which is connected to the main body via a pivot joint that allows the yoke to yaw within the main body.

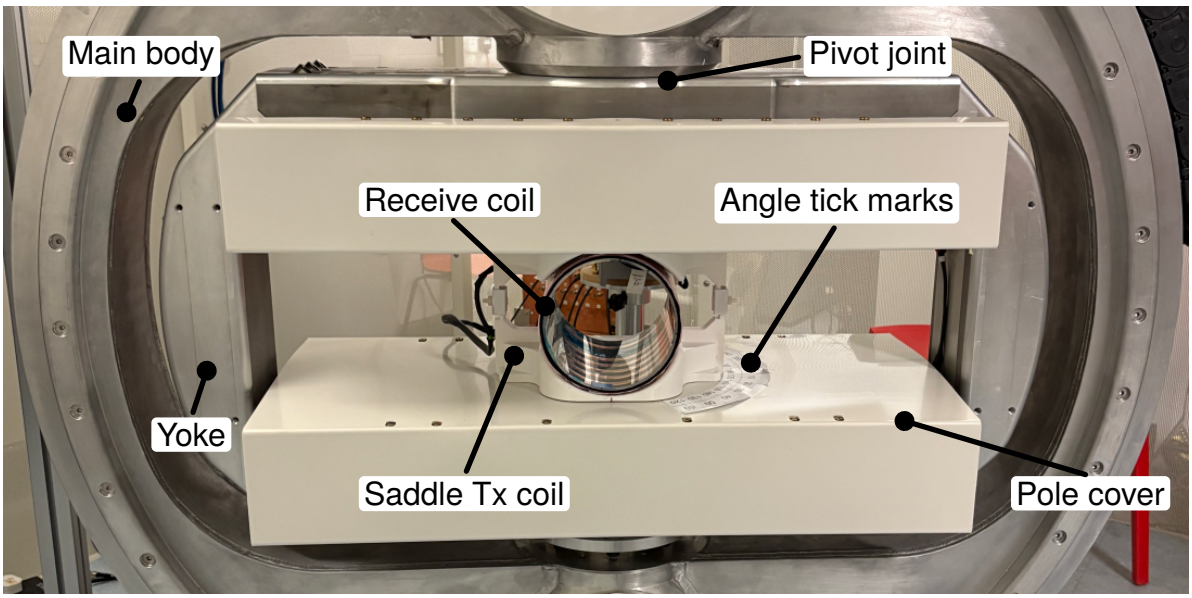


Figure 4.4: Close-up view of central section of low-field rotatable scanner. The receive coil is shown in the centre of the scanner bore, surrounded by the saddle transmit coil. The transmit and receive coils are situated between the two pole faces of the bore, which are enclosed in a PVC covering. The poles are held within the scanner yoke, which is attached to the main body by means of a pivot joint.

The transmit coil housing has protruding pins that seat into holes in each pole cover via a friction fit, constraining translational motion of the housing whilst allowing rotation about the yaw axis of the yoke. This enables the Tx/Rx coils to counter-rotate relative to the yaw of the yoke, ensuring they remain stationary[1] with respect to the subject during yaw rotation. The transmit coil housing forms a friction fit with the pole covers, holding its position once rotated. Proper alignment is verified by angular tick marks on the pole covers. The left side of Figure 4.5 shows a cross-sectional schematic of the magnet poles and the pole-mounted components, including the gradient, transmit, and receive coils, as well as the saddle coil housing and its pin interface with the pole covers.

Figure 4.5 presents a block diagram of the control architecture for the low-field rotatable MR system. The scan control computer interfaces with the spectrometer, which controls the RF and gradient amplifiers. These amplifiers draw power from the supercapacitor bank shown in Figure 4.3b. During RF transmission, the spectrometer detunes the receive coil via a PIN diode circuit to prevent signal contamination or damage. In receive mode, the MR signal from the receive coil and preamplifiers is digitised by the spectrometer and then processed and stored by the scan control computer. The dedicated motion control computer manages the scanner's orientation by interfacing with the roll and yaw motors, as described in Section 4.3.7.

The spectrometer used in this work is a commercially available EVO unit from MR Solutions (Guildford, UK), which provides a wide range of standard MR functionality and pulse sequences via its associated Clinical Scan software, running on Windows 10.

The pole-mounted gradient coils provide a maximum gradient strength of 20 mT/m and are driven by standard gradient amplifiers from Performance Controls Inc. (Montgomeryville, Pennsylvania, USA).

The RF system is divided into separate transmit and receive chains. Two transmit configurations are available: a pole-mounted Helmholtz pair and a saddle coil housed within a custom 3D-printed support, as shown in Figure 4.4. The housing also serves to guide and centre both the receive coil and the subject within the scanner bore. Throughout this thesis, the saddle transmit coil was used due to its superior performance compared with the pole-mounted transmit coil. Both transmit coil configurations produce linearly polarised B_1 fields. This simplifies operation as $\vec{B}_1$ is always perpendicular to

[1] An operator manually counter rotates the transmit coil housing during scanner yaw motion.

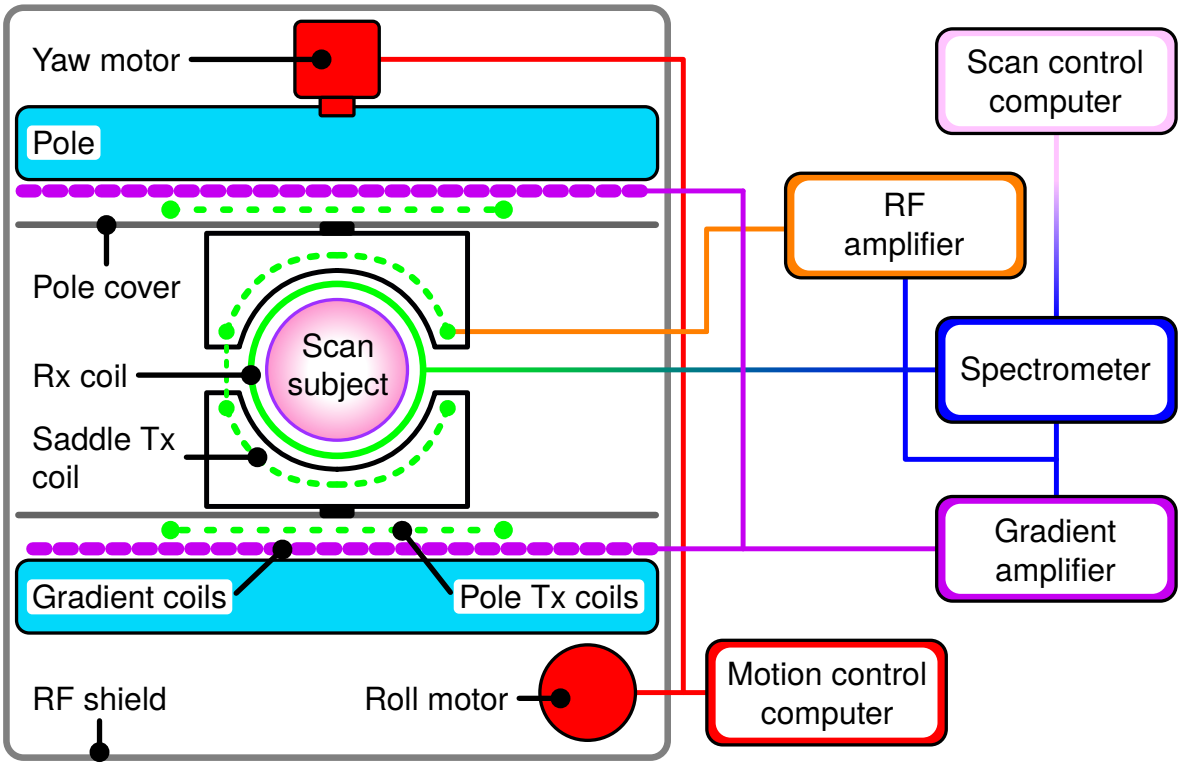


Figure 4.5: Block diagram of low-field MR system. The scan control computer communicates with the spectrometer, which in turn sends control signals to the gradient and RF amplifiers to control the gradient coils and transmit coils. The spectrometer also sends control signals to detune the receive coil during transmit using a PIN diode tuning circuit. The dedicated motion control computer communicates with the roll and yaw motors to orient the scanner. Scan control two-way communication is indicated by dual colour connections. Motion control communication is inherently bidirectional as the motors are monitored using absolute encoders.

$\vec{B}_0$ regardless of any applied yaw rotation. Greater B_1 efficiency could be achieved by employing circular polarisation; however, maintaining the perpendicular orientation between $\vec{B}_1$ and $\vec{B}_0$ under yaw rotation would then require the use of the pole-mounted transmit configuration that would yaw with $\vec{B}_0$, rather than the saddle transmit coil that remains static relative to the patient.

The receive coil is a solenoid comprising five inductively coupled resonant loops, tuned such that the fundamental resonant mode matches the operating frequency of the 0.15T B_0 field of 6.4 MHz. A sixth inductively coupled loop provides the output link [86], transferring the MR signal to the preamplifier and subsequently to the spectrometer. The solenoid design presents limitations arising from signal attenuation when yaw rotation is applied, as $\vec{B}_0$ is no longer perpendicular to the coil axis under yaw rotation; further details of this issue are provided in Section 4.4.

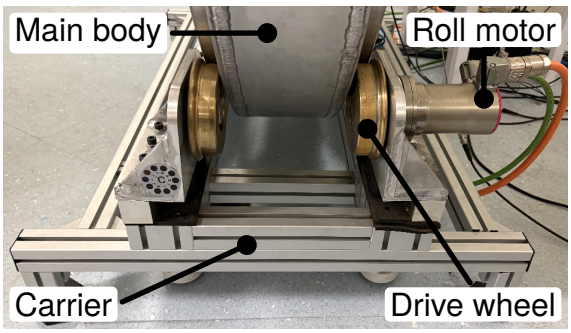
All signal and power cables enter and exit the shielded room via filtered connections in the patch panel shown in Figure 4.2b, with the exception of the motion control signals for the roll and yaw motors. These use a fibre-optic CANBus implementation, which does not require filtering but does necessitate a copper waveguide, as shown in Figure 4.2b, to maintain proper RF shielding.

4.3.3 Motion Platform

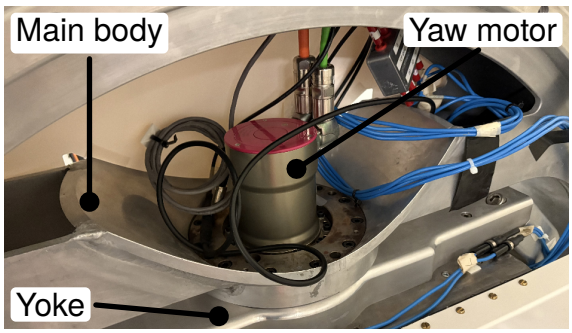
Each pole is mounted to the aluminium yoke, as shown in Figure 4.4, which can yaw within the main body. The main body, in turn, can roll on its carriage. These two degrees of freedom (DOF) allow reorientation of $\vec{B}_0$ about the scanner's isocentre by rotation.

Figure 4.6 shows the roll and yaw motors used to reorient the low-field scanner. The roll motor drives a smooth wheel in direct contact with the scanner main body, which also rests on three additional identical non-driven wheels. Scanner rotation in roll is achieved simply through friction between the drive wheel and the main body. The yaw motor is directly coupled to the yoke via a pivot joint between the yoke and the scanner main body.

The motors are DC strain-wave geared units from Harmonic Drive SE (Limburg an der Lahn, Germany), offering near-zero backlash, high precision, high torque, and low-speed motion. Position is monitored by inbuilt high-resolution absolute encoders.



(a) Close-up view of roll motor and drive wheel interface with scanner main body.



(b) Rear view of top section of scanner main body showing close up of yaw motor interface with yoke.

Figure 4.6: Close up of roll and yaw motors of low-field system. Figure 4.6a shows a close-up of the roll motor and drive wheel that rolls the scanner main body, and Figure 4.6b shows a close-up of the yaw motor and its attachment to the scanner yoke. The roll motor is directly connected to a smooth drive wheel on which the main body rests, and the yaw motor is directly connected to the yoke.

Communication with the motion control computer is via fibre-optic cable through the RF shielded room's patch panel; the signal is converted back to its original electrical form (CANBus protocol) for communication with the motor drivers once inside the shielded room. Motor operation is managed through a vendor-supplied state machine in the motor drivers, which is controlled by the custom motion control software described in Section 4.3.7.

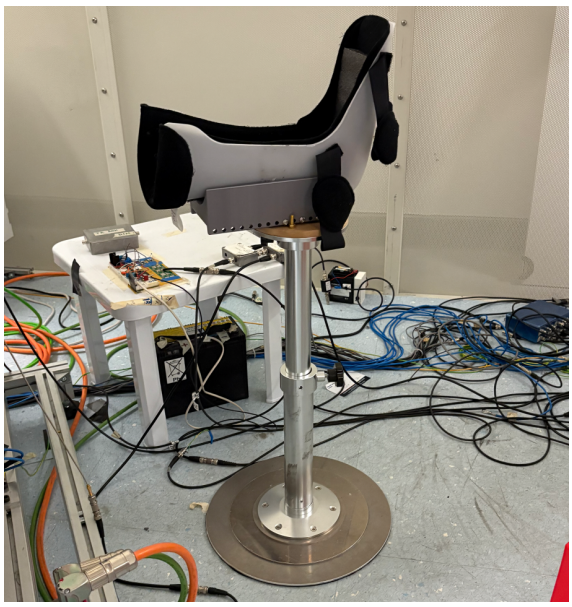
4.3.4 Patient positioning

Figure 4.7 shows a patient volunteer positioned for a scan of their right knee and the immobilising boot used to support their leg on the other side of the scanner.

To position the patient within the scanner as shown in Figure 4.7a, they are first seated in the adjustable patient chair away from the scanner. With their leg extended, the chair is rolled forward by an operator until the leg passes through the transmit and receive coils and into the immobilising boot shown in Figure 4.7b. The chair height and inclination are then adjusted for comfort and precise positioning. The non-scanned leg is free to rest on the scanner carriage or a support stool. The scanner rolls and yaws around the patient, who remains in this position for the duration of the examination. The open design of the scanner allows the patient to read or watch downloaded videos on a tablet and use headphones during the scan.



(a) Patient volunteer seated in the adjustable patient chair.



(b) Boot used to support and immobilise the lower leg of the patient.

Figure 4.7: Patient positioning within the low-field rotatable scanner. Figure 4.7a shows a patient volunteer with their right knee within the scanner bore and their lower leg extending through to the other side of the scanner. Their non-scanned leg is free to rest outside of the scanner. Figure 4.7b shows the boot used to support and immobilise the patient's lower leg. The foot and ankle are held in place with Velcro straps, and the heavy metal base of the support structure prevents movement during scan acquisition.

4.3.5 Magnet Design

The magnetic field strength of the low-field scanner is 0.15T, corresponding to a Larmor frequency of $0.15\text{T} \times 42.58$ MHz/T ≈ 6.4 MHz, which is the operating frequency of the scanner's RF components. The B_0 field is generated by an array of permanent magnets, as shown in Figure 4.8. This configuration was designed by McGinley et al. [16] using iterative finite element simulations to produce $\vec{B}_0$ transverse to the magnet poles with maximum diameter of the spherical volume (DSV) over which the B_0 field is uniform. In this thesis, the term pole refers to a complete permanent magnet array unit, with two such units forming the top and bottom poles, which are mounted within the yoke as shown in Figure 4.4.

Figure 4.8 shows a computer-aided design (CAD) model of the magnet poles. Each pole measures 650mm by 350mm, with a vertical separation of 220mm. The backplates consist of low-carbon steel to provide magnetic flux return, and support an array of Neodymium-Iron-Boron (Nd-Fe-B) permanent magnets. The magnets are arranged such that opposing polarisations face each other on the opposite pole assemblies,

generating a transverse magnetic field parallel to the pole faces and aligned with the major axis of each pole [16]. The different colours of the permanent magnet subunits in the figure indicate variations in height, polarity, and magnet size.

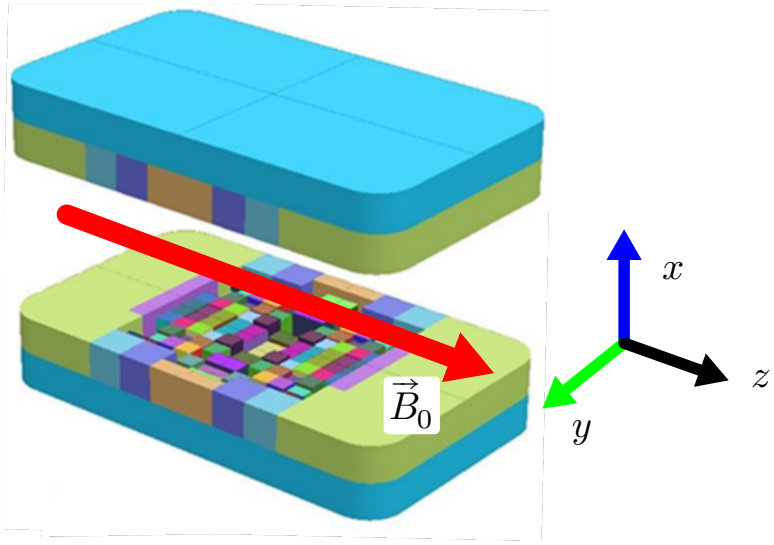


Figure 4.8: CAD model of low-field magnet design composed of an array of permanent magnet subunits. The $\vec{B}_0$ orientation is shown by the red arrow as parallel to the major axis of the pole faces. The z-axis is defined as parallel to $\vec{B}_0$ by convention.

In practice, the usable DSV (where the field strength is 0.15T$\pm$10μT) is reduced to approximately 120-130 mm due to manufacturing tolerances and the passive shimming required for a permanent magnet system [16]. The effective pole gap is similarly limited to 200 mm by the presence of the pole-mounted transmit and gradient coils, and the pole covers. Whilst these dimensions were sufficient for all data acquisition in this thesis, future iterations of the low-field system would benefit from a larger pole gap and DSV to improve ease of use and accommodate a greater range of patient sizes.

4.3.6 Coordinate System

A global Cartesian coordinate system was defined such that any $\vec{B}_0$ orientation can be described by the applied roll and yaw angles. By convention, the z-axis is parallel to $\vec{B}_0$ in the scanner's home orientation[2]. The x-axis is defined as vertical and perpendicular to the z-axis, and the y-axis is perpendicular to both x and z, satisfying $\hat{x} \times \hat{y} = \hat{z}$, where $\hat{x},\ \hat{y},\ \hat{z}$ are the unit vectors along each axis. This definition preserves conventional right-handed geometric identities, ensuring that standard rotation matrices are applicable.

[2]The global coordinate system is static and does not rotate with the scanner.

Figure 4.9 shows the axes of the global coordinate system, the positive rotation directions about the x- and y-axes, and $\vec{B}_0$ in the home orientation of the scanner. The origin of this coordinate system is defined as the isocentre of the scanner. A pure yaw rotation of the scanner corresponds to a rotation of $\vec{B}_0$ about the x-axis, and a pure roll rotation corresponds to a rotation of $\vec{B}_0$ about the y-axis.

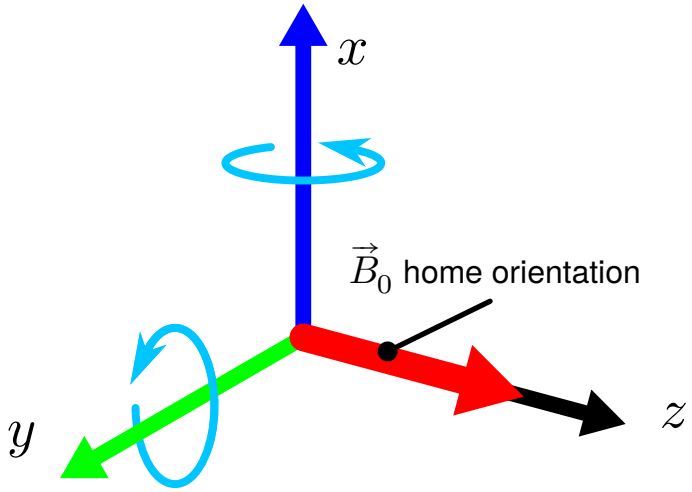


Figure 4.9: Global coordinate system of the rotatable low-field scanner, and the $\vec{B}_0$ orientation when the scanner is in its home configuration. The positive rotation directions (following the right-hand rule) about the x- and y-axes are shown in light blue.

The $\vec{B}_0$ orientation for an applied roll and yaw of θ_R and ϕ_Y degrees respectively is denoted $\hat{B}_0(\theta_R, \phi_Y)$, where $\hat{B}_0 = \vec{B}_0/||\vec{B}_0||_2$. By convention:

$$\hat{B}_0(0,0) = \hat{z} = [0,\ 0,\ 1]^T$$

Any $\hat{B}_0$ orientation is then defined in terms of this initial condition, as described by Equation 4.1.

$$
\begin{aligned}
\hat{B}_0(\theta_R, \phi_Y) &= \mathbf{R}_y(\theta_R)\,\mathbf{R}_x(\phi_Y)\,\hat{z} \\
&= \begin{bmatrix} \cos\theta_R & 0 & \sin\theta_R \\ 0 & 1 & 0 \\ -\sin\theta_R & 0 & \cos\theta_R \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\phi_Y & -\sin\phi_Y \\ 0 & \sin\phi_Y & \cos\phi_Y \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \qquad 4.1 \\
&= \begin{bmatrix} \sin\theta_R\ \cos\phi_Y \\ -\sin\phi_Y \\ \cos\theta_R\ \cos\phi_Y \end{bmatrix}
\end{aligned}
$$

Where $\mathbf{R}_x(\phi_Y)$ and $\mathbf{R}_y(\theta_R)$ are 3D rotation matrices describing rotations of ϕ_Y and θ_R degrees about the x- and y-axes respectively. The transformation $\mathbf{R}_x(\phi_Y)$ is applied before $\mathbf{R}_y(\theta_R)$, as rotation operations in 3D are non-commutative, and this ordering is required to model the physical system. In practice, the yaw axis is mechanically coupled to the roll motion of the main body, meaning that roll and yaw inputs can be applied in any order within the motion control software.

4.3.7 Motion Control Software

Figure 4.10 shows a screenshot of the graphical user interface (GUI) developed to control the position of the rotatable scanner.

The software runs on Windows 10 on the dedicated motion control computer and was implemented in C++ using the Qt open-source GUI framework. The software utilises Harmonic Drive C++ libraries to communicate with the commercially available Yukon Drives (also from Harmonic Drive) via CANBus, which is converted to a fibre-optic signal for transmission through the RF shield and then reconverted to CANBus for drive communication. All drive operations are managed through a state machine, as described in the Yukon Drive Application Manual [87], the status of which is displayed to the user in the roll and yaw status readouts.

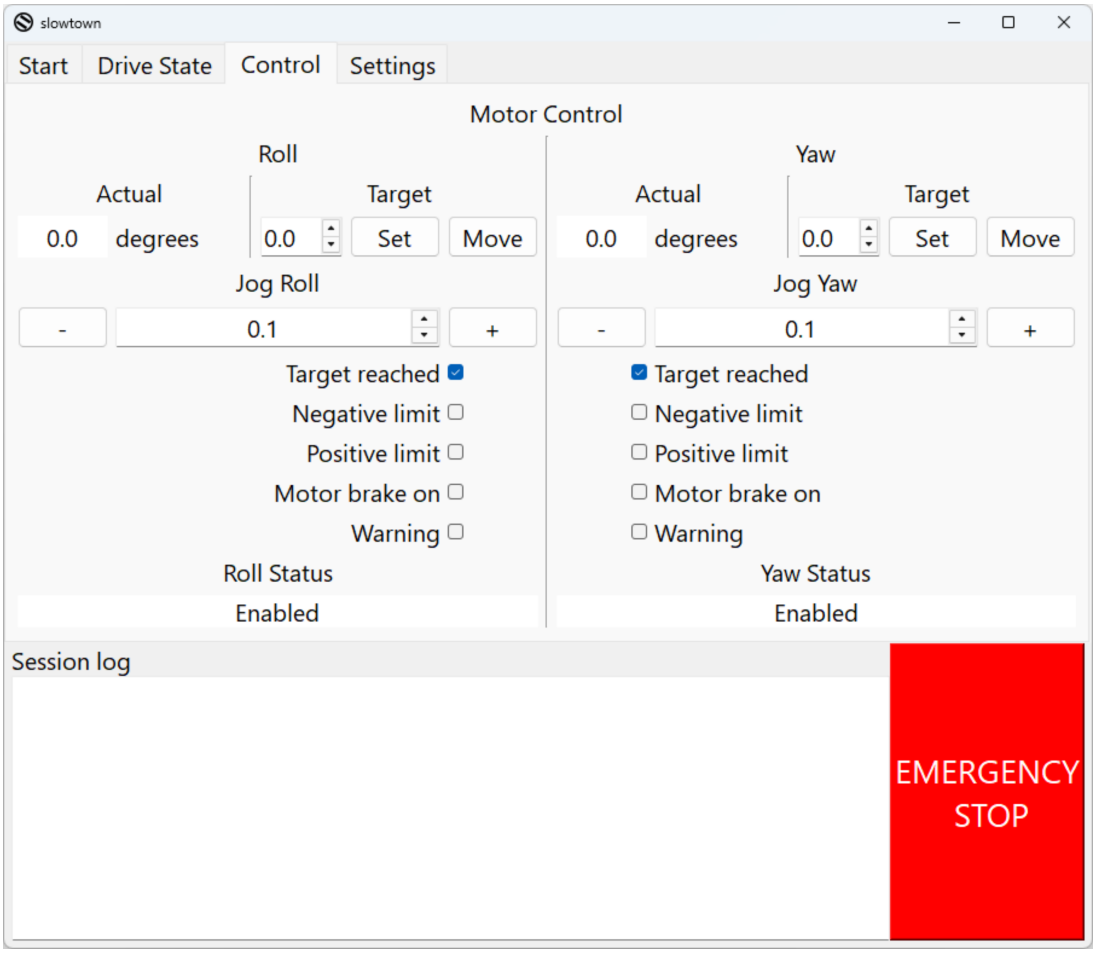
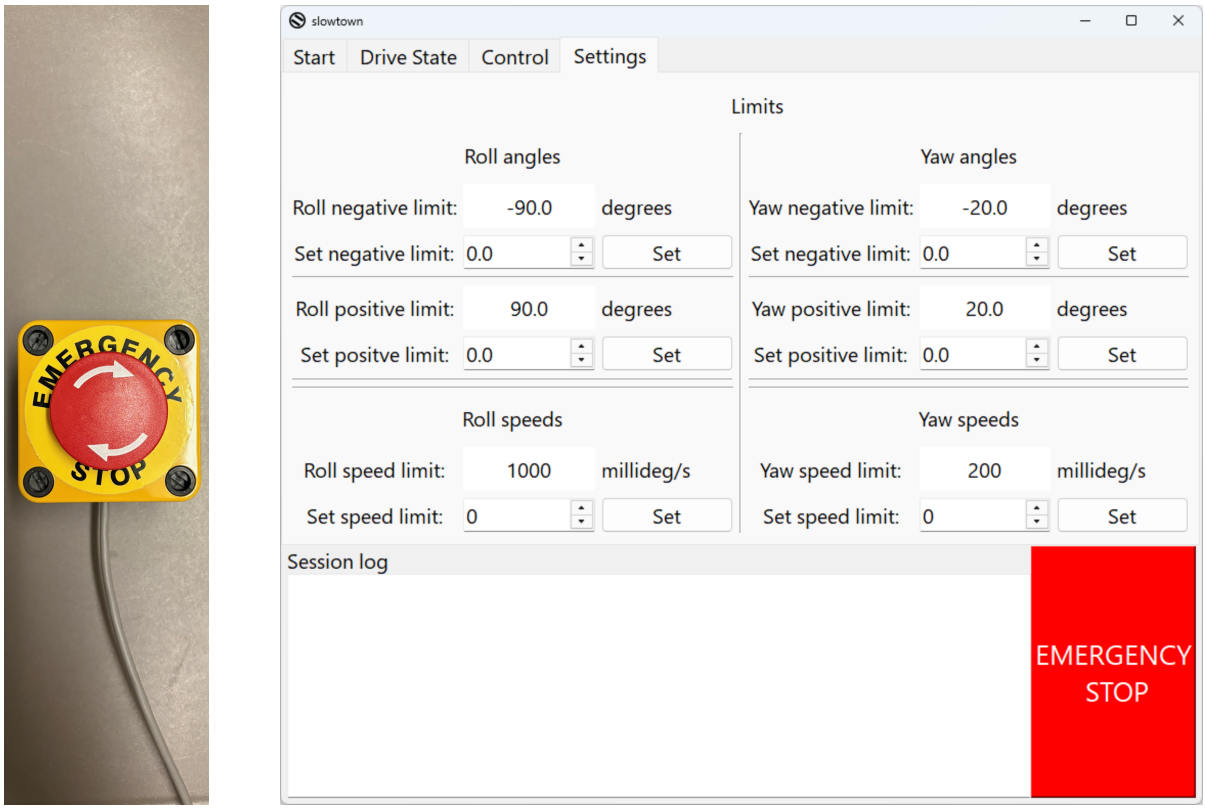


Figure 4.10: Screenshot of motion control software used to set scanner roll and yaw angles. Target values are specified independently of initiating scanner movement. The interface provides live numerical readouts of roll and yaw angles, checkboxes indicating limits and binary states, and textual status readouts for motor state display. A session log records all user inputs and motor responses, and a large emergency stop button immediately halts all scanner motion.

4.3.8 Patient Safety

Figure 4.11 shows the motion control software's settings page, which includes safety features such as position and speed limits, a session log, and an emergency stop button.



(a) Emergency stop button.

(b) Scanner motion settings page in the control software, showing configurable position and speed limits, and software-based emergency stop button.

Figure 4.11: Physical emergency stop button and scanner control software settings page with software emergency stop button.

The emergency stop button shown in Figure 4.11a is mounted on the patient chair, remaining accessible throughout the examination, and will immediately halt scanner movement if pressed. The operator has an additional software-based emergency stop button within the control software, as shown in Figure 4.11b, which also stops all scanner motion. The patient can also easily alert the operator to any issues verbally, as the RF shield is perforated and located in the same room as the operator station. If necessary, the operator can also cut power to the motors directly at the wall sockets located beside the operator station. The control software shown in Figures 4.10 and 4.11b includes a session log that records and displays all software inputs and motor drive status outputs, providing the operator with a direct record of the system's state during motion.

Figure 4.11b shows the roll and yaw limits set at $\theta_R \in [-90^\circ, 90^\circ]$, $\phi_Y \in [-20^\circ, 20^\circ]$. Mechanically, the scanner can roll nearly $\pm 180^\circ$ on its carriage and yaw the yoke $\pm 180^\circ$. As discussed in Section 2.3, only $\pm 90^\circ$ of roll and yaw rotation is required to orient $\vec{B}_0$ anywhere on a hemispherical maximum necessary operating envelope. However, for patient safety and comfort, yaw was limited to $\pm 30^\circ$ to prevent the yoke from interfering with the patient or their other leg, thus reducing the operating envelope along the y-axis. The scanner is also operated at low speeds, 1000 millidegree/s for roll and 200 millidegree/s for yaw, allowing both the patient and operator sufficient time to respond to any issues.

The low-field strength of the system makes it inherently safer than conventional MRI. For example, the 5 Gauss line, indicated by the dotted line on the floor in Figure 4.2a, lies entirely within the shielded room. Low-field operation also reduces projectile hazards and mitigates SAR concerns.

4.3.9 Pulse Sequence

In the present work, all MR images were acquired using a 3D Fast Low Angle SHot (FLASH) pulse sequence, the parameters for which are given in Table 4.1. 3D FLASH is a spoiled gradient echo sequence [88, 89] that produces predominantly T1-weighted contrast through the use of short TR and a small flip angle. However, as a gradient echo sequence it is inherently sensitive to T2* decay, a property that has been exploited to sample the magic angle effect in tendon, ligament, and menisci [14, 34, 90].

To capture the magic angle effect, a sequence with a short TE is necessary [11, 14, 34, 90] due to the intrinsically short T2 / T2* of anisotropic collagen-rich tissues. Bydder et al. [8] noted the T2 of tendon and ligament as 2–8 ms when the collagen fibres are parallel to $\vec{B}_0$. However, the T2 of anisotropic collagen-rich tissues is strongly dependent on fibre orientation relative to $\vec{B}_0$, reaching a maximum at the magic angle. For example, T2 in the human Achilles tendon increases from 0.2 ms at 0° to 22 ms at 54.7° relative to $\vec{B}_0$ [25], and from 7 ms to 23 ms in the bovine Achilles tendon [12]. It is therefore desirable to use a TE < 10 ms, ensuring that sufficient signal is retained from tissues with short T2 at unfavourable orientations relative to $\vec{B}_0$.

The 3D FLASH pulse sequence and its parameters were chosen based on pulse sequence optimisation work carried out by Chappell [34]. Chappell assessed the sampling

Table 4.1: Parameters for the 3D FLASH pulse sequence used for all acquisitions on the low-field rotatable scanner.

Parameter	**Value**
TE	5.2 ms
TR	13 ms
Flip angle	30°
FOV	$150 \times 150 \times 150$ mm
Resolution	$1 \times 1 \times 1$ mm
No. of averages	2
Total scan time	10 minutes

of the magic angle effect in a caprine knee patellar tendon by rotating a knee specimen within the bore of a conventional 3T scanner. Chappell analysed both proton density turbo spin echo (PD TSE) and 3D FLASH sequences. These sequences were selected for evaluation on the basis of their (future) availability on the low-field rotatable scanner, whose hardware was still under development at the time of the pulse sequence optimisation study. Chappell [34] found that in the caprine knee patellar tendon at 0° and 54.7° relative to $\vec{B}_0$ T1 FLASH had a 7.8-fold increase in signal intensity and PD TSE had only a 2.7-fold increase in signal intensity. Overall, Chappell [34] concluded that 3D FLASH yielded a greater magic angle effect than the PD TSE and was "much quicker to perform in a time that may be tolerable to research participants".

The 3D FLASH sequence used was a spoiled gradient echo technique which is specifically designed to disrupt steady-state T2 coherences and produce T1-weighted images [89]. However, as a gradient echo sequence it is also sensitive to T2* decay, which is the relevant contrast mechanism for detecting the magic angle effect in short-T2* collagen-rich tissues. It should be noted that the T2* sensitivity (rather than T2) of the FLASH sequence does not compromise the MADI technique. Firstly, T2* incorporates the underlying T2 of the tissue via Equation 2.5; thus, the orientation-dependent dipolar anisotropy that drives the magic angle effect manifests in T2* as well as T2, as discussed in Sections 2.2.5 and 2.3. Secondly, at low field, T2* converges towards T2 due to the lower susceptibility inhomogeneities as discussed in Section 2.2.5. The

sequence therefore provides a practical means of sampling $\vec{B}_0$ orientation-dependent contrast in collagen-rich tissues, despite producing images that would conventionally be characterised as T1-weighted.

Alternative pulse sequences such as balanced steady-state free precession (bSSFP), TSE, and ultrashort echo time (UTE) may offer SNR advantages over 3D FLASH for sampling the magic angle effect in collagen-rich tissues. bSSFP offers superior SNR efficiency per unit time and reduced off-resonance sensitivity at low field [91]; TSE provides improved acquisition efficiency by collecting multiple lines of k-space per TR [92]; and UTE sequences, with TE on the order of 100–500 μs, are particularly promising for high SNR imaging of collagen-rich tissues whilst also sampling the magic angle effect [93].

However, the availability of pulse sequences on the prototype low-field rotatable scanner used in this work was determined by the scanner pulse sequence library (proprietary from MR Solutions). The 3D FLASH sequence employed in the present work was selected on the basis of optimisation work conducted by Chappell [34] that targeted sequences available on the prototype low-field scanner as previously mentioned. Systematic evaluation of alternative pulse sequences as scanner capabilities develop is identified as a direction for future work in Section 9.5.2.

4.3.10 Examination Protocol

A standard MADI exam requires two operators present. The motor drives are first powered on at the wall socket, after which the scanner orientation is set using the motion control software, whilst the first operator monitors the patient's positioning and holds the receive and transmit coils in place within the scanner during yaw rotations to keep the coils stationary in the lab frame of reference. The patient or second operator can halt the motion at any time with the hardware and software emergency stop buttons. Once positioned, the sliding doors of the RF-shielded room are closed and the motor drives are powered off to prevent interference during imaging. This process is repeated for each scanner orientation, resulting in a total examination time of approximately 1.5 hours for a typical protocol with $N = 7$ orientations.

All datasets were acquired as isotropic volumes with a voxel size of $1 \times 1 \times 1$ mm^3 and volume dimensions of $150 \times 150 \times 150$ mm. The total scan time for each volume was 10 minutes, comprising two excitations that were subsequently averaged. Data from the scanner were stored as directories of 150 2D DICOM slices by default, each with in-plane dimensions of 150×150 voxels, which were then collated into 3D volumes during processing.

4.4 MADI Data Preprocessing

This section outlines the preprocessing steps performed prior to the core stages of the MADI methodology. Firstly, intensity rescaling is required to standardise the intensity range of each scan. The system records data in the DICOM format, which stores scaled intensities, from which the raw values must be recovered. The measured signal is also attenuated when yaw is applied to the scanner due to the solenoid receive coil design. Both effects must be compensated for to allow valid intensity comparisons across scans acquired at different scanner orientations. Secondly, denoising is applied to improve SNR and enhance the interpretability of the acquired MR images, although the MADI technique is robust to low SNR.

4.4.1 DICOM Intensity Rescaling

The DICOM format is a standard medical imaging format used across a wide range of modalities, including X-ray, CT, ultrasound, and MRI. To represent the variety of units and dynamic ranges of these modalities using a single format, the raw image data are rescaled into an unsigned integer range determined by the bit depth used, which can be 8, 12, or 16 bits. The rescaling is automatic, and is a function of the intensity values in the data. Without accounting for this rescaling, the intensities from one image in a set of N cannot be directly compared to another from the same set, thereby undermining the fundamental principles of the MADI technique. Prior to processing, the raw values for each image are recovered using the “Rescale Slope” and “Rescale Intercept” parameters stored in the DICOM header of each image, as defined by Equation 4.2.

$$I_{\text{RAW}} = R_S \cdot I_{\text{DICOM}} + R_I \tag{4.2}$$

Where I_{RAW} are the recovered raw intensities, I_{DICOM} are the intensities stored in the DICOM files, and R_S and R_I are the rescale slope and intercept from the DICOM headers. Rescaling each image in this manner ensures that any intensity variation across the N images arises from the magic angle effect, rather than from differences in DICOM-assigned rescaling parameters.

4.4.2 Yaw Intensity Rescaling

The sensitivity of the solenoid receive coil design used to acquire the data in this thesis is maximised when $\vec{B}_0$ is perpendicular to the coil axis. This orientation maximises the flux induced by the net magnetisation and thus yields the highest signal. When $\vec{B}_0$ is yawed away from perpendicular alignment, the signal is attenuated by a factor of cos ϕ_Y. The corrected intensity for data acquired with yaw applied is therefore given by Equation 4.3.

$$I_{\text{YC}} = \frac{I_{\text{YAW}}}{\cos \phi_Y} \tag{4.3}$$

Where I_{YC} are the intensities corrected for yaw attenuation, and I_{YAW} are the yaw attenuated intensities. This correction was effective up to angles of 30°, or an attenuation of $\approx 14\%$, beyond which image quality was substantially degraded by the associated loss in SNR. However, as noted previously, yaw was restricted to $\phi_Y \in [-30^\circ, 30^\circ]$ for patient safety and comfort, and thus this limitation did not pose a practical issue.

4.4.3 Denoising

Low-field MRI has inherently low SNR per unit time. SNR can be improved by performing multiple excitations and averaging the resulting images, although for m excitations the gain is limited to a factor of $\sqrt{m}$ for an m-fold increase in acquisition time. In this thesis, two excitations per volume image were acquired and averaged as a practical compromise between acquisition time and SNR. The averaging procedure is described in detail in Section 5.4.1.

Block Matching Denoising

Following the averaging step, block matching denoising, a well-known and powerful denoising technique, was implemented as a preprocessing step to improve the SNR

of the data input to the core stages of the MADI methodology. This approach exploits the self-similarity of regions within an image to remove noise whilst preserving coherent structures. For 2D block matching denoising, a set of small square "patches" of the image are identified that have similar structural characteristics. These patches are stacked to form a 3D "block", where signal components are highly correlated along the third dimension due to the similarity of the patches, whereas noise remains uncorrelated. A denoising transform, such as a wavelet transform, is applied to the block, followed by coefficient thresholding to separate signal from noise. The denoised block is then inverse-transformed, the patches are returned to their original positions, with any overlapping patch regions averaged to obtain the final pixel values. This concept extends naturally to volumetric data, where each patch is 3D, and the block is a 4D array. The technique has been shown to provide state-of-the-art denoising performance for MR data [94, 95].

Block matching was applied independently to each volume as part of the MADI pipeline (Section 4.2). This per-volume workflow was a natural choice given the batch nature of MADI data acquisition: each scan at each $\vec{B}_0$ orientation incurred a 10-minute wait, therefore denoising (and registration) could be performed on the most recent acquisition[3] whilst acquiring the next, effectively eliminating any preprocessing overhead.

The block matching volumetric denoising algorithm, BM4D, developed by Foi [94] and Maggioni et al. [95], was integrated into the MADI preprocessing pipeline to denoise the isotropic volume images prior to registration and analysis. The specific variant of the algorithm implemented was tuned by Foi to the Rician noise profile typical of MR magnitude images [96]. Figure 4.12 illustrates the effect of block matching denoising on a central slice of the fixed image from a MADI dataset acquired using the rotatable low-field scanner. This dataset, a left knee scan from a 25-29 male, is referred to as the "representative dataset" throughout this thesis, as it is used to demonstrate multiple stages of the MADI methodology. The ACL was chosen as the main focus of the representative dataset, as its morphology is more straightforward to visualise in 2D slices than the meniscus, and thus serves as a clearer reference for demonstrating each component of the MADI technique throughout this thesis.

[3] Approximately 1.5 minutes per volume on an AMD Ryzen 7 3700X (3.60 GHz, 8 cores), or approximately 3 minutes with registration (if necessary).

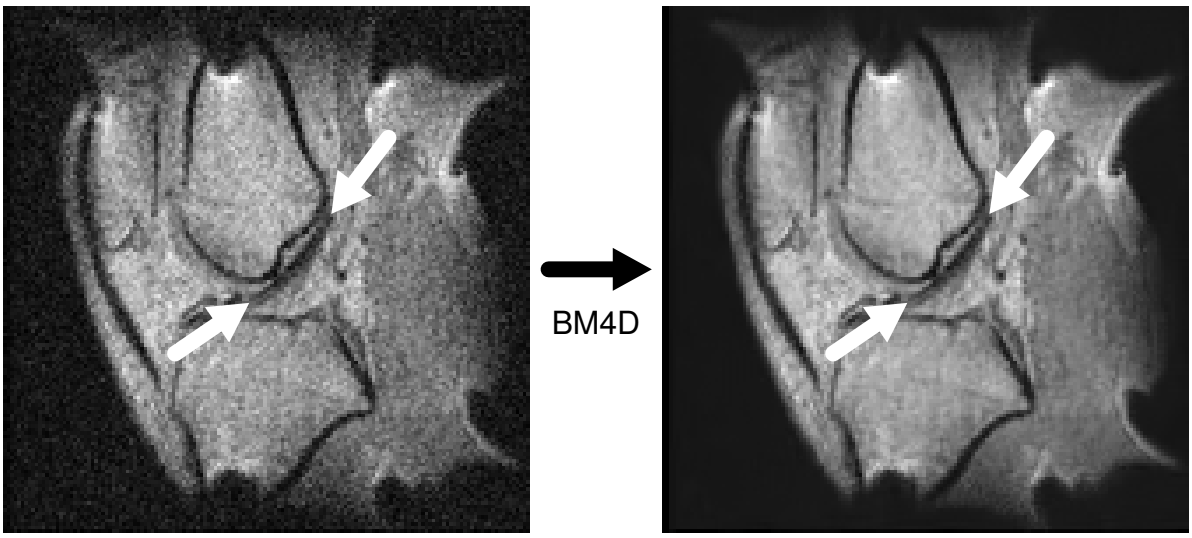


Figure 4.12: Example of block matching denoising applied to a knee scan from the rotatable low-field scanner (sagittal slice 76/150). Left: raw image prior to denoising. Right: image after BM4D denoising applied. The ACL is marked in both images using white arrows.

Qualitatively, Figure 4.12 demonstrates that BM4D substantially reduces noise whilst preserving fine structural details and edge definition. Whilst this is advantageous for standard clinical interpretation of the images, denoising was found to offer little improvement in the accuracy of collagen fibre orientation estimation itself, as discussed in Section B.1. Segmentation, however, benefits from reduced noise, as even with the semi-automatic metric-based approach described in Chapter 6, human input is required to fine-tune the results. Consequently, denoising is primarily valuable for tasks involving human interpretation rather than for the quantitative aspects of the MADI analysis pipeline.

Other Denoising Techniques

An alternative approach to the batch denoising workflow would be to concatenate all N registered volumes along the slice axis to form a single $150 \times 150 \times (150 \times N)$ volume and apply block matching once to the full dataset. This would increase the pool of self-similar patches available to the block matching algorithm, since the repeated anatomy across orientations provides additional repeated structure. However, this would reduce the operational efficiency of the concurrent acquisition and analysis workflow. Batch processing during MADI data acquisition also serves as a quality assurance step to monitor the data being acquired in case any repeat scans are necessary. This is nevertheless an interesting direction for future work.

BM4D was selected for its out-of-the-box state-of-the-art performance and suitability as a per-volume denoising approach. However, the structure of MADI data, N co-registered volumes of the same anatomy acquired at different orientations, presents an opportunity for cross-volume denoising methods that exploit inter-volume redundancy. Techniques such as low-rank and structured low-rank decomposition, principal component analysis (PCA), and tensor PCA are well-established in MRI denoising contexts [97, 98], and may offer enhanced denoising performance when applied to MADI datasets precisely because there is inherent repeated signal. A systematic comparison of denoising approaches in the context of MADI is an important direction for future work, and is discussed further in Section 9.5.3.

4.5 Summary

This chapter provided a complete overview of the MADI methodology and the setup of the rotatable low-field system required to apply the technique in vivo. Section 4.2 presented a high-level overview of the MADI process, illustrated as a flowchart in Figure 4.1, followed by a detailed breakdown of each step in Section 4.2.1, with cross-references to other sections of this thesis for further information.

The MADI technique is concisely summarised here as:

1. Set scanner orientation. Beginning with home orientation.
2. Acquire data.
3. Repeat steps 1 and 2 until N scans are acquired. $N = 7$ typically.
4. Preprocess data to rescale intensities and denoise images.
5. Register moving images to the fixed image.
6. Segment ROI.
7. Perform collagen fibre orientation estimation.
8. Apply geometric distortion compensation.
9. Visualise the data.

Section 4.3 then outlined the hardware and software components necessary to understand the workflow of a MADI examination, beginning with the system installation requirements, followed by an overview of the image acquisition hardware and system

block diagram, the motion platform, patient positioning, magnet design, and the choice of coordinate system based on the MRI convention. It then described the motion control software, patient safety measures, the choice of pulse sequence, and the examination protocol.

Finally, Section 4.4 presented the preprocessing steps specific to the MADI hardware and methodology, including intensity rescaling and denoising, which do not fall within the scope of Chapters 5 to 7 that provide deeper examinations of MADI experimental design or methodology.

5 | Registration

5.1 Chapter Overview

This chapter details all registration components of the MADI methodology. Section 5.2 introduces the concept of image registration and its specific application to MADI, followed by a description of the principles underpinning modern registration techniques. The elastix framework [99], used in this thesis to develop custom, novel registration pipelines, is presented in Section 5.3, along with its adaptation to the requirements of MADI. Subsequently, Section 5.4 describes the details of the custom registration pipeline and its application for excitation averaging with drift compensation, and automatic registration of the datasets acquired during a MADI examination. The chapter concludes with Section 5.5, a description of the distortion compensation methodology, utilising non-rigid registration techniques and a custom 3D printed calibration phantom to correct for systematic geometric distortions.

Registration is the first step in the MADI workflow following image acquisition and preprocessing. Its purpose is to align the N volumetric images acquired at different scanner orientations, enabling subsequent segmentation and analysis of magic angle effect intensity variations. The automatic soft-registration pipeline presented in this chapter represents a substantial improvement in the ease of use and practicality of the earlier MADI implementations, which relied on rigid, manually refined registration methods.

5.2 Registration Overview

In clinical settings it is common for multiple images of a patient to be acquired over time to monitor disease progression or recovery. Additionally, different modalities are often used to acquire multiple images of the same anatomy to obtain modality-specific information on a condition. Spatial alignment of images, whether uni- or multi-modal, facilitates enhanced interpretation of this data by enabling direct visual comparisons between different images. Common examples include overlaying images to highlight differences, viewing them as an animated sequence, or combining modalities to reveal insights that would not be apparent from isolated analysis of a single modality.

Image registration is the process of determining a one-to-one spatial mapping between the voxel domains of two images, thus enabling one image to be transformed so that it is spatially aligned with the other.

In the context of MADI, registration is required to align the $N - 1$ volume images acquired at rotated scanner orientations to the first image acquired in the scanner's home orientation. The reorientation of the scanner introduces substantial spatial misalignment between images, which must be corrected to enable subsequent analysis. A further complication is that the tissues of interest exhibit intensity variations across the dataset due to the magic angle effect, making MADI registration more analogous to multimodal registration, where tissue intensities differ significantly between images, than to unimodal registration, in which intensities are generally assumed to remain constant.

Previous implementations of MADI registration [14, 15] employed single shot rigid registration based on fiducial markers with manual fine-tuning. This approach was sufficient for these early MADI trials, in which excised caprine knee samples were positioned in a spherical jig at the isocentre of a 3T scanner. Under these conditions, geometric distortion was minimal, and simple rigid registration was sufficient. In contrast, the low-field rotatable scanner that enables in vivo MADI introduces geometric distortions that must be corrected, both relative to other images in a MADI dataset and in an absolute sense. Furthermore, the complexity of this registration, combined with limited space within the receive coil precluding the use of fiducial markers, necessitates a fully automatic, telemetry-driven approach.

5.2.1 Registration Principles

In image registration, one image is designated the “fixed” image, and the other the “moving” image. The fixed image is the reference image that remains unchanged during the process, and the moving image is deformed to align with the fixed image. An intuitive approach to registration is to define a transform from the moving image domain to the fixed image domain, and then to apply this transform to the moving image data to produce a registered image.

However, this approach suffers from a critical limitation. Defining the transform from the moving to the fixed image can produce voids in the registered image. This occurs because MR images are defined on discrete voxel grids, and a transform from the moving to fixed image domain cannot ensure that adjacent voxels in the moving image remain adjacent after transformation. Consequently, some voxels in the fixed image domain may receive no assigned intensity under such a transform. Figure 5.1 illustrates this issue with a simple example of registration between two images of an object where the moving image is scaled down along one axis compared to the fixed image.

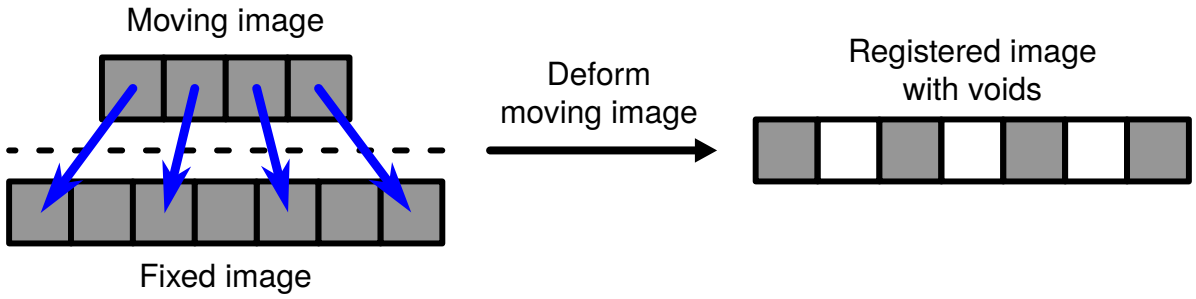


Figure 5.1: Image voids resulting from defining a transform from the moving to the fixed image. The blue arrows represent the transform applied to each voxel centre in the moving image. Upscaling causes previously adjacent moving-image voxels to become non-adjacent in the fixed image domain, leaving unassigned voxels (white) in the registered image.

The blue arrows in Figure 5.1 represent a transform from the moving to the fixed image. The moving image object is scaled down along the horizontal axis relative to the fixed image. Whilst the transform successfully maps corresponding object boundaries, the upscaling causes previously adjacent voxels in the moving image to become non-adjacent in the fixed image domain, leaving voids with no assigned intensity in the registered image. Post-registration interpolation to fill these voids is technically possible, but image quality in regions of high void density would be severely degraded. A more effective solution is to instead define the transform from the fixed image domain to

the moving image domain. A resampled registered image can then be generated by inheriting the fixed image's voxel grid and populating it with intensities interpolated from the moving image. Figure 5.2 illustrates this approach.

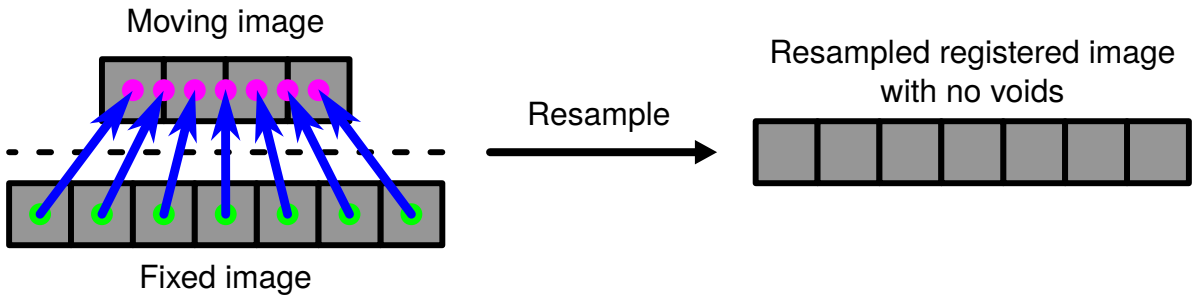


Figure 5.2: Coordinate transform from fixed to moving image domain to avoid voids in the resampled registered image. The blue arrows represent the transform from the fixed to the moving image, originating at the fixed image voxel centres (green dots) and terminating at the corresponding sampling points in the moving image (pink dots). The resampled registered image is generated by interpolating the moving image intensities at the sampling points and assigning them to their corresponding fixed image voxel coordinates, yielding a void-free result.

In Figure 5.2, the fixed image voxel centres (green dots) are mapped into the moving image domain, defining sampling points (pink dots) at which the moving image intensity is interpolated. The resampled registered image is then constructed on the fixed image grid, populated with these interpolated intensities, and is therefore void-free by construction. Figure 5.3 extends this demonstration to a more complex geometry with differing intensities between images, analogous to multimodal acquisition or, in the context of this work, the magic angle effect.

In Figure 5.3, the transform $\mathbf{T}$ (blue arrow) maps each fixed image voxel coordinate into the moving image domain, where the intensity, I, is obtained via interpolation. The interpolated intensity is assigned (pink arrow) to the coordinate C from the fixed image (green arrow) in the resampled registered image. Repeating this for every voxel in the fixed image domain yields a void-free result in which the geometric structure is inherited from the fixed image and all intensity information is interpolated from the moving image. Crucially, $\mathbf{T}$ acts as a sampling map rather than a direct deformation of the moving image and as such the moving image itself is never explicitly transformed.

This formulation is well suited to the MADI pipeline. With known scanner roll and yaw telemetry and appropriate image similarity metrics, $\mathbf{T}$ can be initialised and estimated automatically via conventional parameter optimisation. Different transform models for $\mathbf{T}$ can also be implemented to account for the expected degrees of freedom in misalign-

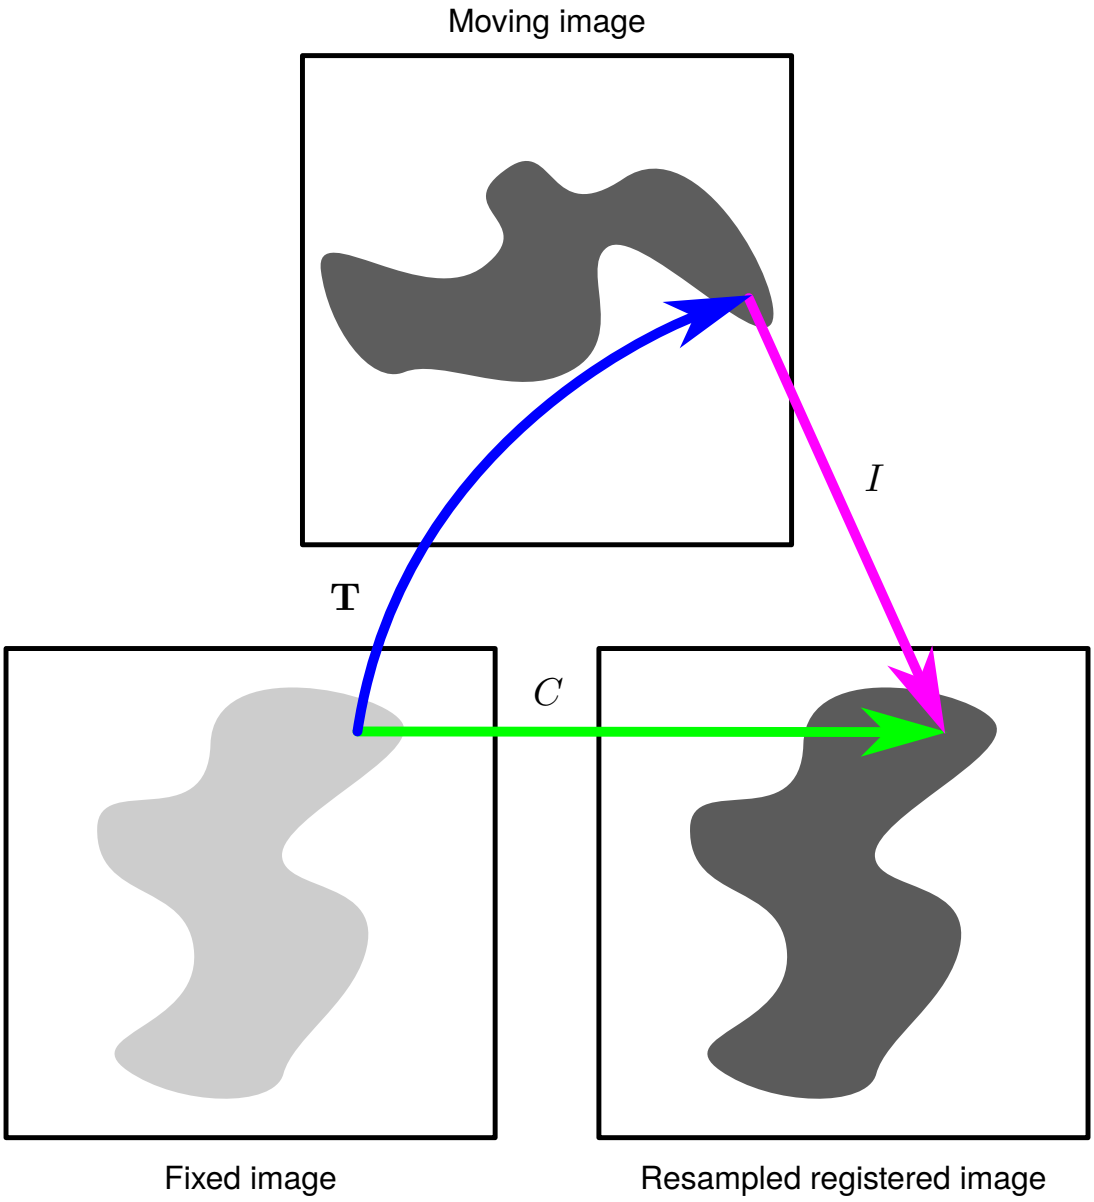


Figure 5.3: Transform-based registration demonstration. The fixed image contains a bright object; the moving image shows the same object with reduced intensity, a different orientation, and geometric distortion. The transform $\mathbf{T}$ (blue arrow) maps a voxel coordinate, C, from the fixed image domain into the moving image domain, where the intensity, I, is obtained via interpolation. I is then assigned to coordinate C in the resampled registered image, which therefore inherits the morphology of the fixed image and the intensity values of the moving image.

ment: a rigid transform encodes translations and rotations only; a similarity transform additionally permits uniform scaling; an affine transform further allows shearing and non-uniform scaling; and a non-rigid B-spline transform supports spatially varying local deformations, as required for geometric distortion correction. Most of these transform types are employed within the MADI registration pipeline developed in this thesis, as described in Section 5.4.

5.3 Registration Framework

In this thesis, the custom MADI registration pipeline was implemented using the elastix toolbox [99], which performs registration according to the principles outlined in Section 5.2.1. This section reviews the mathematical framework for image registration and presents a block diagram illustrating the operation of parametric registration frameworks such as elastix. This provides context for the definitions of rigid, affine, and B-spline transforms used in the novel MADI registration pipeline, as well as for other parameters that control the registration process, including alignment metrics and regularisation settings. A clear understanding of registration fundamentals, and of their formulation and implementation in this framework, was essential to the successful design, implementation, and application of the novel MADI registration pipeline presented in Section 5.4.

5.3.1 Mathematical Formulation of Image Registration

Klein et al. [99] define the intensities of the fixed and moving images as $I_F(\vec{x})$ and $I_M(\vec{x})$ respectively, where $\vec{x}$ is a 3D coordinate vector. $I_F(\vec{x})$ and $I_M(\vec{x})$ have distinct domains, Ω_F and Ω_M, of dimension d, where $d = 3$ for volume images: $\Omega_F \subset \mathbb{R}^d$ and $\Omega_M \subset \mathbb{R}^d$. Mathematically, registration is the task of finding the displacement field $\vec{u}(\vec{x})$ such that $I_M(\vec{x} + \vec{u}(\vec{x}))$ is spatially aligned with $I_F(\vec{x})$.

By defining $\mathbf{T}(\vec{x}) = \vec{x} + \vec{u}(\vec{x})$, registration can equivalently be expressed as the problem of finding a transformation, $\mathbf{T}(\vec{x})$, such that $I_M(\mathbf{T}(\vec{x}))$ is spatially aligned with $I_F(\vec{x})$. This transformation is a mapping from the fixed image domain to the moving image domain, $\mathbf{T} : \Omega_F \subset \mathbb{R}^d \mapsto \Omega_M \subset \mathbb{R}^d$, which corresponds to the blue arrow in Figure 5.3.

The transform $\mathbf{T}$ is denoted $\mathbf{T}_{\vec{\mu}}$ when parametrised by the vector $\vec{\mu}$, which contains all parameters necessary to describe a given transform. For example, for a hypothetical transform restricted to pure translations in three dimensions, $\vec{\mu}$ would comprise three elements representing translations along the x, y, and z axes.

To optimise a transform $\mathbf{T}_{\vec{\mu}}$ over its parameters $\vec{\mu}$, it is necessary to define a cost function, $\mathcal{C}$, as shown in Equation 5.1.

$$\hat{\mu} = \arg\min_{\vec{\mu}} \mathcal{C}\left(\vec{\mu}; I_F, I_M\right) \tag{5.1}$$

Where $\hat{\mu}$ are the optimised transform parameters. For $\mathcal{C}$ to serve as a representative cost function, it must quantify how well the transformation $\mathbf{T}_{\vec{\mu}}$ aligns the images using some metric. The choice of metric depends on the characteristics of the images being registered. Further details on the parameters associated with different transform types, and on the metrics over which the cost function is optimised, are provided in Sections 5.3.3 to 5.3.7.

5.3.2 Registration Framework Block Diagram

Figure 5.4 presents a simplified block diagram of the elastix registration components and their interconnections within the registration workflow, adapted from Klein et al. [99].

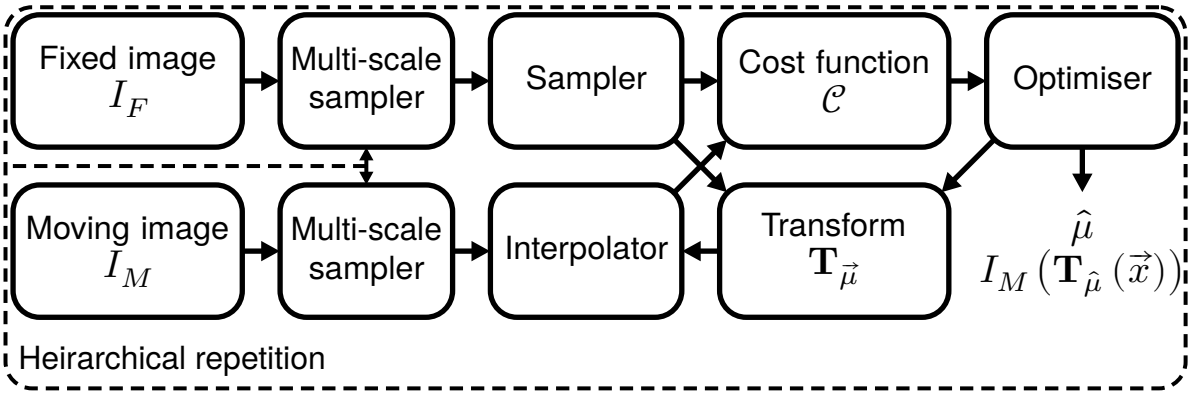


Figure 5.4: Block diagram of the basic registration components of elastix. Fixed and moving image inputs are shown on the left, with the optimised parameters, $\hat{\mu}$, and final optimised transformed image $I_M\left(\mathbf{T}_{\hat{\mu}}\left(\vec{x}\right)\right)$ shown as outputs on the right.

The fixed and moving images shown on the left are the inputs to the workflow. They both initially pass through a multi-resolution sampler block, which applies smoothing and/or downsampling according to a predefined multi-resolution schedule, sometimes referred

to as a “pyramid”. This hierarchical approach enables coarse-to-fine registration: initial low-resolution stages align macroscopic structures, whilst subsequent higher-resolution stages refine the alignment, culminating in a final pass using the original input resolution.

The fixed image is then processed by the sampler component, which selects a subset of voxel coordinates for evaluating the current transformation in the cost function. Using a subset reduces computational cost, as evaluating the entire image at each optimisation iteration would be extremely computationally expensive. The moving image is passed to the interpolator, which retrieves intensities at the sampled coordinates mapped into the moving image domain by the transform component. The cost function then assesses the alignment quality between the sampled fixed image intensities and the corresponding interpolated moving image intensities against a chosen metric. This evaluation is passed to the optimiser, which iteratively updates the transform parameters $\vec{\mu}$ using adaptive stochastic gradient descent to minimise the cost function. The optimisation terminates when the maximum number of iterations is reached or when no further improvement in the cost function is detected, after which the optimised parameters are output along with the final transformed image.

For a detailed description of the internal operation of elastix, the reader is referred to the elastix manual [100] and the original publication by Klein et al. [101].

5.3.3 Rigid Transform

The rigid, or Euler, transform models the image as a rigid body and applies a global transformation to all voxels in the image, as defined in Equation 5.2.

$$\mathbf{T}_{\vec{\mu}}(\vec{x}) = \mathbf{R}(\vec{x} - \vec{c}) + \vec{t} + \vec{c} \tag{5.2}$$

Where $\mathbf{R}$ is a rotation matrix, $\vec{c}$ are the coordinates of the centre of rotation (by default the centre of the image), and $\vec{t}$ is a translation vector. The rotation is parametrised by the Euler rotation angles and the translation vector. In 3D this gives a parameter vector with 6 elements: $\vec{\mu} = \left(\theta_x, \theta_y, \theta_z, t_x, t_y, t_z\right)$.

5.3.4 Affine Transform

The affine transform has a similar formulation to the rigid transform and is also applied globally, as defined in Equation 5.3. However, it allows greater freedom to deform the image through scaling and shearing, in addition to rotation and translation.

$$\mathbf{T}_{\vec{\mu}}\left(\vec{x}\right) = \mathbf{A}\left(\vec{x} - \vec{c}\right) + \vec{t} + \vec{c} \tag{5.3}$$

Where $\mathbf{A}$ is the affine transform matrix, which encodes rotation, non-uniform scaling, and shear (allowing different scaling or shearing along each axis). As defined, the affine transform matrix, $\mathbf{A}$, is a 3×3 matrix for 3D volumes, encoding rotation, scaling, and shear as translations are encoded by $\vec{t}$. The parameter vector for 3D volume images therefore contains 12 elements: $\vec{\mu} = \left(a_{11}, a_{12}, \ldots, a_{32}, a_{33}, t_x, t_y, t_z\right)$.

5.3.5 B-Spline Transform

The B-spline transform [102], defined in Equation 5.4, maps an input coordinate $\vec{x}$ to a locally deformed position using a smoothly varying displacement field parametrised by B-spline basis functions.

$$\mathbf{T}_{\vec{\mu}}\left(\vec{x}\right) = \vec{x} + \sum_{\vec{x}_k \in N_{\vec{x}}} \vec{p}_k \boldsymbol{\beta}^3 \left(\frac{\vec{x} - \vec{x}_k}{\vec{\sigma}}\right) \tag{5.4}$$

Where $\vec{x}_k$ are the spline control points, $N_{\vec{x}}$ is the set of all control points within the compact support at $\vec{x}$, $\boldsymbol{\beta}^3\left(\cdot\right)$ is the cubic multidimensional B-spline polynomial [103], $\vec{p}_k$ are the B-spline coefficient vectors, and $\vec{\sigma}$ is the control point spacing in each dimension.

The control points, $\vec{x}_k$, are arranged in a regular grid overlaid on the fixed image, with spacing $\vec{\sigma}$. Each control point has an associated B-spline coefficient vector, $\vec{p}_k$, which weights its contribution to the local displacement field. B-splines have local support, meaning that the displacement at a given position depends only on a small neighbourhood of control points and their coefficients, rather than the entire grid. For a cubic B-spline (B-spline order $n = 3$) in three dimensions, the compact support spans

$n + 1 = 4$ control points along each axis, resulting in $4^3 = 64$ contributing control points.

The parameter vector, $\vec{\mu}$, for this transform comprises all elements of all the coefficient vectors, $\vec{p}_k$, that parametrise the B-splines. Each vector p_k is associated with a control point and has d elements for an image of dimension d. The total number of control points is determined by the grid spacing, $\vec{\sigma}$, and the image size, $\vec{s}$, with the number of control points along a dimension i given approximately by $P_i \approx s_i/\sigma_i$. The approximation arises because additional control points are placed just outside the image boundaries. The total number of parameters, M, is then $M = P \times d$, where $P$is the total number of control points, $P = \prod_{i=1}^{d} P_i$. For example, a 3D image of size $150 \times 150 \times 150$ voxels with a uniform grid spacing of 10 voxels yields approximately $(150/10)^3 = 3{,}375$ control points, and thus approximately $3{,}375 \times 3 = 10{,}000$ parameters in total.

Owing to the local support of the coefficient vectors $\vec{p}_k$, additional constraints can be applied to the B-spline transform to control its behaviour. Examples include regularisation terms, as described in Section 5.3.7, that prevent excessive deformation, and masking strategies, as described in Section 5.4.4, which can be used to suppress B-spline adjustments near image boundaries that contain no useful signal.

5.3.6 Registration Metrics

To evaluate the quality of a transform and optimise its parameters, $\vec{\mu}$, over successive iterations, the cost function must incorporate a similarity measure that quantifies the degree of alignment between the fixed and moving images. When the two images have similar intensity distributions, simple metrics such as the mean squared difference or the normalised correlation between the fixed and transformed images are sufficient to assess alignment. In MADI, however, the intensity values of collagen-rich tissue structures vary between images by design. Consequently, similarity measures suited to multimodal registration are more appropriate, as these assume structural correspondence across images without requiring intensity values to match.

The metric used for MADI registration was chosen to be mutual information [104–106], which captures structural similarity between images rather than relying solely on intensity correspondence, making it well suited to both multimodal and MADI registration.

The mutual information metric implemented by elastix is defined by Equation 5.5 from Thévenaz et al. [107].

$$\text{MI}\left(\vec{\mu}; I_F, I_M\right) = \sum_{m \in L_M} \sum_{f \in L_F} p(f, m; \vec{\mu}) \log_2 \left(\frac{p(f, m; \vec{\mu})}{p_F(f)\, p_M(m; \vec{\mu})} \right) \quad 5.5$$

Where L_F and L_M denote sets of regularly spaced intensity bin centres for the fixed and moving images respectively, f and m are specific bin centres within these sets, p is the discrete joint probability, and p_F and p_M are the marginal discrete probabilities of the fixed and moving images, obtained by summing p over m and f respectively. During optimisation, only samples from the fixed image domain are transformed at each iteration for computational efficiency, so p must be estimated. This estimation is performed using B-spline Parzen windows [107]. To focus the registration on relevant anatomy and avoid errors caused by artificial edges (e.g. field-of-view boundaries) or by features present in one image but absent in the other, registration masking can be applied. Masking restricts the sampling domain within Ω_F, thereby limiting the region over which the Parzen windows estimate the joint probability. This ensures that the similarity measure is driven by anatomically relevant regions rather than background or artefactual features, as described in Section 5.4.4.

The mutual information metric measures the reduction in uncertainty about the intensity of one image, given knowledge of the intensity in the other. In the context of registration, its value increases as the statistical dependence between corresponding voxel intensities grows, and it is therefore maximised when the images are aligned. Because it is computed from marginal and joint probability distributions rather than direct intensity differences, mutual information assesses structural alignment even when the intensities of a given tissue differ between the fixed and moving images.

5.3.7 Registration Penalties

For non-rigid transformations, the registration problem is inherently ill-posed, and excessive flexibility in the transformation can produce artefacts such as folding, pinching, or overexpansion. Such effects are particularly likely in regions of uniform intensity, including background areas or tissues with homogeneous appearance, such as some muscle regions. To mitigate this, a regularisation term is added to the cost function,

as shown in Equation 5.6, to constrain the transformation and suppress non-physical deformations.

$$\mathcal{C}\left(\vec{\mu}; I_F, I_M\right) = -\mathcal{S}\left(\vec{\mu}; I_F, I_M\right) + \gamma\mathcal{P}\left(\vec{\mu}\right) \tag{5.6}$$

Here, $\mathcal{S}\left(\vec{\mu}; I_F, I_M\right)$ denotes a similarity metric, such as mutual information, and $\mathcal{P}\left(\vec{\mu}\right)$ is a penalty function scaled by a coefficient γ that controls the degree of regularisation. In Equation 5.6, the similarity metric appears with a negative sign because optimisation is formulated as a minimisation problem, so higher similarity corresponds to a lower cost. The penalty term is positive, such that increased B-spline bending results in a higher cost.

In this thesis, excessive curvature of B-spline transforms was penalised using the "transform bending energy penalty" to suppress rapid local displacements, such as adjacent large compressions and expansions. A low weight of $\gamma = 0.01$ was used, as the registration masking described in Section 5.4.4 restricted sampling to regions with sufficient structural detail to prevent excessive deformation of the transform.

5.4 Registration in the Context of MADI

The MADI registration pipeline comprises two interconnected stages. The first is a single-stage registration of the two excitations acquired for each volume image, performed to compensate for image drift arising from magnet heating during operation. The second is a multi-stage registration pipeline, which first corrects for large-scale misalignments introduced by magnet rotation using examination telemetry, and then applies non-rigid registration to correct for geometric distortions between scans, arising from B_0 inhomogeneity within the DSV. Both stages are implemented in elastix, which enables modular composition of transforms and fine-grained control of their behaviour through input parameter files, transform parameter files, and registration masks. The elastix inputs and outputs most relevant to this thesis are summarised in Table 5.1.

The input parameters, I_P, are a set of user-defined values that control the type and behaviour of the transform, its optimisation scheme and rate, the sampling and interpolation strategy, and other configurable settings that tune the behaviour of elastix. In

Table 5.1: Summary of elastix inputs and outputs relevant to the MADI registration pipeline.

Name	I/O Type
Fixed image, I_F	Input
Moving image, I_M	Input
Input parameters, I_P	Input
Registration mask, M_R	Input
Transform parameters, T_P	Both
Registered image, I_R	Output

Equation 5.3 and Figure 5.7, these input parameters are not shown explicitly but are represented by the transform type specified in each transform block. The complete parameter files for all transforms used in this thesis are included in Section C.1.

The registration mask, M_R, is applied to the fixed image and defines the region from which samples are drawn. As discussed in Section 5.3.2, transforming the entire fixed image during optimisation is computationally prohibitive, so random samples are used instead as an approximation until the final registration step, at which point the full image is transformed to produce the registered image, I_R. The registration mask restricts transform evaluation to a specified region, thereby excluding spurious features such as the artificial edge at the boundary of the magnetic field, or features that may be present in one orientation but absent in another. Further detail on the registration masks used in this work is provided in Section 5.4.4.

The optimised transform parameters, T_P, are output by elastix at the end of the optimisation process, as illustrated in Equation 5.3, in the form of a text file. These parameters can also be supplied as input initial conditions to a subsequent transform, enabling transform composition. In this way, successive transforms can be applied using the same fixed and moving image pair, avoiding the repeated resampling that would otherwise occur if each intermediate registered image were generated and then reused. This approach prevents progressive degradation of intensity data due to multiple rounds of interpolation. Figure 5.7 illustrates how successive transforms operate on the same fixed and moving images, whilst propagating the optimised parameters from each stage forward as initial conditions for the next transform in the pipeline.

Once optimisation is complete, the final registered image, I_R, is generated by transforming every point in the fixed image domain into the moving image domain using the specified transform and its optimised parameters. The corresponding intensity values are then interpolated, typically using B-spline interpolation, and resampled to construct the registered image, as illustrated in Figure 5.3. The parameters governing this resampling step are provided in full in Section C.1.

5.4.1 Image Drift Compensation and Averaging

During image acquisition, the gradient coils generate heat that is conducted into the magnet poles on which they are mounted. Permanent magnets such as the Nd-Fe-B poles of the rotatable low-field system used in this thesis exhibit a negative temperature coefficient of remanence of approximately $0.1\%/^{\circ}C$ [108]. As the magnet temperature rises, the field strength decreases, and since the Larmor frequency is directly proportional to the static field strength as per Equation 2.2, heating of the magnet produces a corresponding shift in resonance frequency. For the data in this thesis, two excitations were acquired for signal averaging as described in Section 4.3.10. This drift in resonance frequency produced a misalignment between acquisitions, manifesting as a shift along the frequency-encode axis in the reconstructed image due to the constant frequency offset in k-space.

The temperature of the laboratory itself was also variable (typically $19 - 23^{\circ}$C), and as such the extent of this frequency-induced shift was inconsistent, depending on the laboratory temperature, the duration of the examination, and whether repeated acquisitions were required for any reason. To address this, an automatic compensation step was implemented within the registration pipeline. Figure 5.5 illustrates the affine registration used to align the two excitations prior to averaging[1]. Excitation 1 is the fixed image, I_F, while excitation 2, the moving image I_M, is subject to drift relative to excitation 1 which required correction prior to averaging. The registration outputs are the registered image, I_R, together with the optimised transform parameters T_P.

[1] Since the temperature-induced misalignment was purely translational along the frequency-encode axis, a purely translational transform would have been sufficient. The affine transform, being a superset of a translational transform, was retained as it incurred no meaningful computational overhead for this application and consolidated parameter file management within the registration pipeline.

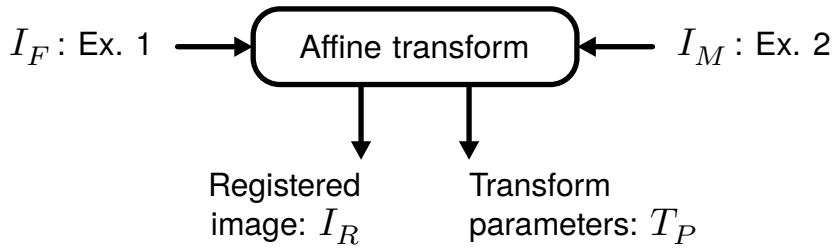


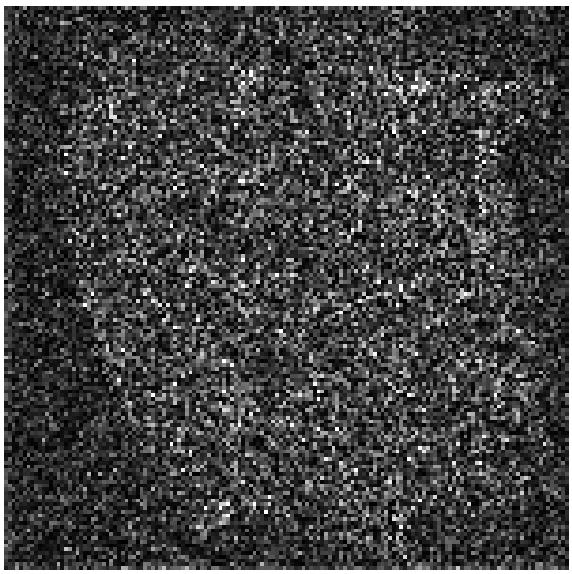
Figure 5.5: Block diagram of affine registration for image drift compensation prior to averaging. The first excitation is the fixed image, I_F, and the second excitation which has drifted relative to the first is the moving image, I_M. The registered image, I_R, is aligned by the affine transform with I_F ready for subsequent averaging.

Figure 5.6 presents difference images of the central slice from the two excitations acquired in the home orientation scan of the representative dataset. Figure 5.6a demonstrates the relative drift between excitations 1 and 2, evident as bright offset edges of the main anatomical features in the difference image. Following registration, Figure 5.6b shows that the two excitations are successfully aligned, with the residual difference image dominated primarily by noise.

The MADI registration pipeline is not strictly continuous between the drift compensation and primary registration stages, as denoising is performed after the averaging step, as described in Section 4.4.3. The denoised images are then supplied to the main registration pipeline detailed in Section 5.4.2.



(a) Difference image of excitations 1 and 2 before affine registration.



(b) Difference image of excitations 1 and 2 after affine registration.

Figure 5.6: Difference images of excitations 1 and 2 before and after affine registration. Figure 5.6a shows the difference image of excitations 1 and 2 from a central sagittal slice (76/150) of the representative dataset. Feature edges are clearly distinguishable as a drift of ≈ 5mm occurred between acquisitions. Figure 5.6b shows the difference image of excitation 1 and registered excitation 2. Figure 5.6b is dominated by noise as features within the two images have been aligned by the affine transform.

5.4.2 Registration Pipeline

Figure 5.7 shows the custom registration pipeline developed for MADI registration using the elastix framework. The inputs to the pipeline are: the N averaged and denoised images, where scan $n = 1$ is the fixed image, $I_{F,n=1}$, and scans $n = i$ for $i = 2 \dots N$ are the moving images, $I_{M,n=i}$; the scanner orientation telemetry, $\theta_{R,n=i}$ and $\phi_{Y,n=i}$, for each acquisition; and the registration mask, M_R, that controls the sample region over which the transform is optimised (see Section 5.4.4). The outputs are: the optimised transform parameters from each stage of the pipeline, T_{Pji} for $j = 0 \dots 3$, and the final registered image, $I_{R,n=i}$, for each rotated acquisition.

The first step is a "dummy" rigid transform call, which generates a template transform parameter file, T_{P0i}, in the correct format for the subsequent rigid registration. This dummy call is restricted to a single optimisation iteration to force immediate output of the template. The rigid transform encodes an Euler transform, with parameters $\vec{\mu}$ consisting only of translations and Euler-angle rotations. The default parameters in the template are replaced with examination telemetry from scan $n = i$ to provide the initial conditions for

the first real rigid transform. Specifically, the Euler angles θ_x and θ_y are set to $-\theta_{R,n=i}$ and $-\phi_{Y,n=i}$ respectively, since reorientation of the subject in the image domain is complementary to reorientation of $\vec{B}_0$ in the laboratory global coordinate system. All remaining elements of $\vec{\mu}$ are set to zero; this populated template is the T_{P1i} input transform parameters file for the first real transform of the pipeline, the rigid transform. This initialisation is required to ensure sufficiently close starting alignment, as otherwise the optimisation may converge to a local minimum or fail altogether.

The second step is a rigid transform with initial conditions provided by T_{P1i}. This accounts for most of the gross misalignment between the fixed and moving images arising from scanner rotation. The optimisation phase of this transform therefore serves only to fine-tune the telemetry-driven alignment. The resulting transform parameter file, T_{P2i}, is then passed as the initial condition to the next stage, the affine transform.

The third step is an affine transform with initial conditions provided by T_{P2i}. This step further refines the alignment prior to non-rigid registration, accommodating scaling or shearing of the moving image that cannot be captured by a purely rigid transform. The output parameter file, T_{P3i}, is supplied as the initial condition for the subsequent non-rigid transform. Importantly, T_{P3i} represents a composition of all preceding transforms, such that each stage only requires the original fixed and moving images as data inputs. No intermediate resampled images are generated at any point in the pipeline, thereby avoiding unnecessary interpolation, resampling, and associated degradation of the original signal.

The final stage is a non-rigid transform, initialised with T_{P3i} and parametrised by B-splines as described in Section 5.3.5. This stage provides the final refinement of each moving image, correcting for geometric distortions that co-rotate with the scanner, as well as soft-tissue deformations in living subjects. At this point the resampled, registered image is generated, as this is necessary to produce the final output of the pipeline.

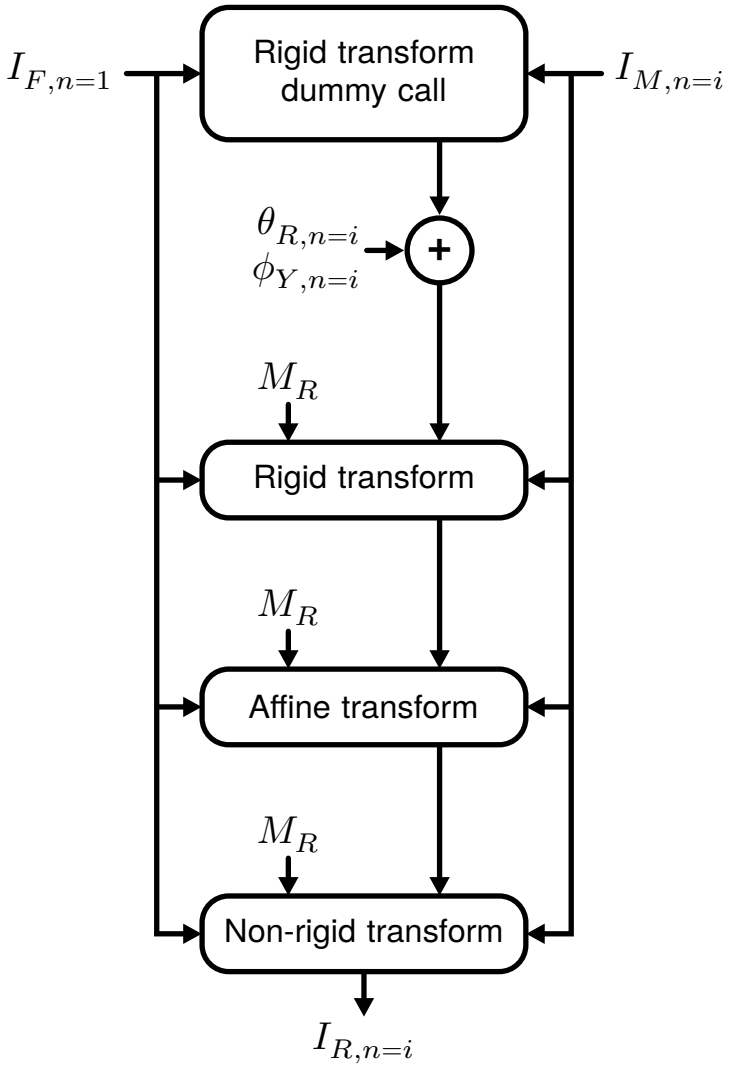


Figure 5.7: MADI registration pipeline. The fixed image, $I_{F,n=1}$, is the first acquisition from the set of N, acquired at the home orientation of the scanner. The moving images, $I_{M,n=i}$ for $i = 2 \ldots N$, are the $N - 1$ acquisitions at rotated orientations of the scanner. $I_{F,n=1}$ is a consistent input to all transform stages of the pipeline as alignment of all moving images with the fixed image is the overall objective of the registration pipeline. $I_{M,n=i}$ is also a consistent input to all transform stages as the optimised transform parameter of each stage are propagated and composed into the next rather than intermediate registered images being generated. The initial transform is parametrised by MADI examination telemetry of the roll and yaw angle used, $\theta_{R,n=i}$ and $\phi_{Y,n=i}$, and all subsequent transforms are optimised over relevant sample regions of the fixed image by the registration mask, M_R.

5.4.3 Registration of Rotated Image Volumes

Figure 5.8 gives an overview of the inputs and output of the registration pipeline from Section 5.4.2 using the representative dataset. The fixed and moving images (Figures 5.8a and 5.8b) exhibit evident anatomical misalignment. Although large-scale bony anatomy (patella, femur, tibia) is identifiable in both, the apparent morphology differs substantially because the joint is viewed from an oblique angle in the moving image due to scanner rotation. The registration pipeline successfully aligned the fixed and moving images as evidenced by the close anatomical correspondence between Figures 5.8a and 5.8c.

The difference image Figure 5.8d reveals two distinct sources of intensity mismatch between the fixed and registered images. Firstly, a bright band is visible at the image centre, which corresponds with the ACL in Figures 5.8a and 5.8c. This difference is due to the magic angle intensity variation caused by the reorientation of $\vec{B}_0$ between the fixed and moving images.

Secondly, bright regions are also visible at the periphery of Figure 5.8d. These arose because the scanner's usable imaging volume boundary and its associated geometric distortions co-rotate with the magnet. Consequently, whilst the apparent orientation of the imaged tissues changes between acquisitions, the usable imaging volume boundary retains the same irregular shape in both the fixed and moving images. Following registration, these irregular usable imaging volume boundaries of the fixed and resampled moving images necessarily no longer coincide, producing an intensity mismatch at the image periphery. This effect is particularly noticeable along the right-hand side of Figure 5.8d where the bright region displays the same outline as the image boundary in Figure 5.8b. These peripheral regions also typically fell outside the registration mask (Section 5.4.4) and therefore received no non-rigid correction, resulting in the discrepancies such as the morphological differences visible at the patella in the top left of Figures 5.8a and 5.8c.



(a) Fixed image. The ACL is visible as a dark band interposed between the femur and tibia. The irregular shape of the usable imaging volume boundary is also visible.



(b) Moving image. The anatomy is significantly displaced from the orientation shown in Figure 5.8a. The same shape of the usable imaging volume boundary as in Figure 5.8a is also visible.



(c) Registered image. The moving image (Figure 5.8b) has been aligned by the registration pipeline to the fixed image (Figure 5.8a). The usable imaging volume boundary shown in this slice is a different irregular shape due to registration.



(d) Difference image between fixed and registered images (Figure 5.8a and Figure 5.8c). Bright regions arise from two sources: magic angle dependent intensity variation at the ACL (centre), and usable imaging volume boundary mismatch at the image periphery.

Figure 5.8: Overview of inputs and output of the MADI registration pipeline, shown for sagittal slice 76/150. Figure 5.8a: fixed image ($\theta_R = 0^\circ$, $\phi_Y = 0^\circ$). Figure 5.8b: moving image ($\theta_R = 30^\circ$, $\phi_Y = 10^\circ$). Figure 5.8c: registered image, following alignment of the moving image to the fixed image coordinate frame. Figure 5.8d: voxel-wise intensity difference between the fixed and registered images.

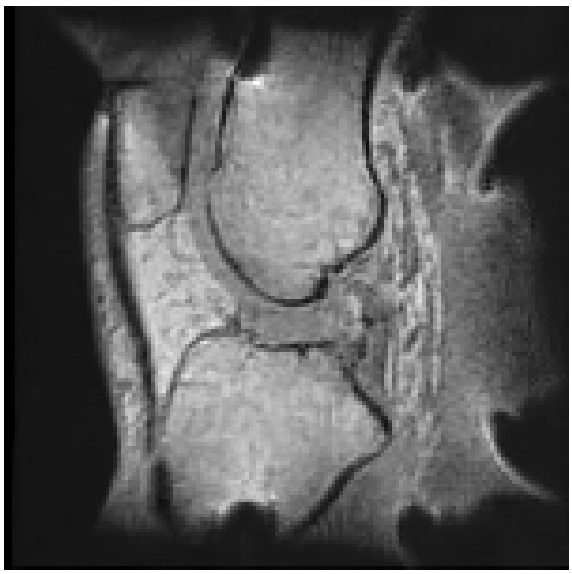
5.4.4 Registration Masking

The registration mask input into elastix is a binary mask with the same dimensions as the input fixed image; for the work in this thesis the dimensions of the mask were therefore $150 \times 150 \times 150$ isotropic volumes. The shape of the binary mask can take any form that fits within this volume, defined in the form of NIfTI (`.nii`) files where values of 1 mark regions to include for domain sampling and 0 mark regions to exclude.

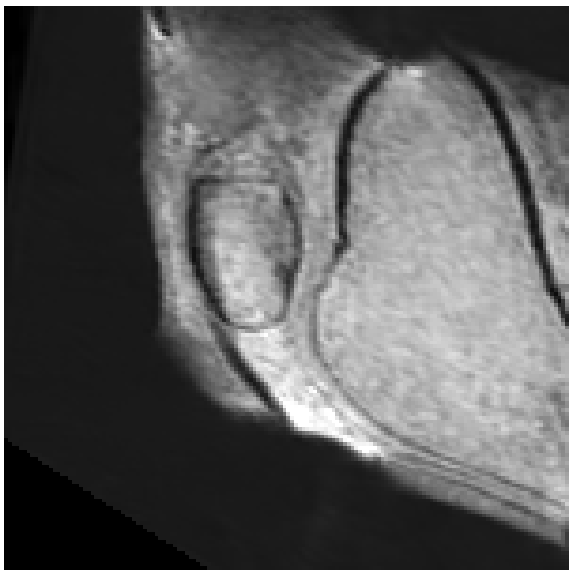
The mask can be generated by any standard segmentation software capable of exporting binary masks, such as 3D Slicer [109]. However, because the scanner rotates about a static patient which in turn co-rotates the irregular magnetic field boundary within the image domain (as shown in Figure 5.8), the most suitable registration mask is often simply a sphere. A spherical mask centred within the image ensures that sampled features will be present across an entire MADI dataset when sized appropriately. The default spherical mask generated for the MADI registration pipeline was of radius 90 mm as this was found to encompass the ACL and meniscus consistently without including artefactual features such as the magnetic field boundary.

5.4.5 Failure Detection

If a patient repositioned themselves between scans to adjust their seating position, this could result in an axial shift of their leg within the receive coil. Important features for successful registration could therefore be moved outside the field of view (FOV) of the scanner. Significant movement would cause transform parameter optimisation to fail with no convergence, resulting in no registered image output. The registrations were typically performed during the acquisition of the next $\vec{B}_0$ orientation of a MADI exam, as the pipeline allows for batch processing. A successful registration became an indication of appropriate patient positioning in the MADI workflow, without necessarily checking the registered image itself, and a failed registration would indicate that the previous orientation would need repeating due to inadequate patient positioning. Occasionally, elastix would instead converge to a local minimum, yielding a spurious registered image result as shown in Figure 5.9. This caused false positives of appropriate patient positioning, only discovered when later processing of the MADI dataset for collagen fibre orientation estimation failed due to the inclusion of the spurious registration result.



(a) Fixed image with correct patient positioning.



(b) Spurious registered image due to incorrect patient positioning.

Figure 5.9: Example of subtle registration failure. Registration of the moving image to Figure 5.9a has converged to a local minimum, yielding obviously incorrect alignment. This occasionally led to a false positive assumption of a completely successful MADI examination.

To prevent this failure mode the operator was automatically presented with an obvious cue to repeat a non-convergent acquisition. Structural similarity index measure (SSIM) checks were implemented between each registered image and the fixed image within the registration masked volume. SSIM < 0.7 triggered a warning in the MADI software that prompted the operator to check the registration and repeat the offending acquisition after adjusting patient positioning.

5.5 Distortion Compensation

Distortion in MRI arises when the actual spatial encoding does not match the assumed linear mapping between frequency and spatial domains, typically due to non-linearities in the main $\vec{B}_0$ field or in the applied gradients. In low-field MRI this problem is often exacerbated by weaker fields, reduced shimming capability, and lower magnetic field homogeneity within the DSV. Methods for distortion compensation include passive shimming (see Section 4.3.5); active shimming using dedicated shim coils; sequence-based estimation of distortions with corrective reconstruction; direct mapping of the $\vec{B}_0$ field to retrospectively correct images; and post-processing methods such as registration of calibration scans to a reference dataset.

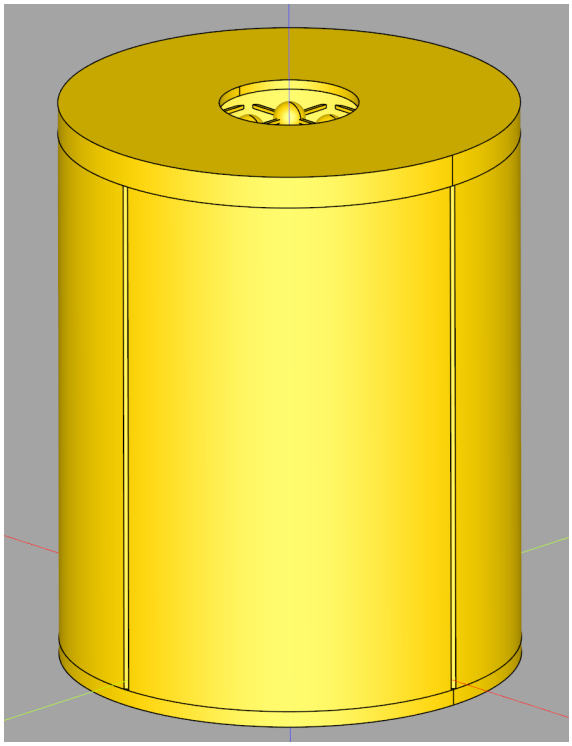
This final approach was adopted for the work in this thesis, as the necessary registration tools were already available within the core MADI pipeline. To implement it, a 3D-printable phantom was designed with a regular internal grid structure. Simulated scans of this phantom were generated to provide a ground truth image, and real scans of the 3D-printed phantom filled with vegetable oil were then acquired. These were registered non-rigidly to the simulated reference, and the resulting transform parameters were stored. Because the geometric distortions corrected by this procedure are systematic to the inherent $\vec{B}_0$ inhomogeneity of the rotatable low-field system, the same transform could subsequently be applied to any scan acquired with the same pulse sequence.

5.5.1 3D Printed Calibration Phantom

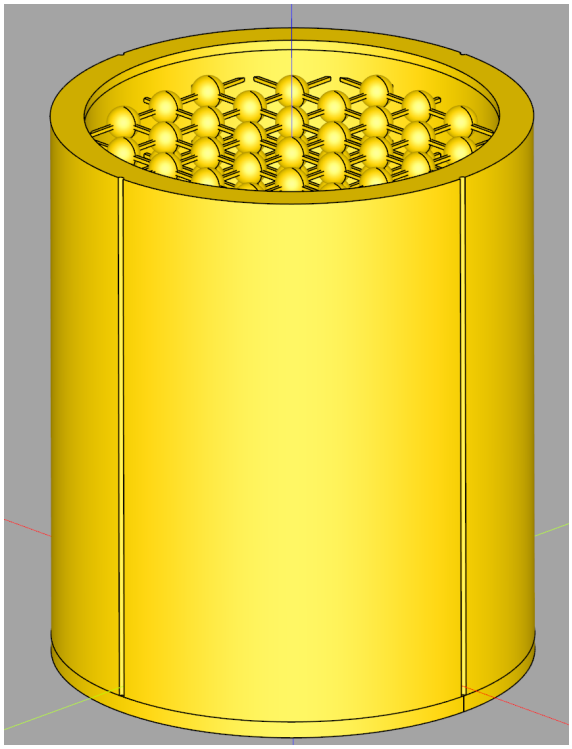
The calibration phantom was designed using the CadQuery [110] module in Python. This facilitated design of the phantom and its regular internal structure using conventional programming paradigms and allowed for rapid design iteration by tuning input variables. Figure 5.10 shows the CAD model within the CadQuery software.

The phantom was designed with 9 layers with a vertical separation between sphere centres of 20 mm. The horizontal separation of sphere centres was also 20 mm. All spheres had a radius of 5 mm, apart from the larger sphere at the centre of the grid which had a radius of 7.5 mm. The cylindrical body of the phantom had an outer radius of 142 mm, and a wall thickness of 12 mm. These dimensions were chosen such that the phantom would fit snugly inside the receive coil which had inner diameter of 143 mm, ensuring that the large central sphere was able to be positioned as close as possible to the centre of the DSV. The rest of the dimensions were chosen such that the grid structure would entirely fill the imaging volume, thus allowing all distortion to be compensated for. The outer wall of the phantom had a debossed notch to allow for square alignment of the grid within the scanner when the phantom was filled with vegetable oil and sealed. Sample code for the calibration phantom generation is available in Section C.2. Figure 5.11 shows the calibration phantom partially glued and also fully assembled, filled, ready to be scanned.

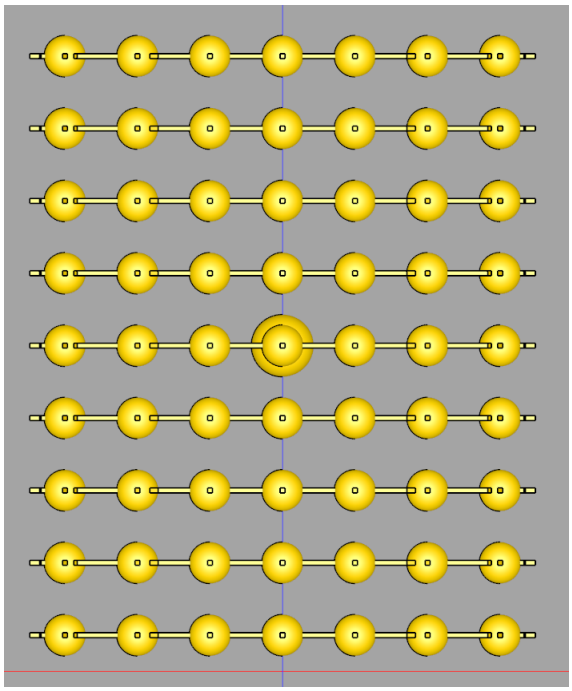
The phantom was 3D printed on a Prusa MK3S (Prague, Czech Republic) using PLA and 0.2 mm layers. It was printed in sections that bisected each layer to ensure accurate printing of the spheres with no overhangs. Each half layer was then glued together



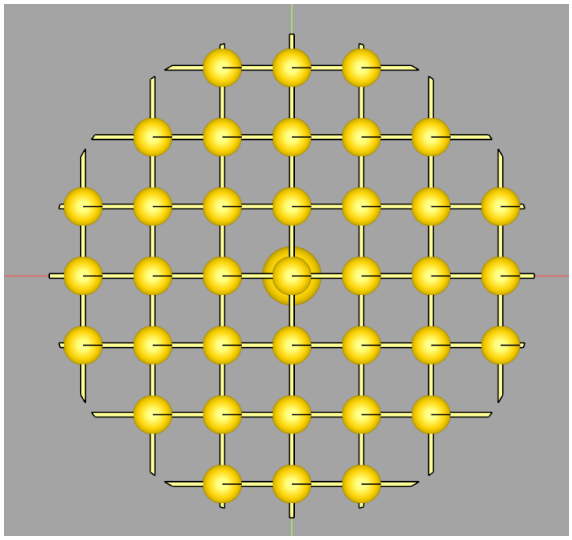
(a) Full outer body of phantom with internal structure visible through top filling hole.



(b) Full outer body of phantom with top section removed to show internal structure clearly.



(c) Side view of internal structure showing regular layers of spheres with a larger sphere in central layer.

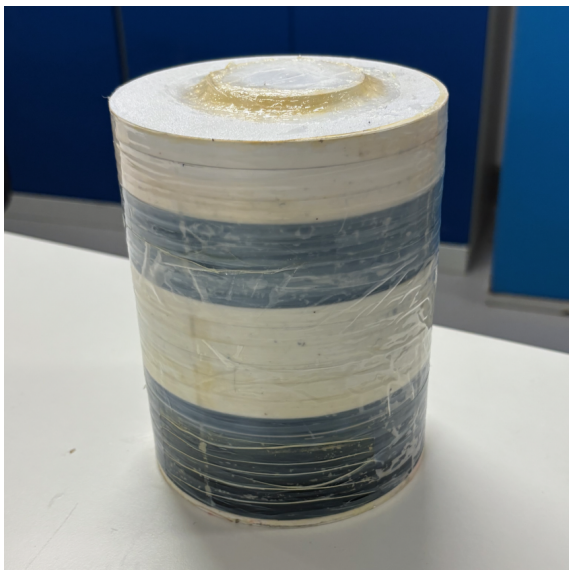


(d) Plan view of internal structure showing regular grid arrangement of spheres with a larger sphere in grid centre.

Figure 5.10: CadQuery model of 3D distortion calibration phantom. The regular internal grid structure provides abundant features for registration and the larger central sphere in the middle of the phantom helps to centre the phantom in the scanner. External notches in the outer body of the phantom provide alignment aids to ensure the grid structure is square within the scanner.



(a) Partially glued layers of the calibration phantom showing the regular internal structure within.



(b) Complete 3D printed calibration phantom filled with vegetable oil and sealed with 3D printed lid and silicone sealant.

Figure 5.11: 3D printed calibration phantom for registration based 3D distortion compensation. Printed in PLA on a Prusa MK3S using 0.2 mm layer heights. Figure 5.11a shows the phantom partially constructed and Figure 5.11b shows the completed, filled, and sealed phantom.

with alignment ensured by the outer notches. The inside of the phantom was coated with a light clear lacquer to prevent leaks through the layered structure of the 3D print, and the outer surface of the phantom was covered in clear tape to both prevent leaks and increase the outer diameter to ≈ 143mm. The phantom was filled completely with vegetable oil and sealed with a 3D printed lid and silicone sealant.

5.5.2 Distortion Correction using Non-rigid Registration

CadQuery was chosen to design the calibration phantom as it allowed for development of a custom scan simulator using Python scripting and CadQuery's boolean operations. To simulate a scan of the CAD model, a $1 \times 1 \times 1$mm voxel was created at the same voxel coordinates as in a regular volume image from the low-field scanner, and boolean operations within CadQuery were used to subtract any overlapping volume between the voxel and the CAD model. The volume of the resulting body was then recorded as an intensity into a simulated image. A volume of 1 represented no overlap (i.e. the voxel was entirely within the vegetable oil), and a volume of 0 represented complete overlap (i.e. the voxel was entirely within the CAD model). Repeating this process for every valid voxel coordinate generated the final simulated image. Figure 5.12 shows an

example of such a boolean operation. The yellow sphere and the green cube overlap, the subtractive boolean operation of the sphere from the cube yields the purple shape. This purple shape has volume V in the range $0 < V < 1$, meaning it would map to a gray intensity in a simulated image with intensity range $0 \dots 1$.

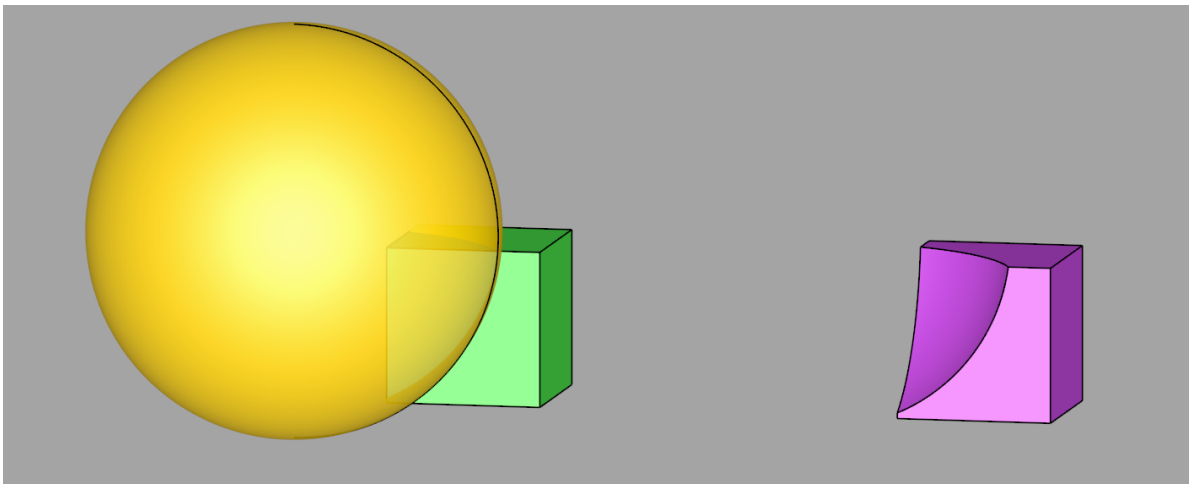
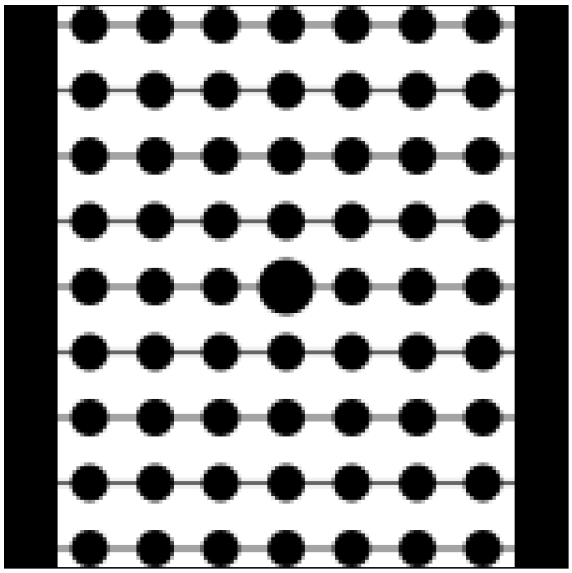


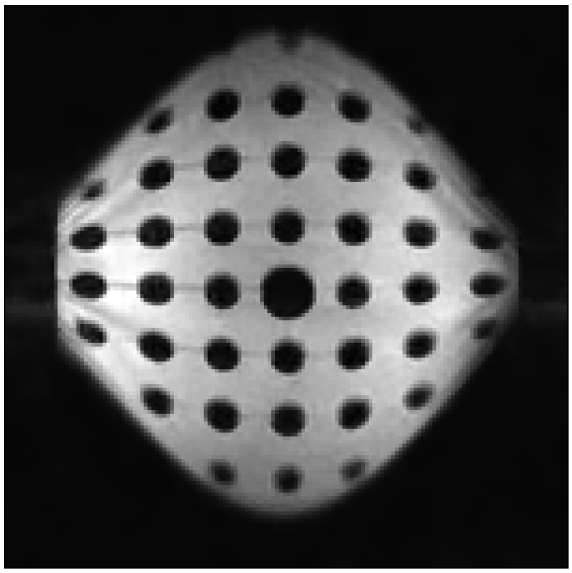
Figure 5.12: Example of a subtractive boolean operation between a part of the calibration phantom CAD model (yellow sphere) and a simulated voxel (green cube), resulting in the new purple shape with reduced internal volume compared to the original cube.

Figure 5.13 shows the central sagittal slice (76/150) of the simulated scan of the CAD model calibration phantom (Figure 5.13a), the real (distorted) scan of the vegetable oil filled phantom (Figure 5.13b), the non-rigid registered phantom scan to the simulated scan (Figure 5.13c), and the difference image between the original and registered phantom scan (Figure 5.13d).

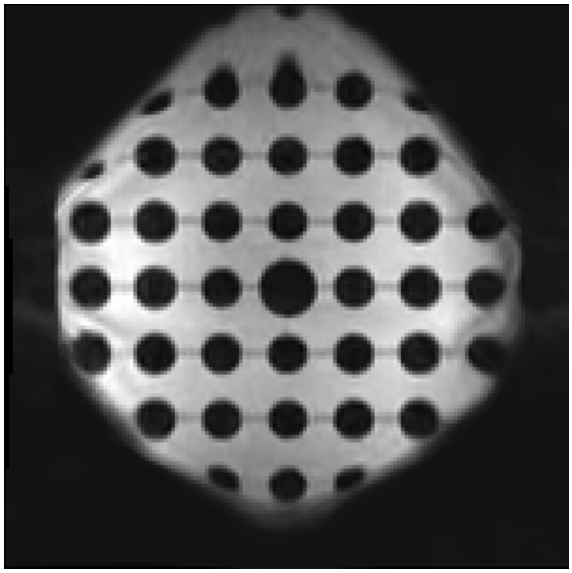
Figure 5.13b shows significant geometric distortion of the calibration phantom, especially towards the boundary of the homogeneous magnetic field. Figure 5.13b was non-rigidly registered to Figure 5.13a using the same non-rigid transform as described in Figure 5.7. The registration result, Figure 5.13c, shows significantly less geometric distortion with a regular grid structure of spheres clearly visible. The non-rigid transform corrects for out of plane distortion as well as can be seen by the recovery of the struts spanning the sphere grid. These struts fade in and out of plane in Figure 5.13b but remain in the full span of the image in Figure 5.13c. The difference image, Figure 5.13d, highlights the extent of the distortions across the whole original image. Significant differences are seen towards the edge of the DSV as previously mentioned, yet it can also be seen that the centre of the image is either scaled down or warped out of plane as the large central spheres have different diameters.



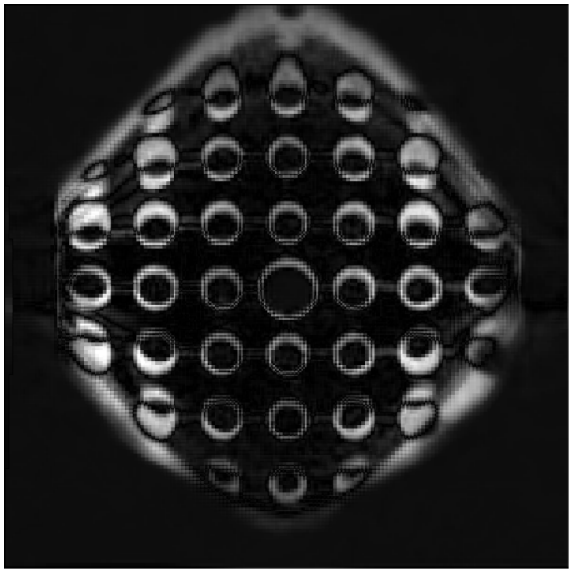
(a) Simulated scan of the calibration phantom using CadQuery.



(b) Real scan of the vegetable oil filled calibration phantom.



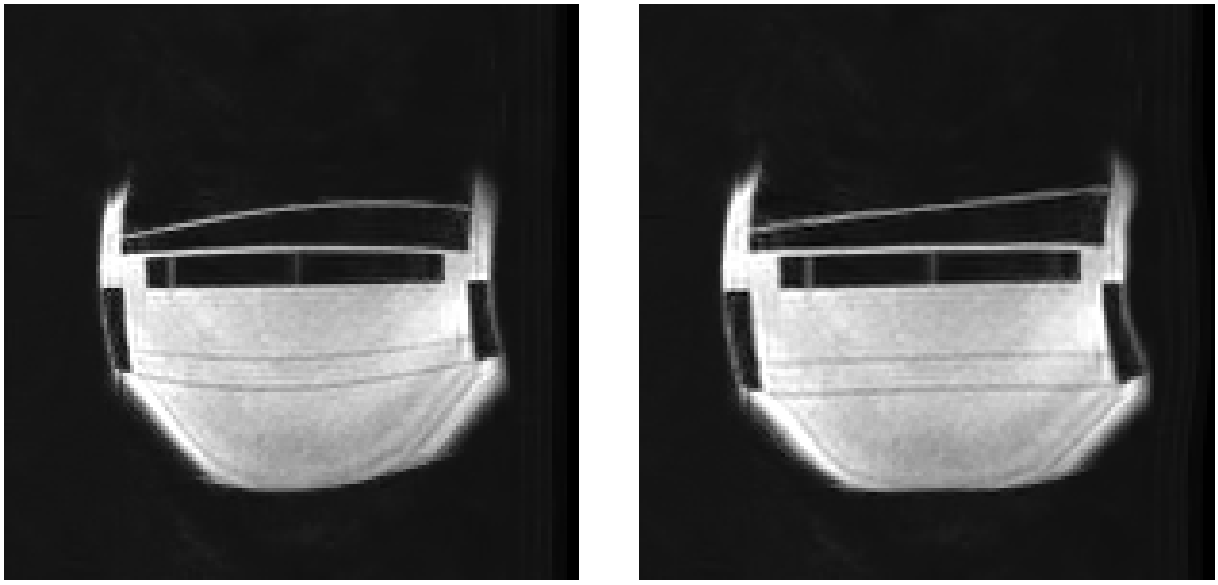
(c) Registered image result of Figure 5.13b non-rigidly registered to Figure 5.13a.



(d) Difference image between Figure 5.13b and Figure 5.13c.

Figure 5.13: Sagittal images (slice 76/150) of the simulated scan of the calibration phantom (Figure 5.13a), real scan of the 3D printed calibration phantom (Figure 5.13b), the registered image of the calibration phantom scan registered to the fixed image (Figure 5.13c), and the difference image between the original calibration phantom scan and the registered calibration phantom scan (Figure 5.13d).

From the registration of the real phantom scan onto the simulated phantom scan, the optimised transform parameters can be applied to any subsequent scan that utilises the same pulse sequence. This can either be done as part of transform composition in the registration pipeline or as an isolated operation on an image using transformix [100], a companion tool to elastix that applies an input transform parameters file to an input image. Figure 5.14 illustrates the application of the stored non-rigid transform parameters to a scan of the American College of Radiology (ACR) standard phantom, acquired with the low-field scanner.



(a) ACR scan with no non-rigid correction transform applied.

(b) ACR scan with non-rigid correction transform applied.

Figure 5.14: ACR phantom scan (axial slice 60/150) before (Figure 5.14a) and after (Figure 5.14b) non-rigid correction transform was applied.

The efficacy of the non-rigid transform is demonstrated by the straightening of the curved features from Figure 5.14a to Figure 5.14b.

5.6 Summary

This chapter presented a complete overview of the registration techniques and frameworks used in the MADI workflow.

Section 5.2 introduced the concept of image registration, its necessity for MADI, and how it had been applied in earlier implementations of the technique. However, these previous registration methods were inappropriate for MADI at low-field and in vivo and thus needed updating to a fully automatic, telemetry-driven, non-rigid registration implementation.

Section 5.2.1 outlined the principles of parametric registration, introducing the counter-intuitive definition of a transform from the fixed to the moving image that is then applied to resample the moving image.

Section 5.3 described the registration framework elastix, which implements this paradigm and was adopted as a suitable and flexible platform for developing the MADI registration pipeline. The various components of elastix were detailed sufficiently to show how they were adapted to MADI and combined into a fully automatic pipeline for entire MADI datasets, with non-rigid registration providing the fine-grained alignment necessary for subsequent segmentation and collagen fibre orientation estimation.

Section 5.4 then described the specific registration applications within MADI, from affine correction of excitation drift due to magnet heating through to the implementation of the complete automatic pipeline.

Finally, Section 5.5 described the extension of non-rigid registration to distortion compensation, using a custom 3D printed calibration phantom and simulated ground-truth scans. This approach proved highly effective, requiring no specialised equipment beyond widely available 3D printing technology, whilst reusing the registration framework already established for MADI. Such distortion compensation will be essential if MADI data are to be applied in contexts such as surgical planning, where geometric accuracy is of critical importance.

6 | Segmentation

6.1 Chapter Overview

This chapter describes the segmentation techniques and framework developed and applied for MADI in this thesis. Section 6.2 introduces the concept of metric-assisted segmentation, which leverages the intrinsic signal variations of $\vec{B}_0$ orientation-sensitive tissues to guide segmentation. Section 6.3 then details the custom segmentation GUI built on this concept, which enables regions of interest to be defined using a bounding box in combination with metric assistance, thereby facilitating precise and data-driven segmentation of the tissues of interest.

In medical imaging, segmentation is the process of dividing an image into regions that correspond to specific anatomical structures or tissues of interest. It is commonly used to isolate regions of interest for further analysis or to generate labelled regions for morphological assessment. In the context of MADI, segmentation is applied to each image in a registered dataset to extract an identical ROI from each of them, across which intensity variations can be analysed for collagen fibre orientation estimation. In conventional workflows, segmentation is performed manually, requiring the operator to delineate ROIs slice by slice, which is time-consuming and labour-intensive. Metric-assisted segmentation, as presented here, reduces this burden by using data-driven analysis rather than human interpretation, enabling large ROIs to be segmented rapidly and effectively.

6.2 Metric-Assisted Segmentation

By the segmentation stage of the MADI pipeline (Figure 4.1), the $N-1$ moving images acquired at rotated scanner orientations have been registered to the fixed image, as described in Section 5.4.2. These N 3D volumes are collated into a 4D image array, where the first three dimensions represent spatial coordinates and the fourth indexes the image set. Intensity data can therefore be analysed at each spatial coordinate along this fourth dimension, allowing standard signal and statistical metrics such as minimum, mean, maximum, range, and standard deviation to be computed. Of particular interest are metrics that quantify variation, as the MADI technique is fundamentally based on intensity changes induced by the magic angle effect. In particular, the standard deviation across the dataset serves as a sensitive marker for identifying tissues that exhibit the magic angle effect.

Figure 6.1 illustrates the formulation of metric-assisted segmentation with a simplified example dataset. A hypothetical set of $N = 5$ registered images is shown, where $I_{R,n=i}$ for $i = 2 \dots 5$ are the resulting images from registration of their respective moving images to $I_{F,n=1}$. Collagen-rich tissues display intensity variation across the dataset, whereas orientation-insensitive tissues such as bone or muscle exhibit consistent intensity. A single slice from each volume image is presented, with green and blue arrows marking corresponding spatial coordinates across the dataset. The standard deviation of intensities measured at the green-arrow locations is significantly higher than at the blue-arrow locations, reflecting the orientation sensitivity of collagen. Standard deviation therefore serves as a useful metric for identifying tissues exhibiting the magic angle effect in a MADI dataset. Repeating this analysis across all voxel locations in the fixed image produces a three-dimensional standard deviation volume, which can then be used as a guide to segment the ROI for the entire MADI dataset.

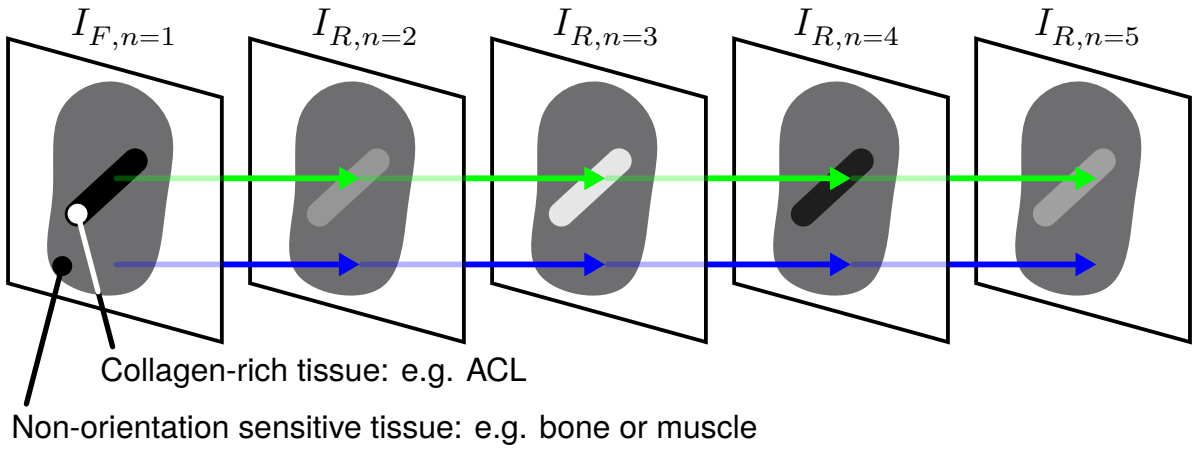


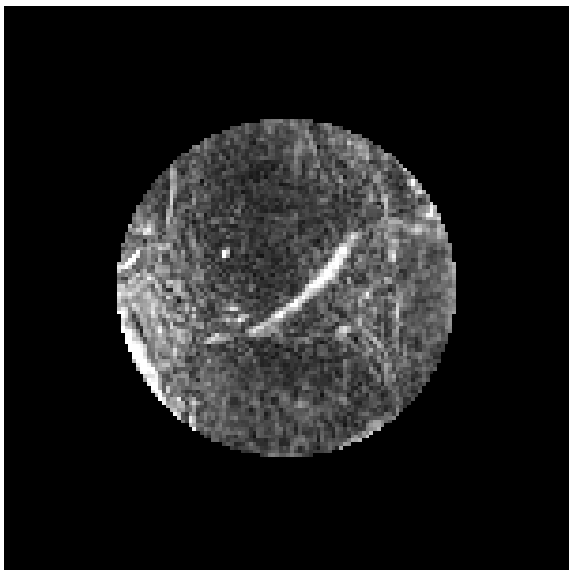
Figure 6.1: Simple example of intensity variations across registered images due to the magic angle effect. 2D slices from a set of $N = 5$ hypothetical registered images are shown; the intensity variations measured along the green arrow through corresponding collagen-rich regions would yield a higher standard deviation than intensities measured along the blue arrow through bone or muscle for example. This standard deviation metric could then serve as a segmentation guide.

6.2.1 Standard Deviation Metric

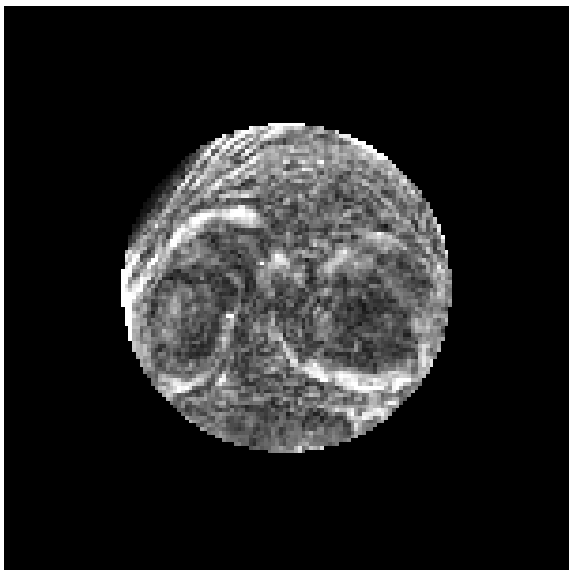
Figure 6.2 shows the standard deviation image volumes for the representative dataset. In Figure 6.2a, the ACL appears bright, reflecting its high standard deviation across the dataset, consistent with the subtraction image shown in Figure 5.8d in Section 5.4.3. Similarly, the meniscus appears as a bright figure-of-eight structure in Figure 6.2b, highlighting its orientation sensitivity.

The full representative dataset is presented, along with derived metric images, in Section D.1.

Previous studies have also employed standard deviation as a metric for segmenting $\vec{B}_0$ orientation-sensitive tissues [14, 15]. Earlier implementations of MADI [14] used a relatively simple segmentation workflow, based on a non-rotatable bounding box applied directly to the raw standard deviation volume, and controlled through a text-based terminal interface. In contrast, the framework presented in Section 6.3 represents a substantial evolution. It enables finer user control of bounding box translation and rotation, introduces more advanced thresholding parameters, and improves usability through interactivity and real-time visual feedback, allowing iterative refinement of the segmentation. This extensible platform for metric-assisted segmentation enhances both the precision and the practicality of the MADI technique.



(a) Standard deviation image highlighting the ACL.



(b) Standard deviation image highlighting the meniscus.

Figure 6.2: Standard deviation images from the representative dataset showing bright ACL (Figure 6.2a) and meniscus (Figure 6.2b) due to magic angle effect intensity variations across the dataset. The registration mask has been applied to these images to exclude regions of high standard deviation due to registration artefacts, including overlapping magnetic field boundaries from unregistered peripheral regions, which produce artificially elevated standard deviation at the volume periphery.

6.3 Segmentation Graphical User Interface

Figure 6.3 illustrates the segmentation GUI developed for metric-assisted segmentation in this thesis. Figure 6.3a shows the software in use for segmenting the ACL from the representative dataset. The left panel displays a central sagittal slice of the fixed image (slice 76/150) with a bounding box defining the ROI limits, while the right panel presents a 3D point cloud of voxel coordinates that satisfy the segmentation criteria. Figure 6.3b shows the same setup with the bounding box hidden and the axial and coronal slices displayed along with the sagittal slice. Voxels meeting the segmentation criteria are highlighted in white and overlaid on the three-axis slice viewer, clearly marking the position of the ACL within the fixed image.

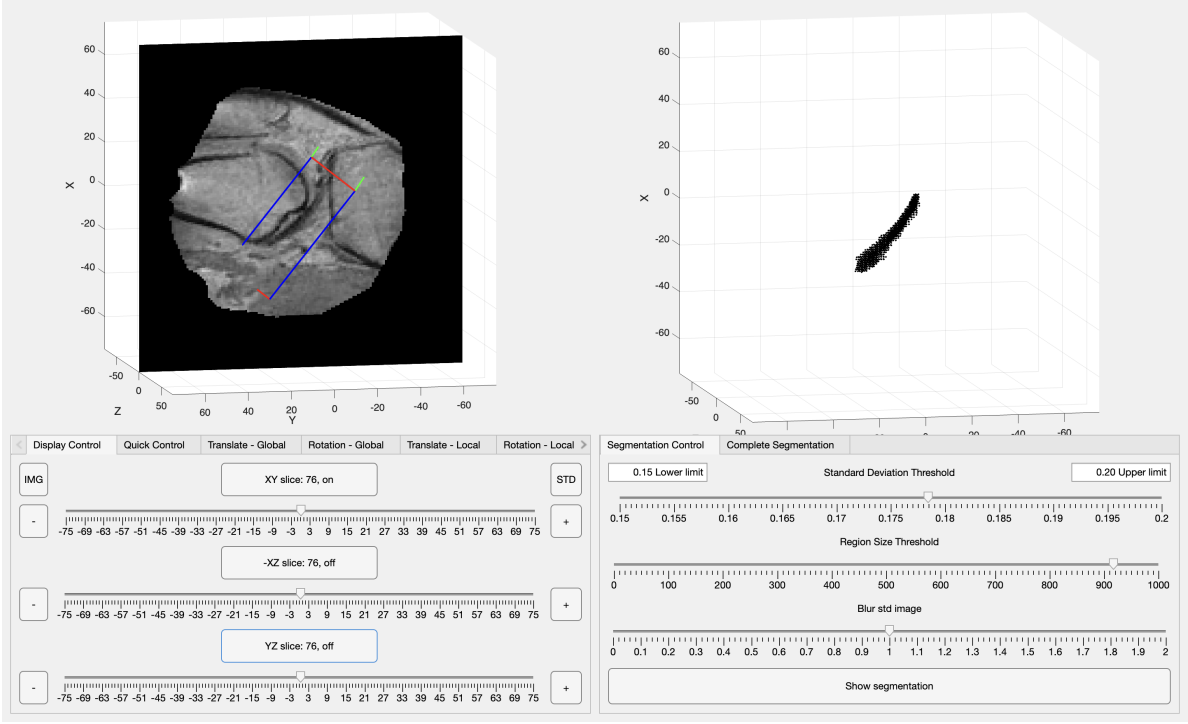
The left panel of the interactive segmentation GUI shown in Figure 6.3 allows precise control of a 3D bounding box within the image domain. The bounding box can be manipulated in both the global and local reference frames: global translations and rotations are applied relative to the basis vectors of the global coordinate system defined in Section 4.3.6, whereas local translations and rotations are applied relative to the basis

vectors defining the three orthogonal axes of the bounding box itself. Its scale can also be independently adjusted along each axis. To facilitate rapid bounding box placement, several quick control options are provided. These include centroid placement, where a user click on a slice immediately sets the bounding box centroid in three dimensions, and face placement, where two bounding box opposing face centroids can be placed to quickly delineate an ROI at an arbitrary orientation. Figure 6.3a illustrates bounding box placement for the ACL, where the user selected both ends of the ligament using face placement to efficiently position the bounding box. This approach enables the user to encompass the anatomy of interest as an initial rough segmentation, which can then be refined using the standard deviation metric.

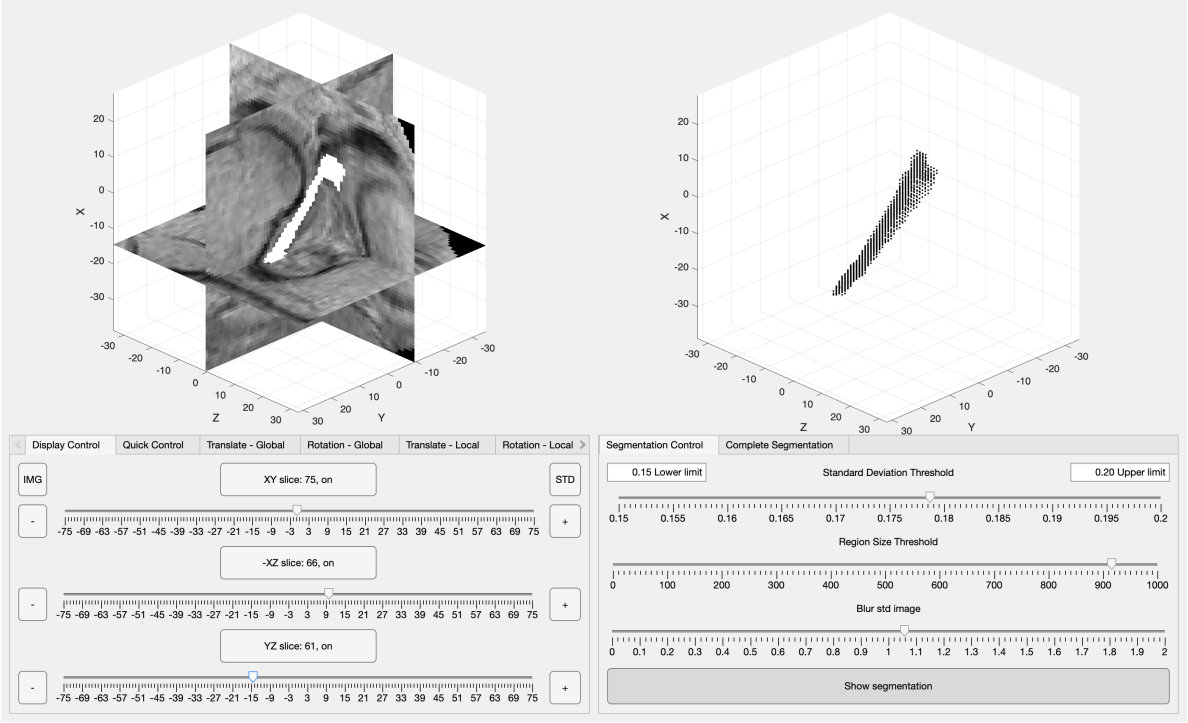
The right panel of the segmentation GUI shown in Figure 6.3 enables fine-tuning of tissue segmentation within the bounding box using standard deviation as a metric. A threshold is applied to the standard deviation volume (as shown in Figure 6.2), normalised to a dynamic range of 0–1, and is adjusted with a slider of controllable resolution[1]. Voxels above the selected threshold are marked and visualised as a 3D point cloud representing the segmentation result. Each point is placed at its associated voxel coordinate, at the centre of each voxel. Additional refinement is provided by the region size threshold and blur factor controls. The region size threshold specifies a minimum allowable region size, with voxels grouped according to face connectivity. This suppresses isolated single-voxel inclusions caused by noise, as well as small off-target regions (portions of the PCL for example) that may fall within the bounding box. The blur factor sets the σ of a 3D Gaussian blur applied to the standard deviation volume: increasing the blur reduces spurious small regions of high standard deviation, whereas decreasing it sharpens the boundaries of the segmented region.

The user can also switch to view the standard deviation volume, instead of the fixed image, which can assist in bounding box placement and threshold selection, as the metric image highlights $\vec{B}_0$ orientation-sensitive tissues of interest as shown in Figure 6.2. Once segmentation is complete, the software generates a binary image that can be exported and further refined in external medical imaging software such as 3D Slicer [109], allowing manual addition or removal of segmented regions if required.

[1]The slider's upper and lower limits can be changed to increase the resolution of the slider graduations.



(a) Segmentation GUI showing bounding box and 3D point cloud of segmented voxels for the ACL.



(b) Segmentation GUI showing highlighted ACL segmentation on three-axis slice image viewer (zoomed in to joint centre), and 3D point cloud of segmented voxels for the ACL.

Figure 6.3: Demonstration of metric-assisted segmentation software. Figure 6.3a shows the bounding box and standard deviation threshold to segment the ACL from the representative dataset, and Figure 6.3b shows the ACL segmentation highlighted on the fixed image of the representative dataset within the 3D slice viewer.

6.3.1 Manual Segmentation

The MADI pipeline represents all masks and segmentations internally as binary images in the NIfTI (`.nii`) format. Consequently, any software capable of importing, editing, and exporting NIfTI files can be used to generate or modify masks and segmentations for MADI. For example, 3D Slicer [109] may be used to produce a fully manual segmentation if required. A more efficient workflow, however, is to begin with a slightly liberal metric-assisted segmentation using the software presented here, and subsequently refine the result in 3D Slicer. This ensures that the segmented voxels exhibit sufficient intensity variation for accurate collagen fibre orientation estimation, whilst alleviating the operator burden of performing a fully manual, slice-by-slice segmentation.

6.4 Summary

This chapter has provided an overview of the segmentation paradigm and framework developed for the MADI workflow. Section 6.2 introduced the concept of metric-assisted segmentation and demonstrated the use of standard deviation as an appropriate metric, computed across the entire MADI dataset, to aid segmentation of $\vec{B}_0$ orientation-sensitive tissues. Section 6.3 described the custom GUI developed in this thesis, which enables rapid delineation of a rough ROI using a flexible bounding box, and its subsequent refinement using the standard deviation metric. The use of standardised file formats ensures compatibility with conventional imaging software, allowing further refinement of metric-assisted segmentations where required.

Recent advances in AI- and computer vision-based segmentation [111–113] demonstrate the potential of these methods, although their reliance on large training datasets and the need for scanner- and sequence-specific adaptation remain limitations. Metric-assisted segmentation, being data-driven rather than operator-driven, offers a means of generating reliable ground truth or labelled datasets to support such algorithms. Furthermore, future integration of metric-assisted segmentation with AI-based methods could provide a hybrid workflow, where the metric serves as a guide for the algorithms and the AI refines the result, thereby achieving robust, precise, and data-driven segmentation in a single package.

7 | Collagen Fibre Orientation Estimation

7.1 Chapter Overview

This chapter addresses the fundamental experimental design questions underpinning the MADI technique: how many scans are required to achieve accurate estimation of collagen fibre orientations? and, what should these scan orientations be? These questions are approached by first outlining the theoretical basis of the collagen fibre orientation estimation algorithm (CFOEA) in Section 7.2, tracing its development and presenting the updated implementation developed in this thesis. Collagen fibre estimation results for the ACL from the representative dataset are also presented. Finally, in Section 7.3, the updated CFOEA is used to produce simulated collagen fibre estimation accuracy maps for any input set of scan orientations. These accuracy metrics are incorporated into a fitness function for a genetic algorithm that generates scan orientation sets designed to maximise the theoretical accuracy of collagen fibre orientation estimation subject to the practical constraints of MADI.

7.2 Collagen Fibre Orientation Estimation Algorithm

In all multi-orientation MRI techniques that aim to exploit the magic angle effect, a common data structure emerges. As discussed in Chapter 6, a MADI dataset consists of a fixed image to which all subsequent moving images (obtained either through sample reorientation or scanner rotation) are registered. Collating these images forms a four dimensional dataset, where the first three dimensions represent spatial coordinates and the fourth indexes the individual scans.

The 3D ROI is typically defined on the fixed image, and then applied to all aligned scans, meaning the extracted data are also inherently four dimensional. Each voxel *location* in

the ROI, i.e. the coordinate associated with each voxel centre, indexes an N-element intensity profile at that spatial coordinate, across the 4th dimension of the dataset. An example of such an intensity profile could be from the green arrow path through the dataset in Figure 6.1.

7.2.1 Previous Algorithms

Two methods for collagen fibre orientation estimation are prominent in prior studies as outlined in Section 1.3. A summary of these studies is provided here with a focus on their choice of scan orientations and collagen fibre orientation estimation algorithms.

Szeverenyi et al. [11] investigated collagen fibre orientation in an excised caprine meniscus, which they scanned at 15 different orientations by rotating the sample in an epoxy sphere within the static B_0 field of a 3T clinical MRI scanner. Fibre orientations were estimated using a simulation-based approach: a bouquet of 900 candidate fibre orientations spread over a hemisphere was generated, yielding an angular resolution of $\approx 5^\circ$. Intensity profiles for each candidate fibre orientation were computed under the same 15 scan orientations as used in the real acquisitions. The intensities were simulated using the orientation-intensity relationship in Equation 2.7. The measured intensity profiles were correlated with the simulated profiles associated with the candidate fibre orientations. The candidate fibre orientation with the highest correlation was then assigned to each voxel in the ROI. This yielded fibre orientation estimations consistent with the expected meniscal microstructure, with fibres aligning tangentially to the meniscus. For their study, Szeverenyi et al. distributed the 15 scan orientations at 30° intervals, constrained to three orthogonal planes. This configuration was chosen based on the heuristic assumption that such a distribution would induce magic angle effect intensity variation in any arbitrarily oriented collagen fibres within the meniscus.

Seidel et al. [10] also used a simulation based approach in their study of the collagen fibre orientations in an excised human iliotibial tract. Similarly to Szeverenyi et al., they held their sample in a jig within a 3T conventional scanner and rotated it within the static B_0 field of the scanner. The rotation scheme Seidel et al. used was to scan the samples at 10° intervals, constrained to two orthogonal planes, totalling 35 scans. This was done to precisely measure the intensity response over a dense set of scan orientations. The maxima of the measured intensity profiles were compared with theoretical profiles

generated using Equation 2.7 to constrain the candidate collagen fibre orientations to at most 4 orientations. The simulated intensity profiles and the profile gradients were then compared to the measured profiles and gradients using either sum of absolute differences, sum of squared differences, or cross-correlation to find the best fitting candidate orientation for each voxel in the ROI.

Performance-wise, Seidel et al. [10] reported a calculation time of ≈ 20 minutes for 332,593 voxels on a 2.1 GHz processor, representing a substantial improvement over the ≈ 60 minutes required for 100,000 voxels on a 3 GHz processor, as reported by Szeverenyi et al. [11]. This improvement arose from reducing the search space to a maximum of four candidate directions per voxel location, compared with the 900 directions evaluated per ROI voxel in Szeverenyi et al.'s method. However, these computational gains are significantly offset by the much larger number of scans required for Seidel et al.'s approach, which limits its potential clinical applicability.

As noted in Section 1.3, these studies served as useful proofs of concept for demonstrating the utility of the magic angle effect as a source of quantitative information on collagen microstructure using MRI. However, both relied on a large number of scan orientations, making them impractical for in vivo application due to prohibitively long examination times, even at high field strength. Nonetheless, our research group adopted the philosophy of Szeverenyi et al. [11], whose approach required fewer scans and offered the potential for further reduction, making it a more feasible foundation for development towards clinical translation.

To this end, Chappell et al. [14] and Brujic et al. [15] adapted the method of Szeverenyi et al. to produce the first iteration of the MADI technique. Their approach differed in two principal ways. Firstly, scan orientations were distributed evenly across a hemisphere rather than being restricted to orthogonal planes; as with Szeverenyi et al., this distribution was based on a heuristic intended to generate magic angle effect intensity variation in arbitrarily oriented collagen fibres, but with the potential to reduce the total number of scans required. Secondly, fewer candidate fibre orientations were simulated, but resolution was improved by refining the initial bouquet-based estimates using a simplex optimisation scheme to obtain the final fibre orientation estimation. In their experiments, sets of 5, 7, 9, 13, and 15 scan orientations were trialled, with 250-500 candidate fibre directions forming the initial bouquet. They concluded that 7 or 9 scans were sufficient for

accurate estimation of collagen fibre orientation, and that larger bouquets of candidate directions improved convergence at the cost of longer computation times.

Brujic et al. [15] also evaluated their implementation of the MADI technique using Monte Carlo simulations with added noise. The "measured" intensity profile of a randomly oriented fibre was simulated under different sets of scan orientations (as above), and synthetic noise was introduced. Promising performance was observed with as few as 7 or 9 acquisitions, even at SNR values as low as 6, with highly consistent results for SNR > 10. However, the simplex optimisation scheme used by Chappell et al. and Brujic et al. to reduce the number of orientations in the candidate bouquet required optimisation of a cost function, defined as the sum of squared differences between measured and simulated intensity profiles. This approach demanded careful manual tuning of the fitting coefficients A and B in Equation 2.7 to achieve stable convergence. Furthermore, the reported performance gains relative to Szeverenyi et al. [11] were, at best, linearly proportional [15] to the reduction in the number of candidate fibres, as the underlying algorithm remained essentially the same, with the simplex optimisation step simply allowing for a more sparse candidate fibre bouquet for the initial correlation step.

This body of previous work does not sufficiently address the practical constraints of multi-orientation imaging in vivo. Whilst the reported low SNR findings provided motivation for pursuing low-field MRI as a platform for MADI, the choice of scan orientations in earlier studies did not consider the ergonomic and hardware limitations inherent to in vivo implementation. In particular, yaw rotations on the low-field rotatable scanner are constrained both by hardware and ergonomic factors as described in Section 4.4: signal attenuation occurs when $\vec{B}_0$ is no longer perpendicular to the receive coil axis, and at certain yaw angles the magnet poles may physically interfere with the patient's legs.

To address these limitations, an updated implementation of the CFOEA was developed, as described in Section 7.2.2. The substantially reduced execution time of the updated CFOEA enabled prospective evaluation of candidate scan orientation sets using simulation, and facilitated non-heuristic orientation selection via a genetic algorithm as described in Section 7.3.3. Moreover, the optimisation step employed by Chappell et al. [14] and Brujic et al. [15] could be replaced with a very dense hemisphere of candidate fibre directions, whilst still achieving faster runtimes than any previous fibre estimation approach.

7.2.2 Updated Collagen Fibre Orientation Estimation Algorithm

The CFOEA, originally developed by Szeverenyi et al. [11] and subsequently refined by Chappell et al. [14] and Brujic et al. [15], is designed to reconstruct 3D fibre orientation information from inherently 1D intensity data. The angle (a 1D scalar) between two lines in 3D determines the intensity expressed by $\vec{B}_0$ orientation sensitive tissues in MRI. The forward model of intensity variation is given by Equation 2.7, which is restated here for convenience.

$$I(\theta) = A \cdot \exp\left(-B \cdot \left(3\cos^2\theta - 1\right)^2\right) \tag{2.7}$$

Where I is the orientation-sensitive intensity, A and B are fitting coefficients, and θ is the angle between $\vec{B}_0$ and the collagen fibres being scanned. All components of Equation 2.7 are 1D, however, the $\cos\theta$ term can be expressed using the dot product to show the 3D basis of the phenomenon, as shown in Equation 7.1.

$$I = A \cdot \exp\left(-B \cdot \left(3 \cdot \left(\left|\hat{B}_0 \cdot \hat{c}\right|\right)^2 - 1\right)^2\right) \tag{7.1}$$

Where $\hat{B}_0$ is the unit vector representing the orientation of $\vec{B}_0$ (as described in Section 4.3.6), and $\hat{c}$ is a unit vector representing the orientation of the collagen fibre being scanned. Since both are unit vectors, the result is simply $\hat{B}_0 \cdot \hat{c} = \cos\theta$. The absolute value of the dot product is taken in Equation 7.1 to restrict θ to $\theta \in [0^\circ, 90^\circ]$, due to the symmetries discussed in Section 2.3.

To recover 3D fibre orientation information from 1D intensity data, the roll and yaw telemetry angles (θ_{roll} and θ_{yaw}) used during acquisition must be recorded in order to compute the corresponding scan orientations using Equation 4.1. These orientations are then used to simulate intensity profiles from a spherical distribution of candidate fibre directions via Equation 7.1. The measured intensity profiles, obtained after registration and segmentation, are correlated with the simulated profiles, and the candidate orientation yielding the highest correlation is assigned to each measured profile. This orientation is then mapped to the spatial location in the ROI corresponding to that measured profile.

A detailed description of the updated CFOEA is given below. By the stage of collagen fibre orientation estimation, the images have already undergone intensity correction (Section 4.4), registration (Chapter 5), and ROI segmentation (Chapter 6). The algorithm then operates on this segmented dataset, proceeding as follows:

1. Generate a hemisphere of K candidate fibre orientations, with $K = 40{,}000$ in the updated CFOEA. This hemisphere represents the complete solution space of possible fibre orientations.

2. Simulate the intensity profiles for each candidate orientation under the N scan orientations (parametrised by θ_R and ϕ_Y telemetry) using Equation 7.1. Store these simulated profiles in an array, forming a dictionary that indexes each profile to its corresponding candidate orientation.

3. For a selected voxel location in the ROI, correlate the associated measured intensity profile with the entire dictionary of simulated profiles.

4. Assign the candidate orientation indexed by the highest correlation between simulated and measured intensity profiles to the selected voxel location.

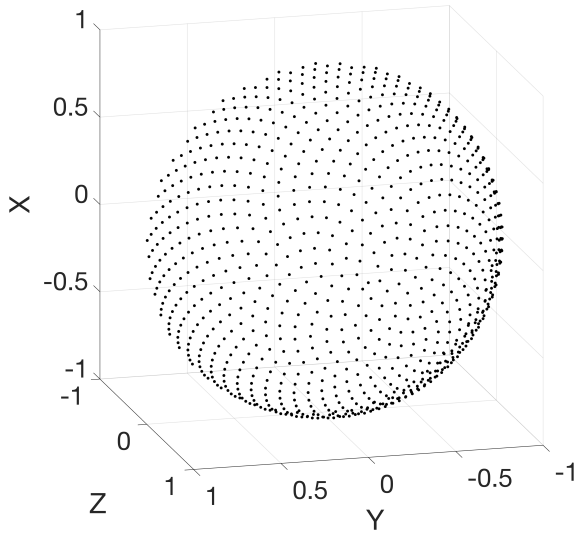
5. Repeat steps 3 and 4 for every voxel location in the ROI.

Efficiency Increase of Updated Collagen Fibre Orientation Estimation Algorithm

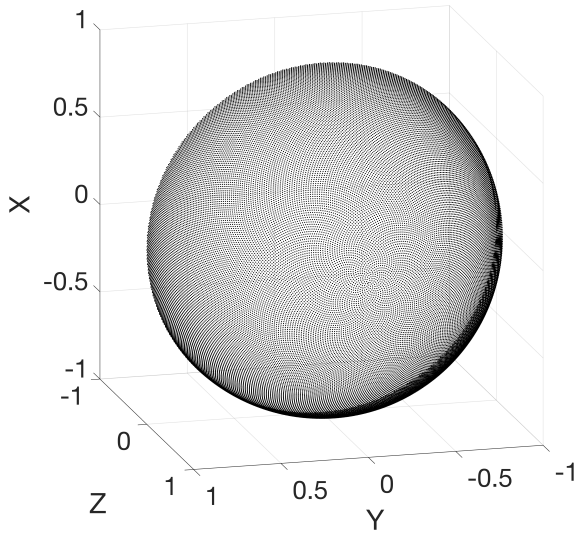
The increased efficiency of the updated CFOEA enabled the number of candidate orientations to be expanded by two orders of magnitude relative to previous implementations [11, 14, 15]. Szeverenyi et al. [11] employed 900 candidate orientations in their initial bouquet, yielding an angular resolution of $\approx 4.6^\circ$. Chappell et al. [14] and Brujic et al. [15] used only 250-500 candidates to speed up fibre orientation computation, but achieved an effective resolution of 0.01° through their optimisation step, provided the fitting coefficients A and B were well tuned.

In contrast, the updated algorithm employed 40,000 candidate orientations, giving a resolution of $\approx 0.7^\circ$, whilst still executing significantly faster than the methods of Szeverenyi et al. [11], Chappell et al. [14], and Brujic et al. [15], and even outperforming Seidel et al. [10], whose approach evaluated at most four candidates per voxel in the worst case scenario. Figure 7.1 shows sets of 900 and 40,000 candidate fibre orientations

evenly distributed over a hemisphere. The orientations are represented as points for clarity but are actually vectors with a common origin at (0,0,0). To generate the dense ball of unit vector candidate fibre orientations, the golden spiral method (see Section E.2) was utilised, which distributes points approximately evenly over a sphere.



(a) 900 candidate fibre orientations distributed evenly over a hemisphere.



(b) 40,000 candidate fibre orientations distributed evenly over a hemisphere.

Figure 7.1: Low and high density sets of candidate fibre orientations generated with the golden spiral method. Figure 7.1a shows 900 candidates and Figure 7.1b shows 40,000 candidates.

Seidel et al. reported the fastest algorithm in prior multi-orientation MRI literature, solving for 332,593 fibres in ≈ 20 minutes. As a benchmark, a dataset of equivalent size (332,593 intensity profiles, each with 35 elements) was simulated and solved using the updated CFOEA, which completed in an average of 45.2 seconds, representing a speed-up of more than 20-fold. While these benchmarks are not directly comparable due to hardware differences (Seidel et al. used a 2.1 GHz processor, whereas this benchmark was performed on a 3.5 GHz processor), the scale of improvement is far greater than can be attributed to hardware alone. Table 7.1 presents data comparing the speeds of collagen fibre orientation estimation algorithms versus the updated CFOEA.

The updated CFOEA is evidently significantly faster than all previous algorithms reported in the literature. For very large numbers of voxels, this translates into substantial time savings. For realistic ROI sizes in current MADI datasets, for example 1245 voxels for the ACL segmented from the representative dataset, the practical gain is smaller due to the limited scale of the segmentation. Nonetheless, the performance improvement is a valuable user experience improvement providing near instantaneous results, and will be highly valuable for high-resolution MADI applications in the future. Critically, the

Table 7.1: Reported speeds of previous collagen fibre orientation estimation algorithms and the updated collagen fibre orientation estimation algorithm. For each source the number of fibres estimated is given (same as number of voxels in the ROI) with the time taken for the algorithm to complete. The data are also presented as a rate (fibres/sec) for comparison. *Chappell et al. and Brujic et al. did not provide a number of fibres estimated, only a relative performance to Szeverenyi et al., so 100,000 was used to match the reported fibre numbers from Szeverenyi et al.

Source	**No. of fibres estimated**	**Time**	**Fibres/sec**
Szeverenyi et al. [11]	100,000	≈ 60 minutes	27.8
Chappell et al. [14] and Brujic et al. [15]	100,000*	$\approx 20 - 30$ minutes	83.3-55.6
Seidel et al. [10]	332,593	≈ 20 minutes	277.2
Present work (benchmark)	332,593	45.2 seconds	7358.3
Present work (representative dataset ACL)	1245	0.2 seconds	6552.6

increased efficiency allowed for integration of the CFOEA into the genetic algorithm for prospective scan orientation selection described in Section 7.3.3, where the algorithm must be executed many times.

The efficiency gains in this implementation of the CFOEA arose from two major optimisations. Firstly, the simulation of candidate fibre intensity profiles under the acquired scan orientations (step 2) was reformulated as a matrix multiplication, as shown in Equation 7.2.

$$\mathbf{T} = \mathbf{CS} \tag{7.2}$$

Where **T** is a $K \times N$ matrix of $\cos\theta$ values for every candidate fibre and every scan orientation combination, obtained from the matrix multiplication on the right-hand side of the equation[1]. **C** is a $K \times 3$ matrix of unit candidate fibre orientations, with each row containing a $\hat{c}$ vector, and **S** is a $3 \times N$ matrix of scan orientations, with each column

[1]This construction essentially performs all necessary dot products as defined in Equation 7.1 as a single operation.

corresponding to a $\hat{B}_0$ vector for a given acquisition. Each element of **T** was then the $\hat{B}_0 \cdot \hat{c}$ component in Equation 7.1, which was used to compute the final simulated intensity profiles.

Secondly, all operations on large data structures, such as the measured and simulated intensity profiles, were implemented using efficient array operations in MATLAB. For instance, the linear algebra in Equation 7.2 is executed without explicitly defined for-loops, instead delegating execution to optimised C/C++ linear algebra/matrix multiplication libraries. Likewise, the correlation step (step 3) is performed using MATLAB's inbuilt `pdist2` function with the correlation metric, ensuring vectorised computation and avoiding unnecessary overhead.

Correlation as a Similarity Metric

For the simulated intensity profiles, the choice of A is not critical, since Pearson's correlation coefficient, Equation 7.3, was used as the similarity metric.

$$r_{xy} = \frac{\sum_{i=1}^{N} (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^{N} (x_i - \bar{x})^2} \sqrt{\sum_{i=1}^{N} (y_i - \bar{y})^2}} \tag{7.3}$$

Where r_{xy} is the Pearson coefficient between -1 and 1 (with 1 representing perfect positive correlation and -1 representing perfect negative correlation), x_i, y_i are pairs of intensities from a measured and simulated profile being correlated against one another, and $\bar{x}, \bar{y}$ are the means of each profile. This measure is invariant under affine transformations, therefore any mismatch in A between simulation and reality does not affect the correlation, as A only scales the relative signal intensity. In contrast, the parameter B is within the exponential term in Equation 7.1, and thus directly modifies the shape of the magic angle effect intensity response in a non-affine manner. B therefore must be matched as closely as possible between simulation and measurement, otherwise spurious correlations could arise and degrade the final collagen fibre orientation estimation.

7.2.3 Collagen Fibre Orientation Estimation Output

The output of the CFOEA is a candidate fibre orientation assigned to each voxel within the ROI. Similarly to visualising the ROI as a 3D point cloud as in Figure 6.3, the fibre orientations are visualised by plotting a 1 mm line segment with the midpoint located at the voxel centre. This produces a vector-field-like representation of the collagen architecture. Figures 7.2 and 7.3 illustrate this output for the ACL segmented from the representative dataset.

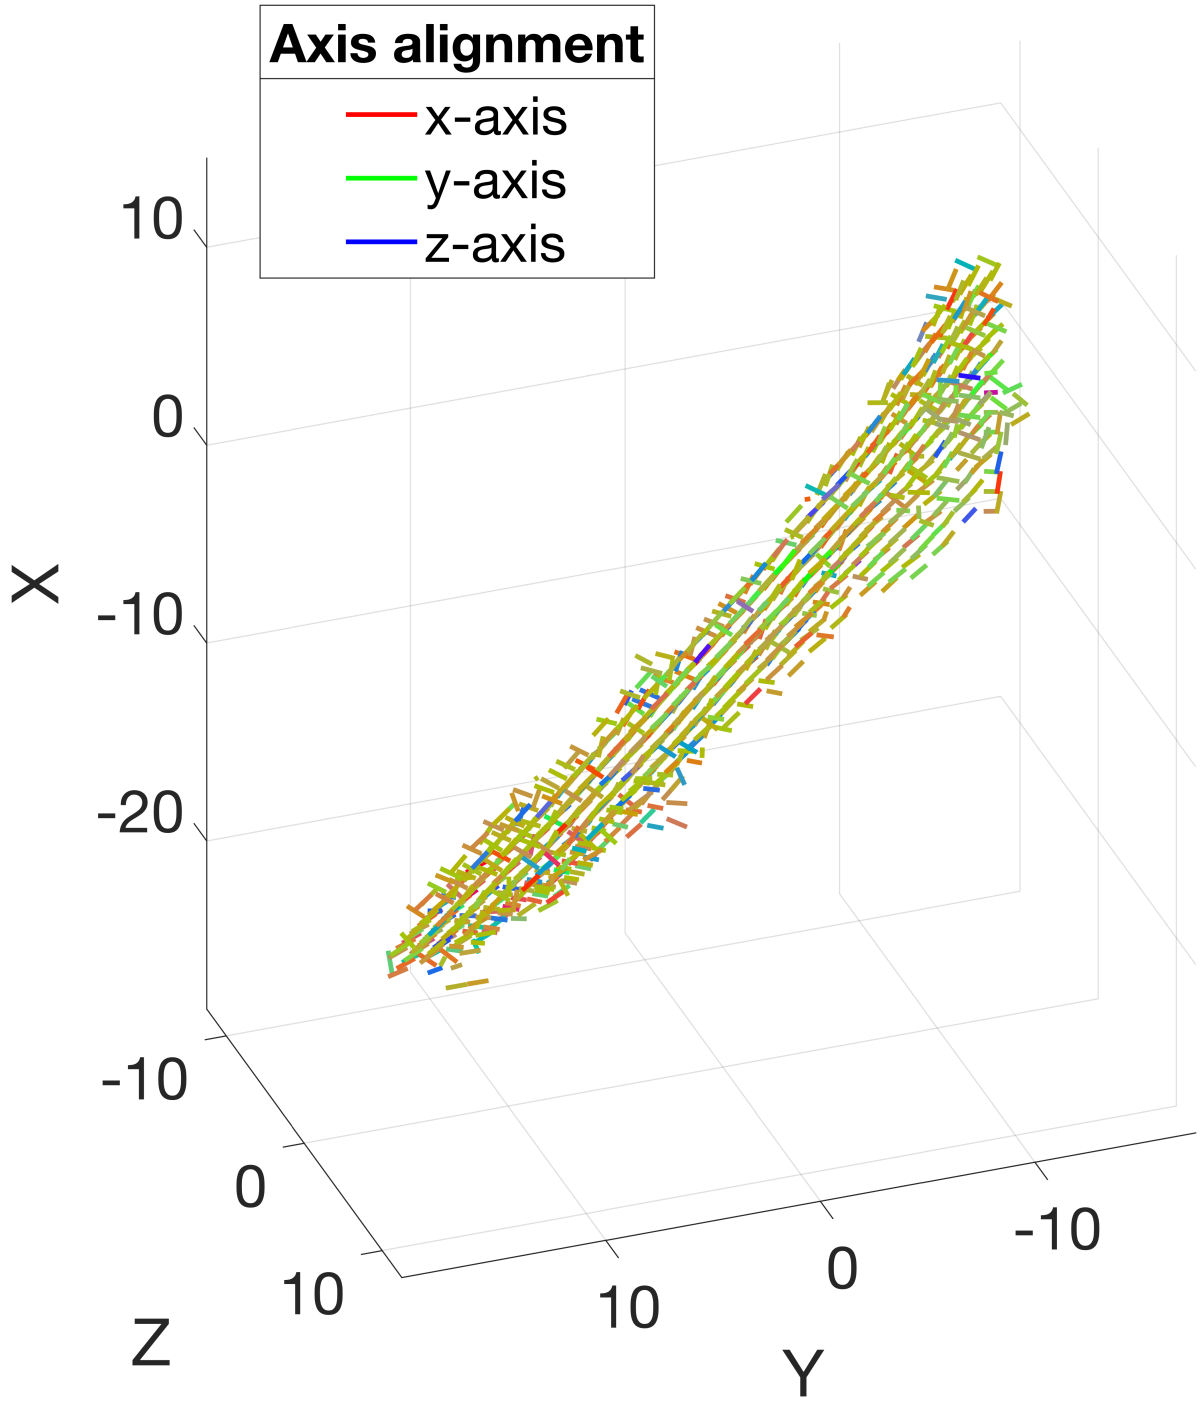
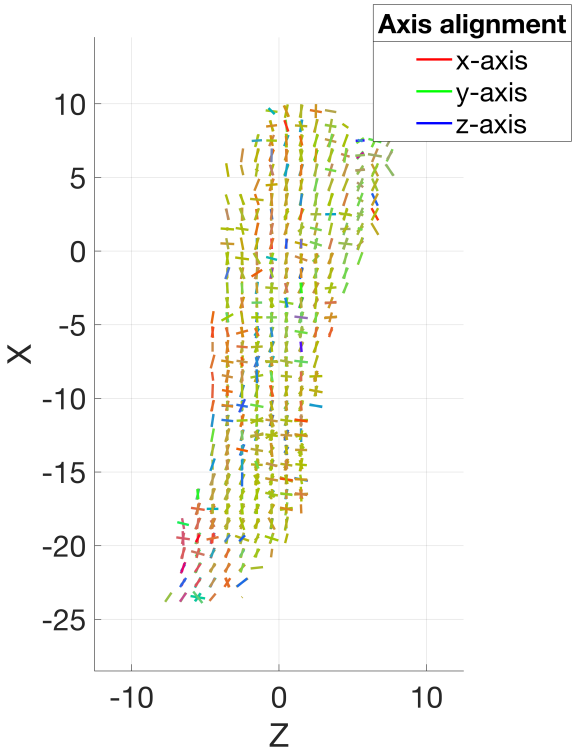
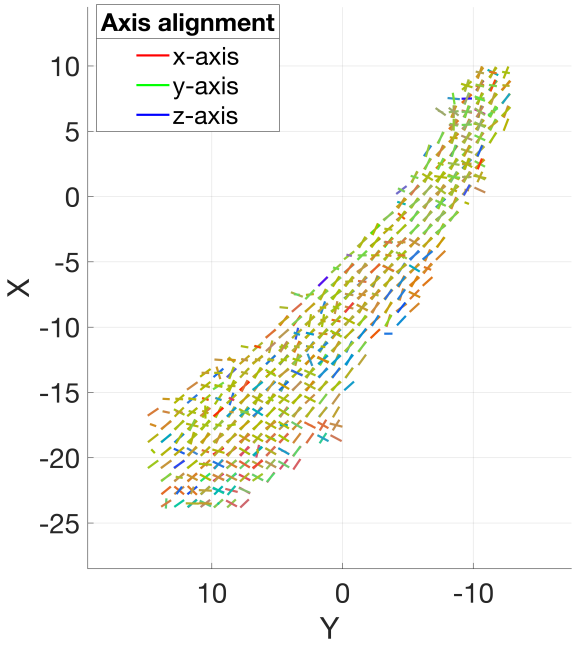


Figure 7.2: Fibre plot of the ACL from the representative dataset viewed from an oblique angle.



(a) Fibre plot of the ACL from the representative dataset viewed looking directly against the y-axis (superior view).



(b) Fibre plot of the ACL from the representative dataset viewed looking directly against the z-axis (lateral view, right to left).

Figure 7.3: Fibre plots of the ACL from the representative dataset viewed directly against the y- and z-axes (Figures 7.3a and 7.3b).

The fibre colours in Figures 7.2 and 7.3 are encoded in RGB according to their orientation. Fibres aligned exactly with the x-, y-, or z-axes appear as pure red, green, or blue, respectively. For all other orientations, colours are generated by interpolating between these primaries, with each RGB channel proportional to the absolute value of the corresponding unit vector component of the fibre orientation. Thus, the components $|x|$, $|y|$, and $|z|$, each ranging from $0 \dots 1$, are linearly mapped onto the RGB scale $0 \dots 255$, producing a smooth and continuous orientation-colour relationship.

Figures 7.2 and 7.3 demonstrate promising collagen fibre estimation results. The predominance of fibres with similar colour indicates coherent alignment within the ACL, consistent with expectations for a healthy ligament. Most fibres appear dark green, corresponding to orientations primarily aligned with the y-axis but with a substantial component along the x-axis. This is in agreement with the established ACL morphology, which forms a tilted, obliquely oriented structure traversing diagonally upwards through the joint. The estimated alignment is visually apparent across all views of the ACL, including the oblique (Figure 7.2), superior (Figure 7.3a), and lateral perspectives (Figure 7.3b). Occasional local deviations from the bulk orientation are evident, though these are predominantly located at the margins of the structure. This suggests that partial volume effects or interpolation artefacts introduced during registration are the likely causes of misalignment. Such fibres constitute only a small fraction of the total estimated set, with the overwhelming majority displaying the expected coherent alignment.

After collagen fibre orientation estimation the parameters A and B were iteratively refined. Each collagen fibre orientation was associated with a set of N measured intensities and a corresponding set of N simulated intensities as described in Section 7.2.2. To choose a suitable B value to match the measured and simulated intensities numerical optimisation was performed: minimising the sum of squared residuals (SSR) between the measured and simulated intensities. Each measured voxel intensity had an associated simulated intensity from the CFOEA as described in Section 7.2.2; thus, SSR between the measured and simulated intensities could be used as a metric to optimise A and B. The CFOEA was run as part of a minimisation optimisation loop where A and B were tuned[2] using the Nelder-Mead simplex algorithm [114]. With this approach a value of

[2]Even though A does not affect the correlation metric it was included as it added negligible computational overhead.

$B = 0.5$ was adopted in this work, which was consistent with values reported in prior studies [11, 14, 15].

7.3 Scan Orientation Selection

7.3.1 Previous Scan Orientation Selection Methods

Previous studies [10, 11, 14, 15] selected scan orientations using heuristic strategies, typically either distributing orientations broadly or acquiring a dense set. These approaches were intended to ensure adequate intensity variation from the magic angle effect for subsequent analysis as outlined in Section 7.2.1. In the early stages of the present work, data from our laboratory's previous study (Chappell et al. [14]), comprising 15 evenly distributed orientations of a caprine knee sample, were re-analysed to investigate the minimum number of scan orientations required for MADI when a priori knowledge of the target anatomy is available. The results from this work were published in *Array*: Lanz et al. [115].

Lanz et al. [115] evaluated subsets of the 15 scan orientations to assess the impact on collagen fibre estimation accuracy, using the full dataset as the reference gold standard. The structure of interest from the caprine knee sample was the patellar tendon, which provided a straight reference to assess the relative orientation of the selected scans against. Thus, each subset dataset could be treated as if scans were being selected with a priori knowledge of the orientation of the patellar tendon. Subsets of 3 to 14 scans were studied, and synthetic noise was also added to simulate SNR levels comparable to those of low-field data. Utilising 4 scan orientations was found to be the minimum viable number of orientations provided a priori knowledge of the patellar tendon orientation was available. Table 7.2 summarises the relative angles between the patellar tendon and the most viable subsets of 3, 4, and 5 scan orientations.

Lanz et al. found that the 3-scan orientation results were highly sensitive to noise, and that increasing the number of scans provided incremental improvements towards the full dataset gold standard. The 4-scan subset was identified as suitable compromise, balancing number of scans, accurate fibre direction estimation, and robustness to noise. A key finding from Lanz et al. was that scans at low relative angles (e.g. 29° and 35°) were the most important to include. These low angle scans contributed significantly to

Table 7.2: Approximate angles between scan orientation subsets and caprine patellar tendon from Lanz et al. [115].

No. of scans	Approximate angles to patellar tendon
3	$35^\circ, 57^\circ, 69^\circ$
4	$29^\circ, 35^\circ, 57^\circ, 69^\circ$
5	$29^\circ, 35^\circ, 57^\circ, 69^\circ, 80^\circ$

the accuracy of fibre direction estimation, whereas although higher angle scans were necessary to provide sufficient intensity variation in the target tissues, there was less sensitivity to their exact values. This result was affirmed by grid search simulation studies assessing fibre estimation convergence with three to five scans at different SNR levels.

This result is intuitive, as low orientation angles correspond to signal intensities uniquely associated with a single θ value in Equation 2.7. As illustrated in Figure 7.4, beyond a critical angle α there is a 2:1 mapping of $\theta \mapsto I$, such that a given I may correspond to two distinct θ values. Consequently, in MADI, where intensities are observed and θ must be inferred as part of fibre orientation estimation, measurements acquired at $\theta \leq \alpha$ impose a stronger constraint on the estimated fibre orientation than those acquired in the range $\alpha \leq \theta \leq 90^\circ$.

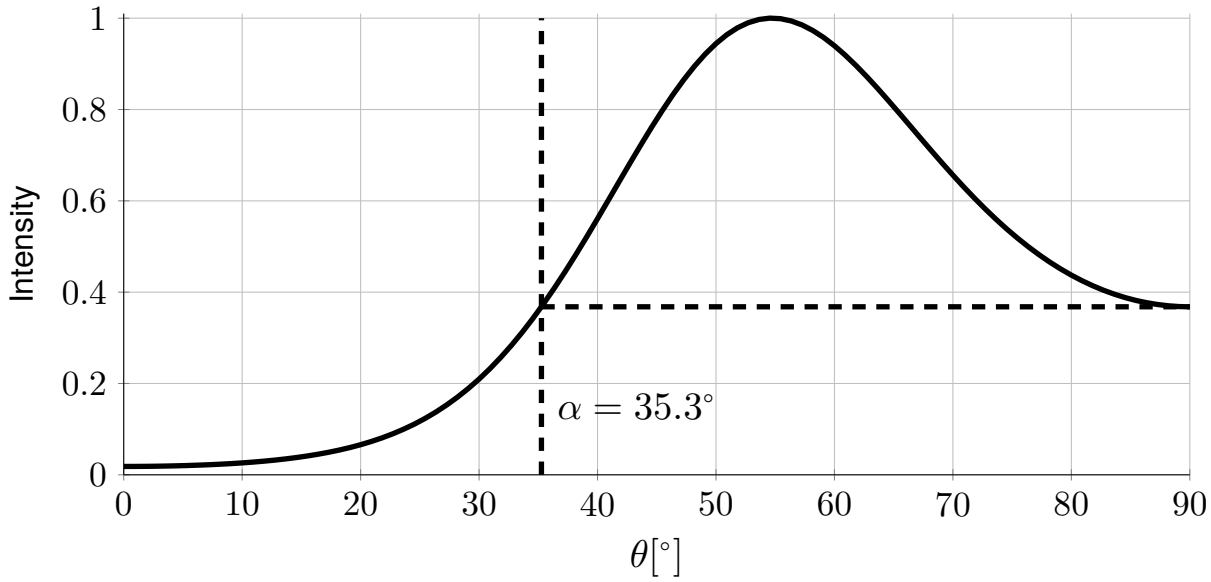


Figure 7.4: Intensity versus θ for $A = 1$ and $B = 1$ with α marked to show the domain where two θ values map to a single intensity value.

From Equation 2.7 the exact value of α is derived as follows:

$$\text{For } \theta = 90^\circ, \quad I(90) = A \cdot \exp(-B \cdot (3\cos^2 90^\circ - 1)^2)$$
$$I(90) = A \cdot \exp(-B \cdot (0-1)^2)$$
$$I(90) = A \cdot \exp(-B)$$

$$\text{Define } \theta = \alpha, \text{ such that } I(\alpha) = I(90)$$
$$I(\alpha) = A \cdot \exp(-B \cdot (3\cos^2 \alpha - 1)^2)$$
$$\therefore A \cdot \exp(-B \cdot (3\cos^2 \alpha - 1)^2) = A \cdot \exp(-B)$$
$$\therefore (3\cos^2 \alpha - 1)^2 = 1$$
$$\therefore \alpha = 35.3^\circ$$

This finding suggests that $\vec{B}_0$ orientations at low angles relative to the collagen fibres are more critical for fibre orientation estimation than sampling strictly within the higher intensity regime associated with $\theta \geq 35.3^\circ$. Consequently, the restricted yaw range of the low-field rotatable scanner may limit estimation accuracy for structures oriented predominantly along the scanner's roll axis (laboratory y-axis) direction, as low-angle orientations relative to these fibres cannot be achieved.

Figure 7.5 illustrates the challenge of sampling fibres aligned with the roll axis of the scanner. A top-down view shows $\vec{B}_0$ in its home orientation, parallel to the laboratory z-axis, with its constrained yaw range of $\phi_Y = \pm 30^\circ$. Since axially aligned fibres are coaxial with the roll axis of $\vec{B}_0$, this yaw constraint is the sole determinant of the achievable angle θ between $\vec{B}_0$ and such fibres, restricting it to $60^\circ \leq \theta \leq 90^\circ$ irrespective of roll. This range lies entirely above the critical angle $\alpha = 35.3^\circ$ derived in the above analysis of Equation 2.7; achieving $\theta \leq \alpha$ would require a yaw of at least 54.7°, far exceeding the current limit. Axially aligned fibres therefore cannot be sampled at or below α, or even at the magic angle. This precludes the low intensities at $\theta \leq \alpha$ and the signal intensity peak at the magic angle from being sampled, rendering them the most challenging case for the present system.

However, evaluating the adequacy of a given set of scan orientations is not straightforward, as collagen fibre orientation estimation accuracy may depend on emergent

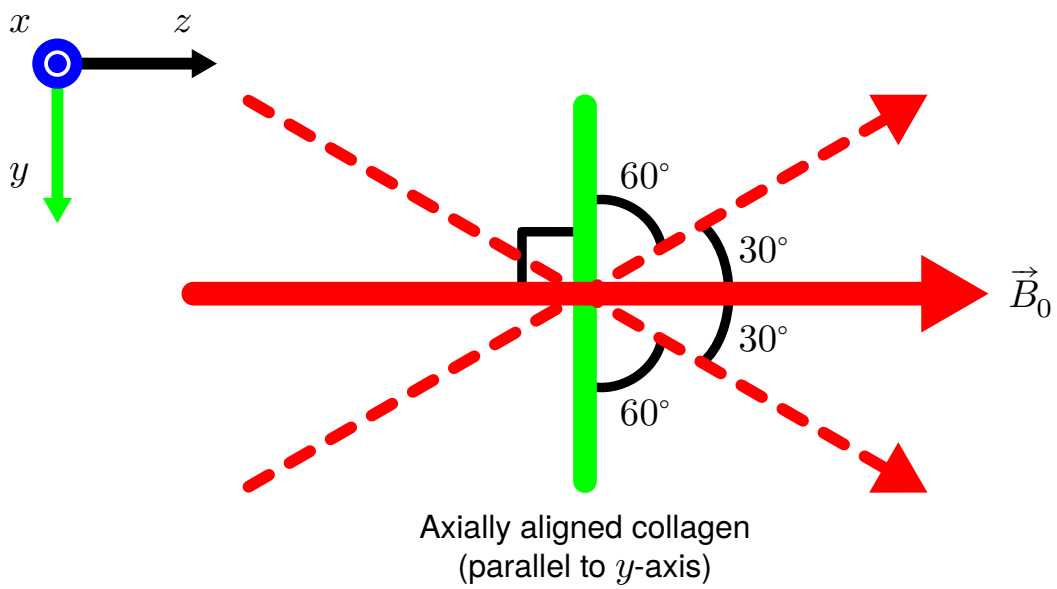


Figure 7.5: Limited angular sampling range relative to an axially aligned collagen fibre due to the scanner's constrained yaw range. The laboratory coordinate system is indicated by the axes in the upper left, with the x-axis directed out of the page. $\vec{B}_0$ is shown in its home orientation as a solid red arrow parallel to the z-axis, with its maximum yaw positions at $\phi_Y = \pm 30^\circ$ indicated by dashed red arrows. The axially aligned collagen fibre is oriented parallel to the y-axis, and the minimum achievable angle between $\vec{B}_0$ and this fibre is 60°, irrespective of roll angle.

properties of their collective three-dimensional distribution rather than the angle each orientation forms relative to any single fibre. This was evident in Lanz et al. [115], where scan orientations were assessed only relative to the patellar tendon, without explicitly accounting for their mutual three-dimensional interrelationships. To enable a more rigorous analysis, accuracy metrics for collagen fibre orientation estimation were developed and incorporated into a genetic algorithm for scan orientation optimisation, as described in Sections 7.3.2 and 7.3.3.

Without hardware constraints, a sufficient number of orientations could in principle be chosen such that any possible fibre orientation has at least one scan with $\theta \leq \alpha$. The yaw constraint renders this impossible for axially aligned fibres as previously discussed. The genetic algorithm was therefore developed to minimise the mean angular deviation of estimated fibre orientations across all possible fibres, potentially exploiting emergent properties of the three-dimensional orientation set that are not accessible via manual design or simple geometric reasoning. Furthermore, the discrete, non-differentiable nature of the deviation metric precludes gradient-based optimisation, making stochastic optimisation with a genetic algorithm a natural methodological choice.

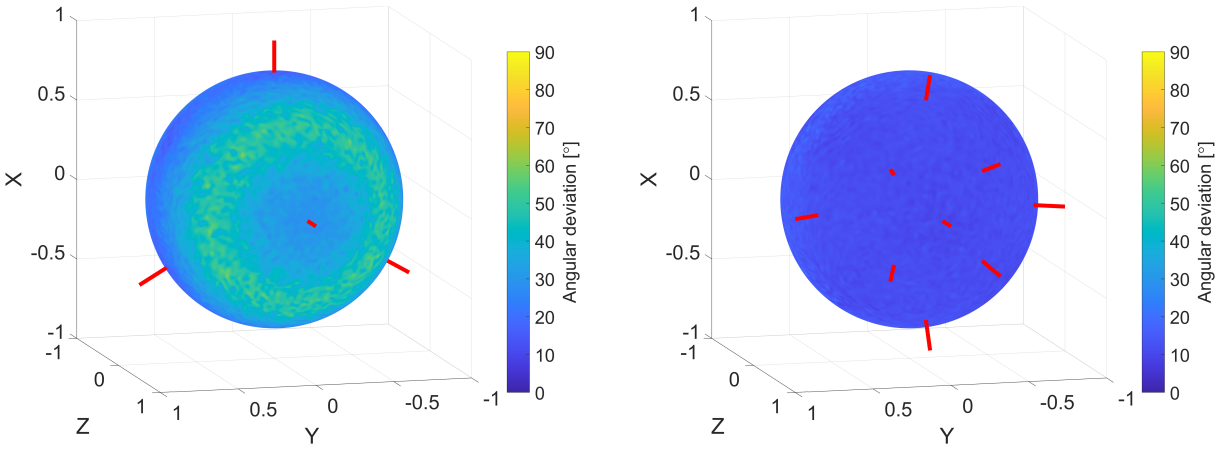
7.3.2 Collagen Fibre Orientation Estimation Accuracy Mapping

The increased speed of the updated CFOEA (Section 7.2.2) allowed for its utilisation as an experimental design tool where previous iterations would have been prohibitively slow. Collagen fibre orientation accuracy mapping was one such tool that was enabled by the updated CFOEA.

To generate a collagen fibre orientation estimation accuracy map for a given set of scan orientations, intensity profiles were simulated for a uniform distribution of 10,000 ground truth fibres over a sphere using Equations 7.1 and 7.2, with parameters $A = 1$ and $B = 0.5$. Rician noise was added with standard deviation, σ, set to $\sigma = 0.2$ for most experiments. This corresponds to an effective SNR of approximately 5 at the peak simulated intensity (since $A = 1$), a lower SNR than observed in the in vivo data, thereby providing a conservative benchmark for evaluating the robustness of candidate scan orientation sets. The noisy profiles were then processed using the CFOEA to generate estimated fibre orientations. This procedure was repeated 50 times, and the mean angular deviation between estimated and ground-truth orientations was mapped onto the surface of a sphere parametrised by the ground-truth directions.

Figure 7.6 shows two such maps of collagen fibre orientation estimation accuracy for sets of 4 and 9 scan orientations. Ground truth fibre orientations located in the brighter regions were more likely to be estimated incorrectly and thus presented a higher angular deviation on average. The 4 scan orientation set shown in Figure 7.6a shows 3 scan orientations distributed at 60° relative to one another with one scan parallel to the z-axis as is always the case for MADI examinations. Low angular deviation is visible close to the scans, as predicted in Section 7.3.1, with high angular deviation in regions far from all scan orientation vectors. Figure 7.6b shows 9 scan orientations distributed in a grid pattern, again with one scan parallel to the z-axis representing the home orientation scan. The angular deviation from the ground truth is low everywhere indicating any arbitrary collagen fibre orientation would be reliably estimated using this set of scan orientations. Unfortunately, due to the large yaw angles required to acquire scans at the orientations shown in Figure 7.6b this set would be impossible to utilise in vivo.

The overall accuracy of a set of scan orientations was measured by taking the total average deviation for all estimated fibres from the corresponding ground truth fibres



(a) Accuracy map for a set of 4 scan orientations.

(b) Accuracy map for a set of 9 scan orientations.

Figure 7.6: Collagen fibre orientation estimation accuracy maps for sets of 4 and 9 scan orientations. Scan orientation vectors are represented by red lines. Figure 7.6a shows an accuracy map with 4 scan orientations and Figure 7.6b shows an accuracy map with 9 scan orientations. $\sigma = 0.2$ for both maps.

(i.e. the mean of all deviation values across the sphere surface). This scalar value was used as the fitness for a set of scan orientations in the genetic algorithm presented in Section 7.3.3. The maximum possible deviation of any given estimated fibre from the ground truth is $90°$ due to the bidirectionality of collagen fibres as discussed in Section 2.3. Thus, the highest possible total average deviation was $90°$ (lowest fitness) and the lowest was $0°$ (highest fitness).

7.3.3 Genetic Algorithm for Scan Orientation Selection

Genetic algorithms are particularly suited to optimisation problems characterised by large design spaces, difficulty or impossibility of computing gradients, and ease of verifying candidate solutions. They are also advantageous for identifying emergent properties that may not be apparent through heuristic reasoning. In many optimisation problems, cost functions can be expressed with explicit gradients, enabling approaches such as adaptive stochastic gradient descent to efficiently minimise error terms; for example, the registration transforms described in Section 5.4.2 were optimised in this way. In contrast, the fitness of a set of scan orientations in MADI depends on noisy simulations as described in Section 7.3.2, making it difficult to define a stable error gradient for conventional gradient-based optimisation. For this reason, a genetic algorithm was selected as an appropriate method to explore the scan orientation design space and to quantitatively optimise scan orientation sets.

A pseudocode implementation of the genetic algorithm is presented in Algorithm 1. Each generation consisted of 400 agents where each agent was a binary array of 1s and 0s that represented its genes. Subsections of the array encoded the two's complement binary representation of the roll and yaw angles for each scan orientation represented by the agent. The first scan roll and yaw angles were never encoded as these were always $\theta_R = 0^\circ$ and $\phi_Y = 0^\circ$. More detail of this encoding scheme is provided in Section E.3. All agents in the first population were initialised with random but valid genes to ensure a diverse gene pool for the algorithm to begin with.

During the generation iterations of the genetic algorithm each agent's genes were decoded into decimal roll and yaw angles and the fitness of the resulting scan orientation set was assessed using the collagen fibre estimation total accuracy as described in Section 7.3.2. For greater execution speed, the density of both the ground truth and solver orientations was reduced by a factor of 100 as this still provided a sufficient fitness metric with lower computational expense.

The next population of agents was produced by breeding with tournament parent selection. Four random agents were selected and the top two fittest would be chosen as parents to produce two offspring. This allowed for a genetically diverse set of parents to be chosen to reproduce with a slight bias towards fitter agents. Breeding was determined by the `breed_rate` where 90% of the time parents were successful in producing offspring and 10% of the time they would pass to the next generation themselves, aiding to preserve genetic diversity. Children were produced by uniform random crossover of genes between parents. All child agents were subject to a 2% chance of gene mutation which served to introduce new genetic material into the next population. The top 5% of each population was also always passed through as "elites" such that the best genes from each generation were always preserved without loss due to breeding or mutation. The genetic algorithm was typically run for 500 generations but convergence was usually observed after < 50.

The genetic algorithm was controlled using several hyperparameters. The number of agents, generations, elites, and the mutation and breeding rates directly affected the algorithm's convergence behaviour. Algorithm 1 shows some representative values for these parameters. To focus the algorithm on a particular number of scan orientations constrained to an appropriate operational window the number of scans N, and the roll

Algorithm 1 High level overview of the genetic algorithm to select scanning orientations for a given number of scans.

Set hyperparameters for genetic algorithm

`num_agents` $\leftarrow 400$

`num_gen` $\leftarrow 500$

`num_elites` $\leftarrow 0.05 \cdot$ `num_agents`

`mut_rate` $\leftarrow 0.02$

`breed_rate` $\leftarrow 0.9$

Set constraints for scanning orientations and simulated intensity data

`num_scans` $\leftarrow N$

`roll_lim_deg` $\leftarrow 90$ ▷ Absolute value, represents angles of $\pm 90^{\circ}$

`yaw_lim_deg` $\leftarrow 30$ ▷ Absolute value, represents angles of $\pm 30^{\circ}$

`noise_sigma` $\leftarrow 0.2$

Initialise population of agents with random but viable starting genes

for $a = 1 \ldots$ `num_agents` **do**

`pop`$[a, :] \leftarrow [10011 \ldots 01]$ ▷ Initial genes constructed to be within constraints

end for

Run the genetic algorithm

for $g = 1 \ldots$ `num_gen` **do**

for $a = 1 \ldots$ `num_agents` **do**

`fitness`$[a] \leftarrow$ `compute_agent_fitness(pop)`$[a, :]$

end for

`max_fitness`$[g] \leftarrow$ `max(fitness)`

`mean_fitness`$[g] \leftarrow$ `mean(fitness)`

`mean_valid_fitness`$[g] \leftarrow$ `mean(fitness` $\forall$ `fitness` > 0`)`

`pop` $=$ `next_pop(pop)` ▷ Next population is bred from current population's agents

end for

and yaw limits were also specified. N controlled the number of genes in each agent, and the roll and yaw limits would ensure that any children generated with scan orientations outside of the permitted operational window would be assigned minimum fitness and would not take part in the next round of breeding.

Figure 7.7 shows the convergence of the minimum and mean total average angular deviation for $N = 6$, $\sigma = 0.2$, and roll and yaw limits set to $\pm 90^{\circ}$ and $\pm 30^{\circ}$ respectively, over 50 generations. The minimum total average deviation (representing the fittest agent from each generation) converges rapidly initially before slowing and finally settling after about 25 generations. The mean population total average deviation shows a similar pattern with smoother overall behaviour with convergence over time as the population's genes become more homogeneous. This convergence profile is a typical example of the convergence behaviour of the genetic algorithm. The converged values of the total average deviation are relatively high in Figure 7.7, however, this is due to utilisation of only 6 scan orientations and the high value of $\sigma = 0.2$.

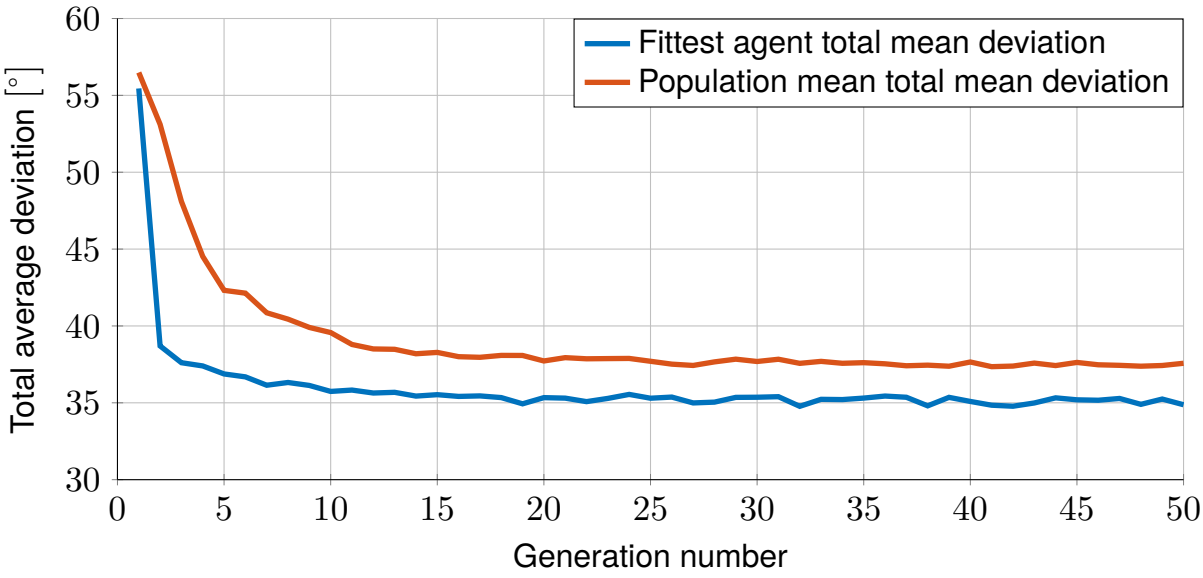
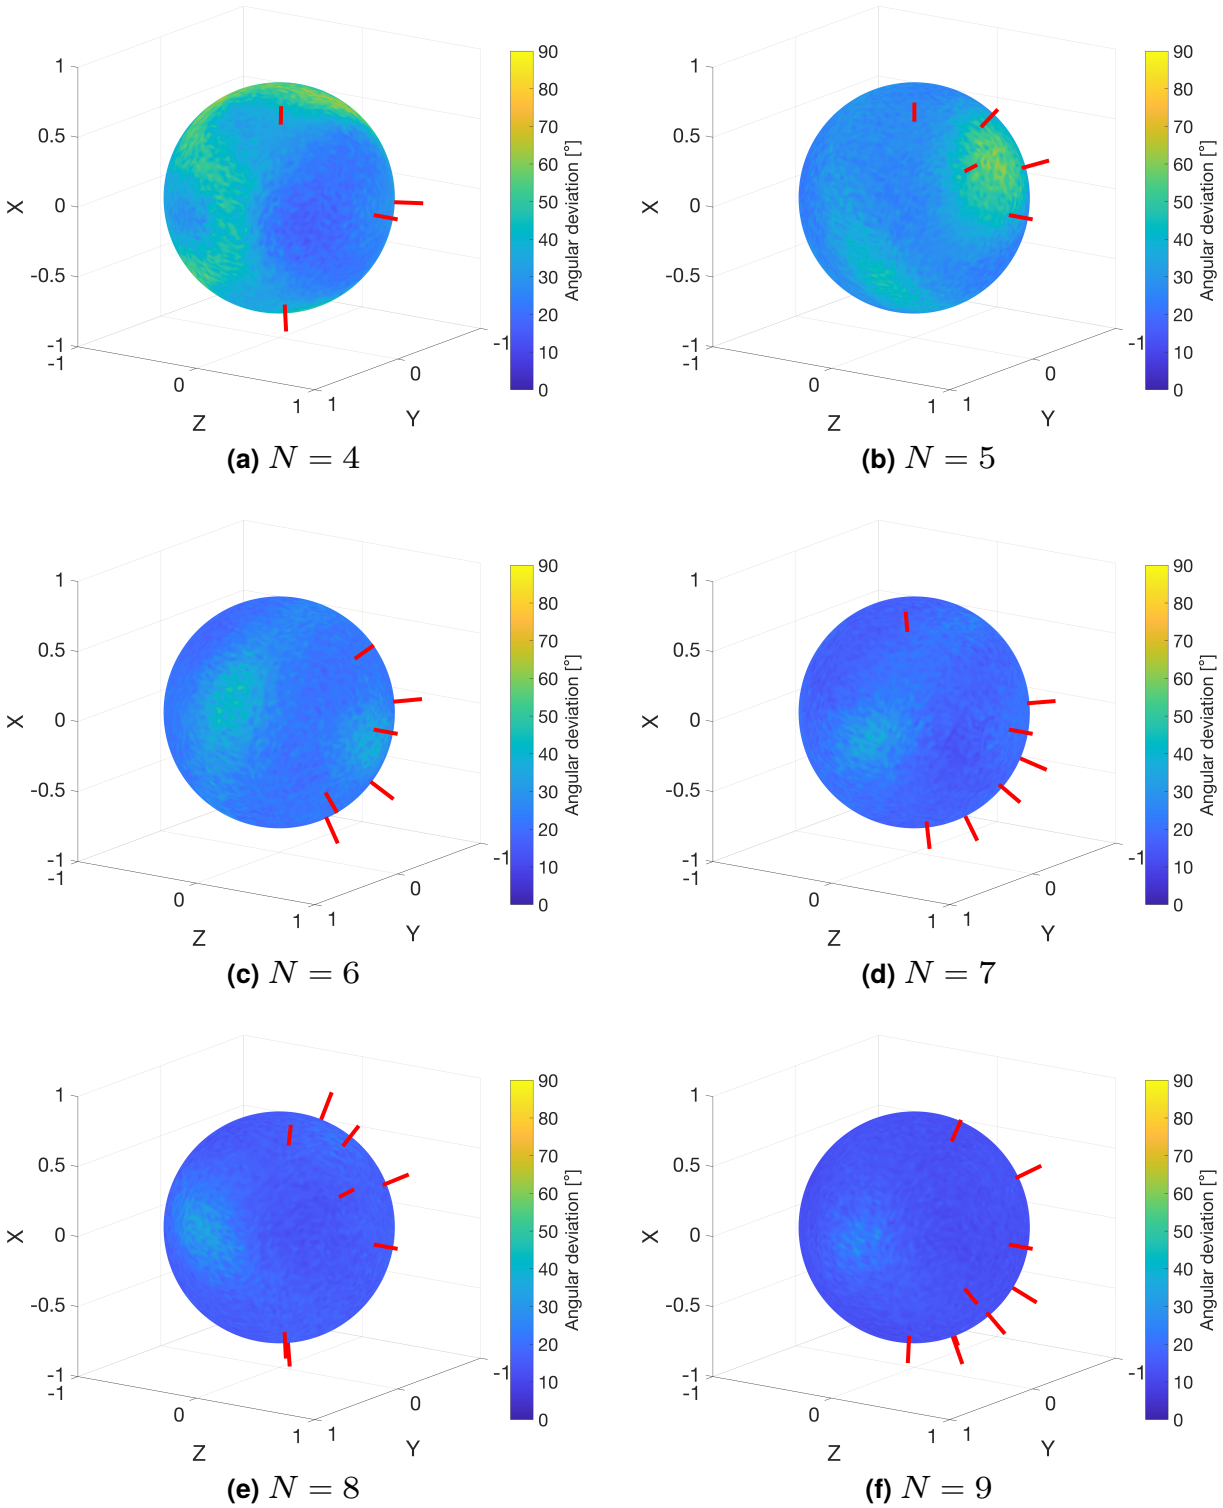


Figure 7.7: Convergence of total average deviation for genetic algorithm with $N = 6$. The blue line shows the convergence of the fittest agent from each generation and the orange line shows the convergence of the mean fitness of the population.

Figure 7.8 shows the collagen fibre orientation estimation accuracy maps for sets of $N = 4, 5, 6, 7, 8, 9$ scan orientation sets, optimised using the genetic algorithm from Algorithm 1. The individual $\vec{B}_0$ orientations from the scan orientation sets are shown in red.



(a) $N = 4$ **(b)** $N = 5$ **(c)** $N = 6$ **(d)** $N = 7$ **(e)** $N = 8$ **(f)** $N = 9$

Figure 7.8: Collagen fibre orientation estimation accuracy maps for $N = 4, 5, 6, 7, 8, 9$ genetic algorithm optimised scan orientations. Bright areas denote regions of high average deviation and dark regions denote areas of low average deviation from the ground truth.

Figure 7.8 shows the progressive improvement in collagen fibre orientation estimation accuracy with increasing N, evidenced by the darkening of the deviation maps from Figure 7.8a to Figure 7.8f. As predicted in Section 7.2.1, the greatest deviation is consistently observed in the region aligned with the y-axis, where the limited yaw range prevents scan orientations from forming small angles with axially aligned fibres. Notably, in Figure 7.8b a localised region of elevated deviation appears in close proximity to four of the five scan orientations. The concentration of orientations in this region yields redundant sampling which is also at low θ, and therefore low I, rendering it particularly susceptible to noise. Both of these effects degrade CFOEA performance. This illustrates a known characteristic of genetic algorithms: a globally fit solution may nonetheless exhibit localised regions of poor performance. In the present work, such solutions were manually reviewed and discarded in favour of alternative genetic algorithm runs; future iterations of the fitness function could mitigate this effect with regularisation for smoothness of the angular deviation.

It should be noted that the plots in Figure 7.8 are illustrative examples of the genetic algorithm output rather than the specific orientation sets used throughout the present work. As genetic algorithms are stochastic optimisers, they do not guarantee convergence to the global minimum; the solutions shown in Figure 7.8 are demonstrative of the qualitative decrease in angular deviation as N increases. As noted previously, a persistent region of elevated deviation aligned with the y-axis is apparent across all subplots which is attributed to the limited yaw range of the scanner: axially oriented fibres cannot be imaged at sufficiently small relative angles to resolve their orientation accurately, representing a current hardware constraint rather than a limitation of the optimisation. The precise roll and yaw angles used for all in vivo acquisitions are tabulated in Tables 8.1 to 8.4.

From simulation studies with the genetic algorithm, it was decided to utilise $N = 7$ scans for most in vivo applications of MADI in this thesis. This choice was a compromise between acquisition time and collagen fibre orientation estimation accuracy. For $N \leq 6$ average deviation was markedly higher than $N \geq 7$, yet generally MADI examinations of $N \geq 8$ would extend acquisition time for marginal accuracy gains thus decreasing patient comfort.

The tabulated scan orientations were not adopted directly from a single genetic algorithm output; rather, candidate sets were drawn from repeated genetic algorithm runs and ultimately selected against practical criteria not captured by a deviation-based fitness function; principally, the angular proximity between successive orientations to minimise patient movement and operator workload during acquisition. Consequently, the scan orientation sets evolved across the studies presented in this thesis, representing a data-informed but operationally constrained selection. Incorporation of such ergonomic criteria into the fitness function to yield a fully automated orientation selection procedure is identified as future work in Section 9.5.4.

7.4 Summary

This chapter presented a detailed account of collagen fibre orientation estimation and how advances in the estimation algorithm informed the experimental design of MADI. It began with a review of existing estimation approaches in Section 7.2, assessing their computational performance before introducing the updated collagen fibre orientation estimation algorithm. The improved efficiency of this implementation enabled its use not only for fibre reconstruction but also as a tool for experimental design and planning. Fibre plots of the ACL from the representative dataset were shown to illustrate the output of the CFOEA.

Section 7.3 described an early study in which scan orientations were selected with a priori knowledge of the target anatomy. This analysis indicated that inclusion of at least one orientation at a low angle ($< 35.3^{\circ}$, as derived from Equation 2.7) was desirable for accurate estimation. The updated CFOEA was subsequently applied to generate orientation accuracy maps, exploiting its speed to incorporate noisy simulations that would have been computationally prohibitive with earlier implementations. These maps provided the foundation for the fitness function of a genetic algorithm used to optimise scan orientation selection under the practical constraints of MADI. This framework allowed both the number of orientations, N, and their distribution to be selected on a data-driven basis rather than by heuristic rules. Moreover, the genetic algorithm permits the inclusion of arbitrary constraints, offering flexibility to accommodate future experimental or technical limitations.

8 | Results

8.1 Chapter Overview

This chapter provides an overview of the types of results obtainable with MADI. Section 8.2 first describes the conversion of collagen fibre unit vectors estimated by MADI into synthetic diffusion tensors, enabling tractography visualisation using standard DTI software. Section 8.3 then presents detailed results for the ACL from the representative dataset, including all stages following fibre orientation estimation and tractography visualisation. Additional analysis of this dataset is used to demonstrate visualisation of the anteromedial and posterolateral bundles, highlighting the substructure of the ACL. Section 8.4 illustrates tractography of a healthy meniscus, shown both independently and overlaid onto the original MR images. Finally, Sections 8.5 and 8.6 present tractography visualisations of a partially torn ACL and a torn meniscus, with correlated arthroscopic findings providing validation of the imaging results.

These results involved both healthy and pathological volunteers who provided informed consent for their data to be published. Ethics approvals for these studies were granted by the Imperial College Research Ethics Committee (ICREC – ref: 22IC7857) and the Integrated Research Application System (IRAS – ref: 301889). In vivo healthy volunteer results from the present work have been published in *MRM*: Ristic et al. [90].

8.2 Tractography Visualisation

The collagen fibre orientation visualisations in Figures 7.2 and 7.3 are similar in nature to the data produced by DTI in the brain. DTI is an MRI technique that characterises the directional dependence of water diffusion within tissue. In DTI, the diffusion process at each voxel is modelled by a second-order tensor, a 3×3 symmetric positive semi-definite matrix that describes the magnitude of diffusion along different spatial directions. The eigenvectors of this tensor represent the principal diffusion orientations,

whilst the corresponding eigenvalues quantify the diffusivities along those axes. This representation captures anisotropic diffusion arising from microstructural organisation, such as axonal fibres in the brain, and provides the basis for tractography methods that reconstruct coherent fibre pathways across voxels.

To enable tractography of MADI data using DTI methods, collagen fibre orientation unit vectors were converted into synthetic diffusion tensors for processing in 3D Slicer [109] via the DTI module [116]. Each voxel's fibre orientation was assigned as the principal eigenvector of the tensor with eigenvalue 1. Two additional orthogonal eigenvectors were generated by taking the cross product of the fibre orientation vector with a random unit vector, followed by a second cross product to ensure orthogonality, and finally computing the third orthogonal vector as the cross product of the first two. Each diffusion tensor was then defined as per Equation 8.1:

$$\mathbf{D} = \mathbf{SMS}^T \tag{8.1}$$

Where **D** denotes the symmetric positive-definite diffusion tensor, **S** is the matrix whose columns comprise the three orthogonal unit eigenvectors, and **M** is the diagonal matrix containing the corresponding eigenvalues, ordered consistently with the columns of **S**. Since the two synthetic eigenvectors do not represent real data, their eigenvalues were set to 0.01 to minimise their influence on the tractography results. The synthetic diffusion tensors were loaded into 3D Slicer by means of a custom Python plugin to interface with the diffusion tensor pipeline at a stage after the tensors would have been generated from a standard diffusion weighted dataset. Once loaded, the tensors could be processed for tractography in the same manner as a standard DTI workflow.

Tractography is a valuable visualisation method as it converts local orientation data into a spatially continuous representation of fibre pathways. Rather than displaying voxel-wise unit vectors, which can be challenging to interpret in aggregate, tractography integrates them into streamlines that resemble the underlying anatomy. This approach facilitates the depiction of large-scale structures, the differentiation of substructures, and the identification of discontinuities, while also providing a visualisation format that is clinically familiar to medical professionals accustomed to DTI-based tractography.

8.3 Healthy ACL

This healthy ACL dataset is from the left knee of a 25-29 year old male, and is also the representative dataset used to showcase the individual capabilities of MADI components throughout this thesis. For a healthy ACL, the expected collagen fibre orientation is in alignment with the macroscopic tendon axis, as described in Section 2.5 and demonstrated in Section 7.2.3. This is further illustrated by the intensity versus θ plot in Figure 8.1, which shows the segmented intensities from each registered image in the representative dataset, plotted at the θ value between the final estimated fibre orientation and each $\vec{B}_0$ orientation, i.e. the intensity profile of each voxel is plotted with a different colour for each point based on what scan it was from. The representative dataset was acquired with $N = 7$ scans, and, given the predominantly straight morphology of the ACL, the data form distinct clusters reflecting the discrete $\vec{B}_0$ orientations at different angles to the ligament. The data also show conformity to the magic angle intensity equation, Equation 2.7, as illustrated by the black line modelled with $A = 3 \times 10^5$ and $B = 0.5$.

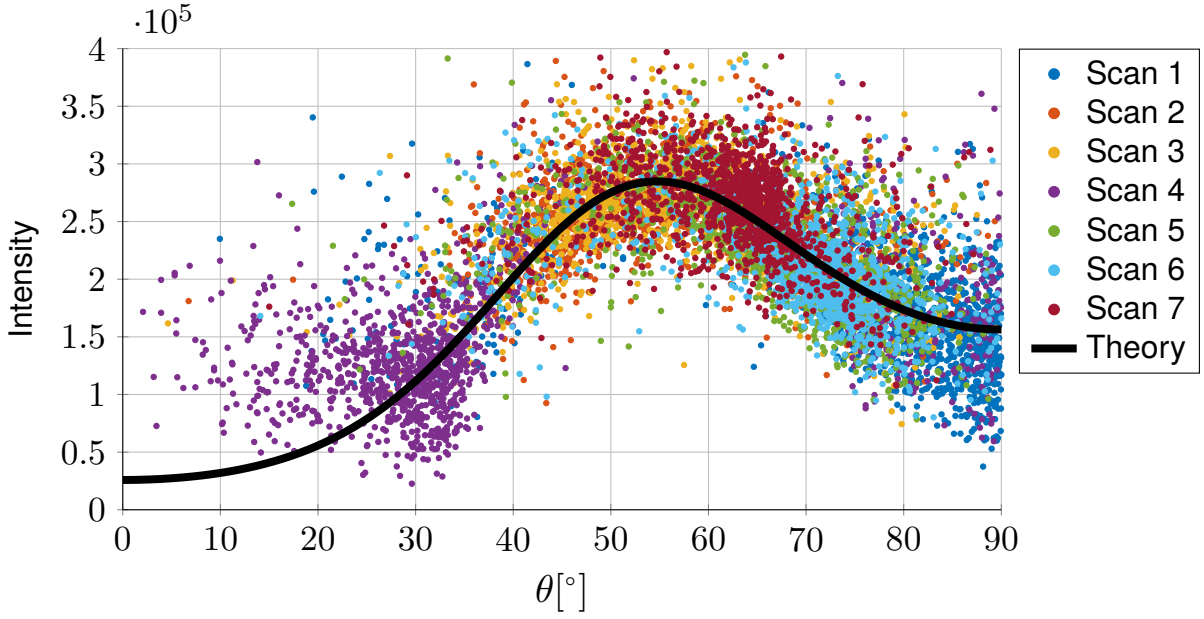


Figure 8.1: Intensity versus θ plot for the representative dataset. Data from each scan are shown as coloured points, the theoretical variation from Equation 2.7 is shown as a black line.

The roll and yaw angles for the representative dataset were chosen using the genetic algorithm from Section 7.3.3 and are given in Table 8.1. Each set of (θ_R, ϕ_Y) was rounded to the nearest 5° as this level of accuracy was sufficient to achieve the distribution

of $\vec{B}_0$ orientations. This dataset used a conservative yaw angle limit of $\pm 15^\circ$ as the limits of the system were still being explored when this dataset was acquired.

Table 8.1: Roll and yaw angles for the representative dataset.

Scan no.	**1**	**2**	**3**	**4**	**5**	**6**	**7**
Roll angle, θ_R	0°	30°	60°	90°	-30°	-45°	-60°
Yaw angle, ϕ_Y	0°	10°	-10°	15°	-10°	10°	10°

8.3.1 ACL Bundles

Figure 8.2 presents an alignment heatmap, where each estimated fibre orientation was mapped to a common origin and assigned a cone of influence of 5°. The surface of the sphere in Figure 8.2 was parametrised by 40,000 evenly spread points. If an estimated fibre orientation was within 5° of a point on the sphere its count was increased. Brighter areas correspond to a higher number of orientations pointing towards that region. The count of fibres pointing towards the points on the sphere was normalised by the highest count. Two prominent bright spots exist in Figure 8.2, potentially indicating the presence of a substructure that is not evident from the initial fibre plots in Figures 7.2 and 7.3.

Anatomically, the ACL is composed of distinct fibre bundles that provide stabilisation across knee flexion and extension, the principal bundles being the anteromedial and posterolateral [117]. To investigate this substructure, the centroids of the bright spots in Figure 8.2 were identified, and fibres from Figure 7.2 were segmented into two subgroups according to their proximity to each centroid. Fibres oriented more than 20° from either centroid were excluded from the analysis. The resulting subgroups are illustrated in Figure 8.3, where two distinct clusters of fibres are clearly visible, showing excellent agreement with anatomical theory.

Fibres in red are denoted as part of the anteromedial bundle, and fibres in blue are denoted as part of the posterolateral bundle as the red subgroup of fibres is positioned more anteriorly (more positive y-axis positioning) compared to that of the blue subgroup. The anteromedial bundle corresponds to the top right bright spot in Figure 8.2 and the posterolateral bundle corresponds to the bottom left bright spot.

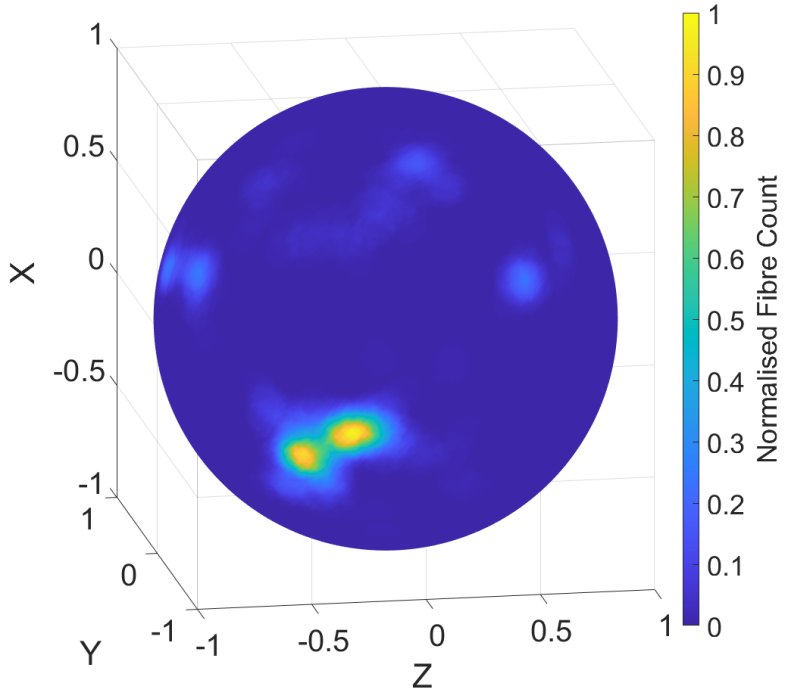


Figure 8.2: Heatmap of estimated collagen fibre orientations from the representative dataset. Brighter areas indicate more estimated collagen fibre orientations pointing towards that region.

The anteromedial and posterolateral fibre bundle substructures were converted to synthetic diffusion tensors using Equation 8.1 and imported into 3D Slicer for visualisation. Tractography was performed by seeding fibre tract growth regions at the centre of each bundle. Figure 8.4 shows the tractography results for the ACL from the representative dataset. The anteromedial bundle is again shown in red and the posterolateral bundle is shown in blue.

For MADI, this result represents an important benchmark as it establishes a conservative lower limit on the achievable angular resolution in vivo. The centroids of the bright spots in Figure 8.2 are separated by $\approx 15^{\circ}$, yet MADI successfully distinguished the two straight, nearly parallel collagen fibre bundles. Following the convention that two populations are resolved when they can be discriminated as distinct, this implies a practical lower angular resolution limit of $\approx 15^{\circ}$. The coherent alignment observed within each bundle further indicates good accuracy of the fibre orientation estimates. Such a result is valuable, as metrics such as resolution and accuracy are otherwise difficult to quantify non-invasively in vivo.

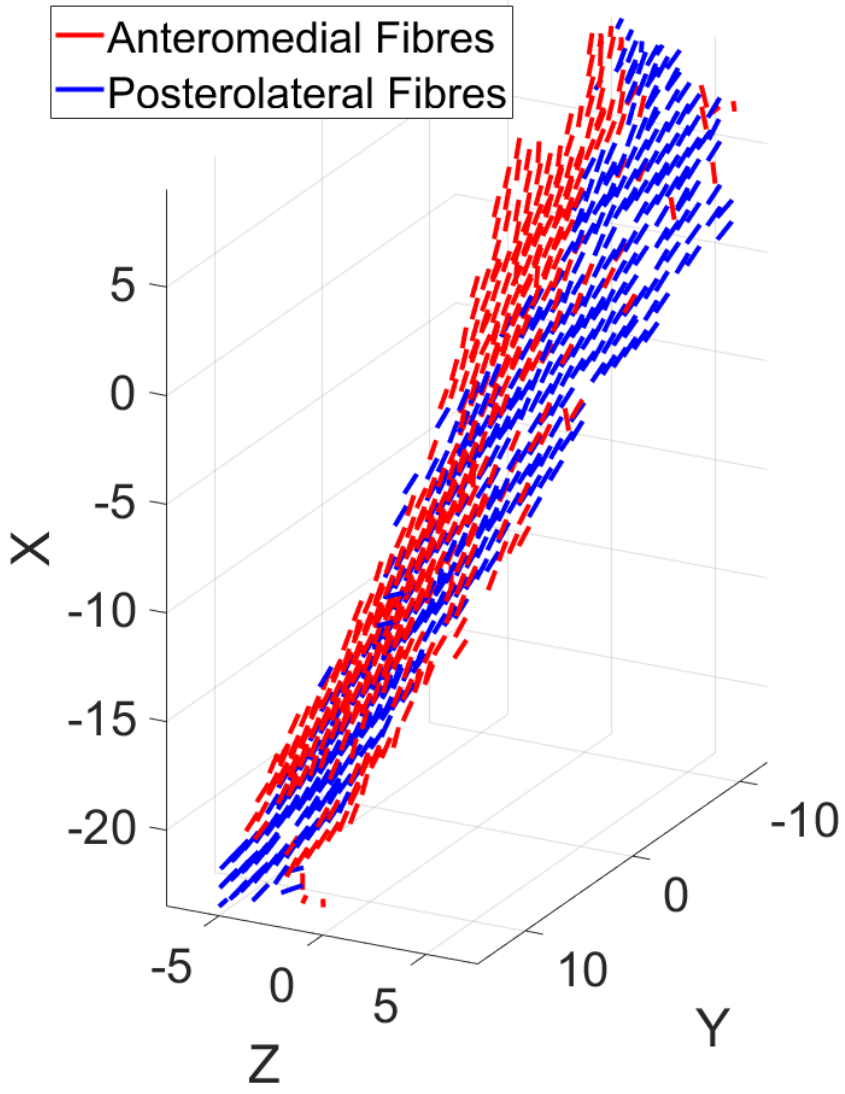
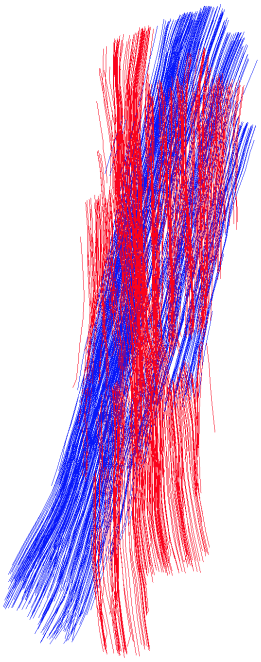
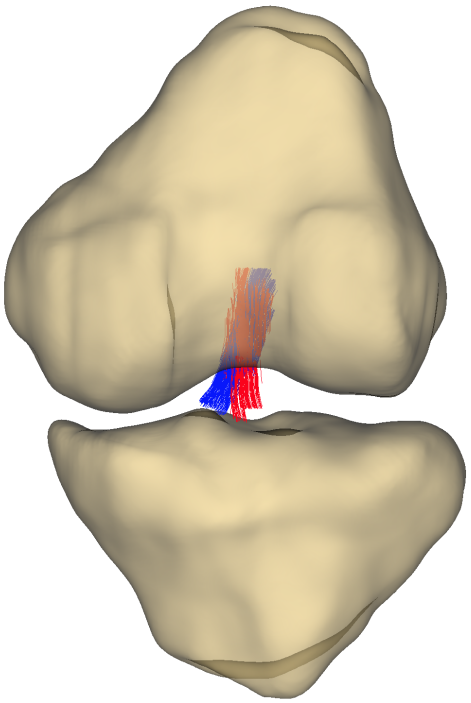


Figure 8.3: Fibre plot of anteromedial and posterolateral bundles of the ACL from the representative dataset. Anteromedial bundle fibres are shown in red, and posterolateral bundle fibres are shown in blue.



(a) Tractography results of the anteromedial (red) and posterolateral (blue) bundles of the ACL.



(b) Tractography results of the anteromedial (red) and posterolateral (blue) bundles of the ACL with bone segmentation to provide anatomical context.

Figure 8.4: Tractography results for the ACL and the anteromedial and posterolateral bundles of the ACL. Figure 8.4a shows the tractography results alone and Figure 8.4b shows the tractography results in the context of the joint with a bone segmentation.

8.4 Healthy Meniscus

This healthy meniscus dataset is from the right knee of a 20-24 year old male healthy volunteer. For a healthy meniscus, the expected collagen fibre orientation is in alignment with the curved structure of the meniscus, with fibres flowing circumferentially around the tissue structure as described in Section 2.5. Figure 8.5 shows the fibre plot of the meniscus, which shows the estimated collagen fibre orientations largely align tangentially to the curvature of the local tissue structure.

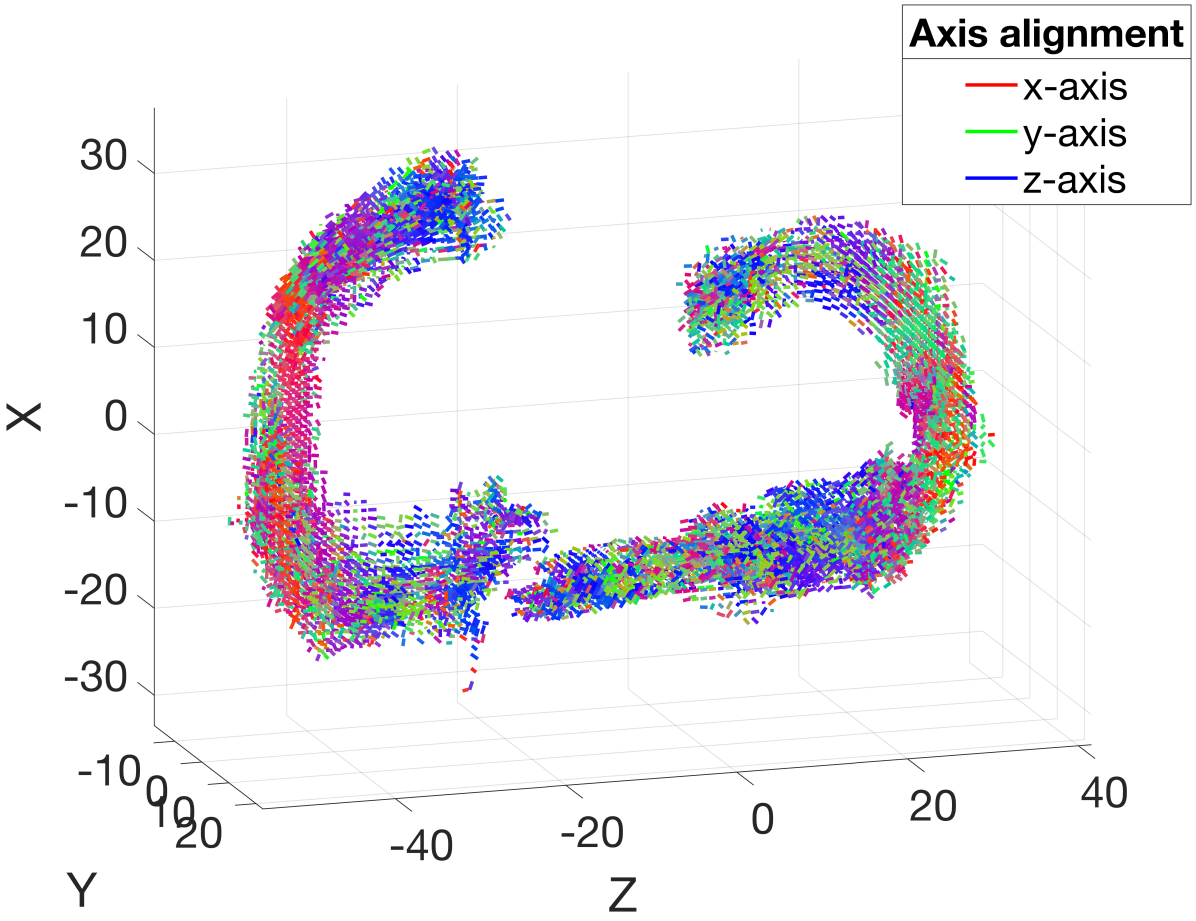


Figure 8.5: Fibre plot of the meniscus from a healthy volunteer.

The fibre colours in Figure 8.5 are encoded as in Figures 7.2 and 7.3, where each collagen fibre unit vector is mapped into RGB space according to its orientation. The colour distribution shows that the majority of fibres follow the curvature of the meniscus, which is predominantly planar in the xz-plane. Most fibres appear red, purple, or blue, reflecting transitions between alignment with the x- and z-axes as they arc around the structure. This pattern is consistent with the expected organisation of a healthy meniscus. Green fibres are also visible, though, as in Figures 7.2 and 7.3, these are likely

largely attributable to over-segmentation or registration-related interpolation artefacts near structural boundaries.

Table 8.2 shows the roll and yaw angles that were chosen using the genetic algorithm from Section 7.3.3. Unlike the representative ACL dataset, these angles were not rounded to the nearest 5° and yaw was restricted to $\pm 20^\circ$. For imaging the healthy meniscus, large yaw angles are not necessarily required for the bulk of the structure as low angle $\vec{B}_0$ orientations relative to the fibres are easily achievable with roll rotations of the scanner alone.

Table 8.2: Roll and yaw angles for the healthy meniscus dataset.

Scan no.	**1**	**2**	**3**	**4**	**5**	**6**	**7**
Roll angle, θ_R	0°	-32°	-67°	-58°	-20°	20°	24°
Yaw angle, ϕ_Y	0°	0°	-7°	16°	20°	15°	-14°

Figure 8.6 shows the tractography of the healthy meniscus dataset, seeded at the end points of each lobe. Fibres are coloured by their mean orientation across their length and exhibit the same overall pattern as in Figure 8.5, appearing predominantly red, blue, and purple as they follow the curved structure of the meniscus within the xz-plane. A small discontinuity is visible on the medial meniscus (left side of Figure 8.6), where fibres appear yellow and orange. This does not reflect a genuine structural disruption; rather, this region lies near the boundary of the magnetic field at certain $\vec{B}_0$ orientations, and the fibres are affected by spurious intensity data that corrupts the local intensity profiles.

This artefactual disruption has implications for the clinical utility of assessing the meniscus in the technique's current form. In a healthy subject, the spurious discontinuity produces a localised deviation in fibre orientation that may be misinterpreted as structural disruption, constituting a potential source of false positives. Conversely, should genuine pathology coincide with the affected region, it would be potentially indistinguishable from the artefact, risking a false negative. Both failure modes are clinically unacceptable for a diagnostic tool. Without a fundamental hardware change to increase the usable imaging volume, the most tractable mitigation is to generate a valid-data label map demarcating the usable field boundary at each $\vec{B}_0$ orientation. In this manner, corrupted voxels could be excluded from collagen fibre orientation estimation, as discussed in Section 9.5.3.



Figure 8.6: Tractography plot of the healthy meniscus dataset. Fibres are coloured according to their mean orientation. A small discontinuity is visible on the left side of the figure in the medial meniscus, yet all other fibres flow smoothly around the meniscal structure.

Figure 8.7 shows the tractography from Figure 8.6 overlaid onto central axial and coronal slices of the mean image from the registered MADI dataset (mean taken across the 4D collated image, similarly to the standard deviation image of Figure 6.2). This overlay provides anatomical context, enabling the tractography to be interpreted relative to surrounding structures such as bone, cartilage, and other connective tissues. Interactively adjusting the slice position relative to the tractography further aids in developing intuition about the trajectory of the meniscus fibres within the joint and the location of its insertion points on the bone.

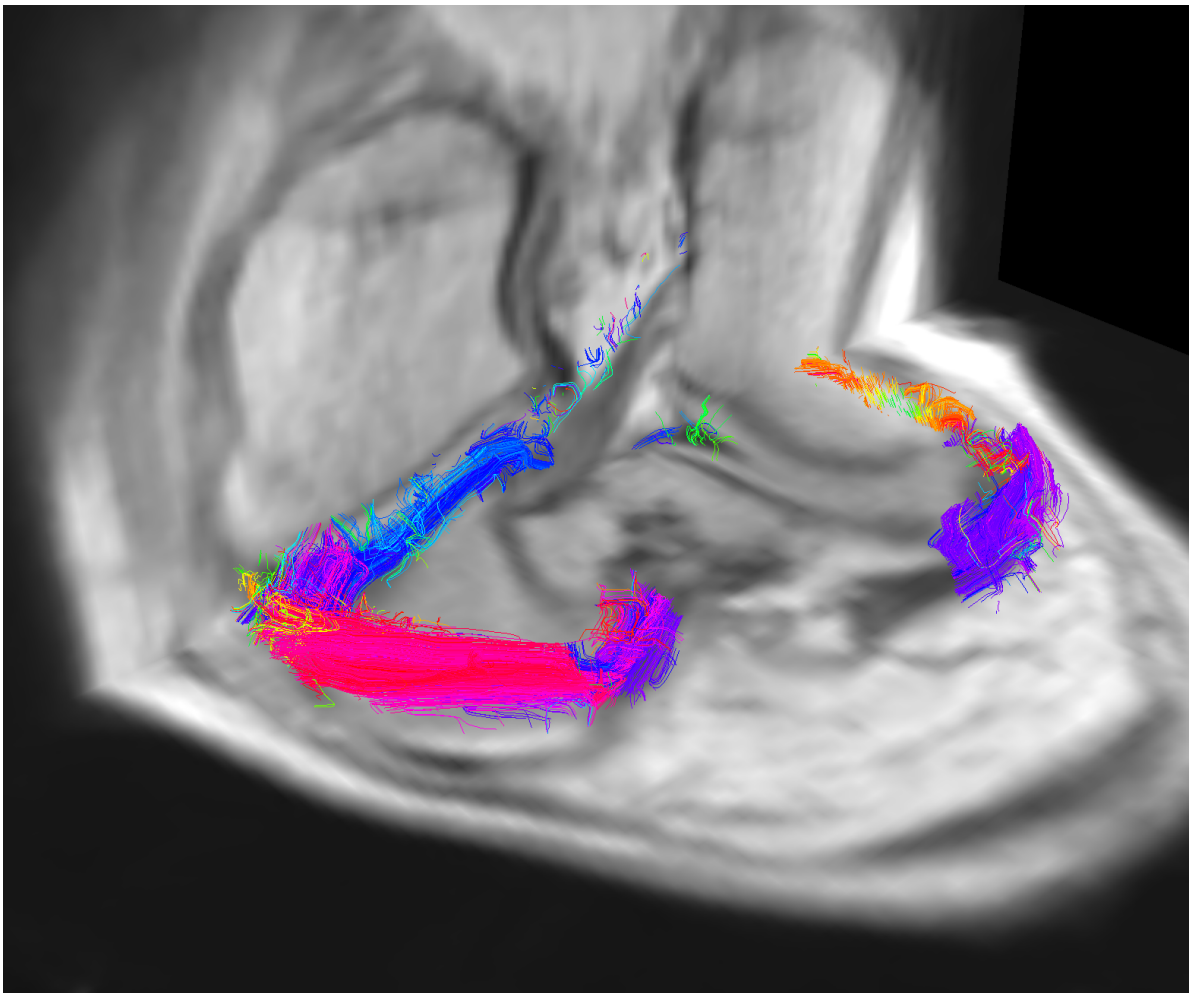


Figure 8.7: Tractography plot of the healthy meniscus dataset, overlaid onto the mean image from the registered MADI images.

8.5 Partially Torn ACL

This partially torn ACL dataset was acquired from the left knee of a 20-24 year-old male patient volunteer. The participant was scheduled for arthroscopy as part of routine clinical management, providing an opportunity to perform a MADI examination beforehand and obtain direct correlation between MADI findings and arthroscopic assessment. Figure 8.8 shows the fibre plot of the torn ACL. The fibres follow the macroscopic structure of the ligament along its length; however, an apparent lesion is visible midway where the ligament bifurcates, suggesting one branch remains attached to the femur whilst the other has become detached from its regular insertion point. At the local scale, the fibre orientations appear coherent along the ligament body, yet the bifurcation visible at the ligament scale indicates partial structural failure with a band of fibres separated from the main ligament body.

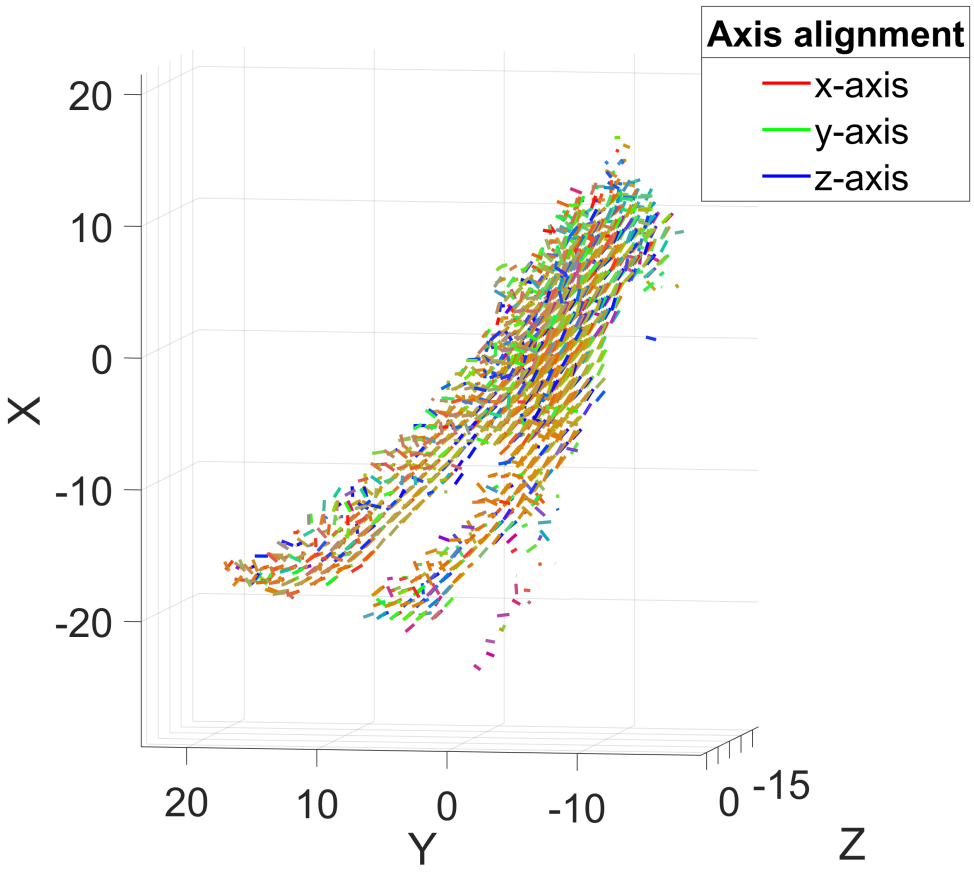


Figure 8.8: Fibre plot of the ACL from a volunteer with a partially torn ACL.

Figure 8.9 shows the tractography of the partially torn ACL overlaid onto the fixed image from the corresponding MADI dataset, oriented consistently with Figure 8.8 which is presented in the laboratory coordinate frame. Tractography was seeded at the ligament

insertions, with streamlines coloured by mean orientation. The intact portion of the ACL manifests as a coherent, uniformly coloured bundle of parallel fibres spanning from tibia to femur. The partial tear is indicated by a discrete bundle of fibres, approximately 3 mm in diameter, that diverges from the main body at the bifurcation and runs parallel to it without reconnecting. This is consistent with the detached branch identified in Figure 8.8. This finding was corroborated by arthroscopy, also overlaid in Figure 8.9, where the white arrow marks the origin of the separated bundle visible in both modalities. All arthroscopic examinations and their interpretation in relation to the MADI data were performed by Mr Chinmay Gupte and Dr Dimitri Amiras.

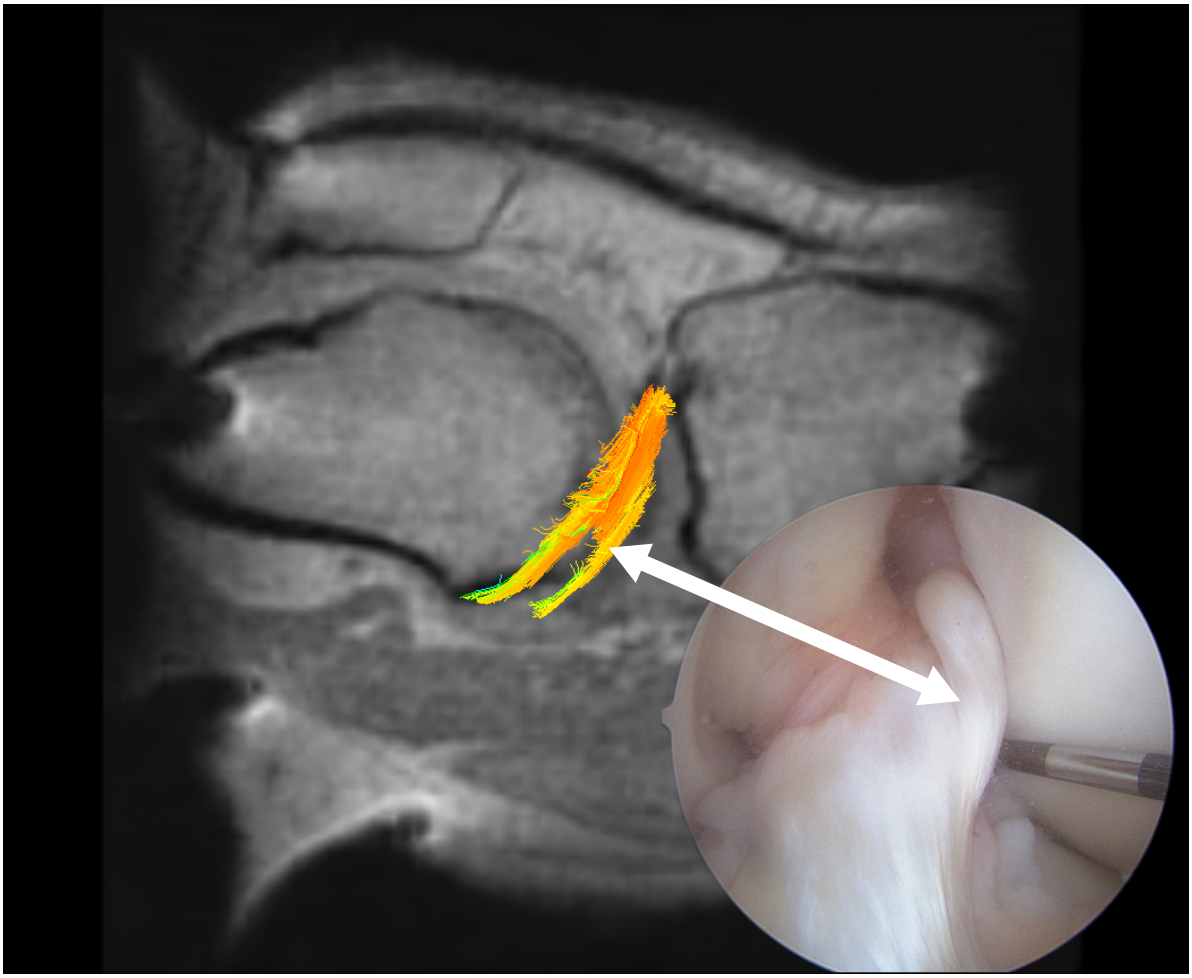


Figure 8.9: Tractography plot of a partially torn ACL, overlaid onto the original fixed image from the MADI dataset. An inset image from arthroscopy also shows the partial tear. A slim bundle of fibres runs parallel from the ACL due to a bifurcation in the ligament. This is also visible in the arthroscopy image as indicated by the white arrow.

This result demonstrates that MADI can detect pathologies also visible under arthroscopy, while providing enhanced visualisation of features that may be difficult to discern on MR images alone. For example, in Figure 8.9 the ACL appears largely intact on MRI, yet the

presence of a stray fibre bundle indicates structural damage. Detection of such loose fibres was further supported by the metric-based segmentation, where the data-driven approach assisted the operator in identifying tissue that might otherwise be overlooked. This establishes a baseline capability for MADI relative to arthroscopy, indicating that pathological features as small as approximately 3 mm in width can be detected.

Table 8.3 shows the roll and yaw angles used for the partially torn ACL dataset. These were the same roll and yaw angles as were selected for the healthy meniscus dataset in Section 8.4.

Table 8.3: Roll and yaw angles for the partially torn ACL dataset.

Scan no.	**1**	**2**	**3**	**4**	**5**	**6**	**7**
Roll angle, θ_R	0°	-32°	-67°	-58°	-20°	20°	24°
Yaw angle, ϕ_Y	0°	0°	-7°	16°	20°	15°	-14°

8.6 Medial Meniscus Tear

This partially torn meniscus dataset was acquired from the right knee of a 20-24 year-old male volunteer, the same participant described in Section 8.5 but imaged in the contralateral knee. The volunteer had sustained injuries to both knees and was scheduled for arthroscopy on the right knee as well, providing another opportunity to obtain direct correlation between MADI and arthroscopic findings, this time focused on the meniscus.

Figure 8.10 shows the fibre plot of the medial meniscus with a suspected medial meniscus tear. It is immediately apparent that the fibre structure is more difficult to assess than that of the ACL in Figure 8.8. A deliberately liberal manual segmentation was employed to ensure the tissues of interest were fully captured, at the cost of including surrounding voxels that do not contain organised collagen fibres. These voxels yield incoherent outputs from the CFOEA, since the CFOEA assigns an orientation to every segmented voxel regardless of whether coherent fibre structure is present; this produces the cloud of disordered fibres visible in Figure 8.10. Nevertheless, coherent fibre structure remains faintly discernible along the outer curvature of the meniscus, where fibres can be seen to run tangentially along the curved structure. The obscuring effect of the over-segmentation renders definitive assessment from the fibre plot alone unreliable. The tractography

in Figure 8.11 therefore provides a more reliable basis for assessing coherent fibre structure, as discussed below.

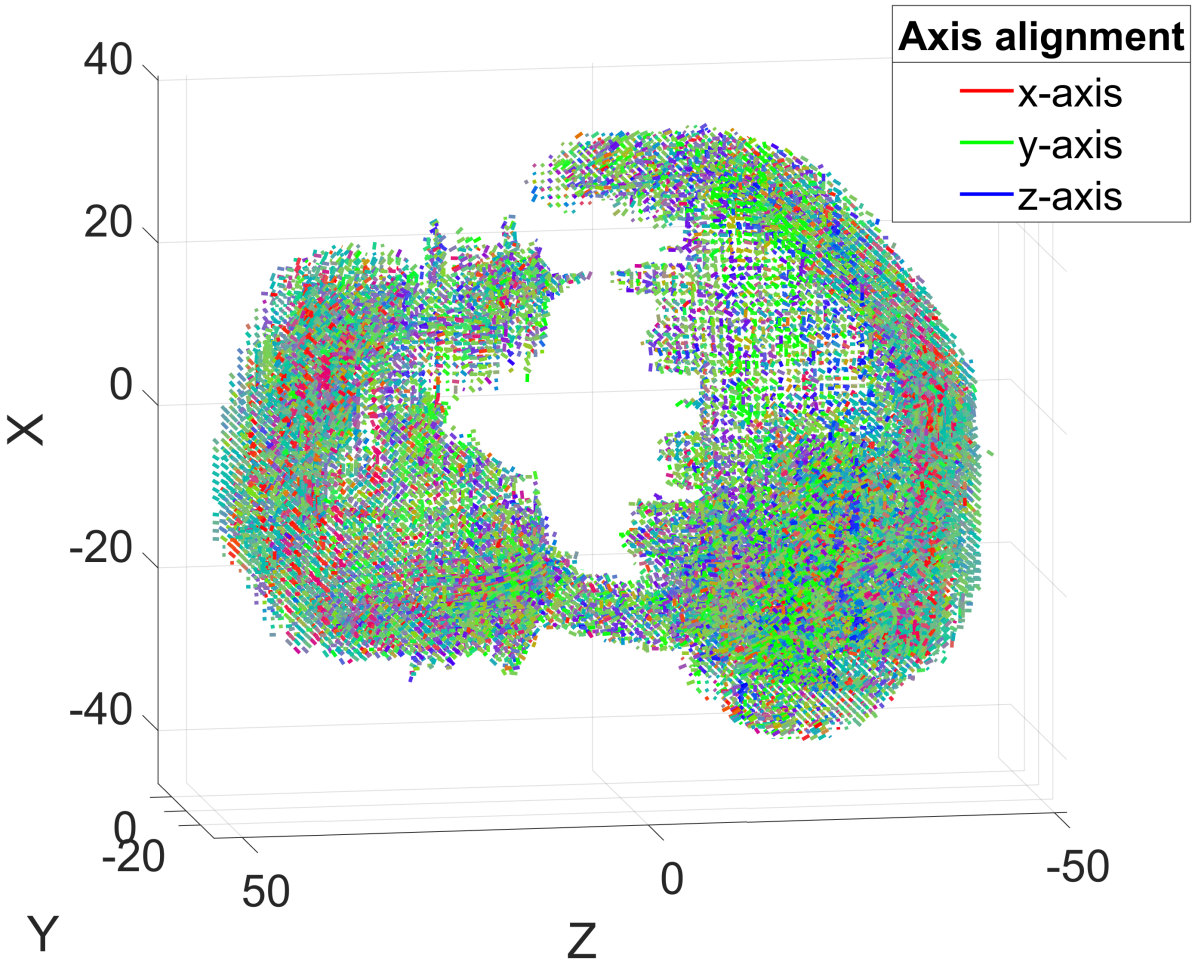


Figure 8.10: Fibre plot of the meniscus from a volunteer with a medial meniscus tear.

Figure 8.11 shows the tractography of the meniscus with a tear in the posterior region of the medial meniscus (lower right-hand portion of the figure). Tractography was seeded from broad regions at the ends of the meniscal lobes to ensure adequate coverage of coherent fibre regions, whose boundaries were ambiguous given the extent of the liberal segmentation. The fibres were coloured by mean orientation, appearing predominantly red, blue, and purple as in Section 8.4. The tractography is displayed against a black background to aid visualisation of the thin and sparse fibre tracts in the torn region, indicated by the white arrows and overlaid arthroscopic findings. The damage is evident in comparison with the intact medial meniscus in Figure 8.6, where the lower portion appears continuous and well-defined; here, the corresponding region exhibits a thin and frayed tract structure. Arthroscopy confirmed a free-edge tear of the medial meniscus, consistent with the region of sparse tractography. Selective visualisation of coherent structure was achieved by imposing minimum tract length and maximum curvature

thresholds, such that tracts originating in incoherent regions, which are too short or too curved to meet these criteria, are excluded from the output. The surviving tracts therefore represent only regions of sufficiently organised fibre architecture, revealing the gross meniscal structures alongside the sparse, near-linear tracts characteristic of the torn region. The tractography parameters used are listed in Appendix F. Once again, all arthroscopic examinations and their interpretation in relation to the MADI data were performed by Mr Chinmay Gupte and Dr Dimitri Amiras.

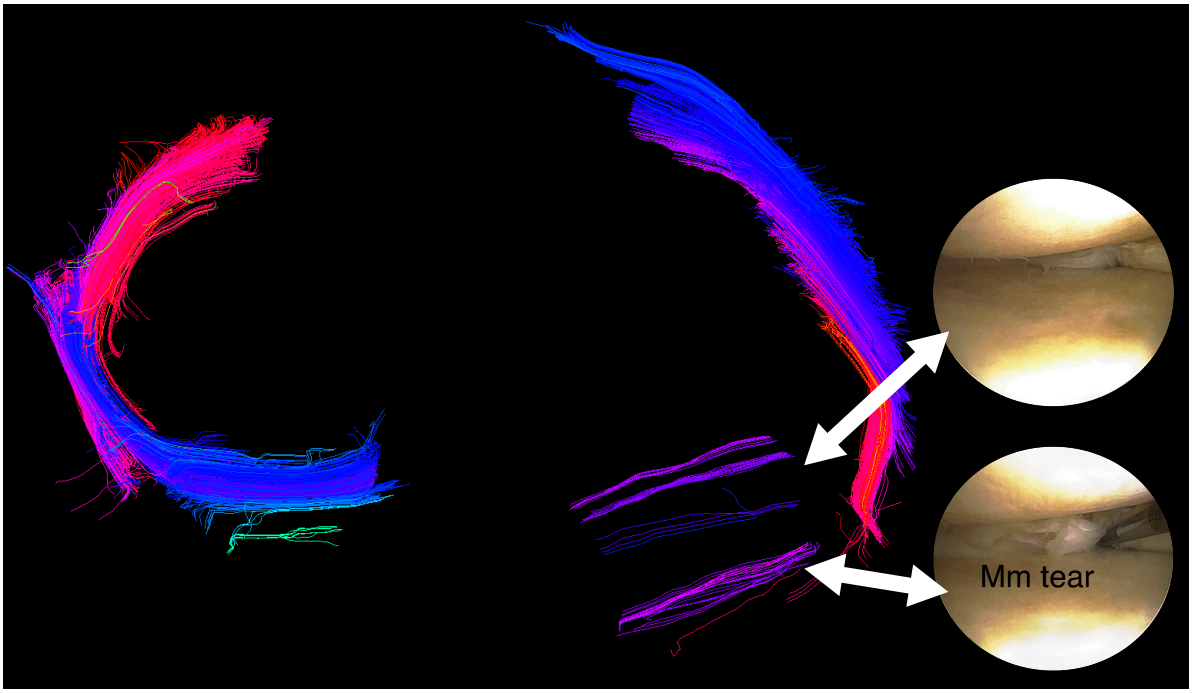


Figure 8.11: Tractography plot of a meniscus with a medial meniscus free edge tear. The tractography results are shown over a black background to better visualise the sparse fibre region in the lower right portion of the figure. Overlaid images from arthroscopy show the medial meniscus (Mm) tear in this region.

This result demonstrates that MADI can detect damage in thin, flat structures such as the inner free edge of the meniscus, as corroborated by the overlaid arthroscopic findings in Figure 8.11. A limitation of MADI, in contrast to arthroscopy, is its sensitivity to tractography seed placement and parameter selection. Appropriate seed regions must be defined to reliably resolve the structures of interest; in this case, the sparse, frayed region associated with the free-edge tear was largely disconnected from the adjacent medial meniscus body, such that insufficient seeding coverage could result in the torn structure remaining unresolved. Whilst an absence of tracts within a segmented region is itself indicative of pathology, it provides limited information about the detailed condition of the tissue. To address this, tractography was complemented with fibre

plots and metric-based volumes (e.g. standard deviation volumes) to ensure that all $\vec{B}_0$ orientation sensitive tissues were systematically examined prior to seeding. Broad seed regions were therefore selected to ensure adequate coverage of coherent fibre architecture, whilst minimum length and maximum curvature thresholds were applied to suppress spurious tracts arising from incoherent voxels.

Table 8.4 lists the roll and yaw angles used for the torn meniscus dataset, which were identical to those selected for the healthy ACL dataset in Section 8.3. The variation in scan orientation sets across different MADI examinations reflects the iterative development and refinement of the MADI technique over the course of this thesis.

Table 8.4: Roll and yaw angles for the medial meniscus tear dataset.

Scan no.	**1**	**2**	**3**	**4**	**5**	**6**	**7**
Roll angle, θ_R	0°	30°	60°	90°	-30°	-45°	-60°
Yaw angle, ϕ_Y	0°	10°	-10°	15°	-10°	10°	10°

8.7 Summary

This chapter presented an overview of the results and visualisations produced by the MADI technique. Section 8.2 introduced the tractography framework, in which synthetic diffusion tensors were generated from collagen fibre orientation unit vectors, enabling tractography plots familiar from DTI to be applied to anisotropic collagen-rich tissues such as the ACL and meniscus.

Section 8.3 provided a comprehensive analysis of the representative ACL dataset. The intensity versus θ plot demonstrated strong agreement between measured signal variation and theoretical predictions, with clustering of intensities reflecting the linear morphology of the ligament. Heat maps of orientation distribution revealed two distinct substructures, which, when segmented, corresponded to the anteromedial and posterolateral bundles of the ACL. These bundles were successfully reconstructed and visualised using tractography. In addition, MADI results were displayed within the context of bone segmentations, highlighting the potential of the technique to support surgical planning by situating fibre architecture relative to joint anatomy.

Section 8.4 presented results from a healthy meniscus. Both fibre plots and tractography demonstrated fibres following the curved meniscal structure, consistent with anatomical expectations. Tractography was also overlaid onto mean MR images in a multi-slice viewer, providing an interactive, anatomically contextualised visualisation that aids interpretation of fibre trajectories within the joint.

Pathological cases were described in Section 8.5 and Section 8.6. In the torn ACL dataset, MADI revealed a loose bundle of fibres separated from the main ligament, corroborated by arthroscopic findings. In the torn meniscus dataset, tractography identified a free-edge tear in the medial meniscus, again consistent with arthroscopy.

A notable characteristic of the MADI tractography workflow is the degree of user interaction required to generate the tracts relative to the generation of the fibre plots. Whilst the fibre plots are a direct voxel-level representation of the estimated fibre orientations, tractography is inherently sensitive to seed placement and parameter selection, introducing a degree of subjectivity into the visualisation. However, this is not necessarily a limitation. The interactive tractography environment in 3D Slicer permits iterative adjustment of seed regions and propagation parameters in real time. This enables a user with appropriate domain knowledge, such as a radiologist or orthopaedic surgeon, to interrogate the data in a targeted and exploratory manner. In this sense, tractography functions less as a fully automated quantitative measurement and more as an interactive visualisation tool, analogous in spirit to the way arthroscopy relies on the skill and experience of the operating clinician to interpret findings. The clinically significant observations demonstrated in this chapter, including the bifurcation of the torn ACL and the free-edge tear of the medial meniscus, were identified through this interactive process. Where precise, voxel-level quantification is required, the fibre plots and associated metric volumes provide a more objective and reproducible basis for analysis.

Together, these findings establish MADI as the first technique to exploit the magic angle effect for in vivo tractography of collagen-rich connective tissues, enabling microstructural fibre architecture to be resolved in tissues that are poorly characterised by conventional MRI. The ACL bundle analysis demonstrated a conservative lower bound of approximately 15° on orientation resolution, providing an empirical benchmark for the technique. The pathological cases confirmed that MADI can reveal clinically relevant disruptions of collagen architecture with direct correlation to arthroscopy, the current gold standard. In

this sense, MADI can be considered a form of virtual arthroscopy, providing structural information beyond that conventionally available in musculoskeletal MRI of the knee. These results therefore position MADI as a powerful and novel imaging method with strong potential for clinical translation, whilst laying the groundwork for future blinded studies to rigorously assess its diagnostic performance relative to arthroscopy.

9 | Conclusions and Future Work

9.1 Chapter Overview

This chapter provides concluding remarks for the thesis as a whole and outlines potential future directions for MADI and rotatable MRI in general. It begins with an overview of the MADI technique and its constituent components, highlighting the key contributions and results from each section in Section 9.2. A summary of the in vivo findings is then presented, with conclusions regarding the performance of both MADI and the low-field rotatable scanner in Section 9.3. Limitations of the technique and the system are discussed in Section 9.4, followed by an exploration of potential avenues for future development in Section 9.5. The chapter and thesis then conclude with final reflections in Section 9.6.

9.2 MADI Technique Summary

This thesis is concerned with the development of the Magic Angle Directional Imaging technique to perform *In Vivo Imaging of Collagen Fibre Structures Using Dedicated Low-Field MRI*.

To this end, MADI was implemented using isotropic volume acquisitions of a static subject at $N = 7$ orientations of a low-field rotatable MRI scanner. The images were registered with a custom pipeline parametrised by scanner telemetry, and regions of interest were segmented semi-automatically with a dedicated framework. Collagen fibre orientations were then estimated using an improved algorithm with substantially higher performance than previous methods, enabling data-driven optimisation of scan orientation sets using a genetic algorithm. This workflow culminated in the first successful in vivo application of MADI, achieving non-invasive visualisation of collagen fibre microstructure using the magic angle effect.

9.2.1 Low-Field Rotatable MRI System

The 0.15 T permanent-magnet rotatable MRI system developed for MADI was successfully used to image connective tissues in vivo, achieving sufficient SNR with clinically tolerable acquisition times. Scanner orientation was controlled by custom motion-control software developed as part of this thesis, with patient safety as a core design principle.

9.2.2 Registration

A custom registration pipeline was developed using the elastix framework, implementing a multi-stage, telemetry-driven approach. The pipeline corrects image misalignment arising from scanner reorientation, and was extended to address systematic geometric distortions using a custom 3D printed calibration phantom and simulated ground truth scans. This geometric correction method is highly extensible and applicable to any imaging device with a sufficiently small FOV to permit 3D printing of a phantom.

9.2.3 Segmentation

A metric-assisted segmentation framework was developed to reduce operator burden by enabling rapid ROI segmentation. An initial bounding box provided coarse segmentation, which was then refined using standard deviation across the MADI dataset as a thresholding metric. The framework was implemented as a GUI to support rapid iteration, with extensibility and interoperability as design goals. Outputs could be further refined in standard medical imaging software such as 3D Slicer [109].

9.2.4 Collagen Fibre Orientation Estimation

Collagen fibre orientation was estimated for each voxel in the ROI using intensity profiles from the 4D registered and segmented MADI dataset. As in previous studies, correlation was used as a similarity measure [11, 14, 15], but the updated CFOEA was significantly more performant than all previous implementations. This efficiency boost enabled the algorithm to be used as an experimental design aid. Fibre orientation estimation accuracy was mapped using the updated algorithm and subsequently this measure was incorporated as the fitness function of a genetic algorithm to optimise scan orientation set selection.

9.3 In Vivo Results

The in vivo results presented here are the first demonstration of non-invasive collagen fibre microstructural assessment using the magic angle effect in MRI. Whilst DTI has previously been applied to non-invasive collagen fibre assessment in research settings, it faces fundamental limitations that have prevented clinical translation, as discussed in Section 3.3. Arthroscopy therefore remains the diagnostic gold standard for connective tissue microstructure, as outlined in Section 3.2. MADI represents a distinct approach that circumvents these limitations, and its successful in vivo demonstration could mark a significant step towards non-invasive MSK assessment beyond the capabilities of both conventional MRI and DTI. Expected healthy tissue results were obtained, showing collagen fibre alignment with the macroscopic structures of the ACL and meniscus. Initial assessments of pathology were also performed, showing good agreement with arthroscopic findings, highlighting how MADI offers a form of "virtual arthroscopy" insight beyond standard MSK MRI examinations.

In total, 30 volunteers were scanned with the low-field rotatable system using the MADI technique. Participants included 10 patients scheduled for arthroscopy, 14 self-selected healthy volunteers, and 6 self-selected volunteers with pathology of varying severity. Whilst this cohort does not provide sufficient statistical power to formally assess the accuracy, reliability, sensitivity, or specificity of MADI, it establishes a strong foundation and clear motivation for future larger-scale validation studies.

9.4 Limitations

The present work faces several limitations spanning cohort composition, data availability, and hardware constraints. The patient cohort was recruited via multiple pathways: self-selected volunteers (both healthy and with suspected pathology), and patients undergoing scheduled arthroscopy at a private orthopaedic practice. Formal diagnoses were unavailable for the majority of participants, and where pathology was present, neither its nature, chronicity, nor mechanism of injury was systematically documented. This precludes robust statistical correlation between MADI estimated fibre orientations and confirmed pathological diagnoses. Taken together, these factors constrain the present study to a proof-of-concept investigation; nonetheless, the results demonstrate physio-

logically consistent fibre orientation maps and preliminary agreement with arthroscopic findings, providing a sound basis for the prospective clinical validation described in Section 9.5.6.

Additional constraints arose from the scanner hardware. The limited DSV meant the knee joint only marginally fit within the homogeneous field volume: at certain scan orientations, meniscal data were occasionally corrupted by B_0 inhomogeneity at the DSV boundary, though ACL imaging was unaffected. The restricted yaw range of the scanner further limited the achievable range of fibre orientations that could be reliably assessed, affecting tissues with predominantly axial fibre orientations, such as the patellar tendon, Achilles tendon, and wrist tendons[1]. Procedural constraints required the first acquisition to always be performed at the home orientation; this could in principle be relaxed, as scanner telemetry provides sufficient pose information to recover a conventional reference frame from any starting orientation. Finally, the inherently low SNR per unit time of the low-field design resulted in extended examination durations; this could be mitigated through imaging acceleration strategies such as parallel imaging or compressed sensing-based undersampling. Future work to address these limitations is discussed in Sections 9.5.1 and 9.5.2.

The low-field design also imposed resolution constraints, with $1 \times 1 \times 1$ mm^3 isotropic voxels representing the practical upper limit for clinically reasonable scan durations. Whilst adequate for the fibre architecture of ligaments and menisci, this resolution is insufficient to resolve the zonal collagen structure of articular cartilage, which is typically 2-4 mm thick and exhibits a continuous variation in fibre orientation from tangential at the articular surface to radial at the osteochondral interface [118]. At $1 \times 1 \times 1$ mm^3, each voxel encompasses multiple collagen zones, such that magic angle signal contributions from zones of differing orientation undergo partial volume averaging, precluding zone-specific orientation estimation. Resolving zonal architecture would require voxel dimensions of approximately $0.3 - 0.5 \times 0.3 - 0.5 \times 0.3 - 0.5$ mm^3 isotropic [118–120], which is infeasible with the current prototype. Extending MADI to articular cartilage therefore represents a longer-term objective contingent on hardware development, and is discussed further in Section 9.5.1.

[1] These tissues were not assessed here but are a natural target for future MADI examinations.

Most of the limitations described above stem from the novelty of the hardware and technique rather than from fundamental physical or methodological barriers, and are therefore addressable through iterative development. The proof-of-concept results presented here demonstrate that the core MADI principle is experimentally validated and clinically translatable, providing a firm foundation for the targeted improvements outlined in Section 9.5.

9.5 Future Work

The work presented in this thesis establishes MADI as a viable technique for non-invasive collagen fibre microstructure assessment in vivo, demonstrating proof-of-concept across both healthy and pathological tissue in a human cohort. Based on this foundation, this section presents a roadmap of future work for the continued development and clinical translation of the MADI technique and the prototype rotatable low-field scanner. This development roadmap spans hardware, acquisition and pulse sequence optimisation, reconstruction and image processing, examination protocol, and broader system utility.

9.5.1 Hardware

As discussed in Sections 4.4, 7.3 and 9.4, the scanner yaw range was restricted to $\phi_Y \in [-30^\circ, 30^\circ]$ by a combination of patient safety and comfort requirements, and the orientation-dependent signal attenuation inherent to a solenoid receive coil. Revised examination protocols, informed by improved characterisation of patient ergonomics, may permit higher yaw angles whilst maintaining safety and comfort, as discussed in Section 9.5.4.

The yaw-dependent signal attenuation can be mitigated with a phased array receive coil. Li et al. [121] describe a design comprising three elliptical receive loops whose containing planes are mutually orthogonal; when wound on a cylindrical former, this geometry places the loop centres at 120° azimuthal offsets around the cylinder axis. The combined sensitivity of the three channels is designed to be invariant with respect to the orientation of $\vec{B}_0$ relative to the coil axis, theoretically extending the usable yaw range to $\phi_Y \in [-90^\circ, 90^\circ]$. A prototype implementation of this design is already in

development; Figure 9.1 shows a custom 3D-printed coil housing with receive loops installed, undergoing preliminary testing at the time of submission.

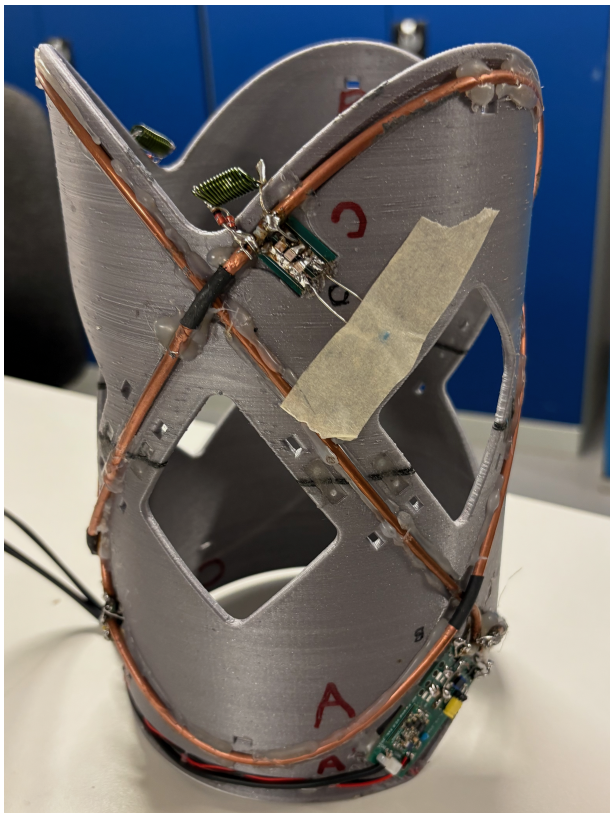


Figure 9.1: 3D printed phased array receive coil whose total signal strength from the three channels combined is invariant with $\vec{B}_0$ orientation relative to the coil central axis.

Whilst ongoing development of the prototype scanner described by McGinley et al. [16] will continue to extend its capabilities, certain limitations are most effectively addressed through a second-generation design iteration of the scanner. A larger DSV would expand the imaging FOV, improve B_0 homogeneity across the joint, accommodate a wider range of patient anatomies, and simplify patient positioning, thus collectively broadening the range of feasible clinical studies beyond the knee. An increase in static field strength would improve SNR per unit time and permit higher spatial resolution, though this must be balanced against the engineering design constraints of a rotatable, open-bore geometry. Further performance improvements could be realised through enhanced gradient hardware and an increased number of receive channels to support more pulse sequences and parallel imaging acceleration. In the longer term, the relatively compact form factor of the rotatable scanner compared to conventional MRI systems presents a credible development pathway towards a portable device or installation in point-of-care settings.

9.5.2 Acquisition and Pulse Sequence Optimisation

A limitation of the acquisition platform was the restricted choice of pulse sequences available on the prototype scanner. The 3D FLASH sequence was selected on the basis of pulse sequence optimisation work by Chappell [34] at 3T, which aimed to maximise the observable magic angle effect; this sequence was subsequently adopted for the low-field system under the assumption that the optimal contrast mechanism would transfer across field strengths. Whilst this proved sufficient to achieve in vivo visualisation of collagen fibre orientations, alternative pulse sequences may offer improvements in SNR efficiency, spatial resolution, or image contrast. Some sequences may enable qualitatively different contrast mechanisms, such as low-field DTI exploiting the magic angle effect to boost T2, following the approach from Wengler et al. [59, 60].

The primary candidate pulse sequences for future assessment are bSSFP, TSE, and UTE, as introduced in Section 4.3.9. bSSFP preserves transverse spin coherence in the steady state, yielding theoretically superior SNR efficiency per unit time [91], a particularly relevant advantage at low field, where SNR is intrinsically reduced relative to clinical field strengths. Since T1 is relatively insensitive to fibre orientation in anisotropic collagen-rich tissues [8], the T2/T1-weighted contrast of bSSFP retains $\vec{B}_0$ orientation dependence, and the magic angle effect should remain detectable. The susceptibility-induced ΔB_0 that drives off-resonance banding artefacts in bSSFP is moreover suppressed at low field [21–23], making it a particularly attractive candidate for this application.

TSE acquires multiple k-space lines per TR via a train of refocusing pulses, reducing scan time relative to single-echo sequences for equivalent resolution and SNR [92]. As a spin echo-based sequence, TSE is sensitised to T2 rather than T2*, which at high field confers a meaningful advantage by refocusing static ΔB_0 induced dephasing; however, at low field this distinction is marginal, since reduced susceptibility effects cause T2 and T2* to converge. A further constraint is that the minimum achievable effective TE in TSE is bounded by the echo train timing, which may compromise signal retention from the short-T2 collagen-rich tissues. TSE nonetheless warrants investigation.

UTE sequences, capable of achieving echo times of the order of 100–500 μs, are well-suited to imaging collagen-rich tissues whose T2 can fall to the order of 1 ms or below, particularly at orientations away from the magic angle [8, 25]. In this regime, FLASH and

TSE sequences suffer substantial signal loss prior to echo acquisition, whereas UTE retains signal from these ultrashort-T2 components. The orientation dependence of T2 in collagen-rich tissues is preserved under UTE acquisition [93], making it a candidate for magic angle-based fibre orientation mapping at higher sensitivity than is achievable with conventional gradient echo sequences.

Implementing new pulse sequences on the prototype low-field scanner requires close collaboration with MR Solutions (Guildford, UK) to deploy them on the EVO spectrometer platform. Such access would enable pulse sequence optimisation to be performed directly on the prototype system, allowing sequence parameters to be tuned against image quality objectives, such as SNR or contrast-to-noise ratio (CNR), as well as against the sensitivity of the B parameter from Equation 2.7, which depends on both tissue properties and acquisition parameters. Broader pulse sequence availability would also facilitate investigation of qualitatively distinct contrast mechanisms; of particular interest is the combination of diffusion-weighted acquisitions with magic angle-enhanced orientation sensitivity using the rotation capabilities of the scanner, following the approach of Wengler et al. [59, 60], which could substantially extend the diagnostic scope of the system.

9.5.3 Reconstruction and Image Processing

Compressed Sensing

The challenges of extended acquisition time and low SNR inherent to the low-field prototype may be mitigated through compressed sensing (CS) and deep learning-based reconstruction, as discussed in Section 3.5 [21–23]. CS exploits the sparsity of MRI signals in appropriate transform domains, such as wavelet or finite difference (for total variation regularisation) representations, to reconstruct complete images from under-sampled k-space data, thereby reducing acquisition time without a proportional SNR penalty. Undersampling is a particularly attractive area of future work as it would directly reduce the examination duration, which is currently one of the principal barriers to clinical translation. Deep learning-based reconstruction offers a complementary approach, learning image priors from training data to further improve reconstruction quality at high undersampling factors [23]. Together, these strategies could reduce examination

duration, enable acquisition of additional orientations within the same time window, or permit higher spatial resolution, all of which would directly benefit the MADI framework.

Denoising

As discussed in Section 4.4.3, denoising represents a further avenue for methodological refinement. Block matching denoising was selected for the present work as a well-validated approach appropriate for a feasibility study; whilst its effect on the CFOEA was shown to be marginal in Section B.1, denoising remains important for image interpretability and for supporting downstream segmentation workflows. A natural near-term extension would be to apply BM4D to the concatenated multi-orientation volume rather than to each volume independently, exploiting inter-volume similarity to improve noise suppression. Comparison against alternative single-volume denoising schemes, including non-local means and deep learning-based approaches, also represents a worthwhile direction for future work.

The more substantive opportunity, however, lies in exploiting the multi-orientation acquisition structure of MADI. Each MADI dataset comprises multiple volumetric acquisitions of the same anatomy, which introduces redundancy across the acquisition dimension that is directly analogous to the redundancy across diffusion-encoding directions in DTI. This structure satisfies the low-rank assumption underlying Marchenko-Pastur PCA (MPPCA) denoising described by Veraart et al. [97]: the magic angle signal model constrains the MADI signal across orientations to a low-dimensional manifold parameterised by a small number of quantities (principally B and θ from Equation 2.7), whilst thermal noise distributes energy across all principal components. MPPCA exploits this separation via the Marchenko-Pastur distribution of noise eigenvalues to perform data-driven rank thresholding, removing noise-dominated components whilst preserving signal. Olesen et al. [98] extend this framework to genuinely multidimensional data, such as a multi-contrast dataset, using higher-order singular value decomposition (HOSVD), retaining the tensor structure of the dataset rather than flattening it to a matrix. This is a natural fit for MADI, where the data is analogous to a multi-contrast dataset due to the magic angle effect intensity variation across the acquisition dimension.

Segmentation and Valid Data Maps

The MADI framework offers several opportunities to advance segmentation methodology beyond the approaches employed in the present work, as discussed in Chapter 6. One avenue is to exploit the voxel-wise standard deviation metric as a segmentation prior. This metric serves as a biomarker for tissues exhibiting the magic angle effect; it could therefore be used to construct tissue-label maps that provide an additional input channel for learning-based segmentation models to be applied to MADI data. A second avenue is tractography-guided segmentation, in which an initial liberal segmentation mask is used to seed fibre tractography, and the resulting tracts are subsequently pruned based on geometric constraints such as minimum and maximum length, and curvature. Voxels through which no valid tract passes would then be excluded from the final segmentation, providing an anatomically informed refinement of the initial mask in an iterative fashion.

A further practical improvement concerns the construction of valid data maps for each scan orientation. In the present work, voxels near the boundary of the DSV are susceptible to B_0 inhomogeneity and gradient nonlinearity, which can corrupt the signal at certain orientations, as observed in the meniscal data discussed in Section 8.4. Rather than excluding an entire orientation acquisition when edge-of-field effects are present, a per-voxel validity mask could be generated for each scan orientation, allowing the CFOEA to proceed using only the valid subset of acquisitions on a voxel-wise basis. For a seven-orientation acquisition in which one orientation is partially corrupted, affected voxels would be estimated using six valid data points, improving clinical utility and reducing artefactual disruption to the estimated fibre orientations. This approach becomes increasingly tractable as the examination protocol matures and the set of acquired orientations becomes sufficiently standardised to permit pre-computation of expected validity maps.

9.5.4 Examination Protocol

Accumulated experience from in vivo scanning has improved understanding of the ergonomic constraints of the rotatable low-field scanner. In particular, it is preferable to yaw the scanner yoke such that its front edge moves away from the contralateral leg, which rests outside the scanner, as this direction of rotation is less ergonomically restrictive for the patient. This asymmetry becomes a more significant design constraint

as the yaw range is extended, since greater yaw rotations bring the scanner yoke into closer proximity with the contralateral leg. Examination protocols for left and right leg imaging therefore require mirrored operational envelopes, with mirrored protocols derived for each leg.

These ergonomic considerations can be incorporated directly into the genetic algorithm optimisation framework described in Section 7.3.3 through refinement of the fitness function. Specifically, the feasible orientation space could be defined by an asymmetric operational envelope that encodes the contralateral leg clearance constraint as a hard boundary on permissible yaw angles. Additional regularisation terms could be introduced to penalise large inter-orientation angular displacements, reducing operator workload during the examination by minimising the inter-orientation rotation of scanner repositioning steps. A further regularisation term could be applied to the angular deviation map to encourage globally smooth fibre estimation accuracy across the imaging volume.

With increasing familiarity with the in vivo examination protocol, patient positioning procedures will become better understood, enabling further protocol refinements. The current requirement to acquire the first scan at the home orientation, imposed to ensure reliable anatomical reference, could in principle be relaxed. Scanner telemetry provides sufficient information to recover a conventional reference frame computationally from any starting orientation. Removing this constraint would decouple the protocol optimisation from a fixed reference acquisition, enabling fully unconstrained, data-driven selection of scan orientations by the genetic algorithm and potentially improving the optimality of the resulting angular sampling. Alternatively, in cases where a priori anatomical knowledge is available, a deliberate choice of initial scan orientation could be made to ensure peak magic angle intensity is sampled in the target tissues [115].

9.5.5 Broader System Utility

Beyond the MADI technique, the low-field rotatable scanner may also have standalone clinical utility as a dedicated system. Its rotation capability enables multi-orientation acquisitions from which derived image representations, unavailable from conventional fixed-orientation scanners, can be computed. By acquiring datasets at multiple orientations, minimum, mean, and maximum projections of the registered images can be computed. The minimum projection could offer insight into fluid or fatty ingress into

collagen-rich structures, serving as a marker of pathology; the mean projection could provide a representation of anatomy largely free from the magic angle effect; and the maximum projection offers a practical realisation of the magic angle contrast imaging concept proposed by Bydder et al. [8] and Wengler et al. [59, 60], which has previously been constrained by the patient positioning limitations of conventional scanners. Figure 9.2 illustrates this concept applied to the representative dataset.

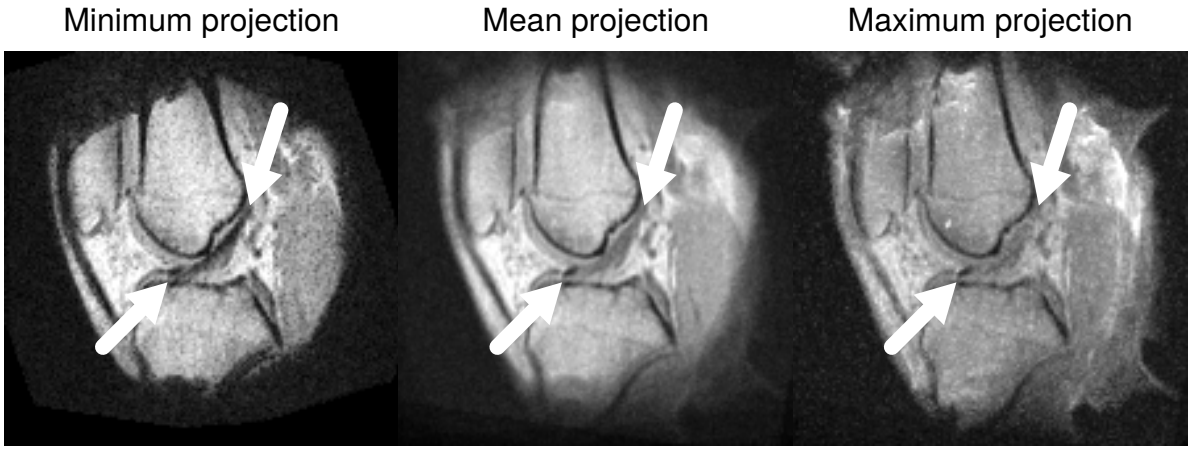


Figure 9.2: Minimum, mean, and maximum intensity projections of a knee scan from the rotatable low-field scanner. The ACL is marked in all images with white arrows and shows markedly different intensity projections in each image.

9.5.6 Clinical Validation

The most important future work for clinical translation of MADI is prospective clinical validation. Validation studies will necessarily employ larger, more homogeneous cohorts than were feasible in the present proof-of-concept work, with strict inclusion criteria encompassing age, injury type, mechanism of injury, chronicity, symptom duration, and prior surgical history. Patients scheduled for arthroscopy or ligament reconstruction surgery provide the most informative cohort, as surgical findings constitute a direct ground truth for tissue pathology. Complementary imaging with conventional 3T MRI should be included to enable direct comparison of diagnostic performance between modalities. Blinding the MADI analyst to arthroscopic and 3T findings, and the reporting radiologist to MADI results, would allow independent assessment by each modality. This would support computation of standard diagnostic accuracy metrics including sensitivity, specificity, and area under the receiver operating characteristic curve. Inter- and intra-rater reliability of MADI estimated fibre orientations, and scan-rescan repeatability, should also be characterised as prerequisites for clinical adoption. The contralateral

limb provides a natural internal control and should be incorporated into the study design where feasible.

Longitudinal studies represent an equally compelling topic of future work, following the approach of Van Dyck et al. [57, 63]. Tracking the structural evolution of ligament repairs and tendon grafts over time would position MADI as a quantitative tool for monitoring tissue remodelling, with direct clinical applications including objective assessment of graft ligamentisation timeline, criteria for return-to-sport based on collagen maturity, and longitudinal monitoring of the stages of tendon, ligament, or meniscal degeneration. These applications would substantially extend the clinical scope of MADI beyond initial diagnosis and towards long-term patient management, representing a compelling case for its adoption as a complementary modality within existing MSK imaging pathways.

9.6 Final Conclusions

Connective tissues composed of anisotropic collagen fibres present a fundamental challenge for conventional MRI: their signal intensity is typically very low, and also strongly orientation-dependent due to the magic angle effect, rendering standard imaging protocols unreliable for microstructural assessment. Arthroscopy remains the diagnostic gold standard for connective tissue pathology, yet it is invasive and costly. This thesis, *In Vivo Imaging of Collagen Fibre Structures Using Dedicated Low-Field MRI*, addressed this gap by developing and demonstrating Magic Angle Directional Imaging, a technique that exploits, rather than suppresses, the magic angle effect to extract quantitative collagen fibre orientation information non-invasively. The result was the first successful in vivo assessment of collagen fibre microstructure using the magic angle effect in MRI.

This achievement rested on four interconnected technical contributions: a telemetry-driven multi-stage image registration pipeline; a metric-assisted segmentation framework; a substantially improved collagen fibre orientation estimation algorithm; and a genetic algorithm for data-driven optimisation of scan orientation sets. Together, these components enabled MADI to be applied to a cohort of 30 participants, yielding physiologically consistent fibre orientation maps of the ACL and meniscus, a conservative orientation resolution limit of approximately 15°, and preliminary agreement with arthroscopic findings in pathological cases. Whilst the present work is characterised as a proof-of-concept,

the results demonstrate that the core physical and algorithmic principles of MADI are experimentally sound and clinically promising.

The future work outlined in Section 9.5, spanning hardware development, pulse sequence optimisation, advanced reconstruction and image processing, protocol refinements, broader system utility, and prospective clinical validation, collectively define a promising roadmap towards clinical translation. The breadth of these opportunities reflects not only the potential of MADI but also the fundamental versatility of rotatable MRI as a platform. MADI offers, in effect, a form of *virtual arthroscopy*: non-invasive, quantitative, and repeatable assessment of collagen microstructure that has the potential to transform the diagnosis and longitudinal monitoring of musculoskeletal pathology.

References

[1] P. C. Lauterbur. “Image Formation by Induced Local Interactions: Examples Employing Nuclear Magnetic Resonance”. In: *Nature* 242.5394 (1973), pp. 190–191. ISSN: 1476-4687. DOI: 10.1038/242190a0. URL: https://www.nature.com/articles/242190a0.

[2] P. Mansfield and I. L. Pykett. “Biological and medical imaging by NMR”. In: *Journal of Magnetic Resonance* 29.2 (1978), pp. 355–373. ISSN: 0022-2364. DOI: 10.1016/j.jmr.2011.09.020.

[3] E. M. Haacke. *Magnetic resonance imaging : physical principles and sequence design*. New York: Wiley, 1999. ISBN: 0471351288.

[4] N. Phelan, P. Rowland, R. Galvin, and J. M. O’Byrne. “A systematic review and meta-analysis of the diagnostic accuracy of MRI for suspected ACL and meniscal tears of the knee”. In: *Knee Surgery, Sports Traumatology, Arthroscopy* 24.5 (May 2016), pp. 1525–1539. ISSN: 1433-7347. DOI: 10.1007/s00167-015-3861-8. URL: https://doi.org/10.1007/s00167-015-3861-8.

[5] X. Li, Q. Hou, X. Zhan, L. Chang, X. Ma, and H. Yuan. “The accuracy of MRI in diagnosing and classifying acute traumatic multiple ligament knee injuries”. In: *BMC Musculoskeletal Disorders* 23.1 (Jan. 2022), p. 43. ISSN: 1471-2474. DOI: 10.1186/s12891-021-04976-1. URL: https://doi.org/10.1186/s12891-021-04976-1.

[6] C. G. Peterfy, D. L. Janzen, P. F. Tirman, C. F. Van Dijke, M. Pollack, and H. K. Genant. “”Magic-angle” phenomenon: a cause of increased signal in the normal lateral meniscus on short-TE MR images of the knee”. In: *AJR. American journal of roentgenology* 163.1 (1994), pp. 149–154. ISSN: 0361-803X. DOI: 10.2214/AJR.163.1.8010202. URL: https://pubmed.ncbi.nlm.nih.gov/8010202/.

[7] S. J. Erickson, R. W. Prost, and M. E. Timins. “The ”magic angle” effect: background physics and clinical relevance.” In: *Radiology* 188.1 (1993). PMID: 7685531, pp. 23–25. DOI: 10.1148/radiology.188.1.7685531. eprint: https://doi.org/

10.1148/radiology.188.1.7685531. URL: https://doi.org/10.1148/radiology.188.1.7685531.

[8] M. Bydder, A. Rahal, G. D. Fullerton, and G. M. Bydder. “The magic angle effect: A source of artifact, determinant of image contrast, and technique for imaging”. In: *Journal of magnetic resonance imaging* 25.2 (2007), pp. 290–300. ISSN: 1053-1807.

[9] K. I. Momot, J. M. Pope, and R. M. Wellard. “Anisotropy of spin relaxation of water protons in cartilage and tendon”. In: *NMR in Biomedicine* 23.3 (2010), pp. 313–324. DOI: https://doi.org/10.1002/nbm.1466. eprint: https://analyticalsciencejournals.onlinelibrary.wiley.com/doi/pdf/10.1002/nbm.1466. URL: https://analyticalsciencejournals.onlinelibrary.wiley.com/doi/abs/10.1002/nbm.1466.

[10] T. Seidel, N. Hammer, N. Garnov, G. Schneider, and H. Steinke. “An algorithm for the calculation of three-dimensional collagen fiber orientation in ligaments using angle-sensitive MRI”. In: *Magnetic Resonance in Medicine* 69.6 (June 2013), pp. 1594–1602. ISSN: 1522-2594. DOI: 10.1002/MRM.24408. URL: https://onlinelibrary.wiley.com/doi/full/10.1002/mrm.24408.

[11] N. M. Szeverenyi and G. M. Bydder. “Dipolar anisotropy fiber imaging in a goat knee meniscus”. In: *Magnetic resonance in medicine* 65.2 (2011), pp. 463–470. ISSN: 1522-2594. DOI: 10.1002/MRM.22645. URL: https://pubmed.ncbi.nlm.nih.gov/20939091/.

[12] R. M. Henkelman, G. J. Stanisz, J. K. Kim, and M. J. Bronskill. “Anisotropy of NMR properties of tissues”. In: *Magnetic resonance in medicine* 32.5 (1994), pp. 592–601. ISSN: 0740-3194.

[13] M. Rea et al. “Devices to facilitate magic angle studies in peripheral MRI”. In: *Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering* 37B.4 (2010), pp. 220–225. DOI: 10.1002/cmr.b.20171. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/cmr.b.20171. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/cmr.b.20171.

[14] K. E. Chappell et al. “Detection of maturity and ligament injury using magic angle directional imaging”. In: *Magnetic resonance in medicine* 82.3 (Sept. 2019), pp. 1041–1054. ISSN: 1522-2594. DOI: 10.1002/MRM.27794. URL: https://pubmed.ncbi.nlm.nih.gov/31081201/.

[15] D. Brujic, K. E. Chappell, and M. Ristic. “Accuracy of collagen fibre estimation under noise using directional MR imaging”. In: *Computerized Medical Imaging and Graphics* 86 (Dec. 2020). ISSN: 1879-0771. DOI: 10.1016/J.COMPMEDIMAG.2020.101796. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7721590/.

[16] J. V. McGinley, M. Ristic, and I. R. Young. “A permanent MRI magnet for magic angle imaging having its field parallel to the poles”. In: *Journal of Magnetic Resonance* 271 (Oct. 2016), pp. 60–67. ISSN: 1090-7807. DOI: 10.1016/J.JMR.2016.08.001.

[17] D. W. McRobbie, E. A. Moore, M. J. Graves, and M. R. Prince. *MRI from Picture to Proton*. 2nd ed. Cambridge University Press, 2006.

[18] H. Mitchell, T. Hamilton, F. Steggerda, and H. Bean. “THE CHEMICAL COMPOSITION OF THE ADULT HUMAN BODY AND ITS BEARING ON THE BIOCHEMISTRY OF GROWTH”. In: *Journal of Biological Chemistry* 158.3 (1945), pp. 625–637. ISSN: 0021-9258. DOI: https://doi.org/10.1016/S0021-9258(19)51339-4. URL: https://www.sciencedirect.com/science/article/pii/S0021925819513394.

[19] M. H. Levitt. *Spin dynamics : basics of nuclear magnetic resonance*. 2nd ed. Hoboken, N.J: Wiley, 2008. ISBN: 0470511184.

[20] J. Larmor. “LXIII. On the theory of the magnetic influence on spectra; and on the radiation from moving ions”. In: *The London, Edinburgh and Dublin philosophical magazine and journal of science* 44.271 (1897), pp. 503–512. ISSN: 1941-5982.

[21] J. P. Marques, F. F. Simonis, and A. G. Webb. “Low-field MRI: An MR physics perspective”. In: *Journal of Magnetic Resonance Imaging* 49.6 (2019), pp. 1528–1542. DOI: https://doi.org/10.1002/jmri.26637. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/jmri.26637. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/jmri.26637.

[22] T. C. Arnold, C. W. Freeman, B. Litt, and J. M. Stein. “Low-field MRI: Clinical promise and challenges”. In: *Journal of Magnetic Resonance Imaging* 57.1 (2023), pp. 25–44. DOI: https://doi.org/10.1002/jmri.28408. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/jmri.28408. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/jmri.28408.

[23] M. Sarracanie and N. Salameh. “Low-Field MRI: How Low Can We Go? A Fresh View on an Old Debate”. In: *Frontiers in Physics* 8 (2020). ISSN: 2296-424X. DOI:

10.3389/fphy.2020.00172. URL: https://www.frontiersin.org/journals/physics/articles/10.3389/fphy.2020.00172.

[24] H. J. C. Berendsen. “Nuclear Magnetic Resonance Study of Collagen Hydration”. In: *The Journal of Chemical Physics* 36.12 (June 1962), pp. 3297–3305. ISSN: 0021-9606. DOI: 10.1063/1.1732460. eprint: https://pubs.aip.org/aip/jcp/article-pdf/36/12/3297/18825209/3297_1_online.pdf. URL: https://doi.org/10.1063/1.1732460.

[25] G. D. Fullerton, I. L. Cameron, and V. A. Ord. “Orientation of tendons in the magnetic field and its effect on T2 relaxation times.” In: *Radiology* 155.2 (1985). PMID: 3983395, pp. 433–435. DOI: 10.1148/radiology.155.2.3983395. eprint: https://doi.org/10.1148/radiology.155.2.3983395. URL: https://doi.org/10.1148/radiology.155.2.3983395.

[26] S. J. Erickson, I. H. Cox, J. S. Hyde, G. F. Carrera, J. A. Strandt, and L. D. Estkowski. “Effect of tendon orientation on MR imaging signal intensity: a manifestation of the ”magic angle” phenomenon.” In: *Radiology* 181.2 (1991). PMID: 1924777, pp. 389–392. DOI: 10.1148/radiology.181.2.1924777. eprint: https://doi.org/10.1148/radiology.181.2.1924777. URL: https://doi.org/10.1148/radiology.181.2.1924777.

[27] C. W. Hayes and J. A. Parellada. “The Magic Angle Effect in Musculoskeletal MR Imaging”. In: *Topics in magnetic resonance imaging* 8.1 (1996), pp. 51–56. ISSN: 0899-3459.

[28] G. D. Fullerton and A. Rahal. “Collagen structure: The molecular source of the tendon magic angle effect”. In: *Journal of Magnetic Resonance Imaging* 25.2 (2007), pp. 345–361. DOI: https://doi.org/10.1002/jmri.20808. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/jmri.20808. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/jmri.20808.

[29] K. I. Momot. “Hydrated Collagen: Where Physical Chemistry, Medical Imaging, and Bioengineering Meet”. In: *The Journal of Physical Chemistry B* 126.49 (Dec. 2022), pp. 10305–10316. ISSN: 1520-6106. DOI: 10.1021/acs.jpcb.2c06217. URL: https://doi.org/10.1021/acs.jpcb.2c06217.

[30] Physics Insights. *The Field of a Small Magnetic Dipole*. Accessed 26 July 2025. 2007. URL: http://www.physicsinsights.org/dipole_field_1.html.

[31] P. Fratzl. *Collagen : structure and mechanics*. New York: Springer, 2010. ISBN: 1441944818.

[32] H. R. SCREEN. “Hierarchical Approaches to Understanding Tendon Mechanics”. In: *Journal of Biomechanical Science and Engineering* 4.4 (2009), pp. 481–499. DOI: 10.1299/jbse.4.481.

[33] S. Tadimalla and K. I. Momot. “Effect of Partial H2O-D2O Replacement on the Anisotropy of Transverse Proton Spin Relaxation in Bovine Articular Cartilage”. In: *PLOS ONE* 9.12 (Dec. 2015), pp. 1–19. DOI: 10.1371/journal.pone.0115288. URL: https://doi.org/10.1371/journal.pone.0115288.

[34] K. E. Chappell. *Optimisation and visualisation of collagen fibre orientation using Magic Angle Directional Imaging (MADI)*. Feb. 2019. DOI: https://doi.org/10.25560/74059. URL: http://hdl.handle.net/10044/1/72873.

[35] D. Amiel, K. K. Ishizue, F. L. Harwood, L. Kitabayashi, and W. H. Akeson. “Injury of the anterior cruciate ligament: The role of collagenase in ligament degeneration”. In: *Journal of Orthopaedic Research* 7.4 (1989), pp. 486–493. DOI: https://doi.org/10.1002/jor.1100070405. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/jor.1100070405. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/jor.1100070405.

[36] R. M. Aspden. “A Model for the Function and Failure of the Meniscus”. In: *Engineering in Medicine* 14.3 (1985), pp. 119–122. DOI: 10.1243/EMED_JOUR_1985_014_028_02. eprint: https://doi.org/10.1243/EMED_JOUR_1985_014_028_02. URL: https://doi.org/10.1243/EMED_JOUR_1985_014_028_02.

[37] A. J. S. Fox, A. Bedi, and S. A. Rodeo. “The basic science of human knee menisci: structure, composition, and function”. In: *Sports Health* 4.4 (July 2012), pp. 340–351.

[38] P. Van Dyck et al. “Three tesla magnetic resonance imaging of the anterior cruciate ligament of the knee: can we differentiate complete from partial tears?” In: *Skeletal Radiology* 40.6 (June 2011), pp. 701–707. ISSN: 1432-2161. DOI: 10.1007/s00256-010-1044-8. URL: https://doi.org/10.1007/s00256-010-1044-8.

[39] E. F. Temponi, L. H. de Carvalho Júnior, B. Sonnery-Cottet, and P. Chambat. “Partial tearing of the anterior cruciate ligament: diagnosis and treatment”. In: *Revista Brasileira de Ortopedia (English Edition)* 50.1 (2015), pp. 9–15. ISSN: 2255-4971. DOI: https://doi.org/10.1016/j.rboe.2015.02.003. URL: https://www.sciencedirect.com/science/article/pii/S2255497115000129.

[40] B. M. Devitt et al. “MRI is not reliable in diagnosing of concomitant anterolateral ligament and anterior cruciate ligament injuries of the knee”. In: *Knee Surgery, Sports Traumatology, Arthroscopy* 25.4 (Apr. 2017), pp. 1345–1351. ISSN: 1433-7347. DOI: 10.1007/s00167-017-4538-2. URL: https://doi.org/10.1007/s00167-017-4538-2.

[41] W. C. Peh and J. H. Chan. “The magic angle phenomenon in tendons: effect of varying the MR echo time”. In: *British journal of radiology* 71.841 (1998), pp. 31–36. ISSN: 0007-1285.

[42] T. Wright, C. Yoon, and B. P. Schmit. “Shoulder MRI refinements: Differentiation of rotator cuff tear from artifacts and tendonosis, and reassessment of normal findings”. In: *Seminars in Ultrasound, CT and MRI* 22.4 (2001). Musculoskeletal MRI and CT, pp. 383–395. ISSN: 0887-2171. DOI: https://doi.org/10.1016/S0887-2171(01)90028-9. URL: https://www.sciencedirect.com/science/article/pii/S0887217101900289.

[43] A. H. Karantanas, A. H. Zibis, and N. Papanikolaou. “Increased signal intensity on fat-suppressed three-dimensional T1-weighted pulse sequences in patellar tendon: magic angle effect?” In: *Skeletal Radiology* 30.2 (Mar. 2001), pp. 67–71. ISSN: 1432-2161. DOI: 10.1007/s002560000307. URL: https://doi.org/10.1007/s002560000307.

[44] R. Crawford, G. Walley, S. Bridgman, and N. Maffulli. “Magnetic resonance imaging versus arthroscopy in the diagnosis of knee pathology, concentrating on meniscal lesions and ACL tears: a systematic review”. In: *British Medical Bulletin* 84.1 (Sept. 2007), pp. 5–23. ISSN: 0007-1420. DOI: 10.1093/bmb/ldm022. eprint: https://academic.oup.com/bmb/article-pdf/84/1/5/1078437/ldm022.pdf. URL: https://doi.org/10.1093/bmb/ldm022.

[45] B. Koo, S. H. Lee, S. J. Yun, and J. G. Song. “Diagnostic Performance of Magnetic Resonance Imaging for Detecting Meniscal Ramp Lesions in Patients With Anterior Cruciate Ligament Tears: A Systematic Review and Meta-analysis”. In: *The American Journal of Sports Medicine* 48.8 (2020). PMID: 31684739, pp. 2051–2059. DOI: 10.1177/0363546519880528. eprint: https://doi.org/10.1177/0363546519880528. URL: https://doi.org/10.1177/0363546519880528.

[46] W. Krampla, M. Roesel, K. Svoboda, A. Nachbagauer, M. Gschwantler, and W. Hruby. “MRI of the knee: how do field strength and radiologist’s experience influence diagnostic accuracy and interobserver correlation in assessing chondral

and meniscal lesions and the integrity of the anterior cruciate ligament?" In: *European Radiology* 19.6 (June 2009), pp. 1519–1528. ISSN: 1432-1084. DOI: 10.1007/s00330-009-1298-5. URL: https://doi.org/10.1007/s00330-009-1298-5.

[47] E. O. Stejskal and J. E. Tanner. "Spin Diffusion Measurements: Spin Echoes in the Presence of a Time-Dependent Field Gradient". In: *The Journal of Chemical Physics* 42.1 (Jan. 1965), pp. 288–292. ISSN: 0021-9606. DOI: 10.1063/1.1695690. eprint: https://pubs.aip.org/aip/jcp/article-pdf/42/1/288/18837047/288_1_online.pdf. URL: https://doi.org/10.1063/1.1695690.

[48] P. J. Basser, J. Mattiello, and D. LeBihan. "MR diffusion tensor spectroscopy and imaging". In: *Biophysical Journal* 66.1 (Jan. 1994), pp. 259–267. ISSN: 0006-3495. DOI: 10.1016/S0006-3495(94)80775-1. URL: https://doi.org/10.1016/S0006-3495(94)80775-1.

[49] S. Mori and P. B. Barker. "Diffusion magnetic resonance imaging: Its principle and applications". In: *The Anatomical Record* 257.3 (1999), pp. 102–109. DOI: https://doi.org/10.1002/(SICI)1097-0185(19990615)257:3<102::AID-AR7>3.0.CO;2-6. eprint: https://anatomypubs.onlinelibrary.wiley.com/doi/pdf/10.1002/%28SICI%291097-0185%2819990615%29257%3A3%3C102%3A%3AAID-AR7%3E3.0.CO%3B2-6. URL: https://anatomypubs.onlinelibrary.wiley.com/doi/abs/10.1002/%28SICI%291097-0185%2819990615%29257%3A3%3C102%3A%3AAID-AR7%3E3.0.CO%3B2-6.

[50] L. J. O'Donnell and C.-F. Westin. "An introduction to diffusion tensor image analysis". In: *Neurosurg. Clin. N. Am.* 22.2 (Apr. 2011), pp. 185–96, viii.

[51] B. P. Witwer et al. "Diffusion-tensor imaging of white matter tracts in patients with cerebral neoplasm". In: *Journal of Neurosurgery* 97.3 (2002), pp. 568–575. DOI: 10.3171/jns.2002.97.3.0568. URL: https://thejns.org/view/journals/j-neurosurg/97/3/article-p568.xml.

[52] K. Kamagata et al. "Advancements in Diffusion MRI Tractography for Neurosurgery". In: *Investigative Radiology* 59.1 (2024). URL: https://journals.lww.com/investigativeradiology/fulltext/2024/01000/advancements_in_diffusion_mri_tractography_for.2.aspx.

[53] S. S. Bhat et al. "Low-Field MRI of Stroke: Challenges and Opportunities". In: *Journal of Magnetic Resonance Imaging* 54.2 (2021), pp. 372–390. DOI: https://doi.org/10.1002/jmri.27324. eprint: https://onlinelibrary.wiley.com/

doi/pdf/10.1002/jmri.27324. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/jmri.27324.

[54] J. G. Raya et al. “Applications of Diffusion-Weighted MRI to the Musculoskeletal System”. In: *Journal of Magnetic Resonance Imaging* 59.2 (2024), pp. 376–396. DOI: https://doi.org/10.1002/jmri.28870. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/jmri.28870. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/jmri.28870.

[55] S. Saran, A. Chhabra, and R. Botchu. “Diffusion weighted imaging in musculoskeletal system: where are we now?” In: *BJR|Open* 7.1 (Jan. 2025), tzaf019. ISSN: 2513-9878. DOI: 10.1093/bjro/tzaf019. eprint: https://academic.oup.com/bjro/article-pdf/7/1/tzaf019/63819801/tzaf019.pdf. URL: https://doi.org/10.1093/bjro/tzaf019.

[56] V. Chianca et al. “Diffusion tensor imaging in the musculoskeletal and peripheral-nerve systems: from experimental to clinical applications”. In: *Eur. Radiol. Exp.* 1.1 (Sept. 2017), p. 12.

[57] P. Van Dyck et al. “Diffusion tensor imaging of the anterior cruciate ligament graft following reconstruction: a longitudinal study”. In: *European Radiology* 30.12 (Dec. 2020), pp. 6673–6684. ISSN: 1432-1084. DOI: 10.1007/s00330-020-07051-w. URL: https://doi.org/10.1007/s00330-020-07051-w.

[58] P. Van Dyck et al. “Diffusion tensor imaging of the anterior cruciate ligament graft”. In: *Journal of Magnetic Resonance Imaging* 46.5 (2017), pp. 1423–1432. DOI: https://doi.org/10.1002/jmri.25666. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/jmri.25666. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/jmri.25666.

[59] K. Wengler et al. “Diffusion tensor imaging of human Achilles tendon by stimulated echo readout-segmented EPI (ste-RS-EPI)”. In: *Magnetic Resonance in Medicine* 80.6 (2018), pp. 2464–2474. DOI: https://doi.org/10.1002/mrm.27220. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.27220. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.27220.

[60] K. Wengler et al. “In vivo evaluation of human patellar tendon microstructure and microcirculation with diffusion MRI”. In: *Journal of Magnetic Resonance Imaging* 51.3 (2020), pp. 780–790. DOI: https://doi.org/10.1002/jmri.26898. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/jmri.26898. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/jmri.26898.

[61] X. Yang et al. “Diffusion tensor imaging for anatomical and quantitative evaluation of the anterior cruciate ligament and ACL grafts: a preliminary study”. In: *J. Comput. Assist. Tomogr.* 38.4 (July 2014), pp. 489–494.

[62] X. Yang, D. Chen, M. Li, D. Shi, B. Zhu, and Q. Jiang. “Diffusion Tensor Imaging of the Anterior Cruciate Ligament Graft After Reconstruction: Repeatability and Diffusion Tensor Imaging Metrics”. In: *Journal of Computer Assisted Tomography* 39.2 (2015). ISSN: 0363-8715. URL: https://journals.lww.com/jcat/fulltext/2015/03000/diffusion_tensor_imaging_of_the_anterior_cruciate.19.aspx.

[63] P. Van Dyck, M. Froeling, C. H. W. Heusdens, J. Sijbers, A. Ribbens, and T. Billiet. “Diffusion tensor imaging of the anterior cruciate ligament following primary repair with internal bracing: A longitudinal study”. In: *Journal of Orthopaedic Research* 39.6 (2021), pp. 1318–1330. DOI: https://doi.org/10.1002/jor.24684. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/jor.24684. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/jor.24684.

[64] S. Liu et al. “Diffusion Tensor Imaging for Quantitative Assessment of Anterior Cruciate Ligament Injury Grades and Graft”. In: *Journal of Magnetic Resonance Imaging* 52.5 (2020), pp. 1475–1484. DOI: https://doi.org/10.1002/jmri.27322. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/jmri.27322. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/jmri.27322.

[65] N. Wang et al. “Magic angle effect on diffusion tensor imaging in ligament and brain”. In: *Magnetic Resonance Imaging* 92 (2022), pp. 243–250. ISSN: 0730-725X. DOI: https://doi.org/10.1016/j.mri.2022.06.008. URL: https://www.sciencedirect.com/science/article/pii/S0730725X22001023.

[66] N. Wang, A. J. Mirando, G. Cofer, Y. Qi, M. J. Hilton, and G. A. Johnson. “Characterization complex collagen fiber architecture in knee joint using high-resolution diffusion imaging”. In: *Magnetic Resonance in Medicine* 84.2 (2020), pp. 908–919. DOI: https://doi.org/10.1002/mrm.28181. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.28181. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.28181.

[67] J. Shen, Q. Zhao, Y. Qi, G. Cofer, G. A. Johnson, and N. Wang. “Tractography of Porcine Meniscus Microstructure Using High-Resolution Diffusion Magnetic Resonance Imaging”. In: *Frontiers in Endocrinology* 13 (2022). ISSN: 1664-2392.

DOI: 10.3389/fendo.2022.876784. URL: https://www.frontiersin.org/journals/endocrinology/articles/10.3389/fendo.2022.876784.

[68] Q. Zhao, A. Holt, C. E. Spritzer, L. E. DeFrate, A. L. McNulty, and N. Wang. "High angular resolution diffusion imaging (HARDI) of porcine menisci: a comparison of diffusion tensor imaging and generalized q-sampling imaging". In: *Quantitative Imaging in Medicine and Surgery* 14.4 (2024). ISSN: 2223-4292. URL: https://journals.lww.com/qims/fulltext/2024/04000/high_angular_resolution_diffusion_imaging__hardi_.1.aspx.

[69] British Medical Association. *NHS diagnostics data analysis*. Accessed 25 June 2025. 2025. URL: https://www.bma.org.uk/advice-and-support/nhs-delivery-and-workforce/pressures/nhs-diagnostics-data-analysis.

[70] Independent report commissioned NHS England. *DIAGNOSTICS: RECOVERY AND RENEWAL; Report of the Independent Review of Diagnostic Services for NHS England*. Available at: https://www.england.nhs.uk/wp-content/uploads/2020/11/diagnostics-recovery-and-renewal-independent-review-of-diagnostic-services-for-nhs-england-2.pdf. Accessed 26 June 2025. Oct. 2020.

[71] Full Fact. *Is the government on track to double the number of CT and MRI scanners in the NHS?* Accessed 26 June 2025. Mar. 2025. URL: https://fullfact.org/government-tracker/double-mri-ct-scanners/.

[72] The Royal College of Radiologists. *Clinical Radiology Workforce Census 2024*. Available at: https://www.rcr.ac.uk/media/4imb5jge/_rcr-2024-clinical-radiology-workforce-census-report.pdf. Accessed 26 June 2025. 2024.

[73] National Health Service. *Diagnostic Imaging Dataset Annual Statistical Release 2022/23*. Accessed 25 June 2025. 2023. URL: https://www.england.nhs.uk/statistics/wp-content/uploads/sites/2/2023/11/Annual-Statistical-Release-2022-23-PDF-1.3MB-1.pdf.

[74] S. G. F. Abram, A. J. Price, A. Judge, and D. J. Beard. "Anterior cruciate ligament (ACL) reconstruction and meniscal repair rates have both increased in the past 20 years in England: hospital statistics from 1997 to 2017". In: *British Journal of Sports Medicine* 54.5 (2020), pp. 286–291. ISSN: 0306-3674. DOI: 10.1136/bjsports-2018-100195. eprint: https://bjsm.bmj.com/content/54/5/286.full.pdf. URL: https://bjsm.bmj.com/content/54/5/286.

[75] NHS Tayside. *Freedom of Information Request: MRI Examinations, Wait Times, and Expenditure 2023*. Available at: https://www.nhstaysidecdn.scot.nhs.uk/NHSTaysideWeb/idcplg?IdcService=GET_SECURE_FILE&dDocName=PROD_374588&Rendition=web&RevisionSelectionMethod=LatestReleased&noSaveAs=1. FOI Reference: IGTFOISA24052. Accessed 25 June 2025. 2024.

[76] NHS University Hospitals of North Midlands. *Freedom of Information Request: MRI Examinations, Wait Times, and Expenditure 2023*. Available at: https://www.uhnm.nhs.uk/media/1pngkx2d/20240221-mri-freedom-of-information-ref-696-2324-2-of-2.pdf. FOI Reference: 696-2324. Accessed 25 June 2025. 2024.

[77] NHS The Rotherham. *Freedom of Information Request: MRI Examinations, Wait Times, and Expenditure 2021*. Available at: https://www.therotherhamft.nhs.uk/sites/default/files/2023-03/6176%20FOI%20-%20MRI%20Services.docx. FOI Reference: 6176. Accessed 25 June 2025. 2022.

[78] NHS Royal Devon University Healthcare. *Freedom of Information Request: MRI Examinations, Wait Times, and Expenditure 2023*. Available at: https://www.royaldevon.nhs.uk/media/44nhjjti/foi-rdf2246-24-mri-imaging-scans-2023-mobile-units-total-machines.pdf. FOI Reference: RDF2246-24. Accessed 25 June 2025. 2024.

[79] NHS Royal Cornwall Hospitals. *Freedom of Information Request: MRI Examinations, Wait Times, and Expenditure 2023*. Available at: https://doclibrary-rcht.cornwall.nhs.uk/DocumentsLibrary/RoyalCornwallHospitalsTrust/Websites/Internet/OurOrganisation/FreedomOfInformation/DisclosureLog/2024/202402/FOIRef23352MRIINFO.pdf. FOI Reference: 23352. Accessed 25 June 2025. 2024.

[80] NHS Doncaster and Bassetlaw Teaching Hospitals. *Freedom of Information Request: MRI Examinations, Wait Times, and Expenditure 2023*. Available at: https://www.dbth.nhs.uk/wp-content/uploads/2024/03/FOI-49-2024-Final.pdf. FOI Reference: -. Accessed 25 June 2025. 2024.

[81] NHS Maidstone and Tunbridge Wells. *Freedom of Information Request: MRI Examinations, Wait Times, and Expenditure 2023*. Available at: https://www.mtw.nhs.uk/wp-content/uploads/2024/02/MRI-imaging.300124.pdf. FOI Reference: FOI/GS/ID 8807. Accessed 25 June 2025. 2024.

[82] NHS The Dudley Group. *Freedom of Information Request: MRI Examinations, Wait Times, and Expenditure 2023*. Available at: https://www.dgft.nhs.uk/wp-content/uploads/2024/03/FOI-012024-000579.pdf. FOI Reference: FOI-012024-000579. Accessed 25 June 2025. 2024.

[83] J. Mallard, J. Hutchison, W. Edelstein, R. Ling, and M. Foster. "Imaging by nuclear magnetic resonance and its bio-medical implications". In: *Journal of Biomedical Engineering* 1.3 (1979), pp. 153–160. ISSN: 0141-5425. DOI: https://doi.org/10.1016/0141-5425(79)90036-0. URL: https://www.sciencedirect.com/science/article/pii/0141542579900360.

[84] W. A. Edelstein, G. H. Glover, C. J. Hardy, and R. W. Redington. "The intrinsic signal-to-noise ratio in NMR imaging". In: *Magnetic Resonance in Medicine* 3.4 (1986), pp. 604–618. DOI: https://doi.org/10.1002/mrm.1910030413. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.1910030413. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.1910030413.

[85] M. Ristic, Y. Gryska, J. V. M. McGinley, and V. Yufit. "Supercapacitor Energy Storage for Magnetic Resonance Imaging Systems". In: *IEEE Transactions on Industrial Electronics* 61.8 (2014), pp. 4255–4264. DOI: 10.1109/TIE.2013.2282596.

[86] J. Mispelter, M. Lupu, and A. Briguet. *NMR probeheads for biophysical and biomedical experiments : theoretical principles & practical guidelines*. Second edition. London: Imperial College Press, 2015. ISBN: 1848166621.

[87] Harmonic Drive AG. *Application Manual AC Servo Controller YukonDrive*. Accessed 25 June 2025. 2017. URL: https://harmonicdrive.de/fileadmin/Downloads/Produkte/Projektierungsanleitungen/YukonDrive_HB_AH_E_1003374_04_2017_V01_1.pdf.

[88] A. Haase, J. Frahm, D. Matthaei, W. Hanicke, and K.-D. Merboldt. "FLASH imaging. Rapid NMR imaging using low flip-angle pulses". In: *Journal of Magnetic Resonance (1969)* 67.2 (1986), pp. 258–266. ISSN: 0022-2364. DOI: https://doi.org/10.1016/0022-2364(86)90433-6. URL: https://www.sciencedirect.com/science/article/pii/0022236486904336.

[89] Y. Zur, M. L. Wood, and L. J. Neuringer. "Spoiling of transverse magnetization in steady-state sequences". In: *Magnetic Resonance in Medicine* 21.2 (1991), pp. 251–263. DOI: https://doi.org/10.1002/mrm.1910210210. eprint: https:

//onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.1910210210. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.1910210210.

[90] M. Ristic, K. E. Chappell, H. Lanz, J. McGinley, C. Gupte, and D. Amiras. "First in-vivo magic angle directional imaging using dedicated low-field MRI". In: *Magnetic Resonance in Medicine* 93.3 (2025), pp. 1077–1089. DOI: https://doi.org/10.1002/mrm.30332. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.30332. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.30332.

[91] K. Scheffler and S. Lehnhardt. "Principles and applications of balanced SSFP techniques". In: *European Radiology* 13.11 (Nov. 2003), pp. 2409–2418. ISSN: 1432-1084. DOI: 10.1007/s00330-003-1957-x. URL: https://doi.org/10.1007/s00330-003-1957-x.

[92] J. Hennig, A. Nauerth, and H. Friedburg. "RARE imaging: A fast imaging method for clinical MR". In: *Magnetic Resonance in Medicine* 3.6 (1986), pp. 823–833. DOI: https://doi.org/10.1002/mrm.1910030602. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.1910030602. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.1910030602.

[93] B. S. Malhi, H. Jang, M. S. Malhi, D. B. Berry, and S. Jerban. "Tendon evaluation with ultrashort echo time (UTE) MRI: a systematic review". In: *Frontiers in Musculoskeletal Disorders* 2 (2024). ISSN: 2813-883X.

[94] A. Foi. "Noise estimation and removal in MR imaging: The variance-stabilization approach". In: *2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. 2011, pp. 1809–1814. DOI: 10.1109/ISBI.2011.5872758.

[95] M. Maggioni, V. Katkovnik, K. Egiazarian, and A. Foi. "Nonlocal Transform-Domain Filter for Volumetric Data Denoising and Reconstruction". In: *IEEE Transactions on Image Processing* 22.1 (2013), pp. 119–133. DOI: 10.1109/TIP.2012.2210725.

[96] A. Foi. *Variance stabilization for Rician-distributed data and its application to noise estimation and removal in MR imaging*. Accessed 25 June 2025. 2016. URL: https://webpages.tuni.fi/foi/RiceOptVST/.

[97] J. Veraart, D. S. Novikov, D. Christiaens, B. Ades-aron, J. Sijbers, and E. Fieremans. "Denoising of diffusion MRI using random matrix theory". In: *NeuroImage* 142 (2016), pp. 394–406. ISSN: 1053-8119. DOI: https://doi.org/10.1016/

j.neuroimage.2016.08.016. URL: https://www.sciencedirect.com/science/article/pii/S1053811916303949.

[98] J. L. Olesen, A. Ianus, L. Østergaard, N. Shemesh, and S. N. Jespersen. “Tensor denoising of multidimensional MRI data”. In: *Magnetic Resonance in Medicine* 89.3 (2023), pp. 1160–1172. DOI: https://doi.org/10.1002/mrm.29478. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.29478. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.29478.

[99] S. Klein, M. Staring, K. Murphy, M. Viergever, and J. Pluim. “elastix: A Toolbox for Intensity-Based Medical Image Registration”. In: *IEEE transactions on medical imaging* 29.1 (2010), pp. 196–205. ISSN: 0278-0062.

[100] S. Klein and M. Staring. *elastix the manual*. Accessed 25 June 2025. 2024. URL: https://courses.compute.dtu.dk/02502/docs/elastix-5.2.0-manual.pdf.

[101] S. Klein, J. Pluim, M. Staring, and M. Viergever. “Adaptive stochastic gradient descent optimisation for image registration”. In: *International journal of computer vision* 81.3 (2009), pp. 227–239. ISSN: 0920-5691. DOI: 10.1007/s11263-008-0168-y.

[102] D. Rueckert, L. Sonoda, C. Hayes, D. Hill, M. Leach, and D. Hawkes. “Nonrigid registration using free-form deformations: application to breast MR images”. In: *IEEE transactions on medical imaging* 18.8 (1999), pp. 712–721. ISSN: 0278-0062. DOI: 10.1109/42.796284.

[103] M. Unser. “Splines: a perfect fit for signal and image processing”. In: *IEEE signal processing magazine* 16.6 (1999), pp. 22–38. ISSN: 1053-5888. DOI: 10.1109/79.799930.

[104] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens. “Multimodality image registration by maximization of mutual information”. In: *IEEE transactions on medical imaging* 16.2 (1997), pp. 187–198. ISSN: 0278-0062. DOI: 10.1109/42.563664.

[105] P. Viola and W. M. Wells. “Alignment by maximization of mutual information”. In: *International journal of computer vision* 24.2 (1997), pp. 137–154. ISSN: 0920-5691. DOI: 10.1023/A:1007958904918.

[106] D. Mattes, D. Haynor, H. Vesselle, T. Lewellen, and W. Eubank. “PET-CT image registration in the chest using free-form deformations”. In: *IEEE transactions on medical imaging* 22.1 (2003), pp. 120–128. ISSN: 0278-0062. DOI: 10.1109/TMI.2003.809072.

[107] P. Thévenaz and M. Unser. “Optimization of mutual information for multiresolution image registration”. In: *IEEE Transactions on Image Processing* 9.12 (2000), pp. 2083–2099. DOI: 10.1109/83.887976.

[108] Stanford Magnets. *Neodymium Iron Boron Temperature Ratings*. Accessed 25 June 2025. 2025. URL: https://www.stanfordmagnets.com/ndfeb-magnets-temperature-ratings.html.

[109] R. Kikinis, S. D. Pieper, and K. G. Vosburgh. “3D Slicer: A Platform for Subject-Specific Image Analysis, Visualization, and Clinical Support”. In: *Intraoperative Imaging and Image-Guided Therapy*. Ed. by F. A. Jolesz. New York, NY: Springer New York, 2014, pp. 277–289. ISBN: 978-1-4614-7657-3. DOI: 10.1007/978-1-4614-7657-3_19. URL: https://doi.org/10.1007/978-1-4614-7657-3_19.

[110] CadQuery contributors. *CadQuery*. DOI: 10.5281/zenodo.10513848. URL: https://github.com/CadQuery/cadquery.

[111] M. Perslev, A. Pai, J. Runhaar, C. Igel, and E. B. Dam. “Cross-Cohort Automatic Knee MRI Segmentation With Multi-Planar U-Nets”. In: *Journal of Magnetic Resonance Imaging* 55.6 (2022), pp. 1650–1663. DOI: https://doi.org/10.1002/jmri.27978. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/jmri.27978. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/jmri.27978.

[112] X. He, Y. Hu, Z. Zhou, M. Jarraya, and F. Liu. *Few-Shot Adaptation of Training-Free Foundation Model for 3D Medical Image Segmentation*. 2025. arXiv: 2501.09138 [cs.CV]. URL: https://arxiv.org/abs/2501.09138.

[113] T.-W. Wang et al. “Artificial Intelligence Detection and Segmentation Models: A Systematic Review and Meta-Analysis of Brain Tumors in Magnetic Resonance Imaging”. In: *Mayo Clinic Proceedings: Digital Health* 2.1 (2024), pp. 75–91. ISSN: 2949-7612. DOI: https://doi.org/10.1016/j.mcpdig.2024.01.002. URL: https://www.sciencedirect.com/science/article/pii/S2949761224000038.

[114] J. A. Nelder and R. Mead. “A Simplex Method for Function Minimization”. In: *The Computer Journal* 7.4 (Jan. 1965), pp. 308–313. ISSN: 0010-4620. DOI: 10.1093/comjnl/7.4.308. eprint: https://academic.oup.com/comjnl/article-pdf/7/4/308/1013182/7-4-308.pdf. URL: https://doi.org/10.1093/comjnl/7.4.308.

[115] H. Lanz, M. Ristic, K. E. Chappell, and J. V. McGinley. “Minimum number of scans for collagen fibre direction estimation using Magic Angle Directional Imaging (MADI) with a priori information”. In: *Array* 17 (Mar. 2023), p. 100273. ISSN: 2590-0056. DOI: 10.1016/J.ARRAY.2022.100273.

[116] I. Norton et al. “SlicerDMRI: Open Source Diffusion MRI Software for Brain Cancer Research”. In: *Cancer research (Chicago, Ill.)* 77.21 (2017), e101–e103. ISSN: 0008-5472.

[117] T. Zantop, M. Herbort, M. J. Raschke, F. H. Fu, and W. Petersen. “The Role of the Anteromedial and Posterolateral Bundles of the Anterior Cruciate Ligament in Anterior Tibial Translation and Internal Rotation”. In: *The American Journal of Sports Medicine* 35.2 (2007). PMID: 17158275, pp. 223–227. DOI: 10.1177/0363546506294571. eprint: https://doi.org/10.1177/0363546506294571. URL: https://doi.org/10.1177/0363546506294571.

[118] J. G. Raya. “Techniques and applications of in vivo diffusion imaging of articular cartilage”. In: *J. Magn. Reson. Imaging* 41.6 (June 2015), pp. 1487–1504.

[119] U. Ferizi et al. “Diffusion tensor imaging of articular cartilage at 3T correlates with histology and biomechanics in a mechanical injury model”. In: *Magnetic Resonance in Medicine* 78.1 (2017), pp. 69–78. DOI: https://doi.org/10.1002/mrm.26336. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.26336. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.26336.

[120] A. Guha, C. Wyatt, D. C. Karampinos, L. Nardo, T. M. Link, and S. Majumdar. “Spatial variations in magnetic resonance-based diffusion of articular cartilage in knee osteoarthritis”. In: *Magnetic Resonance Imaging* 33.9 (2015), pp. 1051–1058. ISSN: 0730-725X. DOI: https://doi.org/10.1016/j.mri.2015.06.004. URL: https://www.sciencedirect.com/science/article/pii/S0730725X15001496.

[121] B. K. Li, H. Wang, A. Trakic, C. Engstrom, E. Weber, and S. Crozier. “An orthogonal-based decoupling method for MRI phased array coil design”. In: *NMR in biomedicine* 25.6 (June 2012), pp. 835–842. ISSN: 1099-1492. DOI: 10.1002/NBM.1799. URL: https://pubmed.ncbi.nlm.nih.gov/22134869/.

Appendix A: Appendices Introduction

A.1 Detailed List of Contributions

This section provides a representative, but by no means exhaustive, list of contributions to the work in this thesis.

This work was partly supported by the Wellcome Trust Innovator Award WT215908/Z/19/Z and MedTechOne PSR085_MEME PSR079_MEME.

Harry Lanz: software development of motion control software and GUI (C++), registration pipeline (MATLAB with elastix subroutines), calibration phantom with programmatic CAD using CadQuery (Python), segmentation framework and GUI (MATLAB), updated collagen fibre orientation estimation algorithm and genetic algorithm (MATLAB), MADI data visualisation 3D Slicer plugin (Python), other small software components/scripts generally developed using MATLAB (e.g. intensity correction and BM4D denoising); hardware development and construction: motion control electronics mounting within electronics box within shielded room; 3D printing of: calibration phantom, immobilisation boot mount, transmit coil mount, phased array coil body.

Dr Mike Ristic: hardware development: electronics, RF toolchain, motion platform design and commission, gradient design, supercapacitor design, magnet design; project principal investigator: ethics approvals, funding acquisition, supervision, project administration.

Dr Karyn Chappell: MR pulse sequence optimisation and investigation: tuning of parameters for magic angle effect detection, tuning of parameters for clinically relevant image contrast and quality; scan planning and acquisition: scanner operation, scan protocol design, patient positioning and management, patient recruitment, ethics approvals.

Dr John McGinley: hardware development: electronics, RF toolchain, motion platform design and commission, gradient design, supercapacitor design, magnet design, transmit and receive coil design and construction.

Dr Dimitri Amiras and Mr Chinmay Gupte: medical expertise: image interpretation, dictation of desired system requirements for clinical translation, performance of arthroscopy and interpretation of arthroscopic findings, ethics approvals, patient recruitment.

Dr Ruslan Garipov: liaison at MR Solutions who provided the spectrometer for the low-field rotatable MRI system: provided continued support throughout scanner development.

Appendix B: MADI Appendices

B.1 Effect of Denoising on Collagen Fibre Orientation Estimation

Figures B.1 to B.4 show histograms of the deviation between corresponding fibres when collagen fibre orientation estimation is performed with and without application of block matching denoising. To produce these plots the full MADI pipeline was applied to each dataset of the healthy ACL (representative dataset, Section 8.3), the healthy meniscus (Section 8.4), the partially torn ACL (Section 8.5), and the medial meniscus tear (Section 8.6), with and without block matching denoising applied. The same segmentations were used in both runs to ensure that both the raw and denoised outputs contained fibres at corresponding locations. The angular deviation between these pairs of corresponding fibres was then computed (90° being the maximum as collagen fibres are bidirectional) and binned into histograms with bin widths of 2°.

Figures B.1 to B.4 show that denoising has a marginal effect on collagen fibre orientation estimation. Most fibre orientations are within 2° of their counterpart, with a vast majority within 10°. Each dataset has a small widely distributed portion of fibres with deviations between 70°–90°. These may relate to fibres where noise had a significant impact on the orientation estimation, or potentially where registration differences arose due to application on noisy or denoised data. It is interesting to note the bimodal distribution of the angular deviations present in each dataset, with a majority of deviations grouping at $< 10^\circ$ with a peak in the 0°–2° bin, and a wider, flatter minority grouped between 70°–90°. Furthermore, the lack of a clear continuum of deviations from 0°–90° merits future investigation. Nonetheless, the angular deviations between the raw and denoised datasets were not large or prevalent enough to have any meaningful effect on the MADI output of these datasets. Denoising was useful, however, in segmentation workflows which were semi-automatic and hence the denoised images facilitated operator image interpretation throughout these workflows.

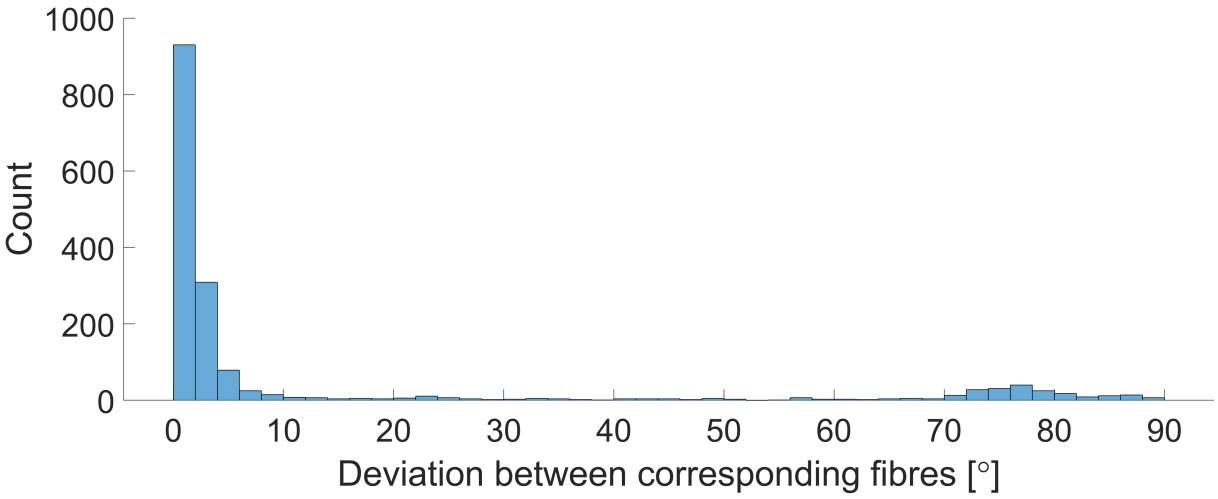


Figure B.1: Histogram showing the deviation between corresponding fibres for the healthy ACL (representative) dataset with and without block matching denoising applied.

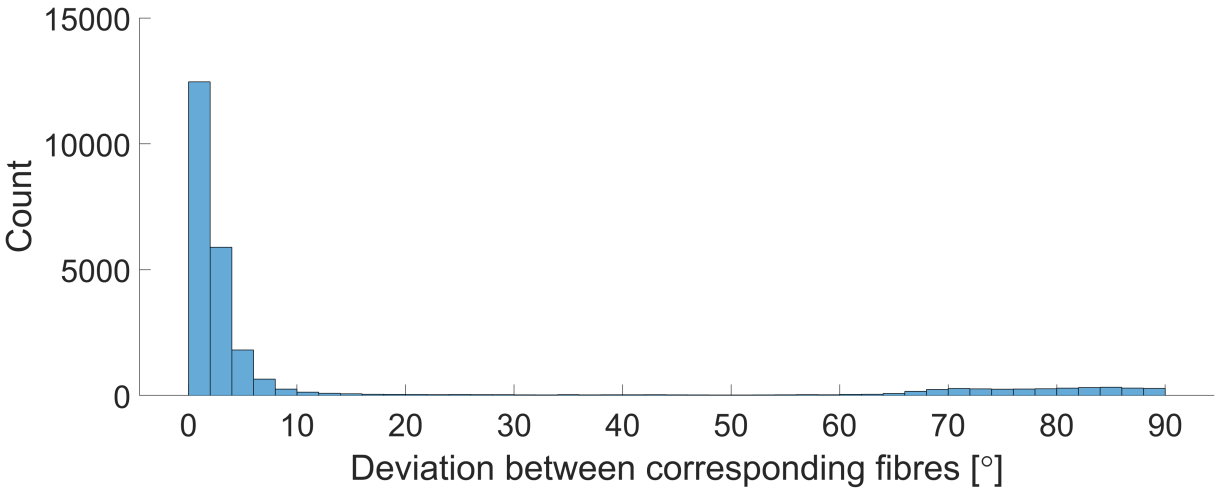


Figure B.2: Histogram showing the deviation between corresponding fibres for the healthy meniscus dataset with and without block matching denoising applied.

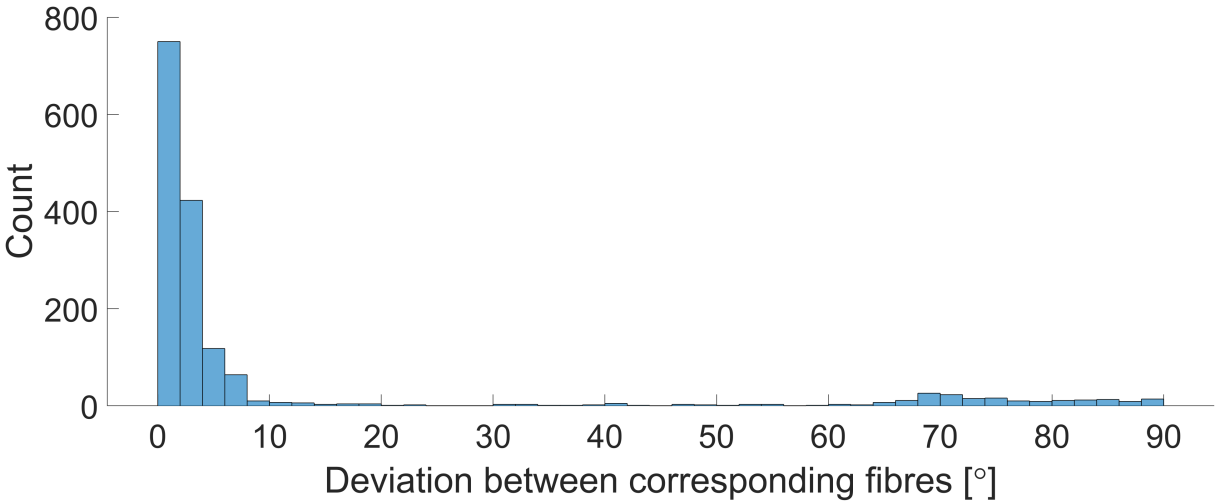


Figure B.3: Histogram showing the deviation between corresponding fibres for the pathological ACL dataset with and without block matching denoising applied.

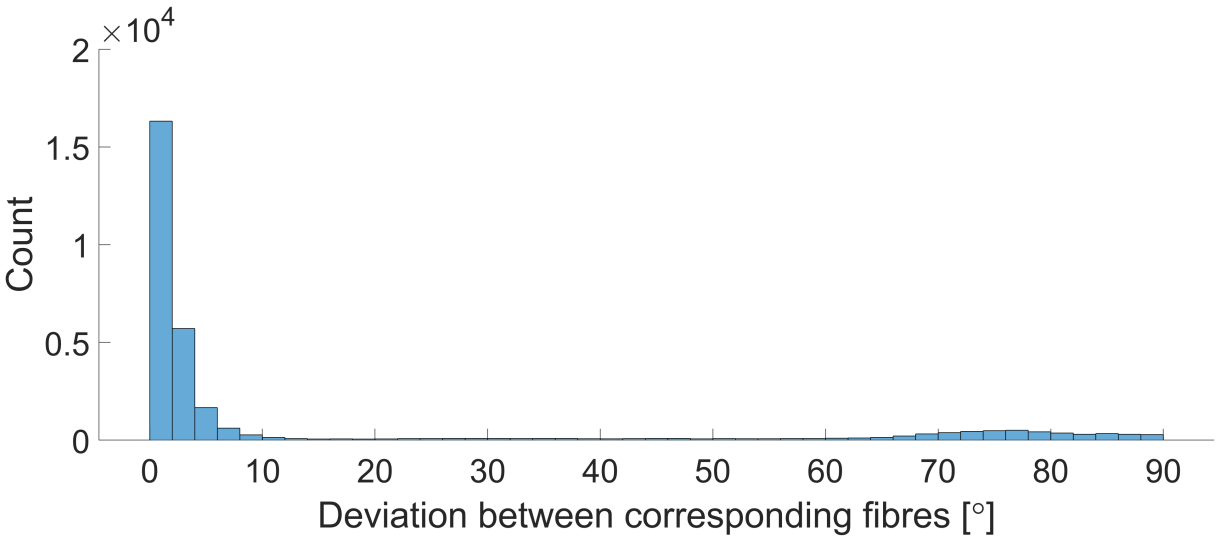


Figure B.4: Histogram showing the deviation between corresponding fibres for the pathological meniscus dataset with and without block matching denoising applied.

Appendix C: Registration Appendices

C.1 Registration: elastix Parameters

C.1.1 Rigid Registration Dummy Call Parameters

```
// INIT RIGID REGISTRATION parameter file, for dummy calls of elastix to
   generate telemetry files of correct format

// Image metadata: information about the image goes here
(FixedInternalImagePixelType "float")
(MovingInternalImagePixelType "float")
(FixedImageDimension 3) // 3D images
(MovingImageDimension 3)
(UseDirectionCosines "true")
(DefaultPixelValue 0) // when image is deformed "new" pixels should be
   initialised to 0
(WriteResultImage "true")
(ResultImagePixelType "double")
(ResultImageFormat "nii") // use nifti format to prevent conversions to and
   from dicom incurring rounding errors

// Registration: registration type and interpolation settings go here
(Registration "MultiResolutionRegistration") // Use multiple resolutions but
   only one metric
(Interpolator "BSplineInterpolator") // Interpolator used during registration
(BSplineInterpolationOrder 1)
(ResampleInterpolator "FinalBSplineInterpolator") // Interpolator used to
   generate final image
(FinalBSplineInterpolationOrder 3) // used to create final image, can create
   bright spots
(Resampler "DefaultResampler") // Use default (other options require
   compilation from source)
(ImageSampler "RandomCoordinate") // sample points for registration are
   selected on random coordinates (not on grid)
(FixedImageBSplineInterpolationOrder 1) // for interpolating random coordinates
(NumberOfSpatialSamples 100) // for dummy calls just use a small number for
   speed
(NewSamplesEveryIteration "true") // avoid sampling same points over and over
   to get better overall registration
(UseRandomSampleRegion "false") // probably best to avoid this as the images
   are already constrained by the artificial edge
(CheckNumberOfSamples "true") // check if samples are within valid region
(RequiredRatioOfValidSamples 0.25) // at least a quarter should be valid

// Image resolutions: settings for different resolutions goes here
```

```
(NumberOfResolutions 1) // number of resolutions to refine registration over,
   dummy call just use 1
(FixedImagePyramid "FixedSmoothingImagePyramid") // pyramids refine resolution
   to initially match larger features
(FixedImagePyramidSchedule 1 1 1)
(MovingImagePyramid "MovingSmoothingImagePyramid")
(MovingImagePyramidSchedule 1 1 1)

// Transform: settings for the transform go here
(Transform "EulerTransform") // euler transform for rigid registration only
   allows rotations and translations
(AutomaticScalesEstimation "true") // automatically estimate scales of
   rotations and translations
// centre of rotation can be left as default (geometric centre)
(AutomaticTransformInitialization "false") // no need for automatic
   initialisation as image centres should be fairly close in general
(AutomaticTransformInitializationMethod "GeometricalCenter")
(HowToCombineTransforms "Compose") // transforms feed into each other by
   composition T(x) = T_1(T_0(x)) rather than by addition

// Optimiser: settings for the optimiser go here
(Optimizer "AdaptiveStochasticGradientDescent")
(MaximumNumberOfIterations 1) // use low number for dummy call
(MaximumNumberOfSamplingAttempts 10) // how many times samples will be
   attempted to be drawn
(AutomaticParameterEstimation "false") // prevent automatic parameter
   estimation as it can be too aggressive
(UseAdaptiveStepSizes "true") // set to true for convenience, when using random
    sample coordinates this may be overwritten
(MaximumStepLength 1.0) // when automatic parameter estimation is false this
   has no effect
(SP_a 100) // set to make step sizes more conservative, smaller as registration
    becomes more refined
(SP_A 20) // set to make step sizes more conservative
(SP_alpha 1.0) // set to make step sizes more conservative
(SigmoidMax 1.0) // set to make step sizes more conservative
(SigmoidMin -0.8) // set to make step sizes more conservative
(SigmoidScale 0.00000001) // set to make step sizes more conservative
(SigmoidInitialTime 1.0) // larger numbers give smaller steps, 0 is default
(NumberOfGradientMeasurements 0) // default 0 gives good automatic estimation
// number of jacobian measurements is left to be initialised by elastix
(NumberOfSamplesForExactGradient 100000) // default
(ASGDParameterEstimationMethod "DisplacementDistribution") // used to estimate
   parameters faster than "Original" method
(MaximumDisplacementEstimationMethod "2sigma")
(NoiseCompensation "true")

// Metric: settings for the metric (or metrics) go here
(Metric "AdvancedMattesMutualInformation") // mutual information is robust when
    intensities may vary (e.g. multi modal or in our case magic angle effect)
(NumberOfHistogramBins 32) // default
(NumberOfFixedHistogramBins 32) // default
(NumberOfMovingHistogramBins 32) // default
```

```
(FixedKernelBSplineOrder 0) // default
(MovingKernelBSplineOrder 3) // default
(FixedLimitRangeRatio 0.01) // default
(MovingLimitRangeRatio 0.01) // default
(FiniteDifferenceDerivative "false") // experimental do not use
(UseFastAndLowMemoryVersion "true")

// Mask: settings for masks (fixed and moving) go here
(ErodeMask "true") // true prevents artificial edge from corrupting
   registration
(ErodeFixedMask "true")
(ErodeMovingMask "true")

// Misc
// There are several warnings that appear in the logs for undocumented
   parameters, best to let elastix initialise these to default values
```

C.1.2 Rigid Registration Parameters

```
// RIGID REGISTRATION parameter file

// Image metadata: information about the image goes here
(FixedInternalImagePixelType "float")
(MovingInternalImagePixelType "float")
(FixedImageDimension 3) // 3D images
(MovingImageDimension 3)
(UseDirectionCosines "true")
(DefaultPixelValue 0) // when image is deformed "new" pixels should be
   initialised to 0
(WriteResultImage "true")
(ResultImagePixelType "double")
(ResultImageFormat "nii") // use nifti format to prevent conversions to and
   from dicom incurring rounding errors

// Registration: registration type and interpolation settings go here
(Registration "MultiResolutionRegistration") // Use multiple resolutions but
   only one metric
(Interpolator "BSplineInterpolator") // Interpolator used during registration
(BSplineInterpolationOrder 1)
(ResampleInterpolator "FinalBSplineInterpolator") // Interpolator used to
   generate final image
(FinalBSplineInterpolationOrder 3) // used to create final image, can create
   bright spots
(Resampler "DefaultResampler") // Use default (other options require
   compilation from source)
(ImageSampler "RandomCoordinate") // sample points for registration are
   selected on random coordinates (not on grid)
(FixedImageBSplineInterpolationOrder 1) // for interpolating random coordinates
(NumberOfSpatialSamples 5000) // low number of samples possible, less than 2000
    is best avoided
(NewSamplesEveryIteration "true") // avoid sampling same points over and over
   to get better overall registration
```

```
(UseRandomSampleRegion "false") // probably best to avoid this as the images
   are already constrained by the artificial edge
(CheckNumberOfSamples "true") // check if samples are within valid region
(RequiredRatioOfValidSamples 0.25) // at least a quarter should be valid

// Image resolutions: settings for different resolutions goes here
(NumberOfResolutions 3) // number of resolutions to refine registration over
(FixedImagePyramid "FixedSmoothingImagePyramid") // pyramids refine resolution
   to initially match larger features
(FixedImagePyramidSchedule 4 4 4 2 2 2 1 1 1)
(MovingImagePyramid "MovingSmoothingImagePyramid")
(MovingImagePyramidSchedule 4 4 4 2 2 2 1 1 1)

// Transform: settings for the transform go here
(Transform "EulerTransform") // euler transform for rigid registration only
   allows rotations and translations
(AutomaticScalesEstimation "true") // automatically estimate scales of
   rotations and translations
// centre of rotation can be left as default (geometric centre)
(AutomaticTransformInitialization "false") // no need for automatic
   initialisation as image centres should be fairly close in general
(AutomaticTransformInitializationMethod "GeometricalCenter")
(HowToCombineTransforms "Compose") // transforms feed into each other by
   composition T(x) = T_1(T_0(x)) rather than by addition

// Optimiser: settings for the optimiser go here
(Optimizer "AdaptiveStochasticGradientDescent")
(MaximumNumberOfIterations 500 750 1000)
(MaximumNumberOfSamplingAttempts 10) // how many times samples will be
   attempted to be drawn
(AutomaticParameterEstimation "false") // prevent automatic parameter
   estimation as it can be too aggressive
(UseAdaptiveStepSizes "true") // set to true for convenience, when using random
    sample coordinates this may be overwritten
(MaximumStepLength 1.0) // when automatic parameter estimation is false this
   has no effect
(SP_a 500 250 100) // set to make step sizes more conservative, smaller as
   registration becomes more refined
(SP_A 20) // set to make step sizes more conservative
(SP_alpha 1.0) // set to make step sizes more conservative
(SigmoidMax 1.0) // set to make step sizes more conservative
(SigmoidMin -0.8) // set to make step sizes more conservative
(SigmoidScale 0.00000001) // set to make step sizes more conservative
(SigmoidInitialTime 1.0) // larger numbers give smaller steps, 0 is default
(NumberOfGradientMeasurements 0) // default 0 gives good automatic estimation
// number of jacobian measurements is left to be initialised by elastix
(NumberOfSamplesForExactGradient 100000) // default
(ASGDParameterEstimationMethod "DisplacementDistribution") // used to estimate
   parameters faster than "Original" method
(MaximumDisplacementEstimationMethod "2sigma")
(NoiseCompensation "true")

// Metric: settings for the metric (or metrics) go here
```

```
(Metric "AdvancedMattesMutualInformation") // mutual information is robust when
    intensities may vary (e.g. multi modal or in our case magic angle effect)
(NumberOfHistogramBins 32) // default
(NumberOfFixedHistogramBins 32) // default
(NumberOfMovingHistogramBins 32) // default
(FixedKernelBSplineOrder 0) // default
(MovingKernelBSplineOrder 3) // default
(FixedLimitRangeRatio 0.01) // default
(MovingLimitRangeRatio 0.01) // default
(FiniteDifferenceDerivative "false") // experimental do not use
(UseFastAndLowMemoryVersion "true")

// Mask: settings for masks (fixed and moving) go here
(ErodeMask "true") // true prevents artificial edge from corrupting
   registration
(ErodeFixedMask "true")
(ErodeMovingMask "true")

// Misc
// There are several warnings that appear in the logs for undocumented
   parameters, best to let elastix initialise these to default values
```

C.1.3 Affine Registration Parameters

```
// AFFINE REGISTRATION parameter file

// Image metadata: information about the image goes here
(FixedInternalImagePixelType "float") // float as output image will be nifti
   rather than dicom
(MovingInternalImagePixelType "float")
(FixedImageDimension 3) // 3D images
(MovingImageDimension 3)
(UseDirectionCosines "true")
(DefaultPixelValue 0) // when image is deformed "new" pixels should be
   initialised to 0
(WriteResultImage "true")
(ResultImagePixelType "double")
(ResultImageFormat "nii") // use nifti format to prevent conversions to and
   from dicom incurring rounding errors

// Registration: registration type and interpolation settings go here
(Registration "MultiResolutionRegistration") // Use multiple resolutions but
   only one metric
(Interpolator "BSplineInterpolator") // Interpolator used during registration
(BSplineInterpolationOrder 1)
(ResampleInterpolator "FinalBSplineInterpolator") // Interpolator used to
   generate final image
(FinalBSplineInterpolationOrder 3) // used to create final image, can create
   bright spots
(Resampler "DefaultResampler") // Use default (other options require
   compilation from source)
```

```
(ImageSampler "RandomCoordinate") // sample points for registration are
   selected on random coordinates (not on grid)
(FixedImageBSplineInterpolationOrder 1) // for interpolating random coordinates
(NumberOfSpatialSamples 5000) // low number of samples possible, less than 2000
    is best avoided
(NewSamplesEveryIteration "true") // avoid sampling same points over and over
   to get better overall registration
(UseRandomSampleRegion "false") // probably best to avoid this as the images
   are already constrained by the artificial edge
(CheckNumberOfSamples "true") // check if samples are within valid region
(RequiredRatioOfValidSamples 0.25) // at least a quarter should be valid

// Image resolutions: settings for different resolutions goes here
(NumberOfResolutions 3) // number of resolutions to refine registration over
(FixedImagePyramid "FixedSmoothingImagePyramid") // pyramids refine resolution
   to initially match larger features
(FixedImagePyramidSchedule 4 4 4 2 2 2 1 1 1)
(MovingImagePyramid "MovingSmoothingImagePyramid")
(MovingImagePyramidSchedule 4 4 4 2 2 2 1 1 1)

// Transform: settings for the transform go here
(Transform "AffineTransform") // affine transform allows scaling and shearing
   as well as translation and rotation
(AutomaticScalesEstimation "true") // automatically estimate scales of
   rotations and translations
// centre of rotation can be left as default (geometric centre)
(AutomaticTransformInitialization "false") // no need for automatic
   initialisation as image centres should be fairly close in general
(AutomaticTransformInitializationMethod "GeometricalCenter")
(HowToCombineTransforms "Compose") // transforms feed into each other by
   composition T(x) = T_1(T_0(x)) rather than by addition

// Optimiser: settings for the optimiser go here
(Optimizer "AdaptiveStochasticGradientDescent")
(MaximumNumberOfIterations 750 1500 2000)
(MaximumNumberOfSamplingAttempts 10) // how many times samples will be
   attempted to be drawn
(AutomaticParameterEstimation "false") // prevent automatic parameter
   estimation as it can be too aggressive
(UseAdaptiveStepSizes "true") // set to true for convenience, when using random
    sample coordinates this may be overwritten
(MaximumStepLength 1.0) // when automatic parameter estimation is false this
   has no effect
(SP_a 1000 500 250) // set to make step sizes more conservative, smaller as
   registration becomes more refined
(SP_A 20) // set to make step sizes more conservative
(SP_alpha 1.0) // set to make step sizes more conservative
(SigmoidMax 1.0) // set to make step sizes more conservative
(SigmoidMin -0.8) // set to make step sizes more conservative
(SigmoidScale 0.00000001) // set to make step sizes more conservative
(SigmoidInitialTime 1.0) // larger numbers give smaller steps, 0 is default
(NumberOfGradientMeasurements 0) // default 0 gives good automatic estimation
// number of jacobian measurements is left to be initialised by elastix
```

```
(NumberOfSamplesForExactGradient 100000) // default
(ASGDParameterEstimationMethod "DisplacementDistribution") // used to estimate
   parameters faster than "Original" method
(MaximumDisplacementEstimationMethod "2sigma")
(NoiseCompensation "true")

// Metric: settings for the metric (or metrics) go here
(Metric "AdvancedMattesMutualInformation") // mutual information is robust when
    intensities may vary (e.g. multi modal or in our case magic angle effect)
(NumberOfHistogramBins 32) // default
(NumberOfFixedHistogramBins 32) // default
(NumberOfMovingHistogramBins 32) // default
(FixedKernelBSplineOrder 0) // default
(MovingKernelBSplineOrder 3) // default
(FixedLimitRangeRatio 0.01) // default
(MovingLimitRangeRatio 0.01) // default
(FiniteDifferenceDerivative "false") // experimental do not use
(UseFastAndLowMemoryVersion "true")

// Mask: settings for masks (fixed and moving) go here
(ErodeMask "true") // true prevents artificial edge from corrupting
   registration
(ErodeFixedMask "true")
(ErodeMovingMask "true")
```

C.1.4 B-Spline Registration Parameters

```
// SOFT REGISTRATION - BSPLINES parameter file

// Image metadata: information about the image goes here
(FixedInternalImagePixelType "float") // float as output image will be nifti
   rather than dicom
(MovingInternalImagePixelType "float")
(FixedImageDimension 3) // 3D images
(MovingImageDimension 3)
(UseDirectionCosines "true")
(DefaultPixelValue 0) // when image is deformed "new" pixels should be
   initialised to 0
(WriteResultImage "true")
(ResultImagePixelType "double")
(ResultImageFormat "nii") // use nifti format to prevent conversions to and
   from dicom incurring rounding errors


// Registration: registration type and interpolation settings go here
(Registration "MultiMetricMultiResolutionRegistration") // Use multiple
   resolutions and multiple metrics
(Interpolator "BSplineInterpolator") // Interpolator used during registration
(BSplineInterpolationOrder 1)
(ResampleInterpolator "FinalBSplineInterpolator") // Interpolator used to
   generate final image
```

```
(FinalBSplineInterpolationOrder 3) // used to create final image, can create
   bright spots
(Resampler "DefaultResampler") // Use default (other options require
   compilation from source)
(ImageSampler "RandomCoordinate") // sample points for registration are
   selected on random coordinates (not on grid)
(FixedImageBSplineInterpolationOrder 1) // for interpolating random coordinates
(NumberOfSpatialSamples 5000) // low number of samples possible, less than 2000
    is best avoided
(NewSamplesEveryIteration "true") // avoid sampling same points over and over
   to get better overall registration
(UseRandomSampleRegion "false") // probably best to avoid this as the images
   are already constrained by the artificial edge
(CheckNumberOfSamples "true") // check if samples are within valid region
(RequiredRatioOfValidSamples 0.25) // at least a quarter should be valid

// Image resolutions: settings for different resolutions goes here
(NumberOfResolutions 2) // number of resolutions to refine registration over
(FixedImagePyramid "FixedSmoothingImagePyramid") // pyramids refine resolution
   to initially match larger features
(FixedImagePyramidSchedule 2 2 2 1 1 1)
(MovingImagePyramid "MovingSmoothingImagePyramid")
(MovingImagePyramidSchedule 2 2 2 1 1 1)

// Transform: settings for the transform go here
(Transform "RecursiveBSplineTransform") // b spline transform for soft
   registration
(BSplineTransformSplineOrder 3)
(FinalGridSpacingInVoxels 12.0 12.0 12.0) // 12 voxels, order of grid spacing
   vs local deformation should be similar
(GridSpacingSchedule 2.0 2.0 2.0 1.0 1.0 1.0)
(PassiveEdgeWidth 0) // bands control points that are fixed towards the edges
   of the image
(HowToCombineTransforms "Compose") // transforms feed into each other by
   composition T(x) = T_1(T_0(x)) rather than by addition

// Optimiser: settings for the optimiser go here
(Optimizer "AdaptiveStochasticGradientDescent")
(MaximumNumberOfIterations 1500 2500)
(MaximumNumberOfSamplingAttempts 50) // how many times samples will be
   attempted to be drawn
(AutomaticParameterEstimation "false") // prevent automatic parameter
   estimation as it can be too aggressive
(UseAdaptiveStepSizes "true") // set to true for convenience, when using random
    sample coordinates this may be overwritten
(MaximumStepLength 1.0) // when automatic parameter estimation is false this
   has no effect
(SP_a 250 100) // set to make step sizes more conservative, smaller as
   registration becomes more refined
(SP_A 20) // set to make step sizes more conservative
(SP_alpha 1.0) // set to make step sizes more conservative
(SigmoidMax 1.0) // set to make step sizes more conservative
(SigmoidMin -0.8) // set to make step sizes more conservative
```

```
(SigmoidScale 0.00000001) // set to make step sizes more conservative
(SigmoidInitialTime 1.0) // larger numbers give smaller steps, 0 is default
(NumberOfGradientMeasurements 0) // default 0 gives good automatic estimation
// number of jacobian measurements is left to be initialised by elastix
(NumberOfSamplesForExactGradient 100000) // default
(ASGDParameterEstimationMethod "DisplacementDistribution") // used to estimate
   parameters faster than "Original" method
(MaximumDisplacementEstimationMethod "2sigma")
(NoiseCompensation "true")

// Metric: settings for the metric (or metrics) go here
(Metric "AdvancedMattesMutualInformation" "TransformBendingEnergyPenalty") //
   mutual information is robust when intensities may vary (e.g. multi modal or
    in our case magic angle effect), bending energy prevents crazy distortions
(Metric0Weight 1.0) // weight mutual information and bending energy
(Metric1Weight 0.01)
(Metric0Use "true")
(Metric1Use "true")
(NumberOfHistogramBins 32) // default
(NumberOfFixedHistogramBins 32) // default
(NumberOfMovingHistogramBins 32) // default
(FixedKernelBSplineOrder 0) // default
(MovingKernelBSplineOrder 3) // default
(FixedLimitRangeRatio 0.01) // default
(MovingLimitRangeRatio 0.01) // default
(FiniteDifferenceDerivative "false") // experimental do not use
(UseFastAndLowMemoryVersion "true")

// Mask: settings for masks (fixed and moving) go here
(ErodeMask "true") // true prevents artificial edge from corrupting
   registration
(ErodeFixedMask "true")
(ErodeMovingMask "true")


// Misc
(UseCyclicTransform "false")
```

C.2 CadQuery Python Code for Calibration Phantom

```
import cadquery as cq
from math import floor, ceil
import time
from cadquery import exporters
from math import tan, cos, pi

# define cylinder wall parameters
outer_rad = 142/2
wall_thickness = 10
inner_rad = outer_rad-wall_thickness
mid_rad = (outer_rad+inner_rad)/2
```

```
# define sphere parameters
sphere_spacing = 17.5 # from center to center
sphere_rad = 5
central_sphere_rad = 7.5 # for sphere right in the middle of the phantom

# define grate parameters
grate_line_width = 1.2
grate_height = 1.2

semi_points = floor(outer_rad/sphere_spacing)
num_points = 2*semi_points+1
coords_start = -semi_points*sphere_spacing

coords_vec = [coords_start + i*sphere_spacing for i in range(num_points)]

# define phantom structure

phantom_structure = ["slice", "slice", "slice", "slice", "center", "slice", "
    slice", "slice", "slice"]
num_slices = len(phantom_structure)

def make_sphere_grid(z_height= 0, central_slice = False):
    points = []
    # generate points that will be occupied by spheres
    for x_coord in coords_vec:
        for y_coord in coords_vec:
            if (x_coord**2 + y_coord**2)**0.5 + sphere_rad < inner_rad:
                points.append((x_coord, y_coord))

    if central_slice:
        points.remove((0,0))

    gridded_spheres = (cq.Workplane("XY")
                    .workplane(offset=z_height)
                    .pushPoints(points)
                    .sphere(sphere_rad))

    if central_slice:
        gridded_spheres = gridded_spheres.sphere(central_sphere_rad)

    return gridded_spheres

def make_grate(z_height):
    points = [(y, z_height) for y in coords_vec]
    grateYZ = (cq.Workplane("YZ")
            .pushPoints(points)
            .rect(xLen=grate_line_width, yLen=grate_height, centered=True)
            .extrude(outer_rad, both=True))

    grateXZ = (cq.Workplane("XZ")
            .pushPoints(points)
            .rect(xLen=grate_line_width, yLen=grate_height, centered=True)
```

```
            .extrude(outer_rad, both=True))

    intersection_cylinder = (cq.Workplane("XY")
                          .workplane(offset=z_height)
                          .cylinder(grate_height, mid_rad))

    grate = grateYZ.union(grateXZ).intersect(intersection_cylinder)

    return grate

def make_wall(start_height = 0, end_height = sphere_spacing*num_slices):
    phantom_body = (cq.Workplane("XY")
                  .workplane(offset=start_height)
                  .circle(outer_rad)
                  .circle(inner_rad)
                  .extrude(end_height-start_height))


    return phantom_body

phantom_wall = make_wall()

feature_height = sphere_spacing/2 # height of plane that features are centered
    around

grates = []
spheres = []

for i in range(num_slices):
    grates.append(make_grate(feature_height + i*sphere_spacing))
    if i == num_slices//2:
        center_flag = True
    else:
        center_flag = False
    spheres.append(make_sphere_grid(feature_height + i*sphere_spacing,
        center_flag))
    phantom_wall = phantom_wall.union(grates[i]).union(spheres[i])

# add lip for glueing on base
lip_height = 3
base_workplane_offset = 30
base_thickness = 5
inset_tol = 0.2
circle_chamfer_length = 1

phantom_wall = phantom_wall.faces("<Z").circle(outer_rad).circle(inner_rad).
    extrude(-lip_height)

base = (cq.Workplane("XY")
        .workplane(offset=-base_workplane_offset) #
        .circle(outer_rad)
        .extrude(base_thickness)
        .faces(">Z")
```

```
        .circle(inner_rad-inset_tol)
        .extrude(lip_height)
        .faces(">Z")
        .chamfer(circle_chamfer_length))

# add lip for glueing on lip
lid_workplane_offset = 30
lid_thickness = 10
circle_chamfer_length = 1
lid_border = 1
lid_top_thickness = 3
taper_ang = 75
taper_ang_rad = taper_ang*pi/180
top_hole_rad = (inner_rad-inset_tol-circle_chamfer_length-lid_border) - (
    lid_thickness+lip_height-lid_top_thickness)*tan(taper_ang_rad)
marker_rad = 1
joint_height = sphere_spacing/3
joint_offset = 3

phantom_wall = (phantom_wall
                .faces(">Z")
                .circle(outer_rad)
                .circle(inner_rad)
                .extrude(lip_height)
                .faces(">Z")
                .workplane(offset=0)
                .pushPoints([(outer_rad, 0), (0, outer_rad), (-outer_rad, 0), (0,
                    -outer_rad)])
                .circle(marker_rad)
                .cutThruAll())

lid = (cq.Workplane("XY")
       .workplane(offset=sphere_spacing*num_slices+lid_workplane_offset)
       .circle(inner_rad-inset_tol)
       .extrude(lip_height)
       .faces("<Z")
       .chamfer(circle_chamfer_length)
       .faces(">Z")
       .circle(outer_rad)
       .extrude(lid_thickness)
       .faces("<Z")
       .circle(inner_rad-inset_tol-circle_chamfer_length-lid_border)
       .cutBlind((lid_thickness+lip_height-lid_top_thickness)/cos(taper_ang_rad)
           , taper=taper_ang)
       .faces(">Z")
       .hole(top_hole_rad*2))

plug = (lid
        .faces(">Z")
        .workplane(offset=10)
        .circle(top_hole_rad*1.5)
        .extrude(5, combine=False))
```

```
rib = (lid.faces(">Z")
      .workplane(offset=30)
      .circle(outer_rad-joint_offset-0.15)
      .circle(inner_rad+joint_offset+0.15)
      .extrude(joint_height-0.8, combine=False))


split_offset = feature_height + lip_height

for i in range(num_slices):
    if i == 0:
        split_offset = sphere_spacing + lip_height
        split_offset_2 = sphere_spacing/2 + lip_height
    elif i == num_slices-1:
        split_offset = sphere_spacing + lip_height
        split_offset_2 = sphere_spacing/2
    else:
        split_offset = sphere_spacing
        split_offset_2 = sphere_spacing/2

    phantom_wall_top = phantom_wall.faces(">Z").workplane(offset=-split_offset).
        split(keepTop=True) # split at halfway between sphere layers
    phantom_wall = phantom_wall.faces(">Z").workplane(offset=-split_offset).
        split(keepBottom=True) # split at halfway between sphere layers, keep
        bottom for next iteration

    ## TOP HALF OF SLICE
    phantom_wall_top_top = phantom_wall_top.faces(">Z").workplane(offset=-
        split_offset_2).split(keepTop=True) # get top half of slice (split
        halfway up spheres)

    if i != 0:
        phantom_wall_top_top = phantom_wall_top_top.faces(">Z").circle(outer_rad
            -joint_offset).circle(inner_rad+joint_offset).cutBlind(-joint_height
            /2) # create joint ring hole on top of slice

    phantom_wall_top_top = phantom_wall_top_top.faces("<Z").circle(outer_rad-
        joint_offset).circle(inner_rad+joint_offset).cutBlind(joint_height/2) #
        create joint ring hole on bottom of slice

    ## BOTTOM HALF OF SLICE
    phantom_wall_top_bot = phantom_wall_top.faces(">Z").workplane(offset=-
        split_offset_2).split(keepBottom=True) # get bottom half of slice (split
         halfway up spheres)
    phantom_wall_top_bot = phantom_wall_top_bot.faces(">Z").circle(outer_rad-
        joint_offset).circle(inner_rad+joint_offset).cutBlind(-joint_height/2) #
         create joint ring hole on top of slice

    if i != num_slices-1:
        phantom_wall_top_bot = phantom_wall_top_bot.faces("<Z").circle(outer_rad
            -joint_offset).circle(inner_rad+joint_offset).cutBlind(joint_height
            /2) # create joint ring hole on top of slice
```

```
    exporters.export(phantom_wall_top_top, f"./step_slices/slice_{i}_a_top.step
        ")
    exporters.export(phantom_wall_top_bot, f"./step_slices/slice_{i}_b_bot.step
        ")


exporters.export(lid, f"./step_slices/lid.step")
exporters.export(base, f"./step_slices/base.step")
exporters.export(plug, f"./step_slices/plug.step")
exporters.export(rib, f"./step_slices/rib.step")
```

Appendix D: Segmentation Appendices

D.1 Full Representative Dataset and Derived Metric Image Projections



Figure D.1: Representative dataset $n = 1$.



Figure D.2: Representative dataset $n = 2$.

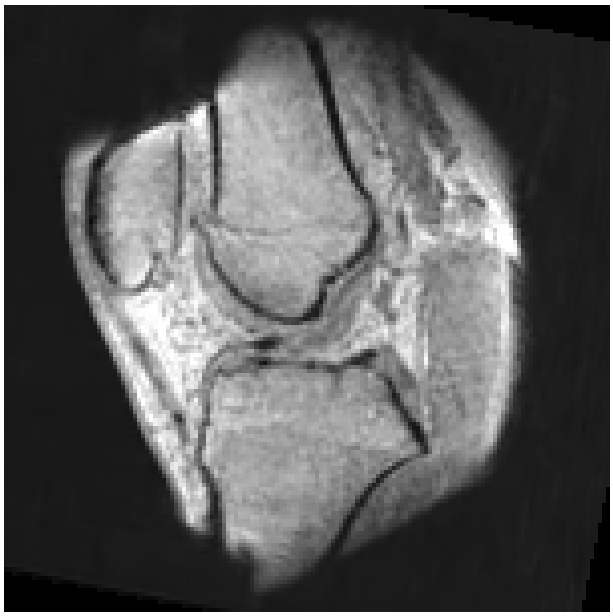


Figure D.3: Representative dataset $n = 3$.

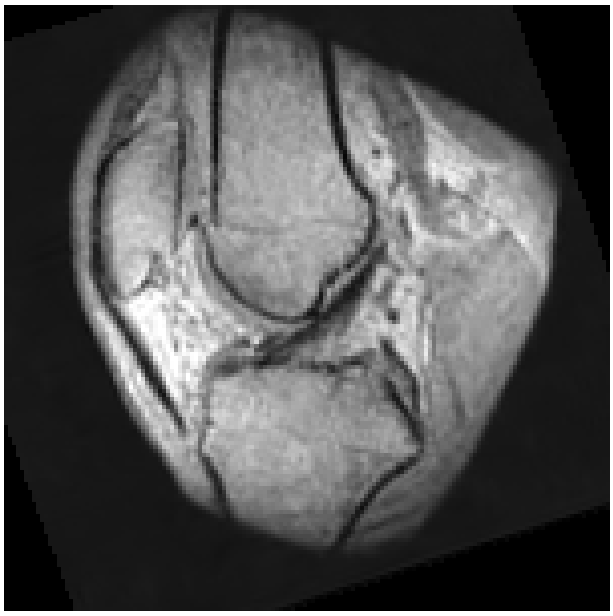


Figure D.4: Representative dataset $n = 4$.



Figure D.5: Representative dataset $n = 5$.



Figure D.6: Representative dataset $n = 6$.



Figure D.7: Representative dataset $n = 7$.



Figure D.8: Representative dataset minimum projection.



Figure D.9: Representative dataset mean projection.

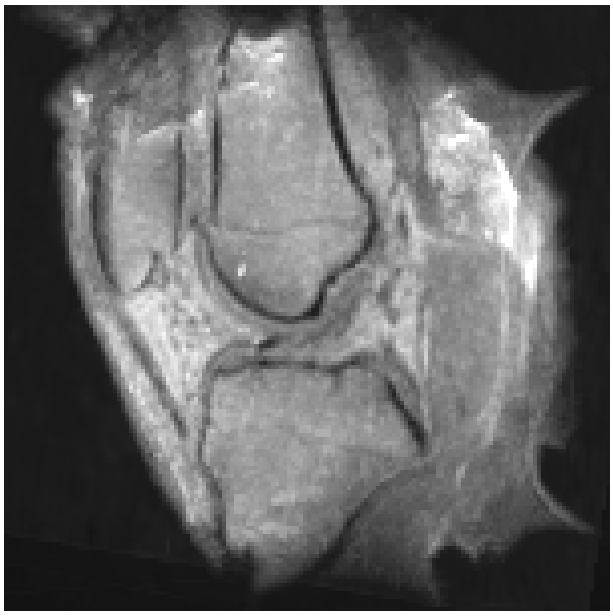


Figure D.10: Representative dataset maximum projection.

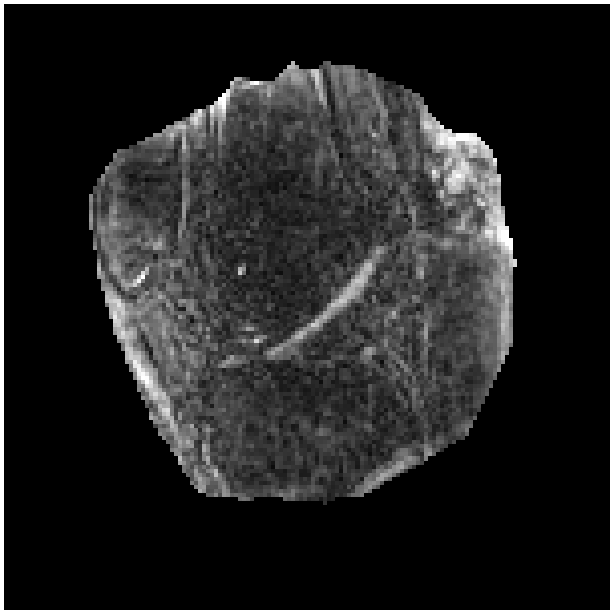


Figure D.11: Representative dataset standard deviation projection.

Appendix E: Collagen Fibre Orientation Estimation Appendices

E.1 Collagen Fibre Solver GUI

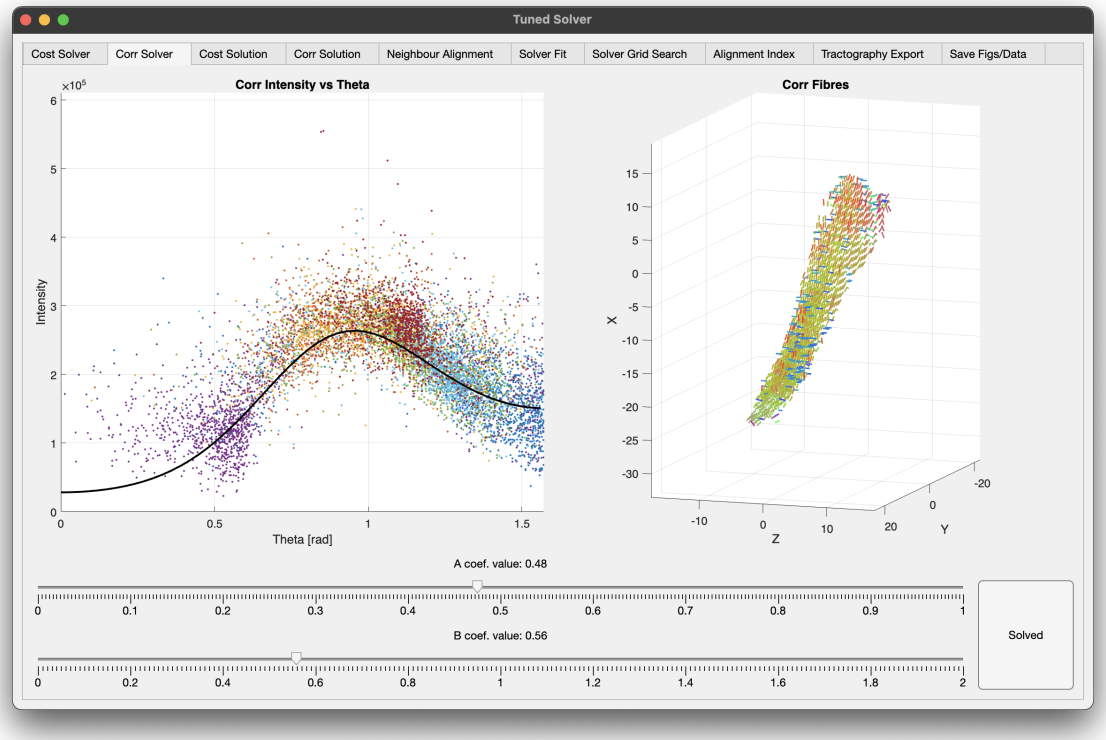


Figure E.1: Solver GUI developed for fine tuning A and B parameters and also generating output figures for MADI process.

E.2 Golden Spiral Points on Sphere

This function will generate points on a sphere that are spread out in steps of the golden angle. This pattern appears in many places in nature such as in sunflower seed heads. It is a simple yet effective way of distributing points roughly evenly.

```
function sphere_points = generate_dirs_golden(num_points)
    % define golden angle
    golden_angle = pi*(3 - sqrt(5));

    % space points evenly in z
    point_indices = linspace(0, (num_points - 1), num_points).';
```

```
    % space the points evenly from +1 to -1
    z_coords = 1 - 2*(point_indices/(num_points-1));

    % get the radius at each z height
    r = sqrt(1 - z_coords.^2);

    % offset each point from its neighbours by the golden angle
    theta = golden_angle*point_indices;

    x_coords = r.*cos(theta);
    y_coords = r.*sin(theta);

    sphere_points = [x_coords y_coords z_coords];

end
```

E.3 Genetic Algorithm Binary Encoding

Each agent in the genetic algorithm consisted of binary genes that encoded θ_R and ϕ_Y in degrees (as integers) for N scan orientations, in the 2's complement format. For example, to represent roll angles from -90° to 90°, and yaw angles from -20° to 20° the encoding was as follows:

$$\texttt{num\_roll\_bits} = \lceil \log_2(90 - (-90) + 1) \rceil = 8$$
$$\texttt{num\_yaw\_bits} = \lceil \log_2(20 - (-20) + 1) \rceil = 6$$

Where the number of bits used was large enough to encode all the positive and negative integers in the desired range and 0 (hence the $+1$). The 2's complement representation of each integer was a sequence of bits that represented coefficients for the following values in these decoding vectors:

$$\texttt{roll\_decode\_vec} = [-2^7 \;\; 2^6 \;\; 2^5 \;\; 2^4 \;\; 2^3 \;\; 2^2 \;\; 2^1 \;\; 2^0]$$
$$\texttt{yaw\_decode\_vec} = [-2^5 \;\; 2^4 \;\; 2^3 \;\; 2^2 \;\; 2^1 \;\; 2^0]$$

Thus allowing for the integers $(-2^{\texttt{num\_bits}-1})$ to $(2^{\texttt{num\_bits}-1}-1)$ to be represented in binary format. Therefore, the complete chromosome for an agent, a, of the population was stored as a row of the, `pop`, population array which was of size [`num_agents`, $(\texttt{num\_roll\_bits + num\_yaw\_bits}) \cdot N$]:

$$\texttt{pop}[a,:] = [[\texttt{roll\_bits\_2}] \ \ldots \ [\texttt{roll\_bits\_N}] \ [\texttt{yaw\_bits\_2}] \ \ldots \ [\texttt{yaw\_bits\_N}]]$$

Where the θ_R and ϕ_Y for scan n are encoded in [`roll_bits_n`] and [`yaw_bits_n`] respectively, for $n \in [2, N]$. Scan 1 was never encoded as it was always $\theta_R = 0^\circ,\ \phi_Y = 0^\circ$ for every set of scan orientations. During initialisation of the 1st generation population in Algorithm 1, each agent was initialised with random roll and yaw angles such that $\theta_R \in [-\texttt{roll\_lim\_deg}, \texttt{roll\_lim\_deg}]$ and $\phi_Y \in [-\texttt{yaw\_lim\_deg}, \texttt{yaw\_lim\_deg}]$. This ensured that the first population consisted entirely of viable agents, therefore increasing the chances of producing viable children.

Appendix F: Results Appendices

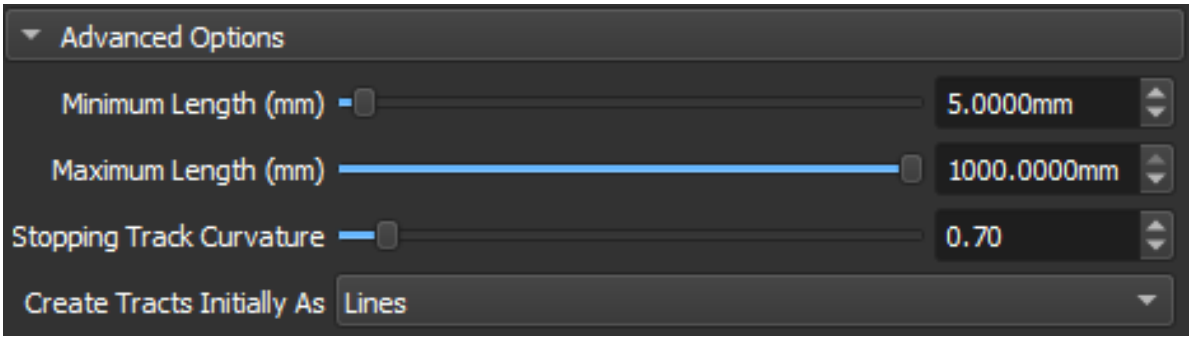


Figure F.1: 3D Slicer DTI tract advanced settings to control tract propagation.

Figure F.1 shows the advanced settings in 3D Slicer that control tract propagation. Minimum length prevents fibres of less than the prescribed length from generating, maximum length is similar but not a concern for the structures of interest in the present work. Stopping tract curvature is a setting to control the radius of curvature (in degrees per mm) that will prevent tract generation. Tracts with too aggressive curvature will therefore not generate, which prevents tract generation in incoherent fibre regions such as those in the liberal segmentation of Figure 8.10.

Appendix G: Volunteer Consent Form

Imperial College London

Consent Form for Participants Able to Give Consent

Full Title of Project: Imaging joints using Magic Angle Directional Imaging (MADI)

Name of Principal Investigator: Mihailo Ristic

	Please initial box
1. I confirm that I have read and understand the participant information sheet version 1.0 dated 08/06/2022 for Imaging joints using Magic Angle Directional Imaging (MADI) and have had the opportunity to ask questions which have been answered fully.	
2. I understand that my participation is voluntary, and I am free to withdraw at any time, without giving any reason and without my legal rights nor treatment / healthcare being affected.	
3. I give/do not give (delete as applicable) consent for information collected about me to be used to support other research or in the development of a new test, medication, medical device or treatment (delete as applicable) by an academic institution or commercial company in the future, including those outside of the United Kingdom (which Imperial has ensured will keep this information secure).	
4. I understand that data collected from me are a gift donated to Imperial College and that I will not personally benefit financially if this research leads to an invention and/or the successful development of a new test, medication treatment, product or service.	
5. I consent to take part in Imaging joints using Magic Angle Directional Imaging.	
6. I consent to my GP being contacted by either the Radiologist or Orthopaedic Surgeon should any incidental findings be found on my joint images during participation in the study.	
7. I give / do not give (delete/mark as applicable) consent to being contacted about the possibility to take part in other research studies.	

_____	_____	_____
Name of participant	Signature	Date
_____	_____	_____
Name of person taking consent (if different from Principal Investigator)	Signature	Date
_____	_____	_____
Principal Investigator	Signature	Date

1 copy for participant; 1 copy for Principal Investigator

To ensure confidence in the process and minimise risk of loss, all consent forms must be printed, presented, and stored in double sided format