Correction of slice misalignment in cardiac MRI studies for the application of patient-specific geometry



Benjamin Villard

Kellogg College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Hillary 2018

Acknowledgements

First and foremost I would like to acknowledge and thank my supervisors, Professor Vicente Grau and Dr. Ernesto Zacur for their constant guidance and support throughout my DPhil. My sincere thanks go to Vicente Grau for his unlimited source of knowledge and insight during our weekly meetings. I appreciate all the help you have provided throughout these years, and your consistent belief in me.

I am extremely grateful to Ernesto Zacur, without who this DPhil would not have been possible. Thank you for all of the help with the last minute late night write-ups and coding. You have truly been a wonderful mentor, friend, and a constant source of motivation. Your teachings have greatly inspired me and have taught me many things, but most importantly, how to properly be proper and think in the most abstract way possible.

I would like to thank Dr. Ramón Casero Cañas for helping me find my initial steps and introducing me to the framework of code development. Thank you for your guidance and help.

I would like to acknowledge all my research group, who made the lab a wonderful environment to work in. I am thankful for the CDT in Healthcare Innovation and the EPSRC for funding this project.

I started and finished my DPhil with Glen Wright Colopy. I am extremely thankful to have shared this journey with you. I have really enjoyed all of the late nights (and seasons) spent with you working in the lab. You have made these years in Oxford the most wonderful. I would like to thank the Colopy family as well. Thank you all so much for your kindness and constant thoughts during this DPhil.

I would like to thank Dr. Konrad Weris. Thank you for being a fun and caring friend. Thank you so much for helping me stay sane and fit during the DPhil.

I am extremely grateful to Dr. Kate Reynolds. Thank you so much for your constant encouragement, and inexhaustible patience. Thank you for taking care of me and providing me with the best environment possible to finish my Thesis. I would also like to thank the Reynolds family for helping me towards the end of my DPhil. Thank you for all of your kindness and support.

Finally, I would like to thank my Mom to whom I am forever indebted. Thank you for making me love learning and allowing me to pursue my dreams.

Abstract

Cardiac Magnetic Resonance (CMR) provides unique anatomical and functional information of the structures of the heart. To quantitatively assess cardiac functions, clinicians typically segment CMR slices in 2D. These can be used for 3D reconstructions allowing to better understand cardiovascular functions as well as predict different cardiac processes. These reconstructions can be used in the area of precision medicine or for shape analysis among populations, as well as serve as the input for biological simulations. However, CMR acquisition times usually necessitate slices to be acquired at different breath holds, which results in potential misalignment of the acquired slices. These artifacts cause any 3D model and subsequent evaluations to become distorted. Correcting for this spatial misalignment is required for accurate 3D reconstruction of the heart.

In order to go from cardiac imaging to the geometric definition of patient-specific anatomies, many steps must be taken to achieve accurate geometric representations. Imaging, segmentation, and meshing challenges make the generation of anatomically realistic volumetric meshes a complex task.

The main aim of this Thesis is to develop tools and methodologies to generate geometrically correct cardiac meshes in the presence of spatial misalignment between slices, to represent the anatomy of patients in a personalized manner.

Three main novel contributions are presented. The first concerns the development of a complete framework for the correction of translational and rotational misalignments between CMR slices, using only the information from the images, allowing the method to be used directly in the clinics. The method aligns the respective slices together by using the line intensities at their intersections. We introduce the normalized cross correlation of local phase vectors as a similarity measure making the method applicable to different CMR protocols, and perform a complete validation using manually traced contours, including the estimation of contouring errors. The second concerns the development of a 3D surface reconstruction method capable of dealing with typical cardiac segmentations which are non-parallel, cross-sectional, sparse, heterogeneous, and non-coincidental contours. An evaluation of the method is performed and a parameter study achieve to tune the method parameters to cardiac applications. Finally a third methodology is presented which includes both prior works and incorporates the individual steps of segmentation, alignment and mesh reconstruction in one method. This method provides a fully automatic way to go from images to 3D personalized models.

Publications

The work found in this thesis has been presented in the following publications:

Journal Articles

1. Villard, B., Grau, V., Zacur, E. Surface Mesh Reconstruction from Cardiac MRI Contours. Journal of Imaging, 4 (1), 16. (2018).

Conference Articles and Proceedings

1. Zacur, E., Minchole., Villard, B., Carapella, V., Ariga, R., Rodriguez, B., Grau, V.: MRI-based heart and torso personalization for computer modeling and simulation of cardiac electrophysiology. BIVPCS 2017. Imaging for Patient-Customized Simulations and Systems for Point-of-Care Ultrasound, pp.61-70 (2017).

2. Villard, B., Zacur, E., Dall'Armellina, E., Grau, V.: Correction of slice misalignment in multi-breath-hold cardiac MRI scans. Statistical Atlases and Computational Models of the Heart. Imaging and Modelling Challenges. STACOM 2016. Lecture Notes in Computer Science (LNCS) 10124 (2017). **Best Paper Award STACCOM 2016**

3. Villard, B., Carapella, V., Ariga, R., Grau, V., Zacur, E.: Cardiac mesh reconstruction from sparse, heterogeneous contours. Medical Image Understanding and Analysis. MIUA 2017. Communications in Computer and Information Science 723 (2017).

4. Minchole, Zacur, Lyon, Ariga, Carapella, Villard, Martinez-Navarro, Bueno-Orovio, Watkins, Grau, Rodriguez.: Investigating the dependency of the QRS complex with the MRI-based heart/torso geometries using personalised computer models. In Computing in Cardiology (CinC), Rennes (France), Sep 2017.

5. Villard, B., Grau, V., Zacur, E. SurFCo: Surface From Contours. In International Symposium on Biomedical Imaging (ISBI), Washington DC, (USA), April 2018.

6. Villard, B., Zacur, E., Grau, V. ISACHI: Integrated Segmentation and Alignment Correction for Heart Images. Engineering in Medicine and Biology Society (EMBC) 2018. Honolulu (Hawai).

In Preparation

1. Villard, B., Zacur, E., Grau, V. ISACHI: Integrated Segmentation and Alignment Correction for Heart Images. In Medical Image Computing and Computer Assisted Intervention (MICCAI) 2018. Granada (Spain).

2. Villard, B., Zacur, E., Ariga, R., Neubauer S., Grau, V. Integration of Segmentation and Alignment Correction of Heart Images towards personalized 3D Models. In Medical Image Analysis (MIA).

Contents

List of Figures **ix**

List of Abbreviations **xii**

1 Introduction **1**
1.1 Motivation . . . 1
1.2 Thesis Contributions . . . 4
1.3 Thesis Overview . . . 6

2 Clinical Background **8**
2.1 Cardiac Physiology and Pathophysiology . . . 8
2.1.1 Cardiac Physiology . . . 8
2.1.2 Cardiac Pathophysiology . . . 16
2.2 Cardiac Tomographic Images . . . 19
2.2.1 Ultrasound . . . 20
2.2.2 Computed Tomography . . . 21
2.2.3 Magnetic Resonance Imaging . . . 22

3 Technical Background **33**
3.1 Image Registration Tools . . . 33
3.1.1 Transformation Models . . . 35
3.1.2 Similarity measures . . . 41
3.1.3 Multilevel image registration . . . 49
3.1.4 Limits of CMRI in image registration . . . 49
3.2 Mesh Reconstruction . . . 51
3.2.1 Common surface construction algorithms . . . 51
3.2.2 Geometric tools . . . 53

4 Literature Review **60**
4.1 Misalignment between CMR slices . . . 60
4.1.1 Motion correction . . . 62
4.2 Surface reconstruction from cardiac contours . . . 68
4.3 Misalignment correction and surface reconstruction integration . . . 70
4.4 Conclusion . . . 71

5 Correction of Slice Misalignment in Cardiac MRI 73
5.1 Introduction . . . 73
5.2 CMR data . . . 77
5.3 Methodology . . . 77
5.3.1 Alignment score and optimization strategy . . . 78
5.3.2 Similarity measure . . . 80
5.3.3 Contour-based alignment . . . 83
5.3.4 Label-based Alignment . . . 83
5.4 Results . . . 85
5.5 Discussion . . . 87
5.6 Conclusion . . . 91

6 Geometrical Representation of Cardiac Anatomy 92
6.1 Introduction . . . 93
6.2 Materials and Methods . . . 95
6.2.1 Synthetic and Real Data . . . 95
6.2.2 Initial Meshes . . . 97
6.2.3 Mesh Deformation . . . 98
6.3 Results . . . 100
6.3.1 Simulation of CMR Acquisitions . . . 101
6.3.2 Synthetic Data . . . 104
6.3.3 CMR Acquired Data . . . 106
6.4 Discussion . . . 114
6.5 Conclusions . . . 118

7 Integrated Segmentation and Alignment Correction of Heart Images 119
7.1 Introduction . . . 120
7.2 Proposed Framework . . . 122
7.2.1 Automatic Segmentation . . . 122
7.2.2 Methodology . . . 126
7.3 Experiments . . . 130
7.3.1 Data . . . 132
7.3.2 Experiment 1: synthetic edgemaps . . . 132
7.3.3 Experiment 2: Atlas template . . . 134
7.3.4 Experiment 3: Real images . . . 136
7.3.5 Experiment 4: Fully automated pipeline . . . 137
7.4 Results . . . 137
7.4.1 Automatic segmentation . . . 137
7.4.2 Experiment 1 . . . 140

7.4.3 Experiment 2 141
7.4.4 Experiment 3 143
7.4.5 Experiment 4 144
7.5 Discussion 147

8 Conclusion and Future Work **156**

Appendices

A Cardiac Surface Meshes **161**
A.1 Simulation of CMR Acquisitions 162
A.2 Application of SurFCo on alternate objects 163

B ISACHI Visual Results **164**
B.1 Experiment 4 164

Bibliography **172**

List of Figures

1.1 Thesis Pipeline 4

2.1 Illustration of the cardiovascular circulation system 10
2.2 Illustration of the cardiac cycle 14
2.3 Different types of cardiac pathologies affecting the morphology of the heart 19
2.4 Illustration of the ultrasound process 21
2.5 Bohr model of the atom 22
2.6 Proton precessing about B_0 24
2.7 Cardiac MRI acquisition terminology 27
2.8 Cardiac MRI views in 2D 28
2.9 Different MRI protocols 30

3.1 Registration example between two images 36
3.2 Non-Rigid transformation examples 40
3.3 Instantaneous phase and amplitude of a signal 44
3.4 The importance of phase in an image 46
3.5 Illustration of the log-Gabor and Riesz filters 47
3.6 Local phase and local amplitude of a cardiac image 48
3.7 Example of multi-scale versus Multi-resolution 50
3.8 The fifteen possible configurations for the 3D marching cube algorithm. 53
3.9 Loop subdivision scheme 55
3.10 Butterfly subdivision scheme 56
3.11 Illustration of various subdivision schemes 57

4.1 Illustration of mesh reconstruction from misaligned slices 61

5.1 MRI slices with their respective inter-sectional Line intensities . . . 79
5.2 Local Phase applied to a late gadolinium image 81
5.3 Local phase effects of varying the wavelength and sigma parameters on a cardiac image. 82
5.4 Labeled-based image alignment on an LGE dataset 84
5.5 Pairwise Energy between slices over several iterations 86

5.6 Binary-based image alignment on an LGE dataset 87
5.7 Motion correction results for an LGE dataset using different strategies. 88
5.8 Example of a segmentation error between two intersecting slices. . . 89
5.9 Surface Reconstruction before and after alignment 90

6.1 Mean statistical shape model of the left ventricle epicardium and right ventricle endocardium fitted to real cardiac contours 96
6.2 SurFCo pipeline 100
6.5 Error measurements between generated mesh and SA contours . . . 105
6.3 Illustration of resulting meshes for different extreme sets of parameters. 108
6.4 Parallel Coordinate graph for the parameter evaluation of the SurFCo method 109
6.6 Error distance between generated mesh and SA contours for an increasing number of SA contours. 109
6.7 Error distance between generated mesh and LA contours for an increasing number of LA contours. 110
6.8 Distance maps between non coincidental, non parallel contours and generated right and left ventricular meshes. 111
6.9 Segmented CMR acquisition 111
6.10 Resulting left and right endo- and epicardium surface meshes for normal and abnormal anatomies. 112
6.11 Contour-to-mesh distances using entire endocardium contour sets compared with contour exclusion. 112
6.12 Leave one out contour to mesh Distances for epi- and endocardial contours for 24 cases. 113

7.1 Training labels for the U-Net. 124
7.3 Automatic Cropping Pipeline. 124
7.2 U-Net CNN architecture. 125
7.4 Representation of ISACHI definitions. 127
7.5 Example of ISACHI edgemaps 130
7.6 Synthetic edgemap obtained by a Gaussian radial basis function . . 133
7.7 Initial data for ISACHI experiment 1. 134
7.8 Isachi second experiment overview 136
7.9 Example CNN segmentation result on the HCM dataset 137
7.10 Example CNN segmentation result on the myocardial infarction dataset 138
7.11 CNN performance validation on HCM data 139
7.12 CNN performance validation on the myocardial infarction dataset . 140
7.13 Visual Result for ISACHI Experiment 2 141

7.14 Atlas alignment validation after ISACHI using the line intensities as a surrogate measure 143
7.15 Visual Result for ISACHI Experiment 2 144
7.16 Example of an abnormal case for Experiment 3 146
7.17 Visual Result for ISACHI Experiment 3 case 1 147
7.18 Example of the CNN segmentation result on an abnormal case for Experiment 4 148
7.19 Visual Result for ISACHI Experiment 4 case 1 149
7.20 Resulting surface meshes before and after ISACHI using expert and automatic contouring 150

A.1 Resulting output meshes for various input shapes 162
A.2 Application of SurFCo on the MIT Bunny 163
A.3 SurFCo application on a pair of lungs 163

B.1 ISACHI result for Experiment 4 on the first normal case. 166
B.2 Meshes generated before and after applying ISACHI on the expert segmented and CNN contours for the first normal case. 166
B.3 ISACHI result for Experiment 4 on the second normal case. 167
B.4 Meshes generated before and after applying ISACHI on the expert segmented and CNN contours for the second normal case. 167
B.5 ISACHI result for Experiment 4 on the third normal case. 168
B.6 Meshes generated before and after applying ISACHI on the expert segmented and CNN contours for the third normal case. 168
B.7 ISACHI result for Experiment 4 on the first abnormal case. 169
B.8 Meshes generated before and after applying ISACHI on the expert segmented and CNN contours for the first abnormal case. 169
B.9 ISACHI result for Experiment 4 on the second abnormal case. . . . 170
B.10 Meshes generated before and after applying ISACHI on the expert segmented and CNN contours for the second abnormal case. 170
B.11 ISACHI result for Experiment 4 on the third abnormal case. 171
B.12 Meshes generated before and after applying ISACHI on the expert segmented and CNN contours for the third abnormal case. 171

List of Abbreviations

1D, 2D, 3D . . One-, two-, or three Dimensional.

4D . . . . . Four dimensional (3D + Time).

CMR . . . . . Cardiovascular Magnetic Resonance Imaging.

MRI . . . . . Magnetic Resonance Imaging.

SAX . . . . . A Short Axis cardiac slice.

LAX . . . . . A Long Axis cardiac slice.

2-4CH . . . . . Two-, three-, four Chambers.

EF . . . . . Ejection Fraction.

SV . . . . . Stroke Volume.

CO . . . . . Cardiac Output.

HR . . . . . Heart Rate.

ED . . . . . End Diastole.

ES . . . . . End Systole.

MI . . . . . Myocardial Infarction.

CAD . . . . . Coronary Artery Disease.

LGE . . . . . Late Gadolinium Enhancement.

ECG . . . . . Electrocardiogram.

HCM . . . . . Hypertrophic Cardiomyopathy.

DCM . . . . . Dilated Cardiomyopathy.

RCM . . . . . Restricted Cardiomyopathy.

DTI . . . . . Diffusion Tensor Imaging.

CT . . . . . Computed Tomography.

US . . . . . Ultrasound.

NMV . . . . . Net Magnetic Vector.

FID . . . . . Free Induction Decay.

HLA . . . . . Horizontal Long Axis.

VLA . . . . . Vertical Long Axis.

LVOT . . . . . Left Ventricle Outflow Track.

RF . . . . . Radio Frequency.

TR . . . . . Repetition Time.

TE . . . . . Time Echo.

SSFP . . . . . Steady State Free Precession.

ROI . . . . . Region of Interest.

DoF . . . . . Degrees of Freedom.

FFDs . . . . . Free Form Deformations.

SAD . . . . . Sum of Absolute Differences.

SSD . . . . . Sum of Squared Differences.

NCC . . . . . Normalized Cross-Correllation.

NMI . . . . . Normalized Mutual Information.

ADI . . . . . Absolute Difference of Intensities.

LP . . . . . Local Phase.

LG . . . . . Log-Gabor.

TPS . . . . . Thin Plate Splines.

SSM . . . . . Statistical Shape Model.

FCN . . . . . Fully Connected Network.

CNN . . . . . Convolutional Neural Network.

BPA . . . . . Ball pivoting Algorithm.

MCF . . . . . Mean Curvature Flow.

RBF . . . . . Radial Basis Functions.

GM . . . . . Global Motion.

ICP . . . . . Iterative Closest Point.

1 Introduction

Contents

1.1 Motivation **1**
1.2 Thesis Contributions **4**
1.3 Thesis Overview **6**

1.1 Motivation

Recently, medical imaging research has seen an increase in the precision medicine field [1–7]. Various initiatives [6, 8] have been introduced to allow for healthcare diagnosis and treatment to be realized on an individual basis. To allow for better visualization, diagnosis, and understanding, medical imaging has seen an increase in 3D applications. The use of 3D geometrical models are increasing in a number of areas, including dentistry, surgical planning, patient risk stratification and evaluation, and device innovation (for virtual and augmented reality purposes) [1]. Recently, 3D models combined with virtual reality have been used in surgical planning to help clinicians anticipate anatomical structures [7]. To quantitatively assess cardiac functions, 3D reconstructions are progressively being used in the precision medicine field [6–8]. Such models allow for global and local 3D analysis, such as volumetric shape analysis, multi-modality fusion, as well as simulations of

biological processes and functions. Depending on the application, either surface or volumetric mesh models can be used. For example, to perform biophysical simulations, a volumetric mesh defining the underlying geometry of the problem is usually required. Computational models for simulations are becoming more complex in order to better understand cardiovascular functions, as well as predict different processes after a clinical event.

In a personalized setting, for example, going from cardiac imaging to the geometric definition of patient-specific anatomies, many steps must be taken to achieve accurate geometric representations. Imaging, segmentation, and meshing challenges make the generation of anatomically realistic volumetric meshes a complex task. In recent years, each of these respective fields has seen an incredible rise in novel techniques.

Although global mortality from cardiovascular diseases is decreasing, morbidity is rising, making cardiovascular diseases one of the most common healthcare disorders [9, 10]. Furthermore, age-specific cardiovascular related deaths are decreasing, signifying a broadening population effect [9]. Enabling early diagnosis and effective treatment is essential to the reduction of the pathological and economical burdens of cardiovascular diseases.

In the past 15 years, several long-term studies have been initiated investigating pathology developments and predispositions such as the UK Biobank study or the Cardiac Atlas Project. These studies have allowed for a significant amount of data to be collected from a wide range of participants (currently $\sim 27,000$ scanned individuals with an aim to image $100,000$ individuals for the UK Biobank imaging study and ~ 3000 individuals for the Cardiac Atlas Project [11, 12]). Using these datasets, population studies can be performed assessing various structural, functional, or genetic factors. Population studies of cardiac anatomy require precise geometric reconstructions as cardiac shape and function are important pathology predictors. In particular myocardial infarctions, characteristically lead to cardiac remodeling in both the left and right ventricles, where the heart undergoes cellular changes, adapting its structural and functional capacities. In order to be able to

capture the different structural or functional changes in the population, accurate personalized models must be generated, allowing for better diagnosis and treatment of the individuals, as well as for better population models.

In cardiac studies, magnetic resonance imaging (MRI) is a befitting imaging tool for its ability to capture, non-invasively and in a non-ionizing manner, the current state as well as 3D changes in bi-ventricular geometry and function. Whilst cardiac computer tomography (CT) images provide high anatomical resolution, segmentations obtained from CMR images can provide additional information, in an non-ionizing manner, about myocardial structure using techniques such as late Gadolinium enhancement or Diffusion Tensor MRI, as well as Cine and tagging MRI.

To capture cardiac structural states as well as perform any biophysical simulations for individual or population studies, an accurate geometrical 3D model must be generated. Typical MRI acquisitions acquire cardiac slices independently, which can cause a loss of 3D consistency between the acquisitions. As such, in order to capture the anatomy, the slices need to be arranged in a spatially consistent manner. These inconsistencies between images are common due to various artifacts including patient motion inside the scanner, breath-hold variability and inherent cardiac motion.

The main aim of this Thesis is to develop tools and methodologies to correct for spatial misalignment in cardiovascular images towards an accurate generation of 3D geometrical meshes representing the anatomy of patients in a personalized manner. Furthermore, in order to be used for high quantities of data from population studies, such as the UK Biobank or the Cardiac Atlas Project, the methodologies presented in this Thesis were developed and tuned for heart like shapes with the intention to be fully automatic. Figure 1.1 represents the overall pipeline considered in this thesis. From typical clinical acquisitions, an initial method to correct for any motion misalignment between slices is introduced in Chapter 5 in a fully automatic way. From the spatially coherent dataset, segmentations are used to generate 3D geometrical meshes using the methodology described in Chapter 6. Whilst these two techniques can be sequentially applied automatically, Chapter 7 combines both approaches in an integrated manner. The personalized generated surface meshes

can be used as inputs for several applications. Population studies can be performed, assessing the structural variability through shape analysis. The surface meshes can also be converted to volumetric meshes, and serve as inputs for biophysical and biophysiological simulations. If temporal data is used, the pipeline can be used on each of the frames, allowing for the evaluation of personalized motion analysis.

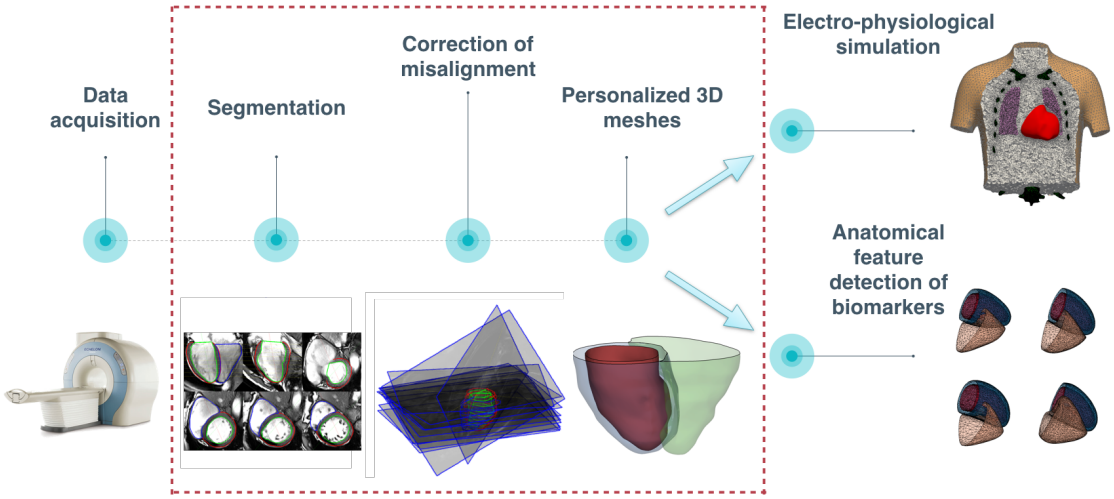

Figure 1.1: Illustration of the pipeline considered in this Thesis. Clinical images are obtained through MR, after which any misalignment is corrected for, before generating 3D geometrical surface meshes. These meshes can be used for biophysical simulations, or population based studies. The dashed box represents the elements of focus in this Thesis.

.

1.2 Thesis Contributions

Three main contributions can be found in this thesis.

The first concerns the correction of misaligned CMR images in 3D by only using image information. This problem can be solved by either using segmentations only, registering the 2D images to a 2D or 3D atlas, registering the 2D images to a low resolution volumetric scan, or by optimizing over the line intersections at the cross sections of spatial data. The latter method, which is used in this Thesis, is the most applicable to the clinic as it can be used on typical acquisitions, without the need for any extra protocols. Using this methodology, the first contribution introduces the use of local phase as a cost function allowing for the use of feature information in the images. While most techniques use Cine acquisitions, the developed method

is applied on late gadolinium enhanced images which focus on highlighting infarcted tissue via the use of contrast, making the use of local phase applicable. Furthermore, a validation methodology is presented which shows the lowest quantitative bounds possible using contour-to-contour distances as a validation measurement.

The second contribution relates to 3D surface reconstruction from cardiac contours, the implementation of which is publicly available at `www.github.com/BenVillard/surfco` and named SurFCo (Surface From Contours). The methodology is presented which enables the creation of surface meshes from planar delineations. Most segmentations are performed independently in 2D, and when placed in 3D, usually result in being non-coincidental, which is exacerbated by any remaining misalignment. Moreover, after the application of 3D motion correction algorithms, the slices loose their parallelism. In order to build a surface mesh from such data, a method is presented that is able to deal with sparse, heterogeneous, non-parallel, cross-sectional, and non-coincidental contours. In an iterative process, an initial mesh is deformed through a composition of smooth deformations, in a diffeomorphic manner, maximizing the data fitting aspect, whilst rendering visually pleasing surfaces meshes. The method developed can be used on left and right ventricle contours, as well as any other shapes given a set of contours. An evaluation methodology is also presented to assess the algorithm's ability to fit the data whilst interpolating sparse areas in the most accurate way. A multi-parameter study is performed to tune the parameter for heart like shapes.

Finally, a third contribution is presented, which combines the two previous works, in order to integrate segmentation, alignment corrections of heart images, and surface mesh generation in one approach, which we introduce as an algorithm called ISACHI (Integrated Segmentation and Alignment Correction for Heart Images). The method integrates an automatic segmentation via a convolutional neural network, to render the whole pipeline fully automatic. By combining geometric and image information, accurate 3D geometrical meshes can be obtained. At the intersection of the images and the mesh, patches are generated from an initial surface mesh. The edgemap information from images, obtained via the local phase, drives the

patches towards areas of high feature information in agreement between the two intersecting planes, aligning the slices, while deforming the contours simultaneously to resolve discrepancies. The contours deform fitting the data more appropriately, correcting for any segmentation error, and becoming spatially consistent. Similarly to the previous contributions, an evaluation methodology is presented, assessing the method at different levels.

The work in this Thesis represents a signifiant improvement over current techniques used in clinics and in the broader scientific community. The methodologies developed as part of this Thesis are fully automatic and tuned for the heart. With the advent of large cohort studies and the need for large cohorts, used for machine learning techniques, having automatic methods is highly beneficial. In particular, the methods developed can be used both in a clinical setting, as well as being incorporated directly in clinical studies such as the UK Biobank.

1.3 Thesis Overview

A brief outline of the different chapters in this thesis is given below:

Chapter 2 provides an introduction to the clinical aspects of cardiovascular imaging. Cardiac physiology and pathophysiology are covered, as well as the different imaging techniques used in cardiac imaging.

Chapter 3 reviews the technical tools used along this Thesis. Notably, these tools are commonly employed in the area of image registration and geometrical mesh processing.

Chapter 4 gives an overview of the current literature. The state-of-the-art research is presented for the correction of cardiac slice misalignment, ranging from geometric to image processing approaches. Surface reconstruction methodologies, and their application to cardiac data are discussed. Finally novel research into the integration of slice misalignment correction and surface reconstruction is presented.

Chapter 5 presents a novel methodology for the correction of slice misalignment, that makes use of a combination of intensity information as well as feature information extracted via local phase. Furthermore, spatial inconsistencies from cardiac segmentations are quantified.

Chapter 6 introduces a new surface reconstruction methodology that is able to work with non-parallel, non-orthogonal, non-spatially consistent, sparse sets of heterogeneous contours. The method is tuned for heart-like shapes and validated against synthetic and clinical data.

Chapter 7 comprises the final research work, incorporating the methodologies developed in the previous chapters. A novel methodology integrating segmentation correction for spatially inconsistent contours as well the correction of slice misalignment is presented. The method is again validated against synthetic and clinical data.

Chapter 8 summarizes the work found in this thesis and discusses its applicability to clinical environments. Some limitations of the methods presented are discussed, as are potential developments which could be implemented in the future to further improve performance.

2

Clinical Background

Contents

2.1 Cardiac Physiology and Pathophysiology 8
2.1.1 Cardiac Physiology 8
2.1.2 Cardiac Pathophysiology 16
2.2 Cardiac Tomographic Images 19
2.2.1 Ultrasound 20
2.2.2 Computed Tomography 21
2.2.3 Magnetic Resonance Imaging 22

This chapter provides a description of cardiac anatomy, physiology and pathophysiology. An account of the heart's remodelling process is also given. Finally an overview of cardiac imaging modalities is provided.

2.1 Cardiac Physiology and Pathophysiology

2.1.1 Cardiac Physiology

2.1.1.1 Basic Cardiac Anatomy and Morphology

The heart is a continuously moving organ, located in the thorax, and is bordered by the lungs, sternum and diaphragm. It has an oblique position within the torso, with its principle axis, usually called the long axis (LA), lying from the right shoulder to the left hip. The top of the heart, where the valves are located is called the *base*, and

the bottom is called the *apex*. It is rotated around 45° in the sagittal and coronal planes [13]. The heart is enclosed inside the *pericardium*, a fibrous sac that contains it within the chest cavity, preventing its over expansion and limiting its movement, whilst providing low friction surroundings enabling easier contractions. The heart is at the center of the cardiovascular system, which has for primary function to deliver, in a rapid and efficient manner, nutrients (such as oxygen, glucose, amino acids) and hormones all over the body, via the arteries. In an adult, its length typically ranges from 11-13 cm long, 8-9 cm wide and 6 cm in thickness, and its weight ranges roughly between 250-300 grams [14, 15]. Heart sizes between healthy individuals can vary depending, among other factors, on individual's exercise aptitude and weight. The heart is made up of four different chamber; the *left atrium* (LAU), the *right atrium* (RAU), the *left ventricle* (LV), and the *right ventricle* (RV). The left and right ventricles act as intermittent pumps. The right side of the heart receives de-oxygenated blood in the RAU via the superior and inferior *vena cava*, and pumps it to the lungs via the pulmonary artery, where it is freshly oxygenated. The LAU of the heart acquires the newly oxygenated blood, and the LV pumps it back to the rest of the body via the artery. In order to regulate the amount of blood coming into and exiting the different chambers, four valves exist. The *triscupid valve* and the *mitral valve* allow for the flow to come into the RV and LV, respectively. The *pulmonary valve* and the *aortic valve* control the blood flowing out into the lungs, and to the rest of the body, respectively. The heart also receives its own dedicated blood supply via the coronary arteries, which form a tree-like structure and are located on top of the right and left ventricles. These vessels ensure the proper delivery of nutrients to the heart, allowing it to fulfill its role. The coronary circulation system acquires its blood supply from the *ascending aorta*, right above the *aortic valve*, where they branch off to surround the two ventricles. An illustration of the cardiovascular system can be seen in Fig. 2.1.

The left and right ventricles are separated by the interatrial and interventricular septum, keeping the blood from merging. These ventricles, as well as the septum, are all part of the heart wall and are made up from different parts, known as the

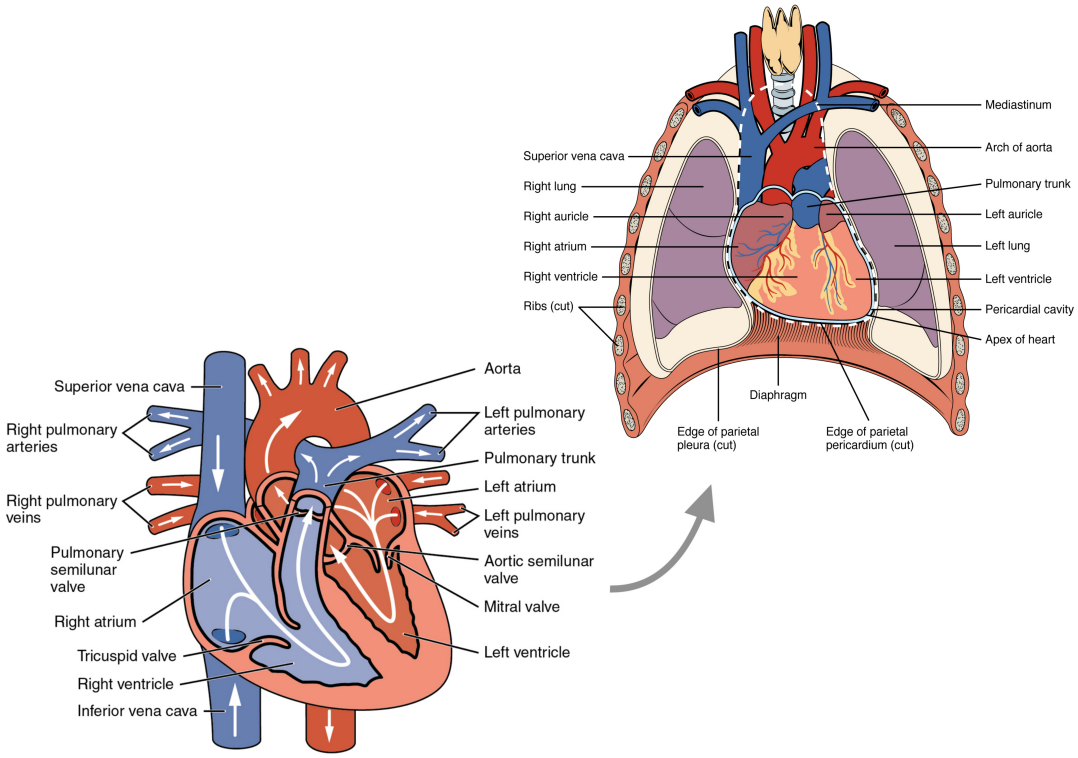

Figure 2.1: Illustration of the cardiovascular circulation system. De-oxygenated blood enters into the right atrium of the heart via the inferior and superior vena cava. The blood flows to the right ventricle, going through the triscupid valce, where it will be pumped out to the right and left lungs, where it is freshly oxygenated. The re-oxygenated blood flows back to the heart, into the left atrium. It then enters the left ventricle, passing through the mitral valve, where it will be pumped out to the rest of the body. Figure altered from [14].

.

epicardium, myocardium and endocardium. The heart wall is made up of connective tissue, endothelial cells, and cardiac muscle. The muscles allow for the beating of the heart, and are arranged in three different, uniquely structured, layers. The outer layer of the heart wall, the epicardium, is primarily made up of connective tissue (elastic and adipose tissue). Its main function is to act as the outer protective layer of the heart. It is also on this layer that the coronary arteries lie. The myocardium is the middle layer, and is composed of cardiac muscle cells, also known as cardiomyocytes. These cells allow the heart to contract, and thus enact its function. Myocardium thickness varies throughout the heart depending on the need. It is typically thickest in the left ventricle, as this muscle needs to pump the blood at a high pressure so that it may flow into the aorta and circulate throughout the

body. The endocardium is the smoothest and thinnest layer, and is constituted by endothelial cells and elastic fibers. It protects the heart valves, and acts as a barrier between the blood pool and the cardiac myocytes. Inside the ventricles can be found papillary muscles as well as endocardial trabeculations. Papillary muscles regulate ventricular valves preventing blood regurgitation and can also help in their contraction. They vary in morphology, and in numbers between the ventricles. The right ventricle possesses three papillary muscles, called the anterior, posterior and septal muscles. The left ventricle contains two papillary muscles, the posterior and anterior muscles. Their electrophysiologically or mechanically induced contractions and relaxations enable the closing and opening of the valves. They can be significant in their involvement in pathological cardiac remodelling. Endocardial trabeculations also play a role in preventing blood regurgitation. The right ventricle tends to have more trabeculations than its ventricular counterpart. Furthermore the right ventricle contains a band of cardiac muscles known as the moderator band. This band plays a significant role in the heart's contraction mechanism as it acts as an electrical conduction pathway from the septum to the right endocardial wall.

The left ventricle's appearance can be generally described as having an oval or prolate ellipsoidal shape. The main structural changes typically occur at the base and apex, whereas the mid-sections transition smoothly in shape. The right ventricle's shape tends to be depicted as pyramidal. With exercise, especially aerobic, the heart undergoes stress induced structural changes. Coronary width increases as well as ventricular myocyte mass, causing the heart to increase in size.

2.1.1.2 The Cardiac Cycle

A two-phase system, generated by electrical signals along the heart, allows the heart to pump out blood around the body. A first phase, *diastole*, occurs when the heart's ventricles fill up with blood and the heart is in a relaxed state. Conversely, the second phase, *systole*, takes place when the heart contracts, pumping blood out. Theses phases can be subdivided into two stages; isovolumetric contraction

and ejection phases for systole and isovolumetric relaxation and atrial contraction phases for diastole. An overview of the cardiac cycle process is given in Fig. 2.2.

Diastole involves the relaxation of the ventricles during which they are filled with blood, a process which lasts around 0.5 s in a normal resting heart. End diastolic volume (EDV) represents the total amount of blood filling the ventricles at end-diastole, and corresponds to around 120 ml. Diastole encompasses the time between the opening of the mitral valve to its closure. During isovolumetric relaxation, ventricular blood pressure falls until it is below the atrial pressure causing the atrioventricular valves to open, refiling the ventricles. Ventricular relaxation can be prolonged in abnormal behavior such as myocardial infarction [13]. The filling process starts via the contraction of the atria. The atrioventricular valves close allowing for the systolic phase to take place.

As a contrast to diastole, systolic contraction is initiated by electrical stimulation, whilst the mitral valve is closed, leading to ejection when ventricular pressure is higher then the adjacent vessels. The duration of isovolumetric contraction lasts around 0.5 s in a normal resting heart. The stored pressure causes the pulmonary and aortic valves to open expelling the blood from the ventricles into the pulmonary and systemic system. End systolic volume (ESV) is considered the point at which the maximum amount of blood has been ejected, and corresponds to around 50 ml. The ejected volume is roughly estimated to be 70-80 ml. This contraction system allows for a ventricular ejection fraction (EF), representing a measure of pumping efficiency, of around 65-70% [13]. This significantly differs for patients with heart problems who tend to have lower ejection fractions, and is also a strong predictor of arrhythmic death [16].

Heart contractions are triggered by electric impulses stimulating cardiac cells. The heart has the ability to generate its own electrical pulses, spreading from cell to cell, which in turn create a mechanical response allowing for muscular contraction. The heart's electrical system is not unlike the brain. Two types of myocardial cells exist: myocardial conducting cells and myocardial contractile cells, which represent 99% and 1% of myocardial cells respectively. Myocardial

conducting cells act similarly to neurons as they allow for the conduction of action potentials throughout the heart. Once the electrical impulses have been generated, the myocardial contractile cells create cardiac contractions. A dense group of myocardial conducting cells located at superior and posterior walls of the right atrium establish the cardiac cycle by generating self induced electric impulses. This group of cells form the *sinoatrial node*, also known as the pacemaker of the heart. The sinoatrial node starts the cardiac cycle by initiating an electrical impulse that makes its way down the heart, passing in a first instance through the atria. This causes both atria to contract, pumping the blood to the ventricles. The impulse then spreads to the atrioventricular node, taking around 50-100 ms to arrive. This time period ensures that the atria have enough time to eject the blood to the ventricle before these start contracting. Afterwards, the electrical impulse travels towards the ventricles through fiber bundles, called atrioventricular bundles. The bundles located in the septum go down towards the apex, feeding into the Purkinje fibers, which split up to go up the left and right myocardial walls. The bundles also split up to go to the right ventricle via the moderator band. The advance of these impulses activates the ventricles to contract, pumping out blood to the rest of the body.

2.1.1.3 Indices of Cardiac Function

Several clinical indices have been developed to evaluate various cardiac functions. One of the most used methods is to evaluate cardiac output (CO), via the computation of ejection fraction, or stroke volume (SV). Depending on the purpose, other metrics can also be evaluated such as wall thickness/thinning. SV relates to the the total output blood flow during a single ventricular contraction, and is given by $SV = EDV - ESV$ (in ml), where EDV and ESV represent the end-diastolic and end-systolic volumes respectively. SV calculations allow to obtain a measure of cardiac output, describing the volume of blood being pumped out of the heart as a function of time, by combining SV with heart rate (HR) such that, $CO = SV \times HR$ (in ml/min). A related index, EF measures ventricular pumping efficiency by computing the fraction of blood output from the ventricle

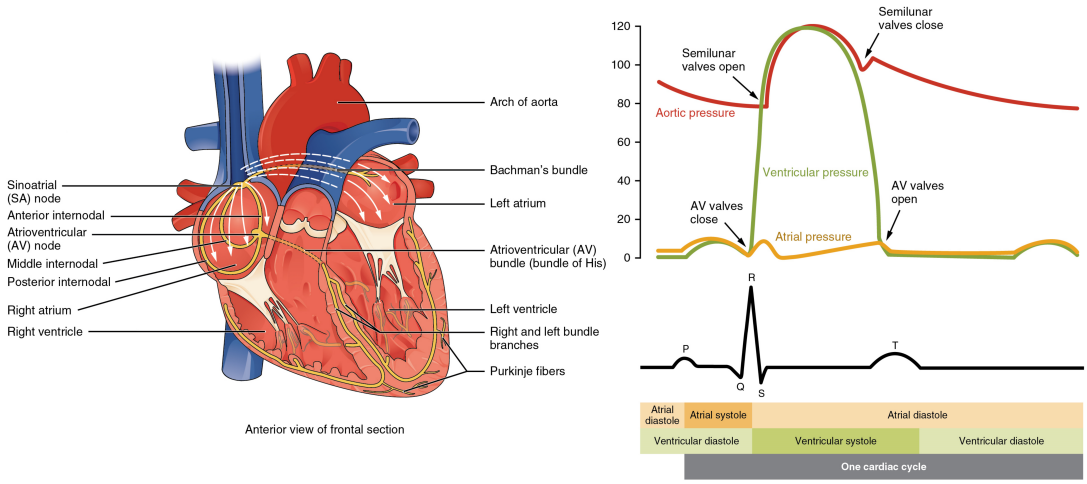

Figure 2.2: Illustration of the cardiac cycle. The electrical impulses of the heart are illustrated in the top left. The top right figure represents the pressure variations during the cardiac cycle, along with the behaviour of the valves. The lower right image represents an electrocardiogram tracing. The **P** wave represents the depolarization of the atria, initiating cardiac contraction. The **QRS** complex represents the depolarization of the ventricles, and the repolarization of the atria, with the R peak signalling ventricular contraction. The **T** wave depicts the repolarization of the ventricles. The associated stages of the cardiac cycle, are linked to the ECG tracing. Figure adapted from [14].

.

during contraction and is given by $EF = \frac{EDV-ESV}{EDV} \times 100 = \frac{SV}{EDV} \times 100$ (in %). EF is clinically used as a primary indicator of heart failure. Normal values range from $60 - 70\%$, which accounts for a $10 - 15\%$ myofiber shortening. As these metrics rely on volumetric calculations, if no 3D volumetric data is available, clinicians currently use methods as Simpson's rule [17] to estimate volumetric assessment from 2D images. Teichholz method [17] used to be the major method of volumetric approximation using 2D area information, however due to shortening and shrinking of the heart in the LA, a better approximation was found using Simpson's rule. Simpson's rule is a numerical integration method, used in ejection fraction calculation to estimate volume measurements by taking into account the short axis (SA) inner diameter for all the slices. Individual 2D slice areas are calculated via segmentation, and multiplied by the slice thickness and gap between slices. This measurement is then summed for all the slices, resulting in an estimated volumetric measurement. These techniques however represent an approximation of the true volumetric measurements, which impacts ejection fraction

quantification. Furthermore, multiplying the area by the slice thickness gives a skewed representation of the volume due to partial volume effects. New methods to compute EF and subsequent metrics are being proposed which make use of 3D reconstructions, motivating the need for accurate 3D surface models [18].

Myocardial mass is another measurement of particular importance, more specifically in pathologies associated with morphological change, such as left ventricular hypertrophy. Similarly to the computation of myocardial volume, myocardial mass is computed by summing the myocardial area times the slice thickness multiplied by the average density of the myocardium (1.05 g/mL), for all slices, where myocardial mass is given by the difference in epicardial and endocardial areas [13, 19]. Typically, myocardial mass tends to be computed solely on the left ventricle as accurate quantification of right ventricular mass can be more complex, due to the increased amount of endocardial trabeculations in the RV [13]. Furthermore, the RV epicardium - endocardium boundary is to thin to be able to distinguish on CMR. To allow for robust and more reliable computation of myocardial mass measurements, the aforementioned premise of using volumetric approaches over geometric assumptions holds.

2.1.1.4 Cardiac Motion

Two main types of motion can occur in the heart. The main kind occurs during the heart's contraction, where it undergoes inherent intrinsic motion. The second type is caused by extrinsic factors, notably other organs, and in particular the lungs during inhalation and exhalation. A study by McLeish *et al* [20] reports the various different types of motion experienced by the heart. Regarding intrinsic motion, non-rigid deformations occur most, as expected, in the left ventricle and right atrium. McLeish *et al* report deformations ranging from 3-4 mm, with some cases ranging up to 7 mm. Extrinsic motion caused by breathing displace the heart in a rigid motion, which occurs primarily within the sagittal direction.

The respiratory cycle, in a normal adult at rest, is estimated around 6 s, during which the heart beats around 6 times. During the respiratory cycle, cardiac motion

is the most restricted during end-expiration, making it the preferred choice for cardiac acquisitions, although it generally is harder to maintain over long periods. During inhalation, the lungs shifts the heart in the craniocaudal direction on average by 12.4 ± 5.9 mm, although [20] reports displacements up to 23.5 mm. Breathing also displaces the heart caudally, where the heart is reported to displace around 4.3 ± 3.7 mm and 2 ± 2.1 mm, in the anterior-posterior and right-left directions respectively. Small rotations are also reported. A full table describing the heart's motion and rotation for patients and normal cohort can be found in [20]. Similar results have also been reported in [21].

During cardiac pump activity, the principal modes of motion occur in the SA plane, as well as longitudinally, causing the heart to shrink as well as twist. Motion is also pertinent in coronary arteries. Coronary artery motion has been shown to be highly variable during the cardiac cycle. Right coronary arteries have been found to move more then their left counterparts, with the most significant motion found in the inferior-superior axis [13, 21].

2.1.2 Cardiac Pathophysiology

As described in the previous section, the heart is a complex and intricate system. For the proper functioning of cardiac activity, different systems must work in unison at various levels. As such, problems stemming from any level render cardiac heart problems common, with myocardial infarctions being the most prevalent cause of death world wide [10, 19]. General cardiac pathologies can be divided into two types; ischemic heart diseases, and muscular diseases.

2.1.2.1 Ischemic Heart diseases

Ischemic heart disease concern the inability of the heart to receive adequate oxygen supply via the blood - a condition characterized as hypoxia. This typically can happen due to coronary atherosclerosis, which involves the thinning of the coronary arteries due to deposits of plaques in the vessels. This reduces the blood circulating to the heart, lowering the incidence of nutrients, which ultimately lead to muscular

infarct, a process known as *myocardial infarction* (MI). Myocardial Infarction, commonly known as heart attacks, is the most severe manifestation of prolonged hypoxia. Due to the lack of oxygen supply, myocardial tissue can become necrotic, preventing the proper functioning of cardiac contractions. Depending on the hypoxic duration, irreversible damage to the tissue can occur, allowing for scar formations.

Myocardial infarction can also have some muscular effects, such as ventricular free wall rupture - usually associated with acute cases of MI. This rupture, although infrequent, can lead to cardiac tamponade or cause a tear between the two ventricles, allowing the passing of blood between them. Cardiac pathology symptoms usually present themselves in form of Angina - chest pain or pressure. This pain is typically associative with coronary artery disease (CAD), which including *atherosclerosis* and *coronary artery embolisms*. The onset of MI can be caused, as mentioned, by the build up of plaque, however it can also occur more acutely by a sudden blockage of the coronary arteries. This can potentially happen if part of the plaque, made up of cholesterol or fatty acids, breaks free and lodges itself in one of the arteries.

Whether the blockage of the coronary arteries takes place suddenly or over time can be distinctly seen on electrocardiogram (ECG) profiles. The most severe case, where complete occlusion of the supply arteries takes place, manifests itself by having the ST segment of the PQRST complex (seen in Fig. 2.2) elevated - usually referred to as STEMI. This is due to tissue necrosis, which causes changes in the polarization of the electrical impulses, causing the ST segment to be elevated on an ECG. Controversially, non ST elevated myocardial infarct (NSTEMI), represents the case were only partial occlusion of the arteries take place.

2.1.2.2 Cardiac myocardial diseases

Cardiac muscular diseases, otherwise know as cardiomyopathies, represents the other most common pathologies. A range of cardiomyopathies exist such as, *hypertrophic, dilated*, or *restrictive cardiomyopathy*, as well as *tachicardia-induced cardiomyopathy*. Most cardiomyopathies are idiopathic, and tend to be genetic or acquired, through cardiac remodelling, for example.

Hypertrophic cardiomyopathy (HCM), is a genetic condition in which the ventricular walls thicken, due to cell growth. In very rare cases, HCM can be acquired. HCM has a distinct impact on the shape and geometry of the heart, as well as cardiac function. Stroke volume tends to be affected, and an increase risk of arrhythmia - irregular heart beats - is common. However, HCM is generally asymptomatic and does not manifest itself unless through some form of exercise.

Cardiac hypertrophy is very similar to HCM, in that cardiac cell size and shape increase as a response to physical training. This cardiac remodelling process can be seen in athletes, in particular ones that perform endurance training or aerobic exercises.

Dilated cardiomyopathy (DCM), is the most common type of cardiomyopathy world wide, and relates to tissue remodelling in an adverse manner. The global shape of the heart typically takes a balloon-like form, whilst the ventricular wall thickness tends to be normal. Due to this enlargement, heart functionality is usually less efficient. Impaired contraction of the ventricles causes reduced cardiac output, and although most persons subject to DCM are asymptomatic, DCM can lead to heart failure.

The least common type of cardiomyopathy, restricted cardiomyopathy (RCM), represents an increased stiffness of the myocardium, impacting cardiac output. This stiffness impairs ventricular filling, causing diastolic dysfunction, and can ultimately cause cardiac failure. As a response to debilitated ventricular function, the atria become disproportionally enlarged causing further dysfunction.

2.1.2.3 Cardiac Remodelling

From the previous section, it is apparent that mechanisms are in place allowing for cardiac remodelling. Cardiac remodelling, or more specifically, ventricular remodelling, is the process in which the heart undergoes morphological and physiological changes. Remodelling can happen as a result of pathological or physiological (exercise) conditions. Cardiac remodelling can be beneficial to the heart, such as is the case with athletes, through ventricular thickening, allowing increased blood

flow to the body. In most cases, however, it causes adverse effects, and is generally accepted as ultimately causing heart failure [22]. The changes associated with cardiac remodelling are still not fully understood. In the case of MI, remodelling can take place from the onset stage, within a few hours, and continue for several months, altering cardiac functionality and shape, as well as generating cellular and molecular changes. Typically, the effect of the remodelling process is roughly proportional to the severity of the infarct.

There has been an increase in research, towards understanding cardiac remodelling in a more detailed manner. Due to the many sub-level changes, multi-compartmental models need to be used. Using different imaging modalities allow for such incorporations, such as obtaining 3D shape via MRI or CT, fiber information via diffusion tensor imaging (DTI), or cellular information via histology [22,23]. Furthermore longitudinal studies can be performed, grasping this remodelling process.

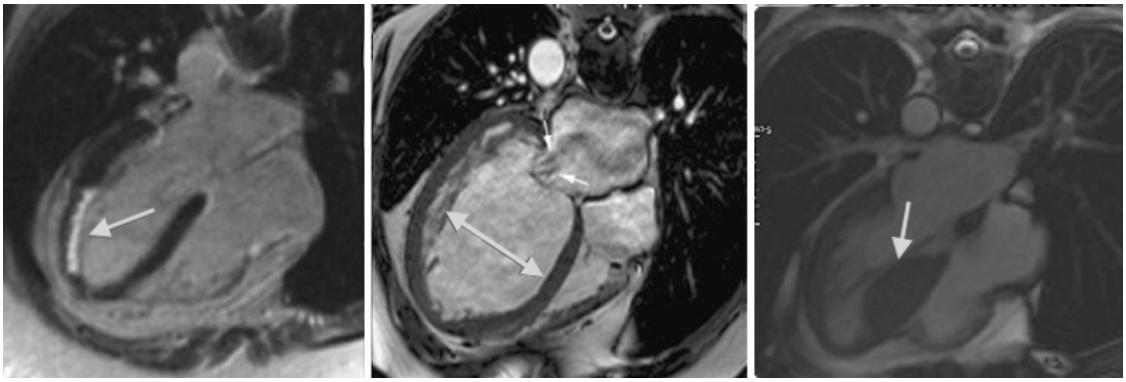

Figure 2.3: Images representing different cardiac pathologies affecting the morphology of the heart. a) Myocardial infarction: the white area in the myocardium represents the infarcted region. b) Dilated cardiomyopathy: the left ventricle takes on a *balloon*-like shape. c) Hypertrophic cardiomyopathy: A thickening of the left ventricle myocardial wall can be observed.

2.2 Cardiac Tomographic Images

The advent of medical imaging technologies has been instrumental in the diagnosis and analysis of numerous pathologies. Nowadays, most imaging tools allow for digital acquisitions, enabling the development of new image processing techniques to tackle the unique issues inherent to each modality. In broad terms, different modalities

can be classified into two categories; ionizing such as computed tomography (CT), and non-ionizing, such as MRI or Ultrasound (US). Due to the nature of the data obtained through our clinical collaborators, this Thesis focuses on MRI acquisitions. This section gives an overview of cardiac MRI, together with a brief introduction to other medical imaging modalities used in cardiovascular medicine.

2.2.1 Ultrasound

Just like the name implies, ultrasound (US), is an imaging method that is based on the reflection of pressure waves. A transducer, is placed on the surface of an object of interest, and allows to visualize the internal structures of that object. Different types of US are available, capturing 1D to 4D data. Generally, the transducer emits an ultrasound wave by converting an electric signal into an acoustic signal, which travels through matter (such as tissue or fat). At surfaces where the density of the matter changes, part of the sound waves gets reflected back to the transducer. The transducer converts the acoustic signal back to an electric signal, which can then be represented as a line within an image. The time between the start of the wave to the reception of the wave (*echo time*) allows the calculation of the distance to the reflecting surface. Due to different levels of reflection of the acoustic signal in different organs, the reflected signals have varying levels of amplitude. These amplitudes are converted to reflect pixel intensities in the US images. Multiple sound waves are produced, typically forming a fan pattern, allowing the formation of an image by composition of the corresponding lines. An illustration can be seen in Fig. 2.4.

Ultrasound has many advantages with respect to other imaging modalities, notably it's accessibility and affordability, as well as being non-ionizing. It also tends to be the initial modality choice for diagnosing heart diseases. However US suffers from poor quality images, with high amounts of noise, compared to CT or MRI. Although 3D probes exist, standard clinical analysis is still performed using a 2D probe and thus depends on the ability of the operator to select the optimal imaging planes.

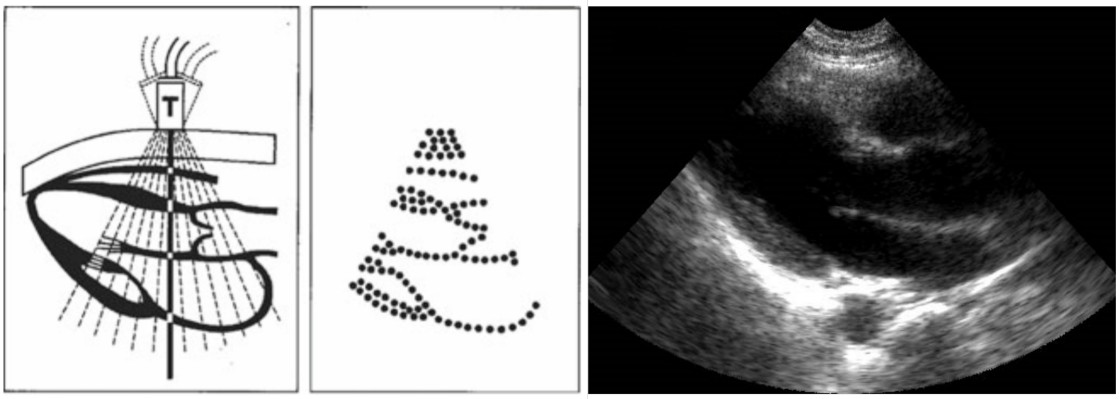

Figure 2.4: Illustration of a cardiac ultrasound scan. Right: diagram representing the transducer scanning the heart. The center black line represents the echo time delineating the various surfaces. Left: the resulting image where the structures can be identified. Figure obtained from [24].

.

2.2.2 Computed Tomography

Computed tomography is an imaging modality based on X-rays that allows capturing cross-sectional images. Typically, a CT machine is composed of an X-Ray tube from which the X-Rays emanate, and an array of detectors facing it. The detectors along with the X-Ray generator are able to rotate fully, allowing a complete cross sectional panoramic acquisition of an object. The generator is composed of an electron beam, which shoot electrons to a respective target material, such as tungsten or molybdenum, generating X-Ray photons. When these photons are generated, they pass trough tissue, interacting with electron shells (see Fig. 2.5). Depending on the interaction, the photons either get absorbed or scatter. The absorption coefficient corresponds to different structures, which reflects into densities in the detector which are used to generate images intensities. As the CT scanner rotates, the absorption coefficient is able to be mapped in 3D. In such a way, after a reconstruction process, a volumetric image is obtained, showing distinct values related to tissue properties.

CT is generally used, in cardiac imaging, for the characterization of coronary arteries. Cardiac CT usually requires the use of intravenous contrast agent to highlight the particular structures of the heart, which also comes with some risks to the patient's health. Furthermore, breath-hold related issues are also a problem,

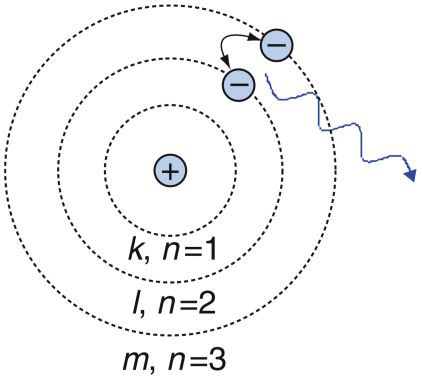

Figure 2.5: Bohr model of the atom. The rings represent the electron shells (given by k), where the amount of electrons per shell is finite (given by n). When an external electron interacts with the atom, there is a change in the characteristics, either through an *jump* between shells, or through interaction with the nucleus. Figure obtained from [24].

potentially causing slice misalignment in the volume.

2.2.3 Magnetic Resonance Imaging

Magnetic Resonance Imaging's ability to generate a wide range of different image types make it one of the gold standard imaging tools for numerous applications. By using electromagnetic and quantum properties to excite nuclei within the tissue, images can be acquired. In the cardiac field, MRI allows for the static and dynamic imaging of the heart, as well as provides functional and structural information. Depending on the type of information wanted, different MRI protocols can be used. Next sections will introduce the basics of MRI, as well as the main techniques used in medical imaging for cardiac applications.

2.2.3.1 Basic Principles

MRI is based on the electromagnetic and quantum properties of atoms, and more specifically hydrogen atoms, although others can be used (such as $^{13}\mathbf{C}$ or $^{23}\mathbf{Na}$). Hydrogen atoms' nuclei are composed of a single proton positively charged. If an external electromagnetic field (B_0) is applied, the protons will align with regards to the applied field. Depending on their energy states, the protons can precess either parallel (low energy protons), or anti-parallel (high energy protons) to B_0. This precession describes the protons spinning on their own axis, which themselves rotate

about B_0, as shown in Fig. 2.6*a*, similar to a gyroscope. The total magnetic force generated by the individual protons cancel out due to their opposing directions, however, a slight excess in low energy protons occurs, proportional to the field strength. This resulting discrepant magnetic field, which points in the direction of B_0, is called the net magnetic vector (NMV) - M_0. The rate of excess protons, and thus the magnitude of NMV follows a Boltzmann distribution, which is a function of temperature and number of spins. It is this distribution that governs the magnetic strength of MRI. MRI with higher magnetic field strength increase the magnitude of NMV, allowing for higher resolution images.

The frequency at which the proton spins about B_0 follows the Larmor equation $\omega = \gamma B_0$, where γ is called the gyromagnetic ratio and ω is the precessional frequency. Thus, ω and NMV are inherently controlled by B_0. Whilst the protons precess about each other, the NMV, M_0, does not. When the protons are in the range of B_0, and no other forces act upon them, M_0 is aligned with B_0 in a steady state. At this point, no signal is measurable. In order to generate an MR signal, a second electromagnetic force is applied, as a radio frequency pulse (RF) - B_1. This new electromagnetic field can manipulate the proton spin direction as per the Larmor theorem. If B_0 is points towards the Z axis, then B_1 is transmitted in the transverse direction ($x - y$ plane). There is an inherent relationship between B_1 and ω, such that if the frequency of B_1 equals ω, a phenomenon known as resonance occurs, causing the protons to start precessing around B_1. As B_1 lies in the transverse plane of B_0, the protons will transition to the $x - y$ plane (shown in Fig. 2.6*b*). As they precess around B_1, M_0 also starts to precess in the transverse plane, forming the basis of the MRI signal. This change in magnetic fields generates electromagnetic radiation which can be received by receiver coils. The extent at which the NMV decreases in the Z axis and increases in the $x - y$ plane is described as the *flip-angle*. Once the RF pulse has been transmitted, the protons begin their recovery to steady state, modifying the NMV. This recovery process occurs due to the dephasing of the previously coherent x-y magnetization, in a process known as *Free Induction Decay* (FID), ultimately leading to the image-generating MR signals. This recovery

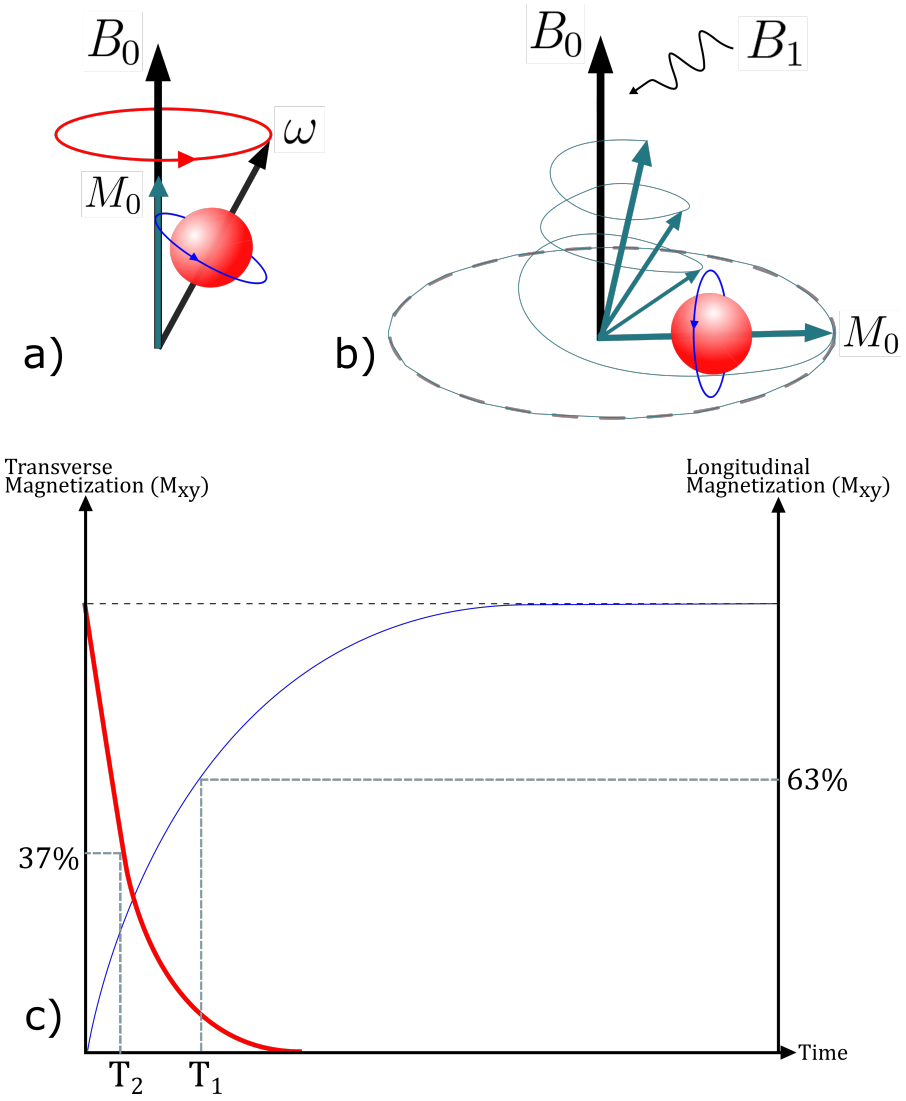

Figure 2.6: a) Illustration of the precession of a proton precessing about an external magnetic field B_0. b) Illustration of the proton precessing to the transverse plane due to the transmission of an RF pulse B_1. The net magnetization vector M_0 aligns with the transverse plane. c) Illustration of the the T1 recovery and T2 decay functions. The T1 signal (in blue) corresponds to the time it takes for the protons to re-align with the external electromagnetic field B_0. The T2 signal (in red) represents the de-phasing time of the protons.

process is tissue dependent due to the distribution of atoms in different tissues. Two main processes, known as *spin-lattice* and *spin-spin*, guide the protons back to steady state. In spin-lattice, the recovery process is given by the time it takes for M_0 to re-align with B_0. This process occurs at an exponential rate, and is dependent on a tissue specific constant **T1**. For example, fat has a short $T1$ recovery, whilst muscle has a longer one. This process also known as longitudinal relaxation (shown

in Fig. 2.6*c*), is given by the equation $M_z(t) = M_0(1 - e^{-\frac{t}{T1}})$, where M_z is the NMV lined up with B_0. When $t = T1$, the NMV M_0, will have regained 63% of its original value. The second process, also known as transverse relaxation, concerns the dephasing of individual proton, in the transverse plane. When B_1 is applied, all the protons spin in phase, however, they begin to dephase shortly after leading to a decrease in M_0. The individual protons start to fan about the transverse plane due to the electromagnetic interactions between themselves. The rate at which the protons dephase, reducing M_0, is given by a constant **T2**, also tissue specific. The following phenomenon follows an exponential decay and is given by the equation $M_{xy}(t) = M_0 e^{-\frac{t}{T2}}$, where M_{xy} is the NMV in the transverse plane. Similarly to $T1$, when $t = T2$, the transverse NMV will have decayed to 37% of it's original value.

By manipulating different RF pulses, different types of information can be acquired, giving distinct information about tissue. Up to here the FID signal represents a sum of all the individual signals generated by the protons. To allow for spatial localization of the signals, three additional gradient fields are applied. If a patient's primary axis lying down is thought of as the Z direction, by applying a gradient varying from negative to positive, spatial information can be obtained as the Larmor frequency of the protons respond differently, in accordance to the gradient field. At the zero crossing, the Larmor frequency ω would equal the specific frequency of the interaction with B_0. By applying a similar approach in the orthogonal directions, specific spatial localization can be obtained [25]. In practice, it must be noted that the application of the gradients results in specific slice thickness as perfect localization is not achievable. This introduces the concept of partial volume effect as the signal obtained represents an average of the protons' excitatory response in that region, in accordance to the slice thickness and pixel size. Furthermore, the magnetic fields are not uniform, causing field inhomogeneities which translate into image artifacts. These can be more or less pronounced as seen in Fig. 2.9*d*. The images are formed through the phase and amplitude of the signal's frequency generated by the excited nuclei. These frequencies are stored in *k-space*, and are then converted to images via the Fourier-transform.

In summary, an MR signal is dependent on the magnitude of proton density, relaxation time, and flip-angle.

2.2.3.2 Cardiac MR Imaging Acquisition

Cardiac MRI has evolved as a major player in the wide variety of cardiac imaging tools, allowing the acquisition of functional, morphological and metabolic information [13]. Current cardiac CMR clinical protocols require patient preparations before the imaging can take place. To compensate for the inherent motion of the heart, the images have to be guided using ECG, in a technique called ECG gating. Due to the physical shifts of the heart in the body due to breathing, patients are instructed to hold their breath at a consistent lung volume between acquisitions. Furthermore, depending on the imaging protocol, contrast might be required.

In a typical clinical acquisition, patients are placed in the MRI scanner lying down in the supine position, with cross-sectional images of the heart taken for a duration of 30-45 minutes. Initial low resolution, fast acquisition time, scout images are acquired in the coronal, sagittal and transversal planes (as shown in Fig. 2.7c), to assess cardiac location. These scout images allow for the cardiac planes to be established after which ECG-gated sequences of long axis slices (LAX) are acquired showing the 2 chamber (CH), 3CH and 4CH of the heart, as can be seen in Fig. 2.8. Once acquired, a stack of short axis (SAX) slices are obtained, typically ranging from the base of the heart, where the cardiac valves lie - in between the atria and the ventricles - to the apex, where the tip of the heart lies. The 2CH view corresponds to the vertical long axis (VLA), 3CH view to the left outflow track long axis (LOVT), and the 4CH to the horizontal long axis (HLA). It should be noted that the oblique orientation of the heart causes the cardiac acquisition planes to be acquired at an angle to the three main acquisition planes. An example can be seen in Fig. 2.7d).

ECG Gating

To compensate for the dynamic motion of the heart, ECG gating (also known as triggering) is typically used, synchronizing the MRI acquisition with the heart contractions. Two different techniques are used: prospective and retrospective

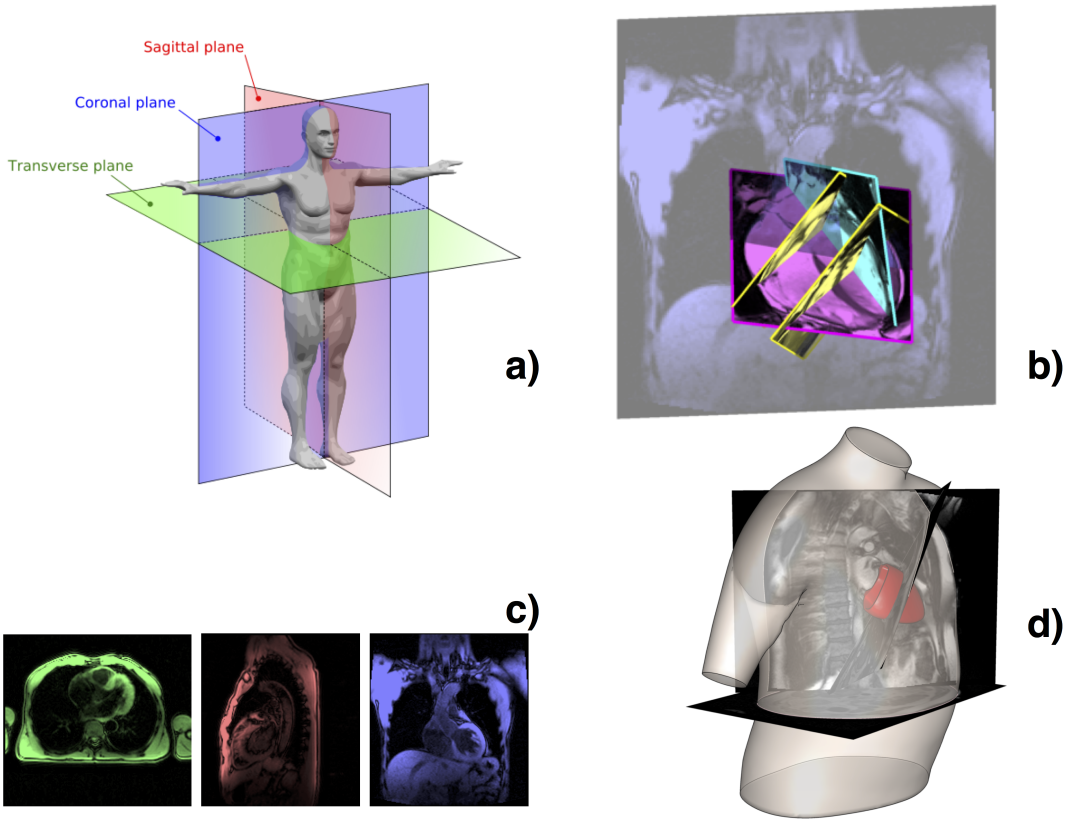

Figure 2.7: Figure describing the different cardiac MRI acquisition terminologies. a) Illustration showing the sagittal, coronal, and transverse planes of a body. b) Figure showing a 3D cardiac stack acquisition in their respective spatial locations. The 3 long axes can be seen intersecting with the stack of short axes. c) The sagittal, coronal, and transverse localisers. d) Illustration representing the location of the heart in the torso. The transverse and sagittal plane image acquisitions can be seen as well as a short axis acquisition. The specific orientation of the SA plane should be noted.

gating. In prospective gating, the R peak of the QRS complex, coinciding to ventricular systole, (shown in Fig. 2.2) is used as a reference point for the acquisition. Different predetermined time delays are used to acquire the heart at various cardiac contraction points during the R-R interval. This interval is related to heart beat (in beats per minute (bpm)) and is typically around 1000-800 ms long (60 - 75 bpm). As such, a specific triggering time must be calculated to account for individual patients' heartbeats, and a waiting time between triggers must be calculated to account for any heartbeat variability. Because of this, prospective gating doesn't permit the acquisition of the full cardiac cycle but around 90%. Cardiac acquisitions are then obtained from the data collected over multiple cardiac cycles. On the

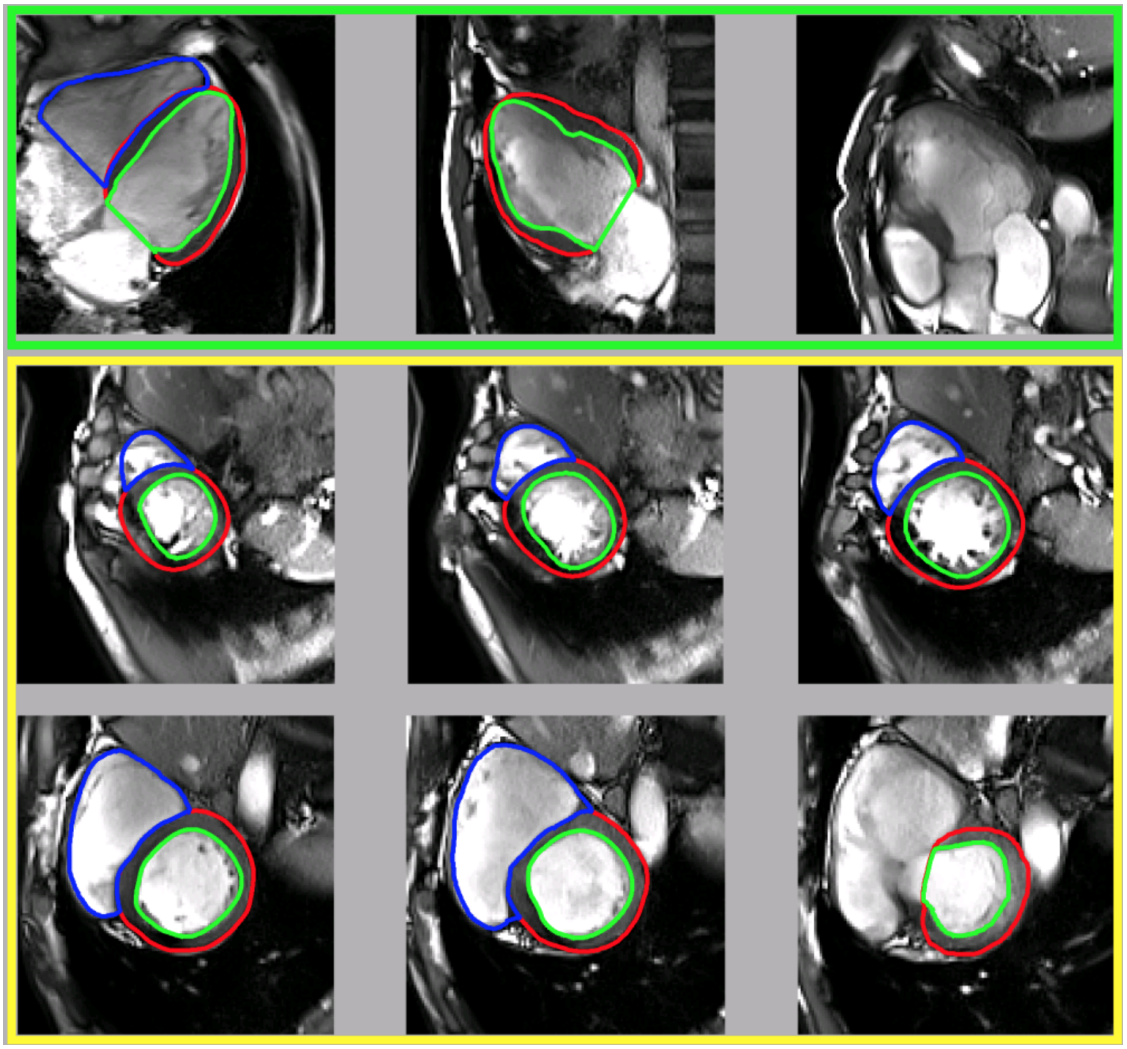

Figure 2.8: Figure showing the different 2D cardiac images acquired at different planes. The first three images, highlighted in the green box represent the 4CH, 2CH, and 3CH, respectively. The rest of the images highlighted in the yellow box represent the short axes, ranging from the apex (2nd row, left), to the base (last row, right). The segmentation of the images represent the different morphological structures of the heart, with the epicardium in red, the endocardium in green, and the right ventricle in blue.

other hand, in retrospective gating, MR scans are continuously acquired and the data is then rearranged to match corresponding cycles.

Breath-Hold Gating

Due to the location of the heart within the torso, the lungs have a significant effect on its spatial location, potentially displacing it several centimeters. To allow for consistent spatial position when acquiring the slices, subjects must perform

multiple breath-holds for periods of 7-15 seconds. Generally images are acquired during end-inhalation or end-expiration breath hold. End-expiration has been found to obtain higher reproducibility between different views [13]. However it is harder to maintain than end-exhalation for longer periods of time, even more so for pathological subjects. Much like ECG-gating, breath hold belts, also known as bellows, exist to compensate for breath hold variability. These belts are fitted to the chest of the individuals and used to monitor the diaphragm motions. The images are then acquired with respect to both the ECG and the breathing bellow. However, they are not typically used for cardiac MRI [26].

2.2.3.3 Limitations of CMR

While image acquisition using CMR allows for various advantages such as soft tissue discrimination or slice acquisition in any spatial plane, CMR acquisitions come with some drawbacks, one of the most relevant to this Thesis being the acquisition time, which brings the necessity, for patients, of cardiac and potentially respiratory gating. Furthermore, the quality of the acquisition is directly related to scanning time, which causes any prolonged scan time to increase motion artifacts and blurring [27]. CMR images do not reflect an absolute value of the signal as the gray values in the reconstructed image depend on a various factors such as the magnitude of the complex signal, coil strengths, and magnetic field intensity among others. Spatial resolution is intrinsically linked to the signal decay and *k-space* sampling. In order to increase the spatial resolution, more samples are needed to be acquired in *k-space*, which will also result in a smaller partial volume effect and thickness slice, however, these come at the cost of longer the acquisition time. Clinical acquisitions have to deal with the trade off between increased resolution and scanning time. CMRI acquisition times have been optimized to maximize any information used for clinical diagnosis as well as minimize scan time. This in turn limits some image data properties as clinical cardiac datasets involve sparse slice acquisition of the heart, causing missing information in in- and out-of-plane resolutions due to the large gap between slices and coarse out of plane resolution.

2.2.3.4 Cardiac MRI Protocols

The manipulation of the MRI parameters such as the RF pulses as well as the MRI coil gradients, can lead to different types of acquisitions enabling the extraction of a wide variety of information. This section describes some of the most used MRI protocols and their typical applications.

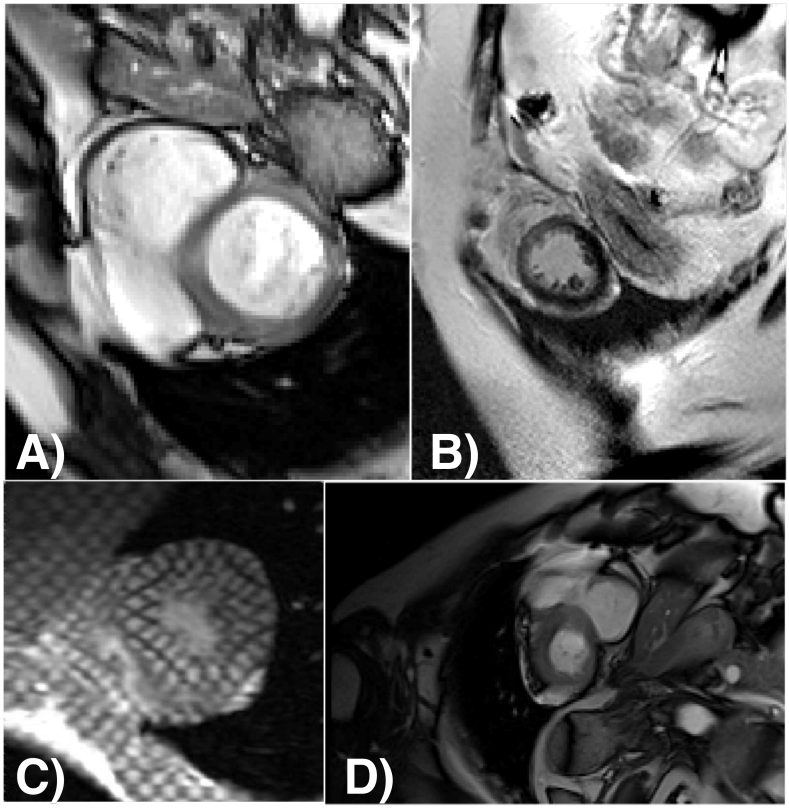

Figure 2.9: Figure representing different cardiac MRI protocols and the effect of field inhomogeneity. a) Cine MRI acquisition of SAX. b) Late gadolinium enhanced MRI acquisition of SAX. c) Tagged MRI of a SAX during systole. d) Example of a Field inhomogeneity caused by the MRI scanner. The left hand side of the image is darker than the right.

T1/T2 weighted

As explained in Section 2.2.3.1, the time constants T1 and T2 are dependent on the structure of each tissue. By applying successive RF pulses, separated by a time interval called *Repetition Time* (TR), and measuring the magnitude of the FID signal after a specific time - know as *echo time* (TE) - an image can be obtained. T1 weighted images have short TR and TE times, allowing to highlight tissue with

fast T1 properties such as fat. T1 weighted images are most commonly used to visualize anatomical details. In contrast to T1 weighted images, T2 weighted images have comparably longer TR and TE times. Water and fluid appear brighter in T2 weighted images, making this sequence optimal for the applications such as the localization of oedema after myocardial infarction. During an acute infarction, a build up of water can ensue in the myocardium. This oedema has been linked with myocardial viability, making T2 imaging an integral part of the imaging protocol in acute myocardial infarctions.

Cine MRI

One of the most common cardiac MRI acquisition protocols is Cinematographic (Cine) MRI. This technique allows for the dynamic acquisition of the heart. High signal intensity of blood flow can be acquired permitting the study of cardiac dynamics, and enabling excellent anatomical details [13]. Cine MRI consists of creating a compilation of acquisitions encompassing a whole cardiac cycle, producing a *movie* of the beating heart. Several scans are taken over time, generating different frames. On average 25 frames are taken making up the whole cardiac cycle. This process is achieved through a pulse sequence known as steady state free precession (SSFP) due to its high acquisition speed, and high contrast to noise and signal to noise ratio. The excellent contrast between myocardium and blood pool further adds to the anatomical information acquired. Cine imaging is predominantly acquired using retrospective ECG gating, allowing for the entire acquisition of the cardiac cycle. A typical Cine acquisition will result in a dataset having around 1.5 mm spatial resolution in the short axis, and around 40 ms temporal resolution. Cine MRI generally results in the acquisition of 2-3 LAX and 8-14 SAX slices, as per Fig. 2.8. Each SAX slices usually consists of an 8 mm slice thickness separated from the other SAX by a 2 mm gap. Having the ability to visualize heart motion is particularly useful for the identification of wall motion irregularities. Furthermore, the acquisitions allow for excellent blood-myocardial differentiation, enhancing cardiac chamber structures.

Tagged MRI

To enable the assessment of cardiac mechanical function, a technique known as *tagged MRI*, allows the placement of grid-like structures on the images. By applying different RF pulses in conjunction with gradients, the protons will experience saturated and equilibrium magnetization spins states at specific locations, producing grid like structures on the images. As the applied magnetization is inherent to specific tissue regions, any cardiac deformation is also applied to the grid permitting the evaluation of motion inside the myocardium. An example can be seen in Fig. 2.9*c*. There exist certain challenges to this technique however, notably its practical restriction to 2D applications, only allowing in plane quantification of deformation. In the clinics, tagged MRI is typically used on SAX, and does not make use of LAX information [28]. Furthermore, changes in the magnetization states of the proton deteriorate over time, causing a loss in the grid-like information on the images.

Gadolinium Enhancement MRI

The use of a contrast agent is an integral part in assessing myocardial viability. This procedure involves injecting a contrast agent into the patient, typically containing gadolinium, and acquiring images straight after the injection, or after a delayed period [29]. Whilst imaging shortly after injection results in the visualization of blood perfusion, delayed imaging allow the diagnosis and quantification of myocardial infarct. Non infarcted regions appear dark, whilst necrosed or fibrotic tissue arise with high intensity (shown in Fig. 2.9b). This is caused by the retention of the contrast in areas of myocardial infarct which experiences a slower wash out rate of the contrast agent. Late gadolinium enhanced (LGE) imaging for cardiac studies have become a standard clinical practice for the assessment of acute and chronic assessment of tissue infarction. LGE images are typically generated using T1 imaging.

3
Technical Background

Contents

3.1 Image Registration Tools 33
3.1.1 Transformation Models 35
3.1.2 Similarity measures 41
3.1.3 Multilevel image registration 49
3.1.4 Limits of CMRI in image registration 49
3.2 Mesh Reconstruction 51
3.2.1 Common surface construction algorithms 51
3.2.2 Geometric tools 53

This Chapter introduces the technical elements found in this Thesis. An overview of image registration tools and surface reconstruction methods is given as well as validation techniques.

3.1 Image Registration Tools

An image can be defined as a spatial distribution of intensities in 2D (known as pixels). A volumetric image is described as an image in a 3D space where each cell (known as voxels) contains the intensity value at that point. This Thesis is focused on medical images, which are usually volumetric images, where each intensity value in the image is represented by its spatial coordinates. Furthermore,

intensity values of medical images typically have an inherent meaning, representing a certain physical property of a biological tissue.

The task of image registration lies in aligning two images of interest which contains landmarks, or points that are mutual to both of the images. The act of registering the two images allows to match these homologous points as to aggregate the information from both images.

Registration is a key tool in recovering the optimal geometric transformation between objects of interest. It is usually posed as an optimization process which aims to either maximize or minimize a certain cost function between the two objects by transforming one of the objects towards the other, given a specific transformation model. The objects are said to be registered when their information, be it their corresponding points, shapes, or visual appearance, are spatially aligned with each other. In such a manner, the common information from both images becomes coincidental, whilst the differing information becomes more apparent. With images, the registration process aims to match the pixels, representative of tissue characteristics, with the corresponding tissue information in the other image. This is achieved by transforming one of the images (often referred to as moving) to match the other image (often referred to as fixed or target) . The most common types of objects used in the registration process are:

- Landmarks (otherwise known as fiducial points): Landmarks are points which represent distinct features in an image, which can be annotated manually or automatically.

- Regions of interest (ROI): Regions of interest are typically biological areas which are delineated (known as segmentation) due to representing a significant area of interest such as a anatomical structures. These segmentations are typically constituted of curves in 2D or 3D with closed boundaries.

- Images: Images are discrete representations of scalar functions with spatially attributed scalar values.

- Significant points: Significant points are point clouds describing a geometrical entity. They are spatially unstructured and can be sparse or dense.

Image registration is used for many different applications, notably to extract the differential information between two similar objects of interest [30, 31]. It is used extensively in the domain of medical imaging, but its application can also be found in face or finger print recognition, remote sensing and video enhancement, among others [32].

Let $\mathcal{I}f$ and $\mathcal{I}m$ be fixed and moving images, or geometrical objects, respectively. The goal of image registration is to find the optimum transformation that maximizes or minimizes an energy function $E : \mathbb{R}^n \mapsto \mathbb{R}$ between the two objects and can be formulated as:

$$\mathcal{R}^* = \underset{R}{\operatorname{argmin}}\, E(\mathcal{I}m^{\mathcal{R}}, \mathcal{I}f)\,, \tag{3.1.0.1}$$

where $\mathcal{R}$ represents the parameters of the transformation model, E the energy function between the objects being registered and $\mathcal{I}m^{\mathcal{R}}$ the transformed image. The energy function E is also known as a *cost function*, an *objective functional*, a *fitness function*, or a *loss function*. $\mathcal{R}^*$ represents the optimum parameters that maximizes or minimizes E. The result of the registration is highly dependent on the choice of the transformation model and energy function, thus an appropriate choice must be considered for each registration problem.

3.1.1 Transformation Models

There are several transformation models which are mainly characterized by their degrees of freedom (DoF). They are generally categorized into two main groups. The first incorporate rigid transformations, including translations and rotations, and affine transformations, including shearing and scaling. The second group represent non-affine, or deformable transformations.

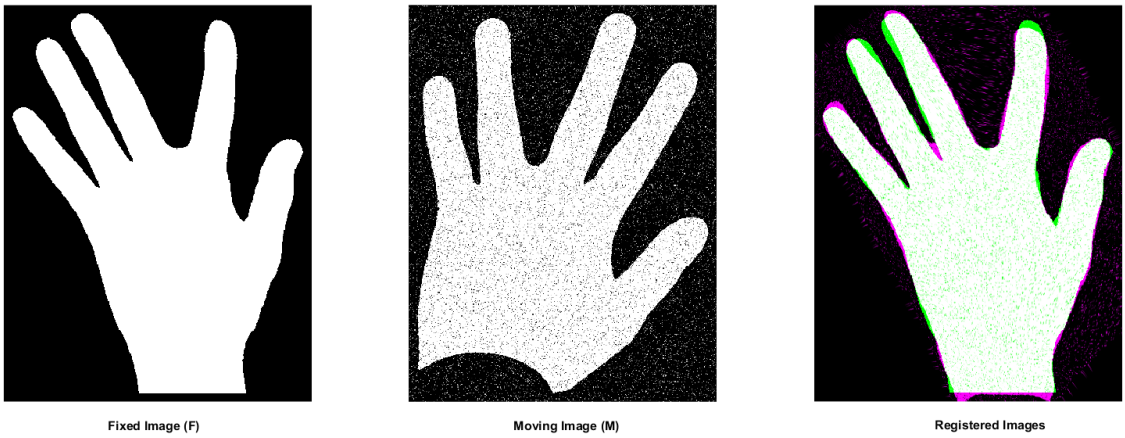

Figure 3.1: Representation of a registration process. The moving image $\mathcal{I}m$, (left) is registered to the Fixed image $\mathcal{I}f$ (center) by applying the inverse of a non-rigid deformation. The final result can be seen (right) where we can see the fingers are matched with each other. The areas where the matching has failed can be seen in purple (representing $\mathcal{I}m$) or green (representing $\mathcal{I}f$).

.

3.1.1.1 Rigid and Affine

Rigid and affine transformations, in the case of image registration, assume that no distorted deformation has occurred between two objects. These transformation models are of particular relevance for this Thesis, as most of the work presented relies on intra-subject data, for which there exists a rigid-body assumption. For example, in the case of correcting for cardiac misalignment, there exists an underlying assumption that the different slices of the heart have been taken at the "same" time point, and can be seen as a "snapshot" of the heart at a specific instance in time. A rigid transformation can be parameterized by six DoFs (in 3D): $\{\alpha, \beta, \gamma, t_x, t_y, t_z\}$, where α, β, and γ represent the rotations about the X, Y, and Z, axis respectively, and t_x, t_y, and t_z represent translations in X, Y, and Z respectively. A rigid transformation can be expressed in homogenous coordinates in matrix form:

$$\mathcal{T}_{rigid}(x, y, z, 1) = \begin{pmatrix} & & & t_x \\ & \mathrm{R} & & t_y \\ & & & t_z \\ 0 & 0 & 0 & 1 \end{pmatrix}, \tag{3.1.1.1}$$

where $\mathbf{R}$ is a matrix encompassing the rotations with respect to each of the axes. The rotation matrix about the X,Y, and Z axes, are respectively defined as:

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos\alpha & \sin\alpha \\ 0 & -\sin\alpha & \cos\alpha \end{pmatrix}, \quad \begin{pmatrix} \cos\beta & 0 & -\sin\beta \\ 0 & 1 & 0 \\ \sin\beta & 0 & \cos\beta \end{pmatrix}, \quad \text{and} \quad \begin{pmatrix} \cos\gamma & -\sin\gamma & 0 \\ \sin\gamma & \cos\gamma & 0 \\ 0 & 0 & 1 \end{pmatrix}. \tag{3.1.1.2}$$

The matrix $\mathbf{R}$ (Eq. (3.1.1.1)) , represents a combination of the three rotation matrices. However, as the product of matrices are not commutative, the order in which the rotation matrices are applied does matter. Furthermore, the center about which the rotation is applied will affect the resulting transformation. All the transformations applied throughout this Thesis are performed, unless otherwise mentioned, about the center of the object undergoing the transformation. This is generally computed as the *centroid* of the object, although other geometric centers are available.

As rotations are not commutative, in some special cases, a specific set of rotation associations will render a rotation about a certain axis arbitrary - a problem known as Gimbal lock. In order to render rotations commutative and unambiguous, they can be parametrized as quaternions, or more specifically, unit quaternions. In the medical imaging field, quaternions are extensively used in image-guided therapy, or tracking systems [25]. Quaternions are useful when dealing with considerable rotations, however, the applications involving rigid transformations found in this Thesis deal with small rotations and as such, the use of Euler angles is justified. In the case of affine transformations, scaling and shearing parameters are also added. In this case, up to twelve DoFs are available.

3.1.1.2 Non-Rigid Transformations

Non-rigid transformations usually involve much higher orders of DoFs, than rigid or affine models and can be classified as parametric or non-parametric. Non-parametric models typically use all the voxel coordinates for which a dense displacement field is stored, such as with Optical-flow, or Demons [33, 34]. In parametric models, the deformation is modeled as a function governed by a set of parameters, typically

substantially less then the amount of voxels. Due to the applications found in this Thesis, only parametric transformations based on spline models will be discussed.

3.1.1.3 Splines

Splines are one of the most widely used tools in image registration [24]. They allow for interpolation or approximation through differentiable numerical models. As described earlier in this Chapter, registration concerns the matching of landmarks, or scalar values, based on the assumption that they are identical in both $\mathcal{I}f$ and $\mathcal{I}m$. Splines provide a way to interpolate displacements over the entire domain of registration (such as the image plane or geometrical landmarks in space), whilst allowing for the matching of these landmarks. Generally, Splines make use of the landmarks as control points of the spline function [24]. There are different types of spline models, of which the most commonly used in image registration are *thin-plate splines* (TPS) and *B-Splines*.

Thin-Plate Splines

Thin-plate splines are a common type of splines that originated from an analogy to the shape of thin metal plates deformed by normal forces at discrete points. They are seen as a 2D generalization of the cubic splines, and are based on radial basis functions in the Euclidean space [24,35–37]. TPS provide a non-rigid transformation field that maps corresponding points in $\mathcal{I}f$ and $\mathcal{I}m$. In its regularized form, TPS can include affine transformations for special cases. TPS are used throughout this Thesis, notably in Chapters 6 and 7. Given a set of source and target points, TPS minimizes an energy functional $\mathcal{J}$ (the bending energy), also known as the biharmonic equation, by a weighted combination of radial basis functions at every point. In 2D, it has the form

$$f(x,y) = a_1 + a_x x + a_y y + \sum_{i=1}^{p} w_i U(\|(x_i,y_i) - (x,y)\|)\,, \tag{3.1.1.3}$$

with the aim of minimizing the energy

$$\mathcal{J}_f = \int\int_{\mathbb{R}^2} (f_{xx}^2 + 2f_{xy}^2 + f_{yy}^2) dxdy\,, \tag{3.1.1.4}$$

where the radial basis function $U = \frac{1}{8\pi} r^2 \log r$ in the 2D case and $U = -\frac{1}{8\pi} r$ in the 3D case. w_i and a_i are unknown coefficients, with w_i relating to the TPS parameters and a_i relating to affine parameters (within the null space of $\mathcal{J}_f$), and p is the number of points. The value r represents the Euclidean distance $r = \sqrt{x^2 + y^2}$. As each individual point is taken into account, the TPS problem can thus be solved as a linear system of equations having the form $Ax = b$:

$$\left[\begin{array}{c|c} K & P \\ \hline P^T & O \end{array}\right] \begin{bmatrix} w \\ a \end{bmatrix} = \begin{bmatrix} F \\ o \end{bmatrix} , \tag{3.1.1.5}$$

where $K_{ij} = U(\left\| \mathcal{M}_{xy_i} - \mathcal{M}_{xy_j} \right\|)$. P represents the moving coordinates $\mathcal{M}_i$ in homogenous coordinates, O is a 3×3 matrix of zeros, o represents a 3×1 vector of zeros, and F represents the target landmarks. Solving this particular system of equations, amounts to solving for interpolation, and will be referred throughout this Thesis as the interpolating TPS. Moreover, the TPS can be easily regularized, which will be referred, throughout this Thesis, as the approximating TPS. In the case of the approximating TPS, the system of linear equations may be re-written as:

$$\left[\begin{array}{c|c} K + \lambda I & P \\ \hline P^T & O \end{array}\right] \begin{bmatrix} w \\ a \end{bmatrix} = \begin{bmatrix} F \\ o \end{bmatrix} , \tag{3.1.1.6}$$

with λ being the regularization parameter and I the identity matrix. This minimizes the function:

$$H[f] = \sum_{i=1}^{p} \left\| f(\mathcal{I}m^i) - \mathcal{I}f^i \right\|^2 + \lambda \mathcal{J}_f \, , \tag{3.1.1.7}$$

where the λ parameter controls the *smoothness* of the bending. When $\lambda = 0$, the TPS function reverts to interpolation. Finally, the coefficients w and a can be solved as

$$\begin{bmatrix} w \\ a \end{bmatrix} = \left[\begin{array}{c|c} K + \lambda I & P \\ \hline P^T & O \end{array}\right]^{-1} \quad \begin{bmatrix} F \\ o \end{bmatrix} . \tag{3.1.1.8}$$

Solving for this matrix inversion, yields an approximating or interpolating transformation displacing the points $\mathcal{I}m_i$ to $\mathcal{I}f_i$. The extension to 3D is trivial.

One of the drawbacks of TPS is the need for a $p \times p$ matrix inversion, where p are the number of points in the dataset. However there exists a variety of methods that reasonably perform TPS transformations using subsampled sets of points [38]. It should be noted that other radial basis functions could be used in the same framework, such as Gaussian or multi-quadratic. Their effects are shown in Fig. 3.2.

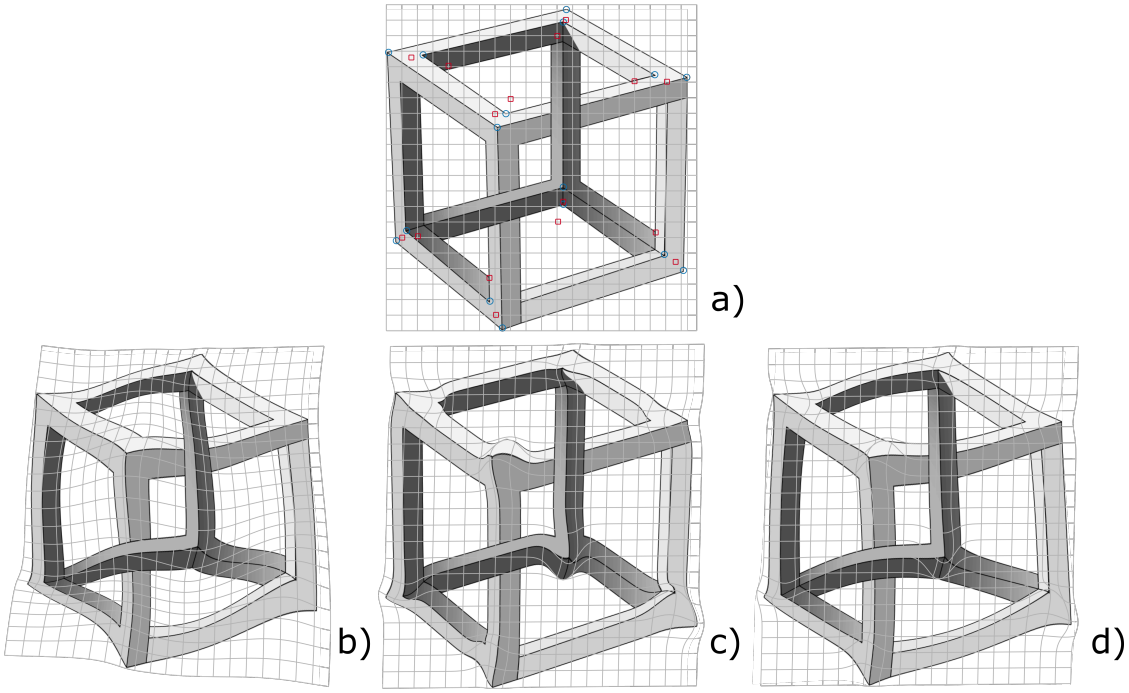

Figure 3.2: Illustration of non-rigid deformations under different operators. a) Original image representing Escher's cube. The blue circle represent the moving points, and the red squares represent the fixed points. b) Thin-plate spline deformation (interpolating) where the radial basis function $U = r^2 \log r$. c) Gaussian radial basis function given by $U = e^{-\frac{r^2}{\sigma}}$. d) Multi-quadratic radial basis function $U = \sqrt{a + r^2}$.

.

B-Splines

B-Splines are a common tool to represent free-form deformations (FFDs). As opposed to TPS, they manipulate deformations via the transformation of control points, and not the landmarks themselves. They are especially used in medical imaging by manipulating a mesh of control points lying on the image domain defined as: $\Omega = \{(x, y) | 0 \leq x \leq X, 0 \leq y \leq Y\}$. The non-rigid transformation matching the set of control points on $\mathcal{I}m$ to the other set of control points on $\mathcal{I}f$ is given by (in 2D):

$$T(x, y) = \sum_{l=0}^{3} \sum_{l=0}^{3} B_l(u) B_m(v) P_{i+l, j+m} \,, \tag{3.1.1.9}$$

where $i = \lfloor x/\delta \rfloor - 1$, $j = \lfloor y/\delta \rfloor$, $u = x/\delta - \lfloor x/\delta \rfloor$, and $v = y/\delta - \lfloor y/\delta \rfloor$, with δ corresponding to the grid size. P represents the set of control points and B_l

corresponds to the lth B-spline basis function with

$$B_0(u) = (1 - u^3)/6 \tag{3.1.1.10}$$

$$B_1(u) = (3u^3 - 6u^2 + 4)/6 \tag{3.1.1.11}$$

$$B_2(u) = (-3u^3 + 3u^2 + 3u + 1)/6 \tag{3.1.1.12}$$

$$B_3(u) = u^3/6. \tag{3.1.1.13}$$

The B-spline transformation can also be seen as a system of linear equations since each control point is associated with a basis function. One advantage of B-splines is that each of the control points is controlled locally according to its neighboring points, making them computationally efficient resulting in a sparse matrix.

3.1.2 Similarity measures

Similarity measures provide a measure of (dis)similarity between two entities, such as intensities in the domain of image registration, or points in the geometrical domain.

3.1.2.1 Image-based measures

Similarity functions are an integral part of image registration. They can be feature-based, which aim at the alignment of specific image features (e.g. edges), or intensity based, which quantifies differences between intensity values. The choice of similarity measure is one of the key factors in the performance of a registration process, and depends on the nature of the data to be registered [39]. In general, the choice of image-based similarity measures in medical imaging will depend on if the data being registered are mono-modal, or multi-modal.

In the case of mono-modal registration, one of the most common measures involves the direct comparison between intensities, using the sum of squared differences (SSD) or sum of absolute differences (SAD):

$$SSD(I_A, I_B) = \frac{1}{N} \sum_{x,y \in \Omega} (I_A(x,y) - I_B(x,y))^2 \text{ , and} \tag{3.1.2.1}$$

$$SAD(I_A, I_B) = \frac{1}{N} \sum_{x,y \in \Omega} \|I_A(x,y) - I_B(x,y)\| \tag{3.1.2.2}$$

where I_A and I_B represent the two image intensities being compared, N the number of pixels, x and y the pixel coordinates and Ω the image domain. However, in cardiac MRI images intensity values might not be equivalent between acquisitions, due to effects such as blood flow or protocols including gadolinium injections which suffer from contrast wash-in/wash-out in the images, increasing the disparity in intensities between slices at different time instants. Furthermore, intensity inhomogeneities due to MRI coils make the use of such measures inapropriate for the comparison between different slices. As such, even though the data is mono-modal, using cross correlation or its normalized version, normalized cross-correlation(NCC) is more appropriate. NCC is a mathematical representation of the similarity between two samples. In image processing, it relates to the relative changes of intensities. NCC is preferably used over cross-correlation due to intensity inconsistencies between images as it compensates for possible differences in intensity scale [40]. It is given by:

$$NCC(x,y) = \frac{1}{N}\sum \frac{(I_A(x,y) - \overline{I_A})(I_B(x,y) - \overline{I_B})}{\sigma I_A \sigma I_B}\,, \tag{3.1.2.3}$$

where N is equal to the amount of pixels in the images, I_A, I_B are the intensity of two images evaluated at pixel coordinates (x, y), $\overline{I_A}, \overline{I_B}$ are their respective means and σI_A and σI_B represent their respective standard deviations. In 1-D Eq. (3.1.2.3) is also known as Pearson's normalized cross-correlation. NCC allows to generalize the assumption of an identical relationship between images, and assumes a linear one.

Regarding multi-modal image registration, one of the most popular similarity measures is mutual information, or its normalized version (NMI), which makes the measure independent to the size of the background region [41–43]. NMI is based on information theory, and relies on the information content, or entropy of the images. In image processing, entropy is defined as a measure of information - or the lack thereof. NMI can be computed by comparing the probability that the images contain a maximal amount of information, which occurs when the images are aligned. The information content of an image can be defined using Shannon-Wienner entropy:

$$\mathcal{H}(\mathcal{I}_\mathcal{N}) = -\sum_{n\in\mathcal{N}} p(n)\,\log p(n)\,, \tag{3.1.2.4}$$

where $\mathcal{N}$ represents the intensity domain of image $\mathcal{I}$, and $p(n)$ represents the probability that a voxel in $\mathcal{I}$ has the intensity value n. A measure of information between the combined images can be computed via their joint entropy, defined as:

$$\mathcal{H}(\mathcal{I}_\mathcal{A}, \mathcal{I}_\mathcal{B}) = -\sum_{a\in\mathcal{A}}\sum_{b\in\mathcal{B}} p(a,b)\,\log p(a,b)\,, \tag{3.1.2.5}$$

where $p(a,b)$ is the joint probability.

Using these criteria for defining the individual and joint entropy of two images, the NMI similarity measure can be defined as:

$$NMI(\mathcal{I}_\mathcal{A}, \mathcal{I}_\mathcal{B}) = \frac{\mathcal{H}(\mathcal{I}_\mathcal{A}) + \mathcal{H}(\mathcal{I}_\mathcal{B})}{\mathcal{H}(\mathcal{I}_\mathcal{A}, \mathcal{I}_\mathcal{B})}\,. \tag{3.1.2.6}$$

As NMI is a measure providing the statistical dependence of two random variables derived from their occurrence probabilities, it cannot be calculated for a single pixel, and necessitates a region to estimate the histogram. The choice of this region is significant since if chosen too small, the estimate is inadequate and if too large, spatial localization is poor.

Practically, the marginal and joint probability distribution of the images can be computed using the histogram of the images, by dividing the frequencies, defined by the intensity values per bins, by the number of voxels. Other, non-parametric, methods can be used to compute NMI, such as by using Parzen windows [24].

Feature-based measures

Feature-based measures are described throughout this Thesis, as measures based on the features derived from the images, as opposed to measures relating to geometrical landmarks such as segmentation points.

One of the most common methods used to derive features from an image is through the use of filters based on spatial derivatives. However, these are more suited to standard images as they depend directly on contrast values as well as do not perform well in medical images due to the high levels of noise. Furthermore, they only detect edges rather than ridges and other structures which might be of interest.

Methods based on local phase (LP) to extract features have gained in popularity, especially with regards to ultrasound imaging [44–48]. Similarity measures based

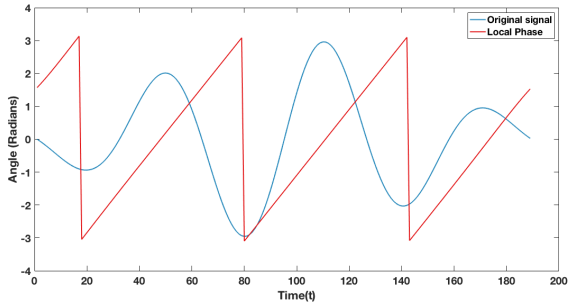

(a) Instantaneous phase of a signal $f(x)$

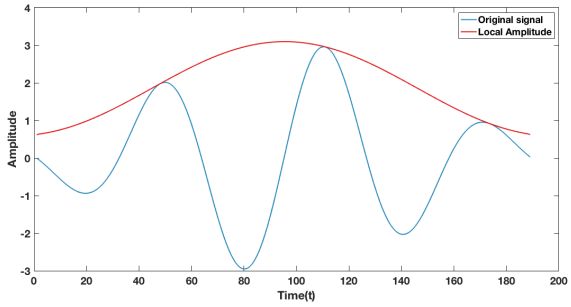

(b) Instantaneous amplitude of a signal $f(x)$

Figure 3.3: Representation of the instantaneous phase and instantaneous amplitude of a signal $f(x)$. **Top**: Sinusoidal function modulated with a Gaussian (in blue), and its phase component (in red). **Bottom**: Amplitude component of the sinusoidal function.

on LP and its variants, such as feature (as)symmetry or phase congruency, allow similarities between extracted image features to be quantified in a robust manner as they are contrast invariant, as well as allow for the detection of features at different scales. The use of local phase can be found in this Thesis in Chapter 5, as it was used as a similarity measure to correct for slice misalignments. Local phase in an image can be computed via the monogenic signal [49], which is an extension of the 1D analytic signal.

The analytic signal of a real-valued signal $f(x)$ corresponds to only positive frequency components. However, the Fourier transform of $f(x)$, $\mathcal{F}(\omega)$ includes both negative and positive frequency components (due to Hermitian symmetry). Thus, if only the positive frequencies of $\mathcal{F}(\omega)$ are taken, the resulting time domain signal will be complex-valued. Furthermore, removing the negative frequency components of a real-valued signal does not result in a loss of information. In order to obtain the analytic signal, the negative components of $\mathcal{F}(\omega)$ must be removed; which can be achieved by adding to the original signal $f(x)$ its anti-symmetric or *odd-symmetric*

version, $\mathcal{F}_h(\omega)$. This is easily achievable through the use of the *sgn* function:

$$sgn(x) = \begin{cases} 1, & x > 0, \\ -1, & x < 0, \\ 0, & x = 0, \end{cases} \text{ giving rise to: } \mathcal{F}_h(\omega) = \begin{cases} \mathcal{F}(\omega), & \omega > 0, \\ -\mathcal{F}(\omega), & \omega < 0, \\ 0, & \omega = 0, \end{cases}, \quad (3.1.2.7)$$

where $\mathcal{F}_h(\omega) = sgn(\omega) \cdot \mathcal{F}(\omega)$ and is the odd symmetric version of the original signal, in the frequency domain. As such, the analytic signal can be described as:

$$\mathcal{F}_a(\omega) = \mathcal{F}(\omega) + \mathcal{F}_h(\omega)\,, \quad (3.1.2.8)$$

where $\mathcal{F}_a(\omega)$ represents the analytic signal in the frequency domain. As mentioned, $\mathcal{F}(\omega)$ is the Fourier transform of the original signal $f(x)$, and $\mathcal{F}_h(\omega)$ corresponds to the negative frequency components with a sign change. The mathematical operation of obtaining $f_h(x)$, is known as taking the *Hilbert transform* of $f(x)$. The analytical signal in the time domain can thus be described as: $f_a(x) = f(x) + f_h(x)$. The analytic signal is useful as it is considered as a function of amplitude and phase, which can both be extracted separately, when represented using polar coordinates:

$$f_a(x) = A(x)e^{i\phi(x)} \quad (3.1.2.9)$$

$$A(x) = \sqrt{f(x)^2 + f_h(x)^2} \quad (3.1.2.10)$$

$$\phi(x) = \text{atan2}\,(f(x), f_h(x)) \quad (3.1.2.11)$$

where $A(x)$ and $\phi(x)$ are the *instantaneous amplitude* and *instantaneous phase* of the signal $f(x)$, respectively. The instantaneous amplitude represents a measure of the energy and corresponds to the envelope of $f(x)$, and the instantaneous phase represents the features of the signal, an example of which can be seen in Fig. 3.3. More specifically, instantaneous phase represents a change in feature; i.e. a value of 0 corresponds to a peak in the signal, and a value of π corresponds to a valley. Being able to extract the phase and amplitude component of a signal is also referred to as the split of identity. Two main properties should be observed: the first, is that the local phase is said to be invariant to the amplitude (it is contrast or illumination invariant) and second, it is also said to be equivariant

to spatial position (it does not contain spatial awareness). A shifted signal will have the same local phase as its original signal.

As the phase and amplitude components are obtained in the frequency domain, they do not contain spatial information. To allow the use of phase as a local descriptor, spatial information must be added. This is achieved through the use of bandpass filters. As structural information occurs at different scales, multiple bandpass filters can be used to extract the wanted frequencies corresponding to the wanted structures. There thus exists a trade off between frequency and spatial information.

The use of log-Gabor filters has proven optimal for the extraction of local phase at specific scale structures for several reasons. Notably, the log-Gabor filter has a zero DC component, as well no negative components. Furthermore, it has an extended tail compared to its non logarithmic counter part, the Gabor filter, which is beneficial as it prevents the higher frequency components from being under represented. The log-Gabor filter, which is defined in Eq. (3.1.2.12), is controlled by two parameters: σ and ω_0, which represent the shape (between 0 and 1), and the center-frequency of the filter, respectively. The center-frequency allows to select the specific frequencies of interest and the shape which allows to control the bandwidth of the filter.

$$LG(\omega) = \exp\left(- \frac{\left(log\left(\frac{\omega}{\omega_0}\right)\right)^2}{2(log(\frac{\sigma}{\omega_0}))^2}\right), \tag{3.1.2.12}$$

where LG stands for log-Gabor, ω_0 is the center frequency and σ is the bandwidth parameter. Figure 3.5 shows the respective influence of the parameters. In order to apply the log-Gabor filter to our signal to obtain the analytic signal, two methods exist. The first is to create the analytic signal by applying the log-Gabor filter to the Fourier transform of the signal to obtain the Hilbert part of the signal as shown in Eq. (3.1.2.13).

$$\mathcal{F}_a(\omega) = LG(\omega)\mathcal{F}(\omega)(1 + sgn(\omega))\,. \tag{3.1.2.13}$$

The second option is to compute the analytical version of the filter, and apply it to the original signal directly, as these two methods are equivalent.

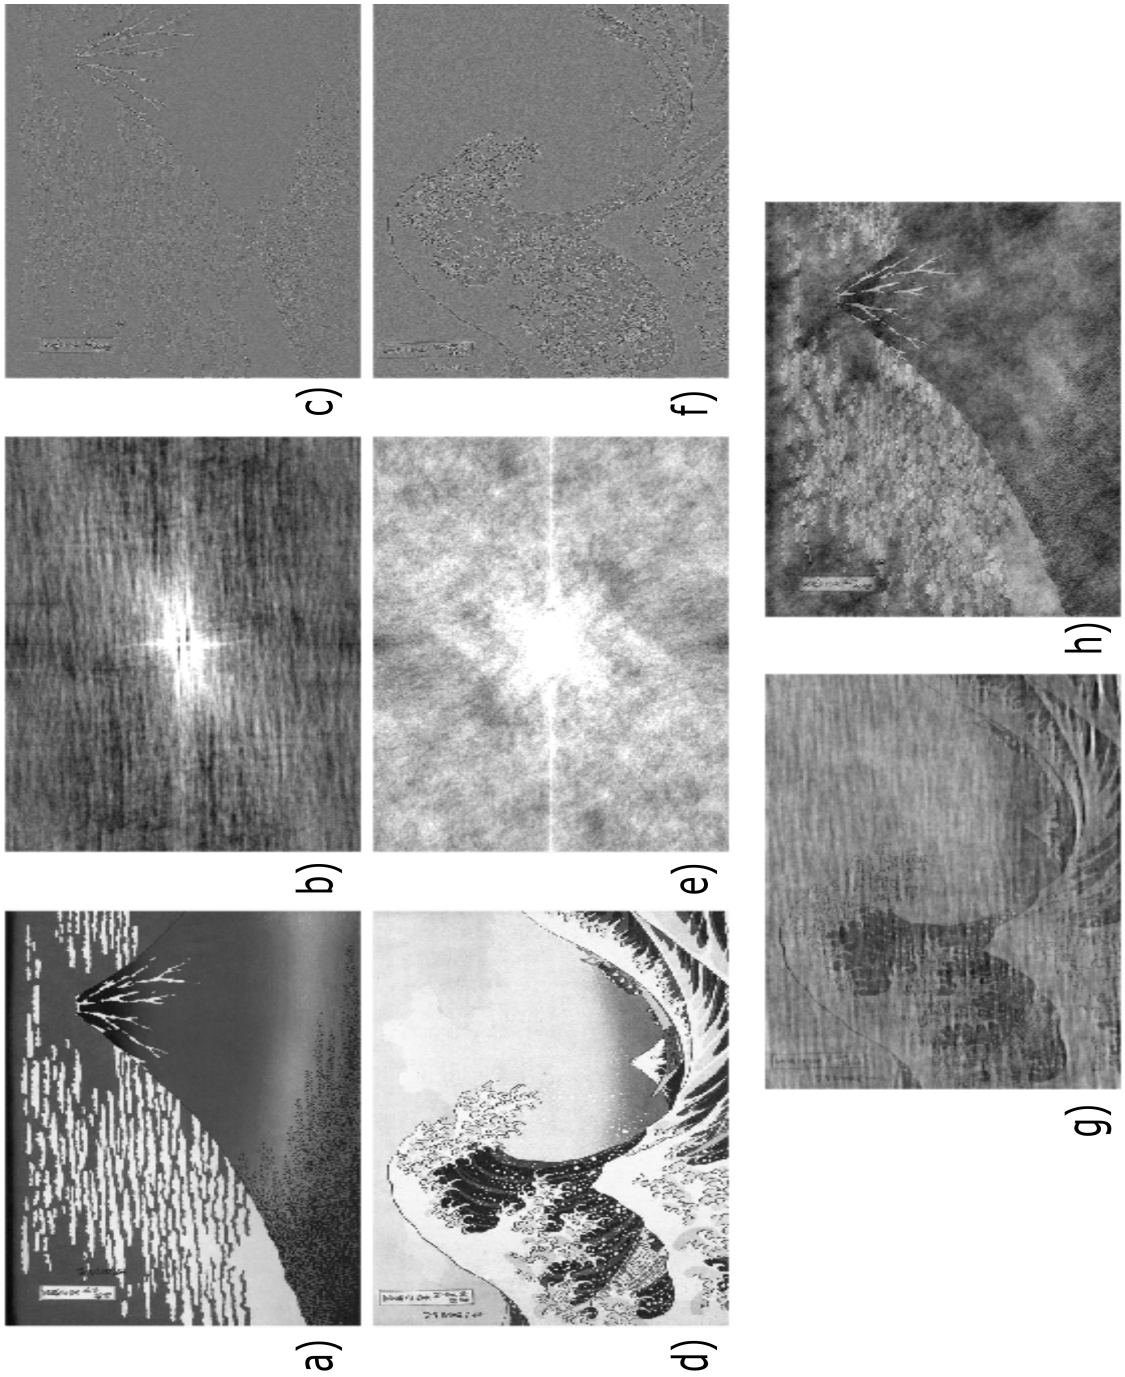

Figure 3.4: Illustration of the importance of phase in images, as first shown in [50]. a) The original image of the Gaifū kaisei painting by Hokusai. b-c) The respective amplitude and phase components. d) The original image of the Kanagawa-oki nami ura painting, also by Hokusai. e-f) The respective amplitude and phase components. g) The image resulting from the substitution of the amplitude of Gaifū kaisei and phase of Kanagawa-oki nami ura. h) Substitution of the amplitude of Kanagawa-oki nami ura and phase of Gaifū kaisei. It can be seen that the phase components dominates the structural information of the images.

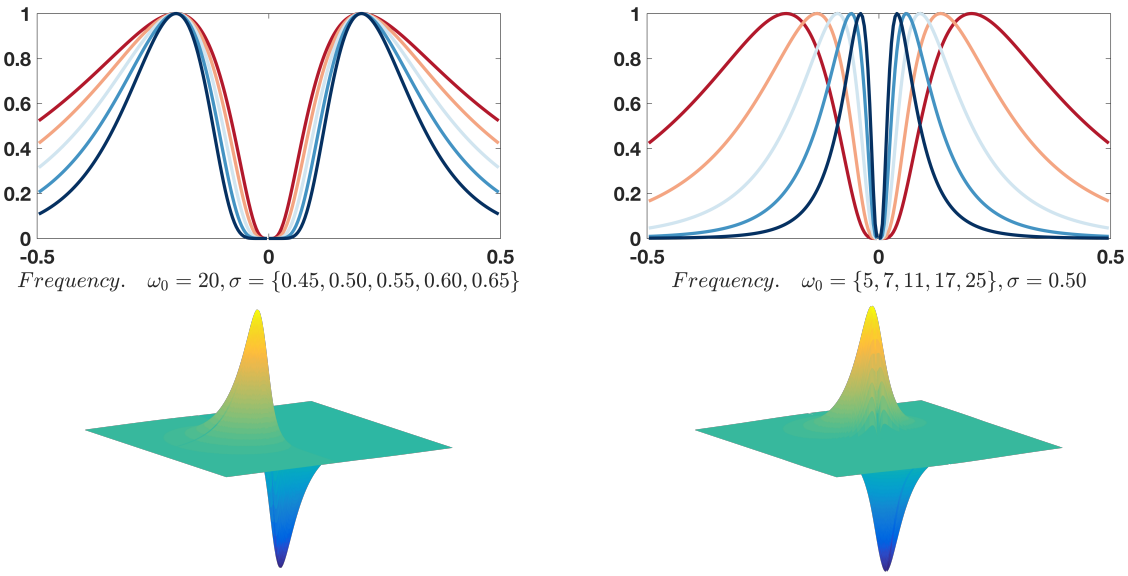

Figure 3.5: Illustration of the log-Gabor and Riesz filters. The **top** row represents different log-Gabor filters according to different σ and ω_0 parameters. The top left figure represents a change in the bandwidth (σ) parameter, whilst having the central frequency (ω_0) fixed. The top right figure represents the log-Gabor filters having σ fixed, and varying ω_0. The **bottom** figures represent the real and imaginary part of the Riesz transform of a log-Gabor filter, respectively.

In image processing, phase is a very important concept, as it was shown in [50] that the phase component, and not the amplitude, encodes the majority of the information in an image. An example of this is shown in Fig. 3.4. However the use of the 1D analytical signal cannot be used in the same way for nD images, as for local phase to be extracted in a useful manner, it needs to be isotropically derived. To that extent, there exists an method to compute a multi dimensional analytic signal, known as the *monogenic* signal.

The monogenic signal can be obtained by substituting the Hilbert transform by the Riesz transform. By using the same prior method, in 2D, the monogenic signal will be composed of three parts. The first part contains the original bandpassed filtered signal, whilst the second and third parts contain the Riesz transform of the bandpass filtered signal, forming a pair of directional filters, as shown in Fig. 3.5. Thus, the application of these three components to an image, form the monogenic signal.

In the same way that the local phase (instantaneous phase in 1D) and local

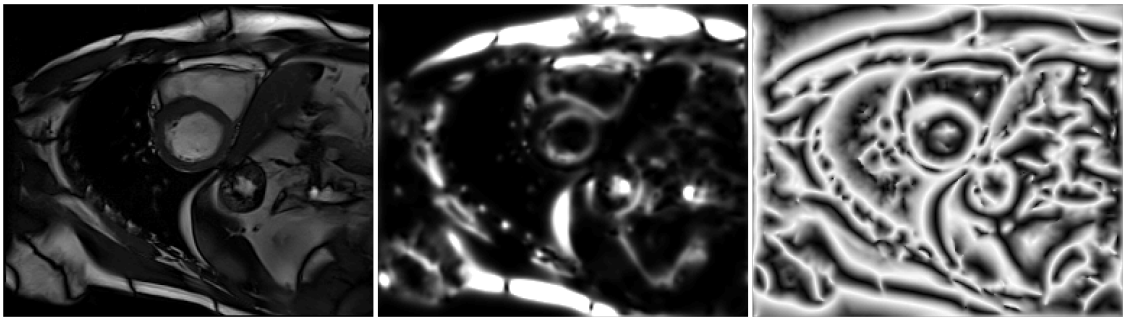

Figure 3.6: Local phase and local amplitude of a SA cardiac image. From left to right: original image, local amplitude, local phase. It can be seen that the local amplitude is dependent on the pixel intensity. The local phase and local amplitude images were obtained by using the monogenic filter with $\sigma = 0.5$ and $\omega_0 = 15$.

amplitude (instantaneous amplitude in 1D) were extracted using polar coordinates, for the 2D case, spherical polar coordinates can be used. The radial coordinate, A, represents the local amplitude (or energy), the elevation angle, ϕ ,represents the local phase, and the azimuth angle, θ, represents the local orientation of the image:

$$A(x,y) = \sqrt{I_{R_1}(x,y)^2 + I_{R_2}(x,y)^2 + I_{R_3}(x,y)^2} \tag{3.1.2.14}$$

$$\phi(x,y) = \text{atan2}\left(\frac{\sqrt{I_{R_2}(x,y)^2 + I_{R_3}(x,y)^2}}{I_{R_1}(x,y)}\right) \tag{3.1.2.15}$$

$$\theta(x,y) = \text{atan2}\left(\frac{I_{R_3}(x,y)}{I_{R_2}(x,y)}\right) \tag{3.1.2.16}$$

where I_{R_1} represents the bandpass filtered image, and I_{R_2} and I_{R_3} represent the directional Riesz component filters applied to the image. A representation of the Riesz transform of the log-Gabor filter can be seen in Fig. 3.5, and the corresponding local phase, and amplitude of a monogenic image can be seen in Fig. 3.6.

It can be seen from Fig. 3.6 that local phase is a very useful tool to extract features from signals or images, which has popularized its use for medical imaging research, notably in the ultrasound community. Different registration algorithms have made use of local phase such as [51], which used the mutual information of local phase as a similarity measure to non rigidly register brain images from MRI and ultrasound. Other uses of local phase include their use for motion estimation instead of traditional approaches such as optical flow [52, 53] .

One of the main limitations of using local phase obtained from monogenic signal, it that the monogenic signal assumes that the image is locally intrinsically 1D. Because of this, it cannot extract corner information. However, in the work presented in Chapter 5, the 1D signal at the line intersections of the two local phase images is taken, and thus the method does not rely on corner information but on the ability to match edge-like structures.

3.1.3 Multilevel image registration

As mentioned in Section 3.1.2.1, structures exist at different levels. In the case of image registration, when optimizing over an energy, or cost-function, there exists the possibility that the energy function contains many local optima, making the method slower, with a small capture range. To that extent, multi-resolution has proven to accelerate the optimization process, whilst multi-scale optimization methods are effective in increasing the capture range of the cost function. An example of both of these techniques can be seen in Fig. 3.7. These methods, notably the multi-scale approach, have been used throughout this Thesis to enable faster convergence. By using a multi-resolution approach, images are decomposed into several images of different resolutions, where higher resolutions depict more localized and finer details of the image, and lower resolutions grasp major, lower dimensional, features of the image. By registering the lower resolution images in a first approach, major features appearing in both images can be matched efficiently. The algorithm can then focus on registering finer, more localized features in the images in each of the stages, making the registration technique more robust.

3.1.4 Limits of CMRI in image registration

As discussed in this Chapter, image registration relies on matching specific features or objects of interest. However, as mentioned in Section 2.2.3.3, clinical acquisition of CMR images comes with certain limiting factors. These factors contribute directly to increasing any difficulty in the image registration problem as it can cause a lack of specific boundaries or features to be identified. Having thicker CMRI slices

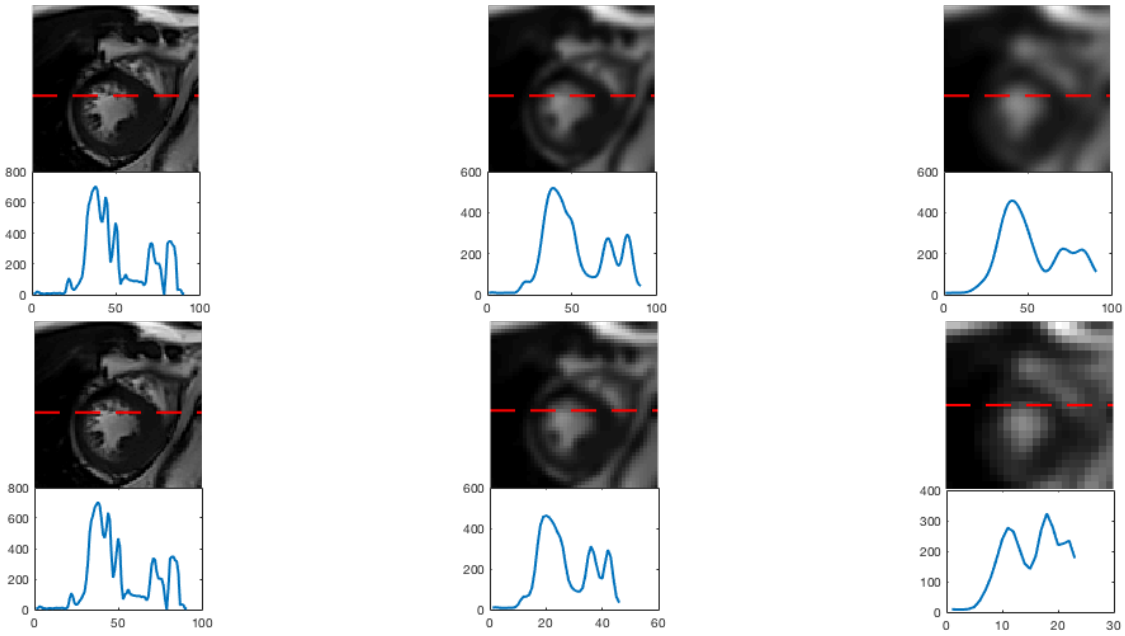

Figure 3.7: Example of multi-scale versus multi-resolution. **Top**: Multi-scale images. The top left is the original image, the middle is the original image filtered with a Gaussian of $\sigma = 2$, and the top right is the original image filtered with a Gaussian of $\sigma = 4$. **Bottom**: Multi-resolution images. The bottom left is the original image, the bottom middle is the half sized downsampled image, and the bottom right is a four times downsampled images. The images have also undergone Gaussian sampling, as per the top row. The representative line profiles, taken at the red-line level, are shown under each image.

leads to a more blurred representation of structures of interests, due to partial volume averaging, loosing any information from smaller structures that could be useful in the registration process, as any differing signal from distinct structures can become merged. As such, the spatial location of a slice could take on multiple locations when aligning intersecting slices. That is, there exist no true optimum spatial location between an intersecting LAX and SAX slices, if the SAX were to be translated out of plane to match information from the LAX. Due to the slice thickness, there exist a certain uncertainty, relatively proportional to the thickness, to the optimum location between the slices, as the alignment would aim to match averaged information between the slices.

To that extent, the methods developed in this Thesis, along with the tools used, contribute to obtaining the most likely position between slices during any registration procedure.

3.2 Mesh Reconstruction

This Section introduces some of the methods used to generate meshes, as well as some common mesh operations such as decimation or subdivision. As the literature in this area is very broad, the main focus lies in the techniques used in this Thesis.

3.2.1 Common surface construction algorithms

A geometric surface can be defined by a set of points, known as vertices in 3D, and their connectivities which make up a collection of connected faces. Their connectivity can be defined, most commonly, as either triangles or squares. These surface meshes describe a surface of a geometric entity, but have no connectivity within the interior of the surface domain. They are intrinsically 2D geometric representations in a 3D environment. To allow for some simulation to occur, the surface mesh must be converted to have 3D connectivity by converting the connectivity shapes from triangles and squares to tetrahedral and hexahedral elements respectively.

Convex Hull

The convex hull is one of the most basic ways to obtain a surface mesh representation of a point set. The convex hull of a point set is the *envelope* of the point set, or more formally, it is the smallest convex set containing the set of points $\{\mathcal{P}\}$. There exist many different incremental algorithms for computing convex hulls, such as Quickhull, Chan's algorithm, Gift wrapping, or Graham scan [54–57]. The main caveat of the convex hull methodology is that only a convex geometry can be obtained, which can be too simplistic.

Poisson Reconstruction

One of the most common methods for constructing a surface from point clouds is the Poisson surface reconstruction algorithm. Given an oriented point cloud (having a normal at each point), an implicit function $\mathcal{X}$ (also known as an indicator function), is defined which represents the zero set of a distance function - where the function is < 0 outside the shape, > 0 inside the shape and is equal to 0 where

the surface lies. From a set of points, their respective normals can be computed by fitting tangent planes to neighboring points. Given the vector field of point normals V, a surface can then be computed by solving the Poisson equation: $\Delta\tilde{\mathcal{X}} = \nabla\overrightarrow{V}$.

One of the main drawbacks of this method is the need for point normals to be defined. Furthermore, this approach is usually applied on dense point clouds, and can fail for sparse data where the normals are ill-defined, which is the case for the data in this Thesis.

Marching Cubes

Marching cubes is a volume oriented algorithm that has been developed specifically for the purposes of surface reconstruction in medical imaging. It is a brute-force method that generates a surface mesh from isodensity volume data [58]. In 2D, an outline from a binary image can be identified by assessing its pixel arrangement in a 2×2 square. There exist only four possible pixel arrangements that define a boundary edge, which can create up to fourteen possible shapes. If the binary image is divided into 2×2 squares, each of the squares can be checked against any of the fourteen possible boundary shapes. In such a way, the edge profile can be obtained. The same process can be applied in 3D. If a 3D data set is given as a volume, every voxel can be represented as a cube having 4 points (corners). The volume can thus be seen as a grid and can be divided into cubic subvolumes of arbitrary size. In 3D, there exist 256 possible boundary configurations, however, if rotations and redundancy are removed, these configurations reduce to 15 unique possibilities, as shown in Fig. 3.8. By increasing or decreasing the size of the subvolumes, finer or coarser surface volumes can be generated [25].

One of the greatest disadvantages of marching cube is its necessity for dense data. Furthermore its generation of large numbers of polygons tend to over tesselate the surface [58].

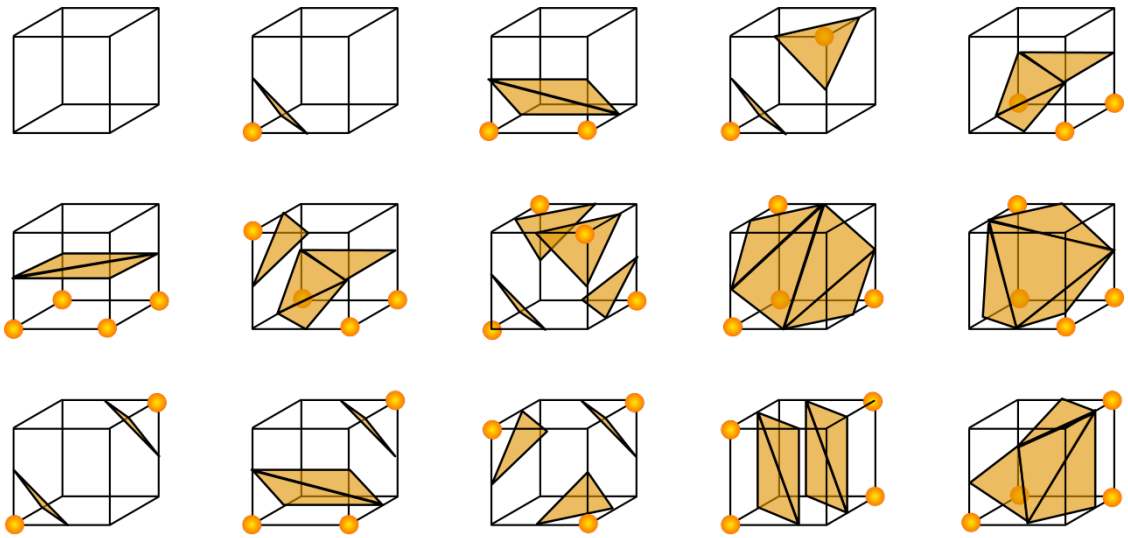

Figure 3.8: The fifteen possible configurations for the 3D marching cube algorithm. Each of these configuration represents a possible edge structure. By comparing a sub-volume against all these combination, a surface volume can be generated.

3.2.2 Geometric tools

Geometric meshes have many degrees of freedom that can alter their shape or geometry, such as how coarse or fine, or how smooth it can be. The complexity and topology of a mesh can be controlled by operations known as subdivision and decimation, and its smoothness can be controlled via smoothing algorithms. Some of the common operation algorithms are described below.

Subdivision and decimation

Given a polygonal mesh, a subdivision algorithm is the process of refining the mesh by subdividing the individual mesh elements (vertices, and faces) into further components, increasing the complexity. Its counterpart, decimation, refers to the process of simplifying the mesh elements, making the mesh coarser.

According to their refinement rules, subdivision algorithms can be divided into two types: *approximating* or *interpolating*. Approximating schemes subdivide the mesh whilst altering the original vertex positions, whereas interpolating schemes maintain them. Some of the most common approximating algorithms include Loop, Catmull-Clark and Doo-sabin. As for interpolating schemes, they include Linear, Butterfly, Modified Butterfly, Kobbelt, $\sqrt{2}$, and $\sqrt{3}$ [59–66] . However, some of the

methods mentioned only deal with quadrilaterals, which are not beneficial for the work described in this Thesis, as it deals with surface triangles. In that regard, the linear, Loop and Modified Butterfly subdivision schemes are the most relevant.

The linear subdivision process is a straight forward iterative algorithm that subdivides triangles into four different parts, where each new vertex lies on the edge, at the mid-point between two vertices. This can be seen in Fig. 3.11. As it is just a linear subdividing scheme, it does not alter the shape of the original mesh, and keeps the original vertices.

The Loop subdivision scheme is a one-to-four face splitting algorithm that works by computing *even* and *odd* vertices which are used to create new triangles (four new triangles for each old one). Even vertices are the original vertices, whereas odd vertices are the new computed vertices; although as the Loop subdivision scheme is an approximating method, both sets of vertices are recomputed, repositioning all the vertices. Each new vertex is computed as the weighted average of its neighboring vertices. For odd vertices, lying on the interior of the mesh, new vertices can be computed such that, $\mathcal{V}_O = \frac{3}{8}(\mathcal{V}_a + \mathcal{V}_b) + \frac{1}{8}(\mathcal{V}_c + \mathcal{V}_d)$. Odd vertices positioned on a boundary can be computed in the following manner: $\mathcal{V}_O = \frac{1}{2}(\mathcal{V}_a + \mathcal{V}_b)$. In the case of even vertices, for the interior vertices, new vertices can be obtained, such that: $\mathcal{V}'_E = \beta\mathcal{V}_E + \alpha\sum_{j=0}^{n-1}\mathcal{V}_{E_j}$, where n is the valence associated to vertex $\mathcal{V}_E$, $\beta = 1 - \alpha$. In the original paper by Loop [62], α could be computed such that:

$$\alpha = \frac{1}{n}\left(\frac{5}{8} - \left(\frac{3}{8} + \frac{1}{4}\cos\frac{2\pi}{n}\right)^2\right). \tag{3.2.2.1}$$

Applying this operation to a vertex of valence 6, for example, will lead to $\alpha = \frac{1}{16}$ and $\beta = \frac{5}{8}$ [67]. However, more recently, Warren *et al* [68] proposed an alternative such that, if $n = 3$, $\beta = \frac{3}{16}$, and in the case that $n > 3$, $\beta = \frac{3}{8n}$. New boundary vertices, in the even case, can be computed as $\mathcal{V}'_E = \frac{1}{8}(\mathcal{V}_a + \mathcal{V}_b) + \frac{3}{4}(\mathcal{V}_E)$. An example of the computation of new even and odd vertices can be seen in Fig. 3.9. All of the operations mentioned above are vector operations on 3D coordinates.

The Butterfly algorithm, initially proposed by Dyn *et al* [65], is also a one-to-four interpolating subdivision scheme, that was later modified by Zorin *et al* [59]

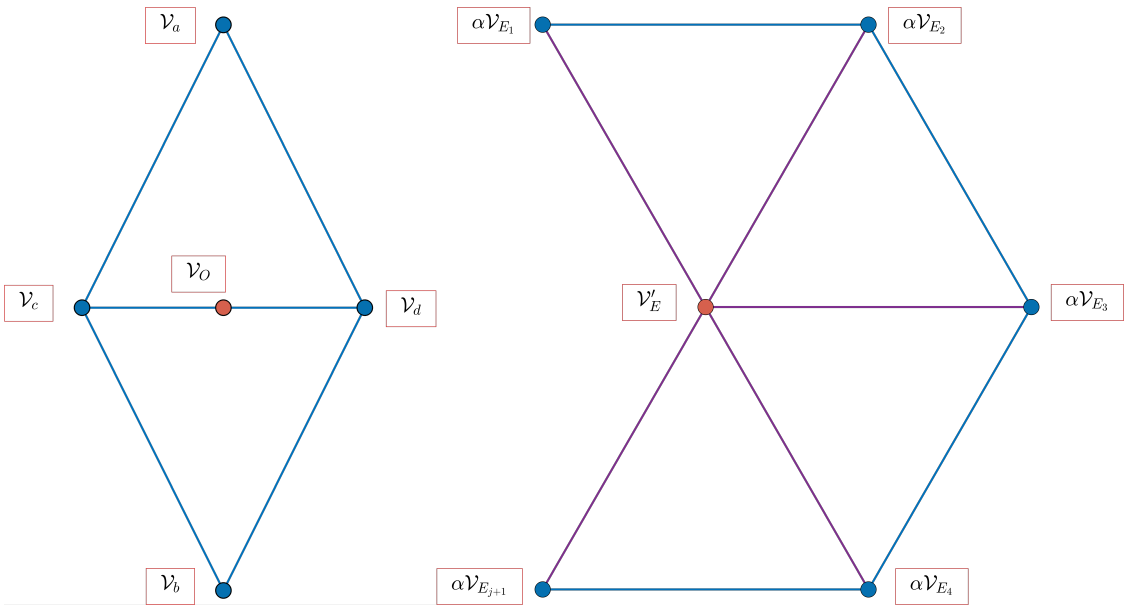

Figure 3.9: Loop subdivision scheme for even and odd vertices. Left: Odd vertices subdivision. Right: Even vertices subdivision scheme.

in order to deal with vertices of valence (number of edges) > 6. Compared to the approximating Loop subdivision scheme, the Butterfly algorithm takes more neighboring vertices into account. In the original Butterfly scheme, an eight point stencil was used to subdivide arbitrary triangles, although Zorin *et al* [59] introduced higher order stencils to work with vertices of valence $\neq 6$, termed *extraordinary points*. Two stencils are shown in Fig. 3.10. Given a surface $\mathcal{M}$, the Butterfly subdivision iteratively subdivides all of the triangle faces into 4 new triangles, retaining the original vertices. New vertices are computed using the edges of the original triangle, using stencils according to their types. Each new vertex is therefore a weighted combination of its surrounding points such that, for a regular interior (show in Fig. 3.10), a new vertex can be calculated as $\mathcal{V}_j = \frac{1}{2}(\mathcal{V}_1 + \mathcal{V}_2) + \frac{1}{8}(\mathcal{V}_3 + \mathcal{V}_4) + -\frac{1}{16}(\mathcal{V}_5 + \mathcal{V}_6 + \mathcal{V}_7 + \mathcal{V}_8)$. For a neighborhood considered extraordinary, the weights associated to find the new vertex can be computed by $\mathcal{S}_j = \frac{1}{n}(\frac{1}{4} + cos\frac{2j\pi}{n} + \frac{1}{2}cos\frac{4j\pi}{n})$, where the number of edges, $n \geq 5$. When $n = 3$, $\mathcal{S}_0 = \frac{5}{12}$, $\mathcal{S}_1 = -\frac{1}{12}$, $\mathcal{S}_2 = -\frac{1}{12}$. When $n = 4$, $\mathcal{S}_0 = \frac{3}{8}$, $\mathcal{S}_{1-3} = 0$.

One of the main aims of subdivision schemes is to be C_1 or C_2 continuous. These relate to the smoothness of the geometry. C_1 and C_2 continuity concerns the

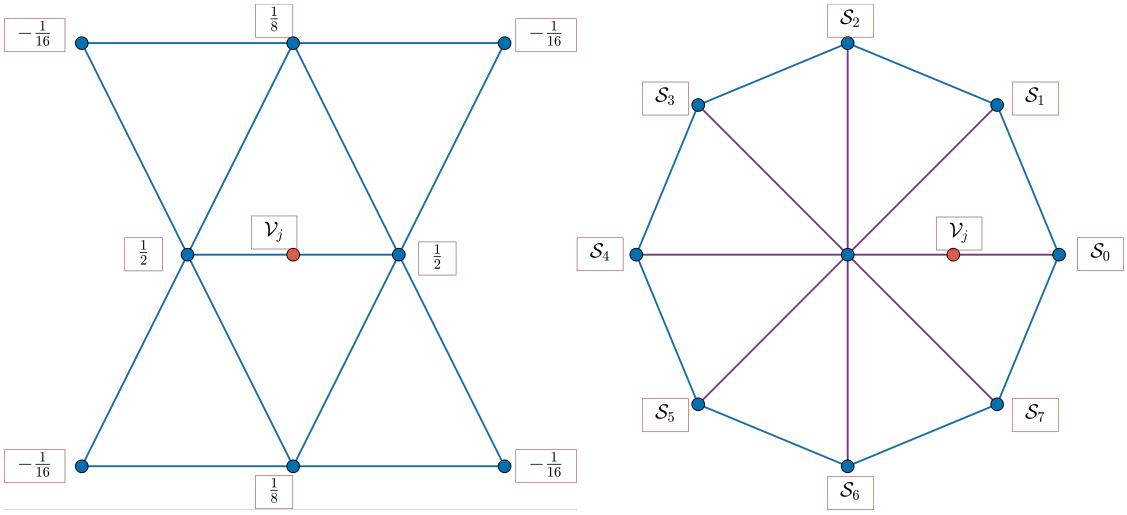

Figure 3.10: Modified Butterfly subdivision scheme for regular and extraordinary vertices. Left: Vertex subdivision for regular interior cases. The Butterfly subdivision scheme gets its name from the shape of the stencil. Right: Vertex subdivision scheme for extraordinary cases (valence > 6). The newly computed vertex is a function of the weighted neighbors.

ability for the parametrization of a curve to be continuous and be differentiable with regards to the first and second derivatives respectively. To that end, the linear subdivision process does not fall under this definition, whereas the Butterfly and Loop schemes do.

Given $\mathcal{M} = (\mathcal{V}, \mathcal{F})$, with $\mathcal{V}$ the set set of vertices, and $\mathcal{F}$ the set of triangle faces, decimation can be described as the process of finding $\mathcal{M}\prime = (\mathcal{V}', \mathcal{F}')$ such that $N_{\mathcal{V}'} \ll N_{\mathcal{V}}$ and the topology change is minimal, where N_m is the number of vertices m. The assessment of decimation algorithms can be made through the computation of some error value. However, there are many possible assessment criteria, which depend on the application of the mesh. Volume or area preservation might be the criterion, or the error might be computed in terms of the geometric distance between the original and the new vertices. In this Thesis, most surface error functions are computed via the latter evaluation, such that:

$$E(\mathcal{M}_1, \mathcal{M}_2) = \frac{1}{k_1 + k_2}\left(\sum_{p \in \mathcal{V}_1} d(p_1, \mathcal{M}_2)^2 + \sum_{q \in \mathcal{V}_2} d(q_1, \mathcal{M}_1)^2\right) \text{ where,} \tag{3.2.2.2}$$

$$d(x, \mathcal{M}) = \min \ \|x - y\| \tag{3.2.2.3}$$

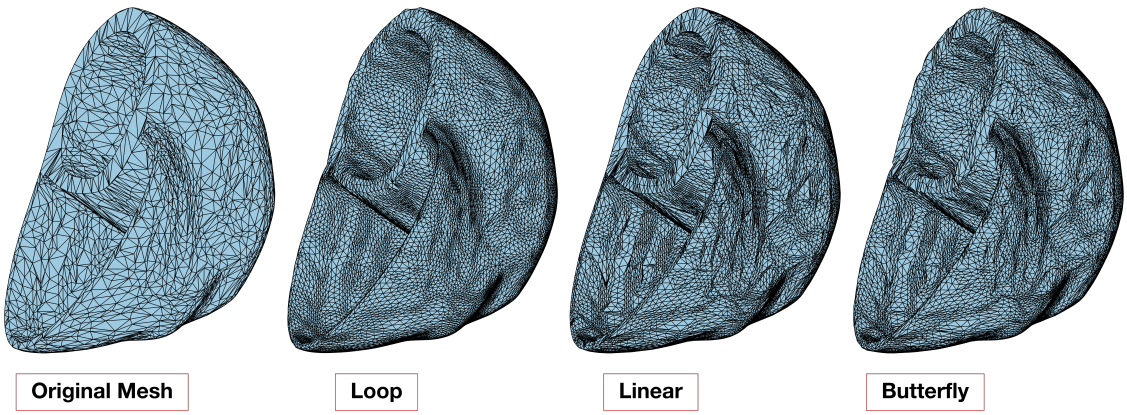

Figure 3.11: Different subdivision schemes. From left to right: original surface, Loop, linear, and modified butterfly subdivision schemes. The Loop subdivision is an approximation method, compared to the linear and modified Butterfly, which are interpolating methods. It can be seen that, although most visually pleasing, the Loop subdivision algorithm can incur some errors. Furthermore, as it is approximating, the original vertices are lost, and the mesh has shrunk. The linear subdivision algorithm subdivides the original mesh triangles, into four using the linear subdivision scheme. It can be seen that all the original features have been preserved. The modified Butterfly scheme also creates four new triangles for each triangle, but additionally smoothes over sharp edges, as demonstrated at the boundaries.

where $\mathcal{V}_i$ is the set of the vertices of the mesh $\mathcal{M}_i$ and k_i its cardinality.

Decimation strategies can be divided into local and global approaches, where local strategies act on a small subset of the mesh, whereas global strategies act on the mesh as a whole. Typically, local approaches are preferred as they allow for higher specificity criterion, such that a specific number of faces must remain after the decimation or that the error measure must not fall below a certain threshold. Local decimation algorithms can be divided into *vertex decimation*, and *edge contraction*.

Given a triangulated surface, Schroeder *et al* [58] proposed a vertex decimation algorithm that aims to remove triangle edges, and perform a re-triangulation that will lead to less components whilst maintaining the topology of the mesh. Given a vertex ($\mathcal{V}_i \in \{\mathcal{V}\}$), a plane can be fitted using the area-weighted average of the triangle normals $\mathbf{n}_i$, areas A_i and centers C_i, between a neighborhood of surrounding vertices [69]. A distance between that plane and the vertex can then be computed as $d = \|\hat{\mathbf{n}} \cdot (\mathcal{V}_i - X)\|$, such that:

$$\mathbf{n} = \frac{\sum_i A_i \mathbf{n}_i}{\sum_i A_i}, \quad \hat{\mathbf{n}} = \frac{\mathbf{n}}{\|\mathbf{n}\|}, \quad C = \frac{\sum_i A_i \mathbf{C}_i}{\sum_i}. \tag{3.2.2.4}$$

A plane well fitted - in terms of the minimum vertex to plane distance - to all the surrounding vertices is indicative of a flat region, and therefore, as it does not contain important information, can be decimated. In contrast, areas with high distances, suggestive of sharp features, should be kept, to preserve the geometry. Once the vertex has been removed, a re-triangulation is performed. The algorithm iterates this process for all vertices.

The other common method, proposed by Gardland *et al* [70], contracts edges by optimizing over a cost function. However, this type of method also allows for non manifold surfaces. By iterating through all vertices in a surface, a cost function can be minimized, such that:

$$E(\mathcal{V}_i) = d(\mathcal{V}_i, P)^2 = \mathcal{V}_i^T A \mathcal{V}_i - 2b^T + c\,, \tag{3.2.2.5}$$

where $d(x, P)$ is the distance between a point to a plane, and the fundamental quadratic $\mathbf{Q}$ is given by $\mathbf{Q} = (A, b, c)$, where $A = (\mathbf{n}\mathbf{n}^T)$, $b = (\mathbf{n}\mathbf{n}^T P)^T$, $c = (P^T \mathbf{n}\mathbf{n}^T P)$. For every triangle that has a vertex, their fundamental quadratic can be added, and minimized, as the sum of square distances between the vertices and each of the planes [69]. As it is a quadratic function, a minimum can be found such that $\mathcal{V}_i = -A^{-1}b$, with the minimum value of $-b^T A^{-1} b + c$. This process makes for a fast and robust way to decimate surfaces. Furthermore, this method is easily implementable, and has been found in practice to have an excellent trade off between computational efficiency, geometric fidelity, and topological generality [69]. A certain control over the decimation can be established by selecting an error threshold, or a decimation percentage. This method, which is used in this Thesis, can be found in Chapter 6.

Smoothing

The goal of smoothing a surface can analogously be thought of as removing high frequencies in a given signal, whilst maintaining the best approximation to the original input. To that extent, the Laplacian smoothing method acts as a low pass

filter by moving the internal vertices to the barycenter of their surrounding vertices in order to smooth a surface [71,72]. The Laplacian smoothing operation is given by:

$$\mathbf{p}_i = \frac{1}{|\mathrm{adj}(i)|} \sum_{j \in \mathrm{adj}(i)} \mathbf{q}_j \,, \tag{3.2.2.6}$$

where p_i is the new position of the vertex i and $\mathbf{q}_j$ are the original coordinates of the points connected to i and $\mathrm{adj}(i)$ is the set of adjacent vertices to i.

Applying this smoothing operation usually shrinks the mesh, however, in the case of Chapter 6, an affine transformation between the original and final positions of the vertices is applied after the smoothing, to maintain the overall size. More recent techniques for mesh smoothing exist such as [73–75], however, one of the benefits of the Laplacian smoothing operator is its simplicity and speed, which still renders it one of the most common tools used.

4
Literature Review

Contents

4.1 Misalignment between CMR slices 60
4.1.1 Motion correction 62
4.2 Surface reconstruction from cardiac contours 68
4.3 Misalignment correction and surface reconstruction integration 70
4.4 Conclusion 71

Motion is an intrinsic part of different organs or organisms, such as the lungs, heart, brain, prostate, liver or fetus. In the field of Medical Imaging, motion related artifacts can cause image acquisitions to appear blurred or suffer from ghosting, ultimately increasing clinical diagnosis difficulties [76]. Correcting for such motion remains a challenging task and depends on the application [77–79]. This Thesis is concerned with motion artifacts relating to the heart that might be caused by breathing motion, patient motion inside the scanner, or gating issues.

4.1 Misalignment between CMR slices

CMR images allow for detailed ventricular anatomical information as well as an accurate representation of myocardial function by using a plethora of available specialized protocols. Acquiring cardiac images is a complex process due to the

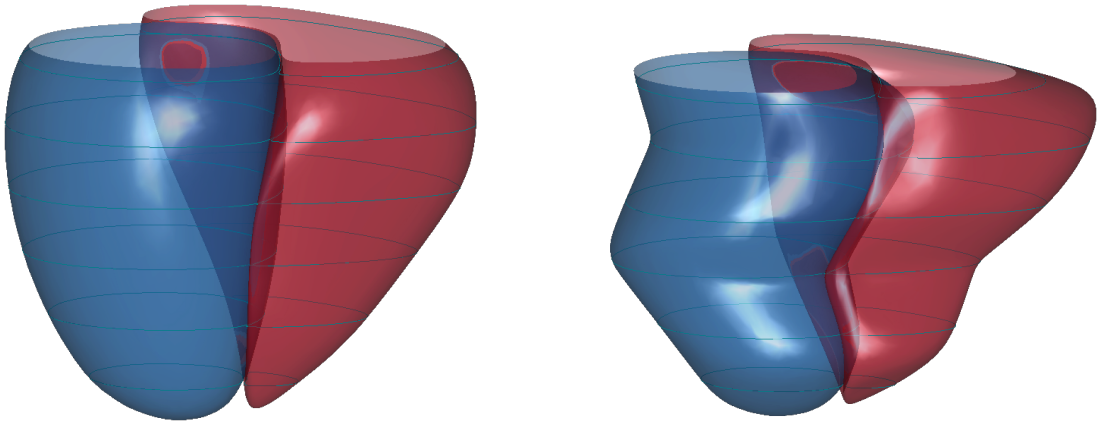

(a) Ground Truth Mesh **(b)** Misaligned Mesh

Figure 4.1: Representation of the effect misaligned CMR slices have on mesh reconstruction. Meshes were generated using SurFCo [87], on short axes contours. **a)** The ground truth mesh. **b)** Mesh reconstructed from misaligned contours.

constant motion of the heart. Whilst 3D acquisitions can be obtained via steady state free procession (SSFP) sequences, they are limited by poor temporal and spatial resolution, making them unreliable for functional analysis [80, 81]. Standard clinical protocols thus typically acquire a collection of 2D slices, oriented either on the short axis or on a long axis plane of the ventricles, each one at a separate breath hold. As mentioned in Chapter 2, ECG gating allows the image acquisitions to compensate for cardiac motion whilst breath holding at the same lung volume for periods ranging from 7 to 15 seconds is used to reduce slice misalignment, although this reduction is highly variable. This alignment distortion may be further enhanced by any patient movement inside the scanner, affecting slice spatial coherence as a 3D dataset [13]. To quantitatively assess cardiac functions, clinicians typically segment CMR slices in 2D, which is then used for 3D reconstructions allowing for global and local evaluations of various parameters. However due to slice misalignments, any 3D reconstruction and subsequent evaluations may become distorted (illustrated in Fig. 4.1). As such the correction of slice misalignment is mandatory for any 3D analysis, such as volumetric ventricle assessment, shape analysis, multi-modality fusion, or mechanical and electrophysiological simulations [40, 82–86].

4.1.1 Motion correction

As the correction of CMR slice misalignment is a necessary pre-processing step for any 3D cardiac analysis, significant amount of research has been dedicated to it. Specifically, when dealing with large number of datasets, the correction of misaligned slices can not be performed manually as per [88], and as such we focus on automated methods.

The literature can be partitioned into two distinct approaches: geometry-based and image-based methodologies. Geometry based methods generally consist of extracting geometric information, such as contours, from the images and using them to realign the slices. Image-based approaches can be further subdivided into 3 main categories: slice-to-volume, volume-to-volume and slice-to-slice alignment.

In contour-to-contour techniques, some studies solely align endo- and epicardial contours with regards to each other. Van Assen *et al* [89] aligned SA centroids with respect to LA centroids, making the assumption that the LV is rotationally symmetric, which can lead to the common banana problem [90, 91]; specifically, cross-sections of an underlying 3D curved object cannot be used to recover the curved shape without using additional information. In a similar approach, Swingen *et al* [83] fitted a polynomial to the epi- and endocardium SA centroids and realigned the slices, only allowing for in-plane movements, by minimizing the difference with respect to the contour centres. These techniques require the availability of accurate contours, which means that in practice, due to the inter/intra variability in expert contouring, the final alignment will depend on the expert; on the other hand, if the application requires the smoothest ventricular shapes possible, methods such as [92, 93] can be used. Su *et al* and Tan *et al* [92, 93], reconstruct a 3D mesh and correct for the misalignment based on the geometric assumption that LV meshes must be smooth. They optimize over the principal curvature components ($\mathcal{K}_1$, defined as the maximal value of normal curvature and $\mathcal{K}_2$, defined as the minimum value of normal curvature) obtained across multiple frames, to align the mesh vertices using in-plane translations. In addition to the drawbacks mentioned above, this method requires access to temporal segmentation which is not routine in clinical

practice as typically only the end-diastole, and end-systole frames are segmented. Other studies make use of the contour points to align the slices, such as Wan *et al* [94], which use classical and generalized iterative closest point (ICP) algorithms to register SA and LA contours points. It should be noted that their data consists of a large amount of segmentations ranging from the apex to the atria, which is not standard clinical practice. Instead of aligning the contour points to each other, Medrano-Garcia *et al* [95], use a shape model of the LV to register the slices together. By placing fiducial landmarks on various slices, they first aligned all the images to the same patient specific coordinate system. The surface model was then fitted to the datasets at the epi- and endocardium SA segmentations, in a *goodness of fit* manner, with the underlying assumption that the best fit between all the contours represent the true position of the slices. The in-plane discrepancy between the individual contours and the model was then used to correct for the slice misalignment. The model was assumed to rely on an implicit stiffness parameter, which controlled the fitting between the contours, and as such is very dependent on this parameter, which can alter the results of the alignment.

Although geometrical approaches focus on obtaining smooth, non convoluted shapes, they have the disadvantage of forfeiting all the additional information available in the image values. To that extent, image registration techniques have been the most widely used approaches. In a slice-to-slice approach, Ector *et al* correct for misalignment using image registration by maximizing mutual information between consecutive SA slices [94, 96]. As previously mentioned in Chapter 2, as slice intensities are averaged over an 8 mm area, this method relies heavily on the assumption that there are no drastic changes between consecutive slices. To compensate for this, other studies use volumetric approaches where individual slices are registered to a 3D volume (slice-to-volume registration). Generally, a 3D volumetric dataset is acquired using SSFP, which is used to register the 2D slices in order to correct for the misalignment. However this type of acquisition is not clinically standard and due to its reduced spatial resolution requires robust image registration techniques. Chandler *et al* [97], use such a volumetric dataset

to which they register a stack of SA slices using a rigid transformation model and NMI as the similarity measure. Zakkaroff *et al* [98, 99], perform a similar study where a high-resolution angiography volume is used for registration with the SA stack. They propose an alignment strategy where the transformation optimization, which consists of only in plane translations and rotations, is a function of the entire stack. Consequently, the cost function is computed with respect to the whole image stack with all slice correction parameters contributing to the result. Their results report that in practice, using such a strategy, increases the accuracy of the metric calculation which in turn increases the robustness of the method. Lötjönen *et al* [100, 101] used a dataset consisting of both an LA and SA stacks to align the datasets together. The transformation of both image stacks was optimized over by maximizing a similarity measure at the intersection of the stacks, using only in- and out-of-plane translations. The choice of slice order can alter the results. Accordingly, [101] compared two approaches; a standard optimization method, and a stochastic one, where the slice being optimized is chosen randomly from the image stack. They report that the stochastic method improves the results. This approach further corroborates the findings in [98, 99]. Furthermore, by having acquired additional LA stacks, the method proved more accurate. In [102], a different data composition is used where the reference volume acquired consists of high resolution isotropic SA and LA single slices. A dataset consisting of a large number of SA images are registered to this dataset, using in-plane translations, to correct for the misalignment with the aim of building a high resolution temporal dataset. The slice-to-volume approach has proven to be popular to other areas such as the brain. For example, [103] used a reference functional MRI (fMRI) volume, acquired in the same scanning session, to align, using a rigid transformation model, an fMRI stack. The slice-to-volume methodology allowed [103] to compensate for any motion that might occur whilst using scanning protocols which require multi-slice acquisitions.

Unfortunately, typical cardiac clinical protocols do not involve the acquisitions of volumetric data, rendering the above approaches impractical. As mentioned in Chapter 2, clinical datasets are commonly composed of a stack of SA slices ranging

from the base to the apex, and 2 - 3 LA slices. To that extent, another approach in the literature uses the intersection line between the different slices to perform the alignment correction. This approach is based on using the image intensities at the slice intersections. By optimizing over the (dis)similarity between intensities on the intersecting line, optimal alignment can be achieved. Barajas *et al* [104], adopt this method to realign the slices with respect to each other by maximizing the mean value of the NMI measure evaluated at all the line intersections at the different planes. Their transformation model consists of in- and out-of-plane translations as well as a simultaneous update of all the slices to not bias the method to a particular slice. They choose to register the slices using only translation as they focus on the left ventricle. However, as the left ventricle can be highly cylindrical, if rotations are not accounted for, the slices could incur even further misalignment by translating out of plane, straying away from the "true" alignment of the slices. A similar approach has also been used for the brain [105], where a simultaneous update of all the slices is performed considering the line intersections between all the slice stacks, and updating collectively all the slices. However the dataset used in [105] is specific to the brain, where a denser acquisition can be obtained. Furthermore, they assume that the MRI contrast between slices are identical, allowing them to use similarity measures such as the MSE (see Section 3), which allows for a direct comparison of the intensities. This is not possible for cardiac applications due to the dynamic aspects of the heart which alter the contrast of the MR images. The method in [104] uses CINE acquisitions, however they only makes use of one of the temporal frames. In order to increase the image information, to help with the registration method, other studies make use of multiple frames [106, 107]. Slomka *et al* [106] use the intensity information from all of the CINE frames, by computing a similarity measure between the intensity of all the frames at their relative slice intersections. The particular method of aligning the slices by using line intersections between planes makes the LA planes the most influential with respect to all of the other planes as they incorporate line intersections with all the other planes, compared to the SA planes which only coincide with other LA planes. To that extent, Slomka *et*

al first realign the LA planes to themselves, which they then use as references to realign the SA planes. This is a sensible approach however one of the LA planes might be severely misaligned, which would bring the other "relatively well" aligned LA planes to another *common* alignment space with the LA plane, misaligning them further from the SA planes. Furthermore, [106] uses a cost function made up of a trade off between the cross correlation and sum of absolute difference of the intensities (SAD), which renders the registration process prone to the choice of the weighting parameters, although they select these weighting parameters empirically. Their transformation model consisted of SA in-plane translations, out-of-plane rotations, and LA in-plane translations, as they assume the LA translations correct the SA longitudinal rotations, although that is true only globally and not for individual slices. As the choice of cost function has a significant impact of the registration result, Ellen *et al* [107] present realignment results using different parts of the cost function used in [106], using in- and out-of-plane translations, as well as a constrained optimization approach. They compare the use of the weighted average of the absolute difference of the intensities and correlation, with each of the two parts individually. They report that the correlation part of the cost function does not add any significant impact to the combination of both similarity measures, and as such, SAD is only needed as a cost function. However, the data used in [107] does not make note of potential clinical scanning procedures which involve multiple continuous acquisition protocols such as CINE followed by LGE acquisitions. In practice, the contrast used in LGE acquisitions can have an impact on the slice intensities which would render the use of SAD as a cost function impractical. The methods in [106, 107] make use of the temporal frames from the CINE acquisitions to increase the robustness of their algorithms. However, a study on this approach's significance was performed in [108], who report that the use of multiple time frames does not meaningfully impact the registration algorithm, but increasing the number of LA planes improves registration and model reconstruction accuracy. These findings further validate that the structural information provided by the LA images is extremely valuable to the alignment results, and that methods relying only on SA

slices render finding the *truest* alignment between slices more difficult. It should be noted that when [108] evaluated the impact of increasing the time frames, they fixed the amount of LA slices at 6. As the acquisition of more than 2-3 LA slices is currently only for research purposes, the use of multiple time frames, using only the standard 2-3 LA planes, might have an effect on the alignment algorithm. Both [107, 108] use methods to obtain a region of interest surrounding the heart so that the registration may be driven only by the information from the heart. However, Carminati *et al* [40], use the global image information without the need of a precomputed ROI. Their method relies on optimizing in-plane translations using the mean cross correlation value between the line intersections intensities, constrained by a search space region of $\pm 10\,\text{mm}$ around the original SA position. They report that the use of an ROI may hinder the alignment process as lung volume differences cause displacement in other areas than the heart. These areas can contribute to the alignment process as the information obtained from the line intersection profiles of these areas may help drive the algorithm to a more true alignment.

The alignment strategy should be carefully taken account as it can alter the resulting alignment significantly. As aforementioned, one method is to use a fixed slice, usually an LA slice, to which all the other slices are aligned to. The drawback to this is the dependence on the choice of the reference slice, which can heavily influence the results if a particularly misaligned slice is chosen as the fixed reference. An alternative method is to alternate the reference slice in an iterative manner and allow the other slices to register to it, as per the approach of [98, 101]. However, trying to optimize the simultaneous transformation of all the slices, as per groupwise registration [109,110], can be costly and less efficient. If a rigid transformation model is used, the optimization space would consist of $6(n-1)$ parameters, with n being the number of slices. A better approach would be to fix all the other slices and allow free transformations of only one slice obtaining the best *global* alignment. Furthermore, using the sum of similarity measures from a slice to several reference slices can lower the influence of outliers, forcing them to converge through an iterative process; also known as alternate optimization [111]. Theoretically, different alignment strategies

could converge to the same minimum, regardless of the optimization parameters used. However, the presence of local minima means that, in practice, the choice of optimization parameters has a substantial effect on the result.

Recently machine learning has been used to tackle the process of slice misalignment [112]. Sinclair *et al* [112], use a fully connected neural network (FCN) to first segment the cardiac planes which they use to correct for the breath-hold motion misalignment. Their method makes use of the alignment strategy developed in Chapter 5, and published in [113].

4.2 Surface reconstruction from cardiac contours

Whilst simple heart models can be used for electrophysiological or mechanical studies, patient-specific meshes require more complex and accurate anatomical models [114]. These models are commonly characterized through 3D meshes, on which shape analysis or finite element methods can be performed. From a set of contours, or a point cloud, a surface mesh can be generated and act as an input to create a volumetric mesh. Depending on the complexity of the surface mesh needed, many parameters need to be taken into account—notably the trade-off between data matching and expected visual appearance of the resulting mesh [114, 115]. The prior art for mesh reconstruction is considerable, and an in depth analysis of current techniques can be found in [116–118].

According to [118, 119], surface reconstruction methods broadly fall under three distinct categories: region growing, computational geometry, and functional methods, all of which can be further divided into interpolating or approximating methods.

Region growing algorithms propagate information from all the points, starting from an initial seed. The most notable of these methods is the ball pivoting algorithm (BPA) [120]. The ball pivoting algorithm is an interpolating technique to generate an alpha-shape surface given an oriented point cloud. In 2D, given an initial seed point, a circle of a chosen radius iterates from point to point connecting them via edges. This process can be extrapolated to higher dimensions. This technique can

tend to fail for complex geometries and as they rely on point normals they are more suited for dense point clouds, with well defined normals.

Computational geometry methods explore all the surrounding neighbourhoods of all the sample points in order to find the best neighbouring candidate points for each of the samples. The most prominent of these techniques is the Delaunay triangulation. Delaunay triangulation is an explicit method that relies on finding a unique triangulation for all points given a set of points [118]. Lim *et al* [121] use Delaunay triangulations in order to reconstruct a cardiac mesh given a set of parallel SAX contours. By computing a connectivity tree between the set of contours, they project adjacent contours onto a 2D plane, where a Delaunay triangulation is computed between the contours, before projecting them back to 3D. Another technique when presented with sparse planar point clouds is to interpolate between the planes. Many algorithms exist to interpolate between contours; however, results are not reproducible between algorithms, which causes discrepancies in further analysis [122]. The method in [123] starts with an initial mesh and proceeds to iteratively refine the mesh. However, equivalently to [115], the input data must have coincidental points at the intersection for the method to work.

Functional methods tend to find a function that best approximates a surface given a set of points, via sets of functions. In [115], the surface mesh is constructed by optimising over the topology, using a level-sets approach, and having the genus as a control parameter. In order for the resulting mesh to represent the most accurate anatomical representation, an image volume is necessary. One of the drawbacks of the method is for the contours needing to have coincidental points at their intersections, which prevents the method from being used on standard CMR acquired data. To obtain uniformly distributed contour points, the authors in [124, 125] first fit a cubic periodic B-spline curve to the contours, allowing them to uniformly sample candidate points. The disadvantages of this approach is that it relies on the the number of control points selected, which in turn, smooths the overall mesh. The use of a template mesh in [126] can also cause the resulting mesh to be biased to the template. When presented with abnormal data, this

bias will cause the resulting mesh to be unable to retrieve the abnormal shape. In the Poisson surface reconstruction technique [127], an implicit function is fitted to the data using provided point normals, after which the reconstructed surface is obtained by extracting an isosurface. Similarly to the BPA algorithm, this method can tend to fail for complex geometries.

4.3 Misalignment correction and surface reconstruction integration

From the previous section, it is clear that there exists a clear need for an alignment method that is robust, contrast invariant and leads to accurate geometrical meshes that describe heart anatomy. Furthermore current methodologies involve subsequent iterative steps of first correcting the misalignment before creating a geometrical 3D model. Paiement *et al* [128, 129], demonstrate a technique that integrates the various steps of registration, segmentation and shape interpolation by using level sets. They make use of level sets to automatically segment and construct a 3D surface of the left ventricle and perform the alignment correction on the level set segmentation. The surface obtained through level sets used mean curvature flow (MCF) and radial basis functions (RBFs) to interpolate between the sparse data. By constraining the level sets to grow by using the combined information found in the intersecting slices, they allow the integration of the registration and level set's segmentation. It should be noted that although they report registration results, no segmentation validation was performed.

As clinicians contour objects of interest independently on each 2D slice, when placed in 3D, the contours are found to not be spatially consistent [113], which is further extended by any intra-inter segmentation variability. Using methods such as [128, 129], refrain from the need for human contouring, which is beneficial for large data set and allow for a global consistent segmentation.

4.4 Conclusion

It is worth noting that the methods described in Section 4.1.1 differ significantly in the registration process, notably the number of degrees of freedom used, as well as the cost functions. When comparing geometry based techniques over methods using image information, several things can be observed; primarily that geometry-based methods suffer from high variability in inter/intra segmentations, and that relying solely on the contours or using template/atlas based methods strays away from the real heart shape. Using techniques based on image intensities at slice intersections relies solely on the data and does not suffer from the aforementioned issues. By optimising over the (dis)similarity between intensities on the intersecting line, optimal alignment can be achieved.

From all the studies trying to correct for motion misalignment between slices, several key points surface. First, the use of LA information is vital to any alignment strategies. Second, a rigid transformation comprising of 6 degrees of freedom (in- and out-of-plane rotations and translations) should be used to try to obtain the best alignment globally and individually. Third, relying on the image information as much as possible will allow for the most *personalized* alignment possible, which renders the slice-to-slice line intersection method the most suitable. Fourth, a cost function should be chosen carefully that makes use of image intensities, but more importantly that allows for the comparison of image features. Fifth, a robust optimisation strategy is needed that captures the problem of slice inter-dependency. Sixth, the use of multiple frames is preferred, although not necessary. However the end-diastole frame should be used for realignment if only one frame is to be chosen as it has the most minimal motion variation. Finally, an *a priori* computed ROI is not necessary for the alignment strategy.

The work in Chapter 5, and published in [113] has been developed with the aforementioned points in mind.

Concerning surface reconstruction, many approaches have been developed in order to produce surface reconstructions from planar contours; however, current techniques do not support data formed by parallel or quasi-parallel noisy contours,

with or without cross-sectional information, that may or may not be coincidental. Furthermore, as CMR data is quite sparse, since each of the SAX contours lie around 10 mm apart, a method is needed that works robustly on non-dense datasets. Furthermore, a method relying solely on the input data without the need for normal calculations or for the contours to be coincidental is needed for clinical integration. The work in Chapter 6 and published in [87, 130, 131] has been developed to that extent.

The literature on slice misalignment correction made apparent that the most accurate and robust methods to obtain personalized models, combined both image information as well as the global geometry of the heart in one integrated manner. As such, the method developed in Chapter 7, combines both works found in Chapter 5 and 6 to integrate slice misalignment correction as well as segmentation correction to obtain smooth, spatially consistent, 3D personalized cardiac models.

5

Correction of Slice Misalignment in Cardiac MRI

Contents

5.1 Introduction 73
5.2 CMR data 77
5.3 Methodology 77
5.3.1 Alignment score and optimization strategy 78
5.3.2 Similarity measure 80
5.3.3 Contour-based alignment 83
5.3.4 Label-based Alignment 83
5.4 Results 85
5.5 Discussion 87
5.6 Conclusion 91

5.1 Introduction

This chapter presents an extension of the work that has been published in [113], and represents the first original research contribution of this thesis. Before generating anatomically realistic meshes of the heart, the first step is to correct for any slice misalignment that might occur, disrupting any spatial coherence between the slices, and ultimately their segmentations. As mentioned in Chapter 4, several aspects must be considered to align cardiac slices. By using the image intensities between the

cross intersecting planes, image registration methodologies can be applied to realign the slices together. This approach is superior to geometrical approaches (described in Chapter 4), which rely on aligning the slices based on the segmentations, as the alignment procedure is image driven. Segmentation based realignment requires the availability of accurate contours, and has the disadvantage of forfeiting all the additional information available in the image values. In practice, as it will be shown in this Chapter, due to the inter/intra variability in expert contouring, the final alignment will depend on the expert; on the other hand, if the application requires the smoothest ventricular shapes possible, methods such as [92] can be used. Furthermore, geometrical approaches can lead to the banana problem [90]. Another method to correct for slice motion misalignment exists by registering the images to either a slice or volumetric template dataset. However this causes a biasing issue, where the choice of the template ultimately guides the registration process, which might perform poorly when dealing with pathological data. In volumetric approaches, a low resolution SSFP dataset of the patient's heart might be obtained prior to the slice-by-slice acquisition and used to correct for any spatial motion. However, this is not standard clinical protocol, and might cause the registration algorithm to be subject to a high amount of local minima. The method presented in this Chapter was developed to makes use of datasets obtained in typical clinical protocols. Furthermore, as the algorithm is fully automatic, it can easily be used in the clinics. By using the line intensities at the slice intersections, information from both SAX and LAX are used. After a clinical onset, such as MI, at least two MRI acquisitions protocols are usually used: Cine and LGE. As LGE uses a contrast agent, highlighting any infarcted area or scar tissue, there may be a high intensity variability between slices. The method should therefore work on both protocols, where the intensity ranges and differences between protocols is significant. As such, a suitable cost function must be chosen. Since the application of the cost function deals with same-modality images, cost functions such as mutual information or NMI are unnecessary. As the main objective of the alignment strategy lies in matching the same information content in two intersecting slices,

the cost function should take into account a measure of similarity between the intensities of the images, and its underlying features. To that extent, our method introduces a new cost function designed to be applied on both Cine and LGE images. The energy function relies on the extraction of image features, via the use of local phase, making the method purely data driven. In the image model presented in [50], it can be seen that images consist of both amplitude and phase information, and that the local phase encodes the main structural information (as shown in Chapter 3, Fig. 3.4). This allows the removal of any intensity inhomogeneities that might cause issues during the registration process, and instead focuses on matching similar structures. Furthermore, through the use of the log-Gabor filter, image features at different scales can be selected, allowing the registration process to match features of different frequency components.

One of the main issues in image registration for the correction of slice misalignment is the choice of optimization strategy. There is an inherent bias with the registration strategy. In sequential approaches, the registration result might be guided differently depending on the choice of the slice being registered to the other images. Techniques such as groupwise registration have been developed which consists of computing the optimal transformation for all slices, and updating them simultaneously, as to avoid any particular slice bias. However, this causes the optimization function to be dependent of a large number of parameters, causing the process to be costly. As typical clinical cardiac MRI datasets contain around 10-15 slices, each optimization step would be covering between $60 - 90$ parameters, increasing the non-convexity of the function. In the specific problem of inter-slice cardiac motion correction, the registration process is inherently biased towards the LAX, as they intersect with all the other slices, whereas the SAX only intersect with the LAX. Some studies [108, 132] use a fixed slice, usually an LAX, to which all the other slices are aligned to. The drawback to this method is the dependence on the choice of the reference slice, which can heavily influence the results if a particularly misaligned slice is chosen as the fixed reference. This can be seen in [108], which uses the information-rich 4CH LAX as the fixed slice, where some

cases still report misalignment of up to 40 mm. An alternative method is to alternate the reference slice in an iterative manner and allow the other slices to register to it individually. This has the negative effect of being highly influenced by outliers. Instead, our method fixes all the other slices and spatially transforms each of the slices individually, obtaining the best global alignment for that slice, with respect to all the other slices. In this Chapter, a specifically designed optimization strategy is presented taking these issues into account. By optimizing a slice's spatial position using the sum of energies between the slice and the other reference slices, the influence of outliers is lowered, forcing the dataset to converge through an iterative process; also known as alternate optimization [111]. Theoretically, different alignment strategies could converge to the same global minimum, regardless of the optimization parameters used. However, the presence of local minima means that, in practice, the choice of optimization parameters has a substantial effect on the result.

In some works that make use of the line intersections to correct for misalignment [106–108], the cost function is composed of an energy evaluated at different frames. This is possible in Cine protocols which acquires different sequences of the beating heart, but not in typical LGE acquisitions, which focus on highlighting an infarcted region. As LGE datasets are used in this Chapter, only one frame was used.

The algorithm presented in this Chapter is applied on real data, and validated using segmentation distances (contour-to-contour distances). However, as segmentations, even from manual experts, suffer from intra- inter- variability, a *true* ground truth validation cannot be provided. To that extent, it will also be shown that contour-to-contour distances suffer from spatial inconsistencies as they are contoured in 2D and do not make use of the 3D coherence of the dataset. The main contributions of this work are:

- The development of a complete framework for the correction of translational and rotational misalignments between CMR slices, based on the combination of short axes (SAXs) and long axes (LAXs) images and applicable to Cine and LGE CMR scans.

- The introduction of the normalized cross correlation of local phase vectors as a similarity measure that makes the method applicable to different CMR protocols.
- A complete validation using manually traced contours, which includes the estimation of contouring errors.

5.2 CMR data

The datasets used in this Chapter is part of the Oxford Study in Acute Myocardial Infarction (OXAMI), taking place at the John Radcliffe Hospital, Oxford, UK. The study has several aims; primarily to better understand the underlying mechanisms causing MI, as well as identifying cardiac biomarkers involved in the remodeling processes of the heart. The study involves scanning patients having ST-elevated MI at multiple time periods: 12-48 hours, 5-7 days, 14-17 days and 6 months. Patients underwent multiple scans using a SIEMENS TrioTim 3 Tesla MRI scanner. Only the first scan, during the acute stage, and the 6 month scan, after significant remodeling occurred, included a full image analysis with both Cine and LGE imaging. The data obtained consists of DICOM files of 2D Cine and 2D LGE sequences of size 216 x 256 pixels (approx: 1.41 mm x 1.41 mm). Each dataset contains between 9 and 13 SA slices from apex to base separated by 8/10 mm and 3 LAXs slices (4CH, 2CH, 3CH). In total 15 LGE datasets were used in this study. For validation purposes, left ventricle epi- and endocardial contours were manually traced by an expert on all SAXs and LAXs using the CMR42 software (Circle Cardiovascular Imaging, Calgary, Canada).

5.3 Methodology

Our method relies on the intensity profiles at the line formed by the intersection between two slices to align the slices together and give spatial coherence to the 3D dataset. This is based on the assumption that two slices will be perfectly aligned when the underlying features of the line profile at their intersection is complementary.

Assuming that slices from a subject are triggered at the same cardiac phase the 3D shape/anatomy of the heart remains fixed among slices. Therefore, rigid-body transformations between slices are considered. As such, we perform rigid registration for each of the slices over 6 degrees of freedom of a rigid-body transformation.

The optimization process needs to be constrained as the problem should avoid slices straying significantly from their initial positions, or merging into one another. For each of the optimization step, the constraints were set to $\pm$ 10 mm for in-plane translations and $\pm$ 3 mm for out-of-plane translations. As for rotations, the constraints where set to $\pm$ 5 degrees on the Euler angles along the X, Y and Z axis. As mentioned in Chapter 2, the study performed by McLeish *et al* [20] described rotations having a low influence on the misalignments, and as such we assessed a 10 degrees rotational freedom to be sufficient. The constraints were set as a function of the optimization, using a hyperbolic transfer function, to allow for smooth constraints. Every step of the optimization was constrained given its current value. In this way, the capture range at every optimization step is constrained, however, given enough iterations, the global spatial transformation of the slices are not constrained. As we assume for slices to be near their optima, having this optimization methodology allows for slices to reach optima in their vicinity, but still reach other far away optima. Even though most datasets are closely aligned, this approach is necessary as we observed in our own work, as well as in other work [108], that in-plane motion could range up to 20 mm, in extreme cases. An overview of the algorithm, given as pseudo-code, can be seen in Alg. 1.

5.3.1 Alignment score and optimization strategy

The global motion (GM) discrepancy of the slices can be measured as the sum of similarity measures E between pair of intersecting slices. Let S_i be the i-th slice and Θ_i the set of its 6 rigid transformation parameters, for example $\Theta = \{t_x, t_y, t_z, \alpha, \beta, \gamma\}$ (in our case 3 translations plus 3 Euler angles). Moreover, let $S_i^{\Theta_i}$ be the transformed version of the slice S_i by the rigid transformation

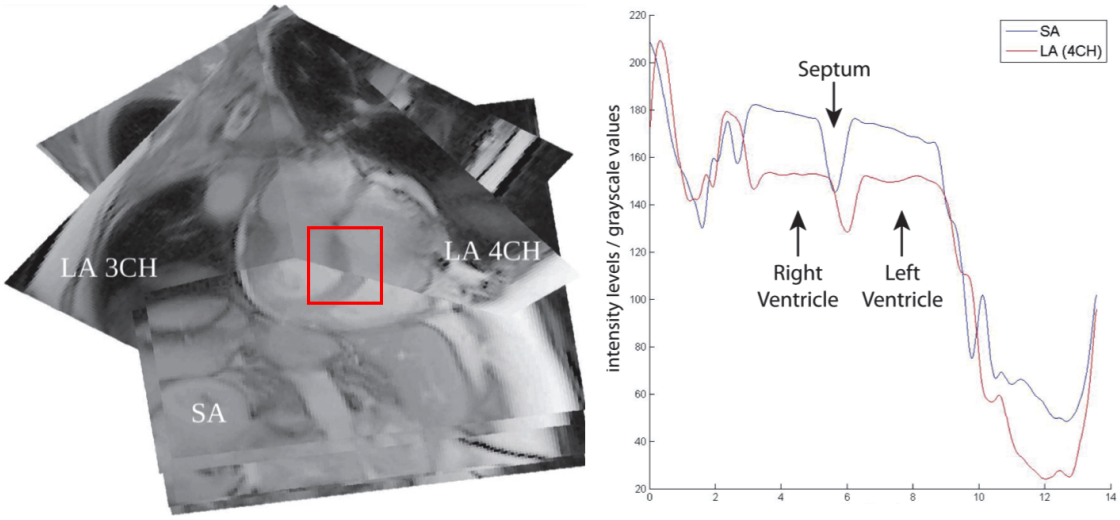

Figure 5.1: MRI slices with their respective inter-sectional Line intensities. Left panel shows some 2D MRI slices in their spatial 3D positions, where a clear misalignment can be shown (red box). Right panel shows the intensity profiles along an intersection line.

defined by parameters Θ_i. By using this notation, the global motion (GM) discrepancy is given by

$$GM\left(\Theta_1, \Theta_2, \ldots, \Theta_n; S_1, S_2, \ldots, S_n\right) = \sum E\left(S_i^{\Theta_i}, S_j^{\Theta_j}\right) , \tag{5.3.1.1}$$

where the summation runs over all pairs S_i, S_j intersecting in a line, *i.e.* a LAX will intersect with all the SAXs and the other LAXs. Minimizing GM is akin to finding the parameters for each slice. In this work, this minimization is performed in an alternate manner, by optimizing the parameters Θ of a single slice whilst leaving the others fixed. As GM is built as a sum of terms, the iterative minimization of partial terms results in the minimization of the global motion discrepancy and the slices match together.

There exists several optimization techniques such as gradient descent, Newton optimization or its approximate versions leading to quasi-Newton methods, several strategies for non-linear conjugate gradient techniques, or Trust regions, where the more representative is the Levenberg–Marquardt algorithm [133]. While in theory, all previous methods will achieve a local minima, or more correctly, a stationary point, they differ in practice, with some finding the global minimum faster for example, or only being used if certain conditions are met such as having a cost

function twice differentiable. All the aforementioned techniques require a gradient, where the optimum transformation is found through an iterative process. One of the most widely used optimization method is gradient descent, however the conjugate gradient methodology is generally accepted as a better alternative as it tends to have a faster convergence rate, since it uses prior information from its previous update calculations. Conjugate gradient is used as the default optimization methodology of choice throughout this Thesis, due to it's good results and easy implementation. The conjugate gradient strategy can be described as finding the optimal transformation parameters $\mathbf{p}^\star$ such that the cost function (or energy) is minimized. This is achieved through a succession of steps such that $\mathbf{p}_{i+1} = \mathbf{p}_i + \alpha \mathbf{d}_i$, where α represents the step size and $\mathbf{d}$ the search direction. The parameter α was automatically calculated using a backtracking line-search, which tries to obtain the best step size such that it is not too small or too large [134]. In the conjugate gradient methodology, $\mathbf{d}_i = -\mathbf{g}_i + \beta \mathbf{d}_{i-1}$, with $-\mathbf{g} = \frac{\partial \mathcal{C}}{\partial \mathbf{p}} \mathbf{p}_i$, β the conjugate gradient parameter and $\mathcal{C}$ the cost function being optimized. There are several formulae to compute β, as defined by Fletcher and Reves, Polak-Ribiére-Polyak, Liu and Storey among others. The optimization used throughout this Thesis uses the Polak-Ribiére-Polyak formulae to evaluate β, as it has been shown to outperform the other methods [135]. It is defined as $\beta_i = \frac{y_i^T \mathbf{g}_{i+1}}{\mathbf{g}_i^T \mathbf{g}_i}$, with $y_i = \mathbf{g}_{i+1} - \mathbf{g}_i$.

5.3.2 Similarity measure

The particularities of image acquisition described in Section 5.2 prevent the use of simple measures based on intensity differences, such as SAD/SSD.

For images with high intensity disparity, feature-based similarity measure can provide a more robust method of assessing similarity between the images at the intersection profiles. As local phase is independent of contrast and not affected by intensity inconsistencies, it is a sensible choice as a similarity measure. The local phase can be obtained through the analytic signal in 1D, and general extensions to higher dimensions have been proposed including the use of oriented filters or the monogenic signal [136]. Although our registration is based on the similarity

between the 1D intersecting line profiles, we compute the local phase for an entire image, and then obtain the line intersection between two local phase images as it results in less noise in the line profiles [136]. The normalized cross correlation between two profiles is used as similarity measure E for an intersecting pair. In 2D, the local frequency information can be obtained by convolving the images with banks of quadrature pairs of log-Gabor filters [48]. A quadrature filter is a complex valued function which transforms a real valued signal to an analytical signal with weighted frequency components. Convolving the image with a filter will result in response vectors encoding phase and amplitude. By using quadrature filters, local phase (Φ) can be estimated by

$$\Phi = \arctan\left(\frac{I_q}{I_p}\right), \tag{5.3.2.1}$$

where I_q represents the magnitude of the odd filters convolved with the image, and I_p the filter response. With respect to 2D, quadrature filters can be generalised through directional formulation [137]. This results in the image shown in Fig. 5.2, which represent an example of an aligned dataset, with its respective local phase images.. Once the local phase images are obtained, the line intersection at both images is taken

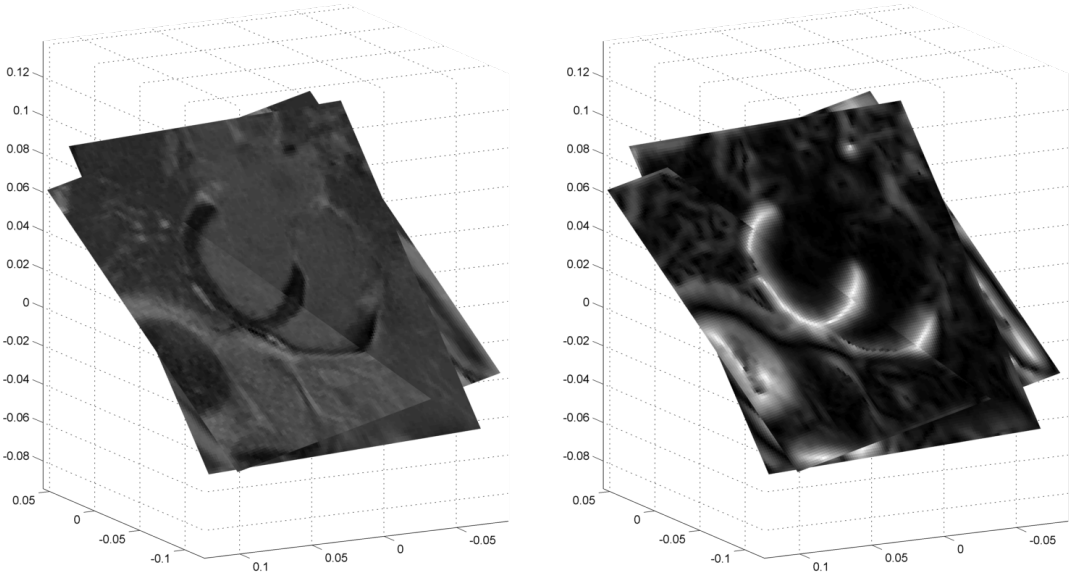

Figure 5.2: Late gadolinium image on the left and local phase image on the right

to account for E measure, using NCC. To allow the registration process to match

features at different scale, the optimization process was performed in a multi-scale approach by calculating local phase images with a different wavelength. We fixed σ to be $\sigma = 0.55$, which we chose empirically, and varied the wavelength $5 \leq w \geq 20$. The effect between the wavelength and sigma choice can be seen in Fig. 5.3.

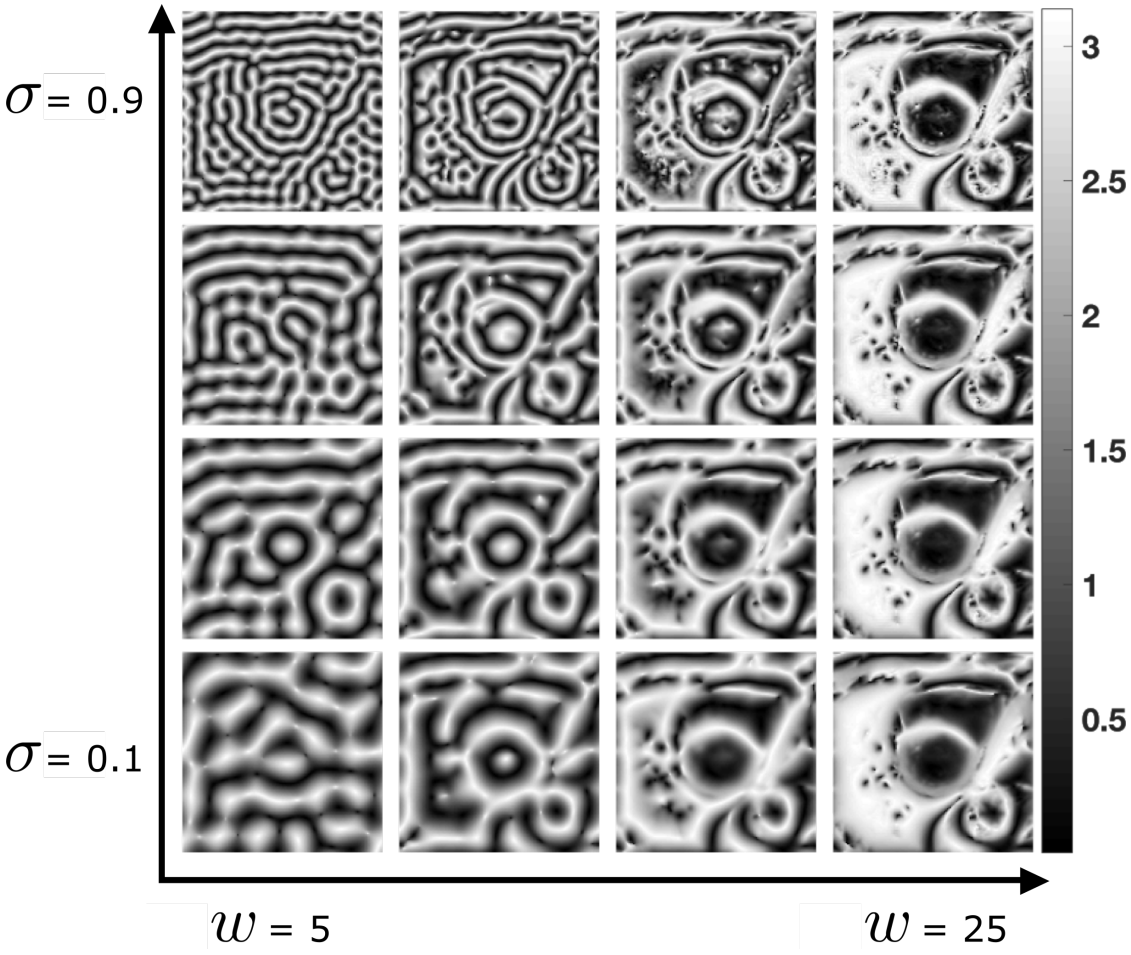

Figure 5.3: Local phase effects of varying the wavelength (w) and sigma (σ) parameters on a cardiac SA image. The colors represent the changes in structure given in radians.

The local phase returns angles representative of features, and as such obtaining the line profiles of the local phase of the images results in a vector of angles. We used the NCC to compare the two vectors as using more traditional methods, such as, the direct comparison of the two vectors would correspond to comparing two angles directly. There exist other approaches for comparing local phase vectors such as the local phase mutual information metric as suggested by [51], however as angles live in a circular metric space; that is, the difference between an angle

of 350° and 10° could either be 20° or 340°, another method would be to use tools developed in the field of circular statistics. A rigorous comparison between angle vectors is performed in [138], using the circular correlation coefficient, but we chose to use the simplified version of the NCC as we estimate the local phase vectors to be somewhat close and therefore comparable to each other. Although the angles live in a circular space, to some extent, locally they can be seen to live in a plane, which makes metrics such as NCC acceptable to use as a local approximation.

5.3.3 Contour-based alignment

Due to inconsistencies in manual contouring of the SAXs and LAXs slices, a "perfect" alignment (i.e. one in which the distances between SAXs and LAXs contour is zero) might not be achievable. In order to establish a baseline error value, we applied a misalignment correction algorithm using the contours from manual images, substituting the energy function E in Eq. (5.3.1.1) by the Euclidean distance (D) between SA and LA contours as shown in Eq. (5.3.3.1):

$$\Theta^{\star} = \text{argmin} \sum D(C_i^{\Theta_i}, C_j^{\Theta_j})^2 \,, \tag{5.3.3.1}$$

where $\Theta^{\star}$ represents the optimum translation and rotation parameters for the contour being spatially optimized. The distance was obtained by taking the respective contour points at the intersection lines between the contour planes. As one intersection between two contours give two contour distances (apart from the LAXs - LAXs interactions), we optimized over the sum of the squares of all of the distances, to allow for a global correction between all contours. We iterated through all the dataset until convergence; which we defined as reached when the global energy between contours was improved no more then $1e^{-5}$. We refer to this algorithm as "contour-based" to differentiate it from the "image-based" algorithm we propose here.

5.3.4 Label-based Alignment

To allow a quantitative evaluation of the methodology with respect the the line intensities (compared to the contour-based alignment strategy), labeled images were

generated from the contours. For each dataset, a labeled image (having the same size as the original image) was generated for each of the contours. Different labels were used for the epicardium and endocardium contours. The myocardium was given a label of 2, and the blood-pool a label of 3. The rest of the image (having no information) was set to zero. An example of a labeled dataset can be seen in Fig. 5.4.

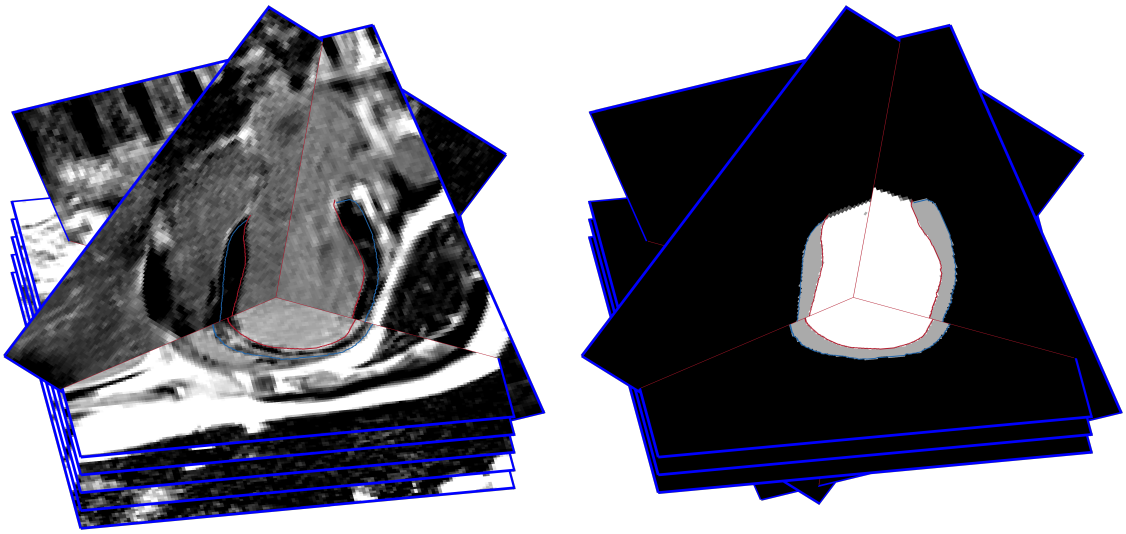

Figure 5.4: Example of a labeled-based dataset. The left image shows the original LGE dataset. The right image shows the labeled image generated from the contours.

This experiment was performed to assess the algorithm's ability to align images, given distinct labels (removing all possible noise). The same alignment process as the image-based one was performed on the labeled images, without using the local phase. As the optimization focuses on matching pixels with the same labels, only the NCC of the line intensities was used. As such, the evaluation criteria E (as per eq. 5.3.1.1) is defined as:

$$E\left(S_i^{\Theta_i}, S_j^{\Theta_j}\right) = \frac{1}{N}\sum \frac{(I^i_{l_{ij}} - \overline{I^i_{l_{ij}}})(I^j_{l_{ij}} - \overline{I^j_{l_{ij}}})}{\sigma I^i_{l_{ij}} \sigma I^j_{l_{ij}}}, \tag{5.3.4.1}$$

where l_{ij} represents the intersection line between the two slices S_i and S_j, such that, $l_{ij} = S_i^{\Theta_i} \bigcap S_j^{\Theta_j}$, and N the total number of points sampled from the line l_{ij}, where the summation is performed across all the samples. I represents the image slice intensities, such that, $I^i_{l_{ij}}$ represents the line profile of slice S_i, evaluated at l_{ij}, at the intersection of the two slices S_i and S_j.

Algorithm 1 Slice alignment correction algorithm

Parameters:

- $\sigma = 0.55$ (bandwidth of log-Gabor filter)
- w (center frequency of log-Gabor filter. Range: $5 \leq w \leq 20$)
- Θ (Transformation parameters)
- N (Number of slices)

1: **if** Cost Function == Local Phase **then**
2: **for** all slices **do** $LocalPhase(Slice)$ with parameters σ, and w for a particular scale
3: **for** $S_i^{\Theta_i}$, with $i = 1$ to $S_N^{\Theta_N}$ **optimize** GM on Θ, (as per Eq. (5.3.1.1)), such that all S^{Θ} are fixed, except $S_i^{\Theta_i}$
4: **if** Cost Function == Local Phase **then** compute E as per Eq. (5.3.4.1), where the intensities I represent the local phase vectors.
5: **if** Cost Function == Label **then** compute E as per Eq. (5.3.4.1), where the intensities I represent the binary labels.

5.4 Results

We investigated the performance of the registration algorithm using the NCC of the local phase signals. The algorithms were used on 15 LGE datasets, each containing 3 LAXs and several SAXs. The image based algorithm was run for 9 iterations, for each dataset, in order to evaluate the convergence of the registration. This number was chosen empirically, as the algorithm appeared to converge by then.

Table 5.1 shows the mean, median, and standard deviation resulting from contour-to-contour distances before the alignment and after, using the image-based, labeled-based and contour-based methods. The upper part of the table represents all of the 682 individual contour-to-contour distances for all the datasets. The lower part reports the values for the 87 contours that were deemed significantly misaligned ($\geq$ 5 mm).

Figure 5.6 shows a example of a labeled-based image alignment. The left image represent the dataset before alignment, and the right image represents the dataset after applying the algorithm using the NCC on the line intensities. It can be seen that after applying the algorithm, the endocardium contours do not coincide fully.

Figure 5.7 shows the different results for the image-based, labeled-based and contour-based algorithms. The resulting meshes built from their respective contours

(in mm)	Before	Image based	Labeled based	Contour based
Median	2.19	1.56	1.09	0.31
Mean	2.82	1.93	1.43	0.46
Std	2.48	1.66	1.30	0.47
Median	7.33	1.96	1.29	0.29
Mean	7.73	2.36	1.68	0.43
Std	2.43	1.91	1.57	0.44

Table 5.1: Table representing the mean, median, and standard deviations of the contour-to-contour distances (in mm) between all the contours for image-based, labeled-based, and contour-based alignment corrections, on 15 LGE datasets. Top: Results before and after the alignment for all of the datasets (682 distances). Bottom: Results before and after the alignment for significantly misaligned slices - ≥ 5 mm (82 distances).

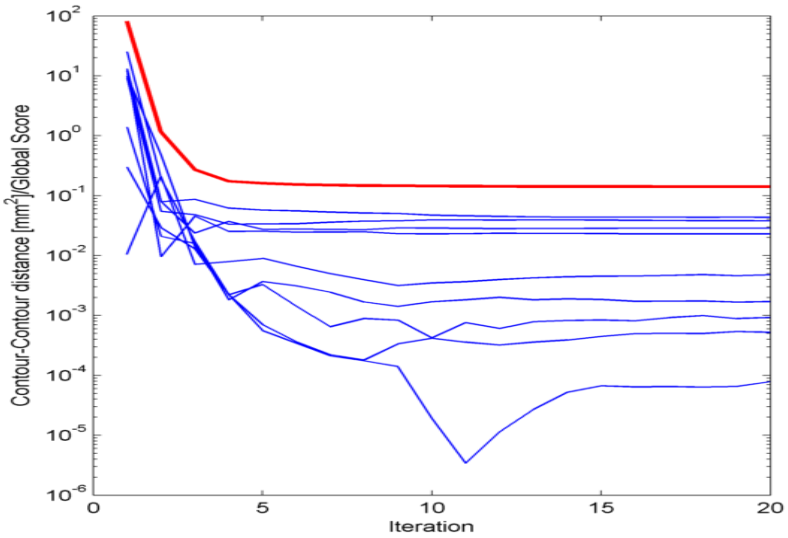

Figure 5.5: Pairwise energy E between slices (shown in blue) over several iterations, for the image based alignment method. The red line shows the global energy GM.

can be seen in the bottom row. For this experiment, the mean, median, and standard deviation of the original misalignment values (calculated from the contour-to-contour distances) was equal to 2.32 mm, 1.88 mm, 1.95 mm, respectively. After applying the image-based local phase algorithm, the mean, median and standard deviation contour-to-contour distances became 1.01 mm, 0.92 mm, 0.72 mm, respectively. The labeled-based alignment resulted in 1.4 mm, 0.87 mm, 0.20 mm, and the contour-based strategy resulted in 0.20 mm, 0.17 mm, 0.18 mm, respectively. It can be seen that the resulting image-based and labeled-based meshes have similar characteristics. The contour-based mesh appears the smoothest of all of the results, however as the

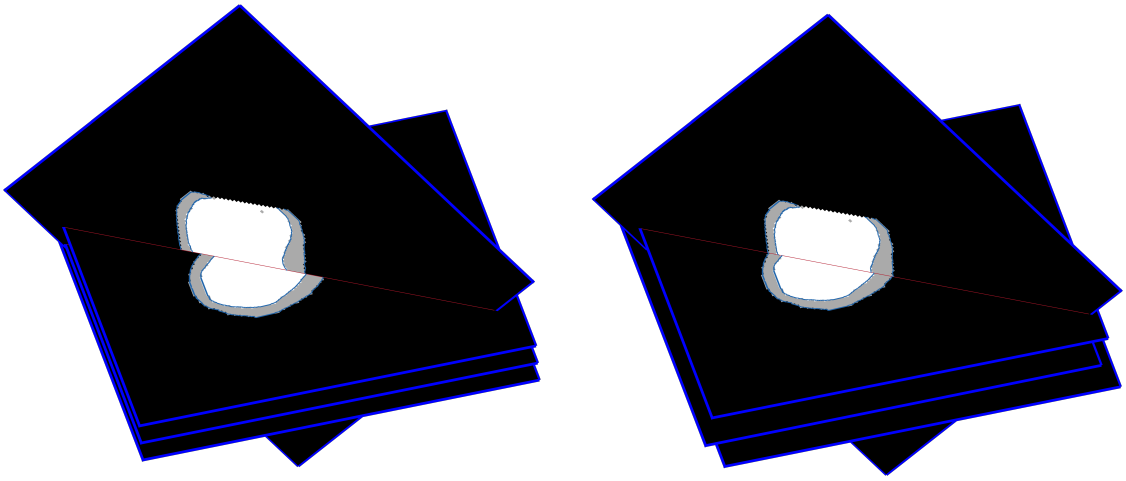

Figure 5.6: Example of a binary-based image alignment on an LGE dataset. The endocardium and epicardium contours were converted to labeled images and aligned using the same image based alignment methodology. The left and right images represent the misaligned and aligned dataset, respectively. It can be seen that after the realignment process, the endocardium contours do not coincide, due to discretization errors in the contouring.

contours have been optimized independently as to obtain the minimum amount of discrepancy, if positioned at the same spatial locations as either the epicardium or endocardium contours, the LGE images do not match, and appear misaligned.

Figure 5.8 shows an example of a segmentation discrepancy between two intersecting images due to the papillary muscles in the LAX. The red, blue and green points on the images represent the points of intersection between the slice and the left ventricle endocardium, left ventricle epicardium and right ventricle endocardium contours of respective slices, respectively. The blue squares highlights a segmentation error between the two images.

5.5 Discussion

Our method relies on using all available information to correct for misalignment by using both the long and short axes slices, in an iterative process. By taking into account all intersections simultaneously, we minimize the effect that individual outliers have on the overall results, forcing these outliers towards the global "consensus" position. By using local phase, we rely on its invariance to changes of intensities in order to focus on obtaining high feature similarity rather than high

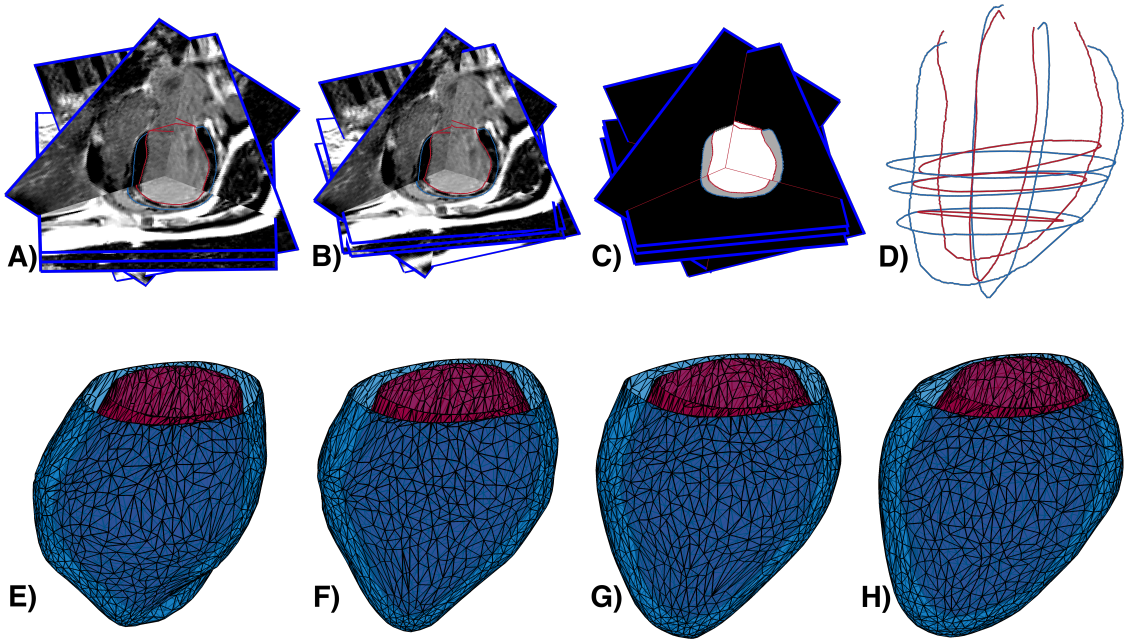

Figure 5.7: Motion correction results for an LGE dataset. A) Misaligned images. B) Alignment result using the image-based method. C) Alignment result using the labeled-based method. D) Alignment result using the contour-based strategy. E) Mesh generated using the contours from A). F) Mesh generated using the contours from B). G) Mesh generated using the contours from C). H) Mesh generated using the contours from E).

intensity profile similarity. Relying on the latter assumes that the images to be registered contain similar intensity profiles which is not always the case.

The slice optimization was constrained to compensate for potential structural symmetry, which might lead to out-of-plane misalignment such as could be the case with midventricle short axis slices. A hard constraint was chosen as an alternative to introducing a regularizer as the latter would induce a bias towards the initial positions of the planes.

The median and mean values for all 15 patients can be observed to decrease notably, however our method distinguishes itself more with regards to the remarkably misaligned slices. We have defined significantly misaligned slices as slices having more than 5 mm between contours. Results show that the median value before the alignment is 7.33 mm (7.73 mm mean). After the alignment, using the image based methodology, a median of 1.96 mm was obtained (2.36 mm mean). When correcting for the motion misalignment using the labeled datasets, the resulting values can be seen to fall under sub-pixel accuracy with a median of 1.09 mm (1.43 mm mean). By

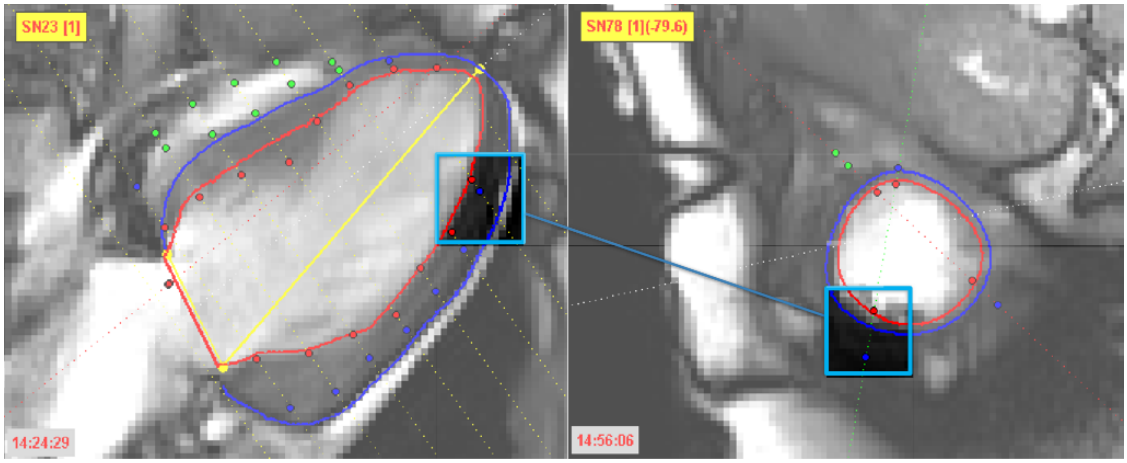

Figure 5.8: Example of a segmentation error between two intersecting slices. The red, blue and green points on the images represent the points of intersection between the slice and the left ventricle endocardium, left ventricle epicardium and right ventricle endocardium contours of respective slices, respectively. The yellow dashes represent the lines of intersection between the image and the other slices. In such a way, the misalignment between the different slices can be observed. The blue squares highlights a segmentation error between the two images.

minimizing the contour-to-contour distances only, results show that zero minimal distance is unobtainable, due to the variability in contouring. As such, the lowest median obtainable with our contours was 0.29 mm (0.43 mm mean), which should be accounted for when looking at the image based results. The aim of using the labeled dataset was to assess the algorithm's ability to realign slices using our optimization procedure. As the line intensities take on step-like characteristics, optimizing over the NCC allows to obtain a good match between the labels. However as the labels are discretized due to the contours having sub-pixel values, a perfect match between labels cannot be obtained. Figure 5.6 shows an example of a labeled dataset before and after alignment. It can be seen that the severe misalignment between slices has been corrected, however, some error between the contours can still be observed. This error can also be seen to occur between the labels. From the results it can be seen that the results from the image-based methodology are quite similar to the labeled-based results. If the contour-based results are used as a lower bound on the possible quantitative measurement, it can be said that the image-based results fall under sub-pixel accuracy. Visually, all the dataset appeared well aligned.

Oscillations along the iterations of the alternate optimization are a concern due to the slice dependencies. This can be observed in the behavior of the global energy at the different iterations (see Fig. 5.5). It can be seen empirically that 10-20 iterations are enough to converge. GM can be seen to decrease and follow a descent path without any jumps occurring, however the evolution of the pairwise energy can result in a non-monotonic (non-descent) path. Even after running the algorithm for more then 100 iterations, some energy pairs will continue oscillating. It can also be seen that for a given pair, the lower the energy, the more it oscillates.

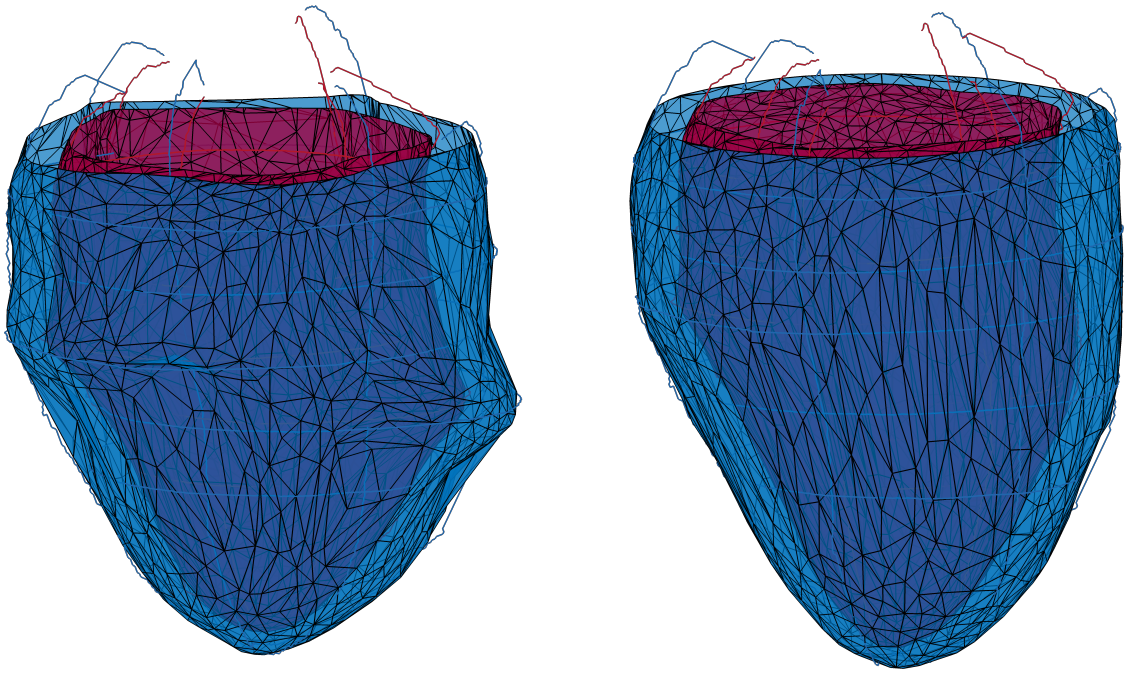

Figure 5.9: Surface reconstruction before and after alignment.

Validation for cardiac image processing comes with some important drawbacks that need to be taken into account. There typically exist no true ground truth due to the nature of the problem. In the majority of cases the ground truth relies on quantification values based on clinical segmentation. However expert segmentation should not be considered as the true ground truth, but more as an approximation, as it suffers from inter-intra expert variability. Furthermore the choice of quantification methods is highly variable and can greatly impact the results. After optimizing the contours individually in a non-constrained manner, the results show that some discrepancy still remains. This value can be seen as the segmentation variance

caused by the discretization of the contouring. Figure 5.8 shows an example of a contouring error between a LAX and a SAX slice. The SAX contour shows a thin myocardium, whereas the LAX contour displays a wide myocardium. It can be seen in this example that if the contours were optimized using the contour-based method, a discrepancy of zero mm could not be achieved. In most of the cases, the papillary muscles being contoured is the reason behind the contouring errors, as they can have similar gray intensity values as the myocardium.

It can be said that correcting for any motion misalignment between slices is indispensable for any cardiac anatomical reconstruction which is clearly shown in Fig. 5.9. Several artifacts can be observed such as non aligned SAX contours causing a "waving" surface. Furthermore spatial discrepancies between SAX and LAX produces depth fissures and ridges of the surfaces and discrepancies between LAX do not allow for a good reconstruction of the apical region.

5.6 Conclusion

We have presented a phase based registration algorithm that corrects for the misalignment of cardiac MRI images. We apply our algorithm to LGE images, as other works mostly focus on Cine acquisitions. A validation methodology was introduced via the use of labeled images, and by minimizing the contour-to-contour distances. We gave a quantitative evaluation of the methodology for 15 different LGE datasets, and showed that the image-based methodology helped correct for drastic misalignment between slices. By optimizing the contour-to-contour distances to obtain the minimum discrepancy range possible, we acquired the minimum quantitative bounds possible, as some error between contours still remained.

6

Geometrical Representation of Cardiac Anatomy

Contents

6.1 Introduction 93
6.2 Materials and Methods 95
6.2.1 Synthetic and Real Data 95
6.2.2 Initial Meshes 97
6.2.3 Mesh Deformation 98
6.3 Results 100
6.3.1 Simulation of CMR Acquisitions 101
6.3.2 Synthetic Data 104
6.3.3 CMR Acquired Data 106
6.4 Discussion 114
6.5 Conclusions 118

This Chapter presents the second major contribution of my research work, and has been published in [130] and [87] and presented in [131]. Once any slice misalignment has been corrected, a 3D representation of the heart can be generated by making use of the segmentations of the images in the dataset. However, any correction algorithm encompassing in- and out-of-plane rotation and translation, will cause the contours to loose their original, well organized, spatial structure. Furthermore, cardiac MRI data, is quite sparse, and their segmentations are generally non-spatially coincident (as discussed in Chapter 5). This Chapter

introduces a method to deal with such data: sparse, non-coincidental, non-parallel, heterogeneous, cross sectional contours.

The tool has been made available on the following GitHub page: `http://github.com/BenVillard/surfco`.

6.1 Introduction

In recent years, much research has looked at creating personalised 3D anatomical models of the heart [115, 122, 139, 140]. These models usually incorporate a geometrical reconstruction of the anatomy in order to understand better cardiovascular functions as well as predict different processes after a clinical event. Besides being fundamental to shape analysis, they are a prerequisite to electrophysiological and mechanical simulations with finite element analysis tools [141]. Furthermore, personalized shape models are increasingly used for surgical planing via 3D printing or device innovation [142, 143]. However, the ability to accurately reconstruct heart anatomy from contoured magnetic resonance images (MRI) in three dimensions commonly comes with several challenges. The main challenge of standard cardiac magnetic resonance (CMR) is the slice misalignments in 3D caused by acquiring the images at separate breath holds. If corrected, the more notable challenge is the trade-off between data fitting and expected visual appearance of the resulting mesh. Most works in the literature tend to over regularize, or bias a data fitting process towards a smooth result [95, 126]. Other methods work based on templates or shape prior, which are usually learned on normal patient anatomies, and tend to fail when presented with abnormal data, or any data outside the learned model [144]. Our approach differs in this, as we maximize the data fitting by composing smooth deformations on a smooth surface. In this Chapter, we propose a process to reconstruct a geometrical 3D surface mesh from a set of contours, partially representing an underlying heart shape anatomy, available from manual delineations or from automatic segmentations.

Whilst cardiac computer tomography (CT) images provide high anatomical resolution, contours obtained from CMR images were used as they provide additional

information, non-invasively, about myocardial structure using techniques such as late Gadolinium enhancement or Diffusion Tensor MRI. In CMR, image acquisitions are typically composed of a stack of SAX ranging from 8–12 slices, around 8 mm thick, with a 2 mm gap between slices, and one or more LAX. Contouring (also know as segmentation or delineation of the anatomical structures) of these typical acquisition protocols result in a spatial configuration of curves delimiting the anatomical structure of interest. These sets of curves are sparse since usually within the SAX stack they are 10 mm apart leaving large gaps to be filled, and cross-sectional since they comprise SAX as well as LAX view contours. Manual or automatic contouring methods are seldom able to account for the 3D environment, which can lead to spatial inconsistencies between long axes and short axes' contours. As discussed in Chapter 5, due to the delineations not being able to grasp perfectly the global 3D shape of the heart, contour-to-contour distances between these contours cannot be minimized to obtain perfect coincidence and therefore, spatial 3D consistency [113].

A good example of this occurs in LAX, where papillary muscles can be hard to differentiate from myocardium, which can result in their inclusion in the myocardial volume, in spite of their exclusion in SAX. Furthermore, due to the image acquisition occurring at different breath holds, as well as any patient movement inside the scanner, the contours might be misaligned. This tends to further increase their discrepancies and extra steps might be needed to correct for this misalignment [113]. This post-acquisition correction of the pose of the images can lead to losing the parallelism in the SAX stack. Finally, since many cardiomyopathies are localized in the left ventricle (LV), many previous efforts were focused on methodologies specifically designed or tuned for the LV. This Chapter is based on the work presented at the MIUA conference [130] and extended into a full journal paper [87]. A thorough evaluation of the reconstruction parameters has been performed. Furthermore, a statistical shape model (SSM) was used to generate realistic anatomical shapes allowing a broader and more comprehensive assessment of the technique.

The main contributions of this work are:

- Development of a surface reconstruction method capable of dealing with non-parallel, cross-sectional, sparse, heterogeneous, and non-coincidental contours,
- Use of our technique on the surfaces of the left ventricle as well as on the right ventricle,
- Evaluation of the method on several realistic shapes generated from a statistical shape model, as well as on real data consisting of normal and abnormal pathologies.

6.2 Materials and Methods

6.2.1 Synthetic and Real Data

Both synthetic and real datasets were used. Synthetic data were generated through the use of a statistical shape model obtained from [145, 146] and is available online (`http://wp.doc.ic.ac.uk/wbai/data/`). Using the modes of variation, different plausible heart-like shapes were generated and used in our experiments. Having the SSM allowed for a quantitative evaluations of our method, since the generated meshes provide ground truth anatomies. Contours were synthesized by slicing the generated meshes at various spatial positions. Parallel and non-parallel contours, simulating a short axis stack, and cross sectional contours, simulating LAX, were produced and used as inputs to our method. We treated the generated shape meshes as ground-truth and compared our results against it. Regarding real data, we synthesised data by fitting the shape model to real contour data by optimizing the contour-to-mesh distances between the shape model and the contours.

The SSM was fitted to the segmented contours, applying Eq. (6.2.1.1). In doing so, we obtained the set of parameters $\mathbf{b}$ and transformation $\mathcal{T}$ that generated the closest possible shape matching the provided contours. As the SSM was acquired from non-abnormal geometries, we used non-abnormal data as the SSM would not be able to generate abnormal shapes if fitted to irregular data. In order to fit the

SSM to the data, we minimized the contour-to-mesh distances, such that:

$$\min_{\mathbf{b},\mathcal{T}} \sum_{k} \sum_{i} \mathrm{d}(C_k^i, \mathcal{T}(\overline{\mathbf{X}} + \Phi\mathbf{b}))^2 \,, \tag{6.2.1.1}$$

where C_k^i is the i-th point on the k-th contour, $\overline{\mathbf{X}}$ the mean shape model, $\Phi = (\phi_1, \phi_2, ...)$ are the SSM modes of variations, $\mathbf{b}$ are a set of parameters that, when applied to Φ along with the rigid transformation model $\mathcal{T}$, generate a new shape. The error metric d is the contour-to-mesh distance. For each of the contours, the shape model with the minimum contour-to-mesh distance was obtained, and a set of spatially consistent contours were generated by slicing the mesh at the same planes as the CMR acquired contour planes. Figure 6.1 shows an example of such a mesh.

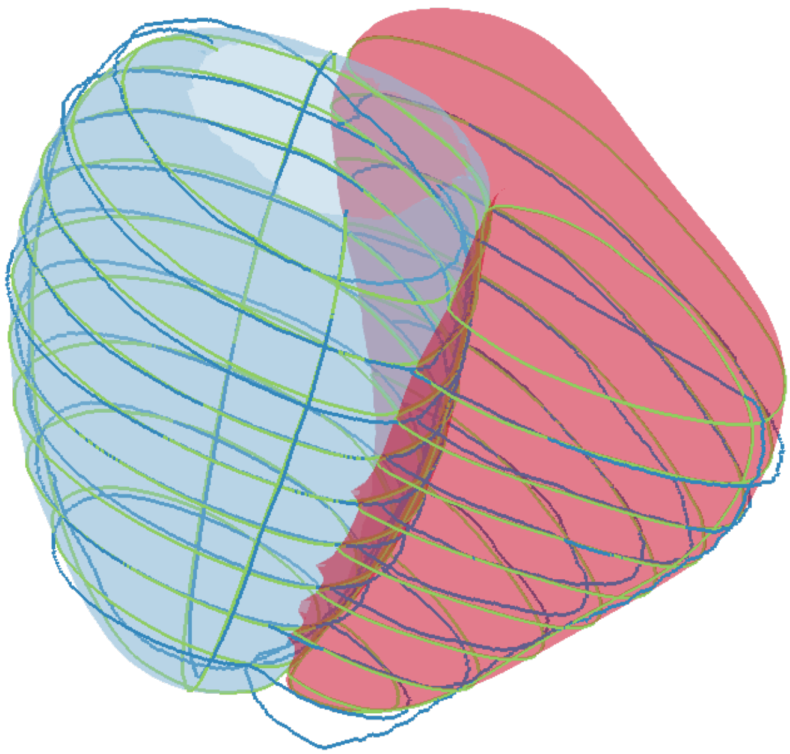

Figure 6.1: An example of the SSM of the left ventricle epicardium and right ventricle endocardium fitted to real contours. The dark blue contours represent the real contours. The green contours represent the synthetic contours generated by slicing the SSM at the same spatial pose of the CMR acquired contour planes.

In such a manner, we used 40 normal subjects to synthesize different meshes.

These meshes were sliced at the same planes of the real contours, thus generating spatially consistent simulated contours, having the same spatial planes as standard CMR acquisitions. This allowed us to have a ground truth mesh with no spatial inconsistencies. Delineated normal and abnormal contours, acquired from cine images at end-diastole were used, which were obtained from a collection of CMR cases from the John Radcliffe Hospital, Oxford, UK (UK ethical approval IRAS project ID: 112814). We define normal and abnormal following the anatomical appearance and clinical condition of the subject. In order to not bias our quantitative assessments, we selected cases that had the same amount of slices, as this is roughly indicative of the hearts having the same height, since the images are taken at regular intervals. Finally, our data consisted of 24 cases, 10 normals, and 14 abnormals randomly selected from the previously mentioned collection. All of the cases were chosen as they have 11 contours, two LAX and nine SAX.

6.2.2 Initial Meshes

Our method requires an initial mesh $\mathcal{M}_0$. This mesh can be computed using various approaches, such as the convex hull of the input data, or by fitting a precomputed mesh to the data such as the meshes found in [86, 95, 121, 126]. In our case, $\mathcal{M}_0$ is computed as a tubular mesh connecting the contours, as explained below. Let $\{C_j\}$ be a set of contours, each lying on a plane with normal n_j, where j is the contour index. A *central* direction is computed among all normal directions by computing the myriad following [147]. All normals lying within a neighbourhood of 20 degrees from that central direction are clustered together and their mean direction is used as main direction representing the long axis. In addition, the contours with normals in belonging to that cluster (which are expected to be the ones in SAX view) are then sorted from the most apical to the most basal (see Fig. 6.2a). Two extra points are added to $\{C_j\}$ providing an upper and lower lid, to fully enclose the initial mesh. We then proceed to create a rough initialization of the mesh by constructing a *tubular* surface going through the contours. For each of the contours, distribute K equally spaced contour points. A ruled surface approach is then used to obtain an

initial mesh, where the initial triangle making up the mesh are composed of the set $\{(C_j(i), C_j(i+1), C_{j+1}(i)), (C_{j+1}(i), C_{j+1}(i+1), C_j(i+1))\}$, with i being the point index of the contours after resampling (see Fig. 6.2b). Depending on the availability of the LAX information, the algorithm initializes the lower bound mesh point from the lowest contour points in that plane, representing the apex, by using the LAX information. Analogously, the same can be done in the upper part of the mesh, where the upper mesh point, representing the most basal slice, is initialized as the mean position of the topmost contour points lying on the plane that contains such point.

6.2.3 Mesh Deformation

As the initial selected points make up the mesh, another set of contour points are selected as the data fitting terms, and act as attractors for the mesh. N points are sampled, using farthest point sampling, allowing for a representation of the global spatial distribution of the contour points. The farthest point sampling method [148] ensures that no contour is sampled more densely than any others. At this point, the initial mesh $\mathcal{M}_0$, the contours, and the attractor points are all scaled to a common reference frame that is irrespective of the units of the provided contours. Once all the preprocessing is achieved, a deformation field is computed by taking the closest point from each of the attractor points on the contours to $\mathcal{M}_0$. The points on $\mathcal{M}$ are pulled towards the corresponding attractor points, and this deformation is extrapolated via the use of approximating thin plate splines (TPS) [149, 150]. This process is applied iteratively in small consecutive steps, to ensure small, smooth deformations as per Eq. (6.2.3.2). Let $\{P_i\}$, lying on the surface $\mathcal{M}_t$, be the set of closest points to the attractor points $\{Q_i\}$. Let $F : \mathbb{R}^3 \to \mathbb{R}^3$ be the deformation field such as

$$\underset{F}{\operatorname{argmin}}\, \lambda \sum_{i=1} (\|F(P_i) - (Q_i')\|)^2 + J_m^d(F), , \tag{6.2.3.1}$$

with

$$Q_i' = P_i + \beta \|(Q_i) - (P_i)\| \frac{(Q_i - P_i)}{\|(Q_i) - (P_i)\|}\,, \tag{6.2.3.2}$$

where J_m^d is the TPS functional using derivatives of order m, and d the image dimensions, as defined in [149, 150]. In our case, we use J_2^3. It should be noted

that λ affects the data matching term. With the computed F, we perform a deformation $\mathcal{M}_{t+1} = F(\mathcal{M}_t)$ iteratively, in a diffeomorphic manner, resulting in a composition of several smooth approximations approaching $\mathcal{M}$ towards the contours. A diffeomorphism is defined as an r-times continuous and smooth (differentiable several times) one to one mapping of a space with r-times differentiable inverse. In our case we define r to be equal to 1 This is is seen in Eq. (6.2.3.2), where Q_i' is the location of the deformation applied on points $\{P\}$, towards $\{Q\}$ as a factor of β. It can be shown that the points $\{P\}$ asymptotically converge to $\{Q\}$. Once the mesh has been deformed, the mesh undergoes subdivision to increase its resolution, allowing for a more appropriate reconstruction of the geometry. We use the interpolating butterfly subdivision algorithm [65], allowing us to increase the number of vertices, in a rounded manner, whilst keeping the position of the original vertices. Afterwards, we smooth our mesh using Laplacian smoothing, which can be seen in Eq. (6.2.3.3) [151], allowing us to clean and fair our mesh. The Laplacian smoothing is performed in the following manner:

$$\mathcal{L}(\mathbf{p}_i) = \frac{1}{\sum\limits_{j \in N_i}} \left(\sum_{j \in N_i} \mathbf{p}_j \right) - \mathbf{p}_i \,, \tag{6.2.3.3}$$

where $\mathbf{p}$ represents the vertices, and $\mathbf{p}_j$ represents the adjacent vertex to $\mathbf{p}_i$. It is well known that this smoothing operation usually shrinks the mesh [71], and as such we follow the smoothing with an affine transformation between the original and final positions of the vertices. As the subdivision process can lead to extremely high amounts of triangles, the mesh is decimated, using quadratic decimation, which allows for some control over the amount of triangles while preserving its geometric characteristics [70].

The latter steps of subdivision, smoothing, and decimation occurs every 20 iterations. This iteration was found empirically, allowing for a fast computational time and an appropriate mesh fitting time. The overall pipeline can be seen in Fig. 6.2. The overall process starting from the farthest point sampling is repeated for a set number of iterations, or until some stopping criterion is achieved, such as the contours-to-mesh distance falls bellow a chosen threshold or a fixed number

of iterations has been reached. In our case, we use as a stopping criterion a set number of iterations equal to 250, which we empirically found to give good quality meshes. That stopping criterion was used along all experiments performed in this work. An overview of the algorithmic process can be found in Alg. 2

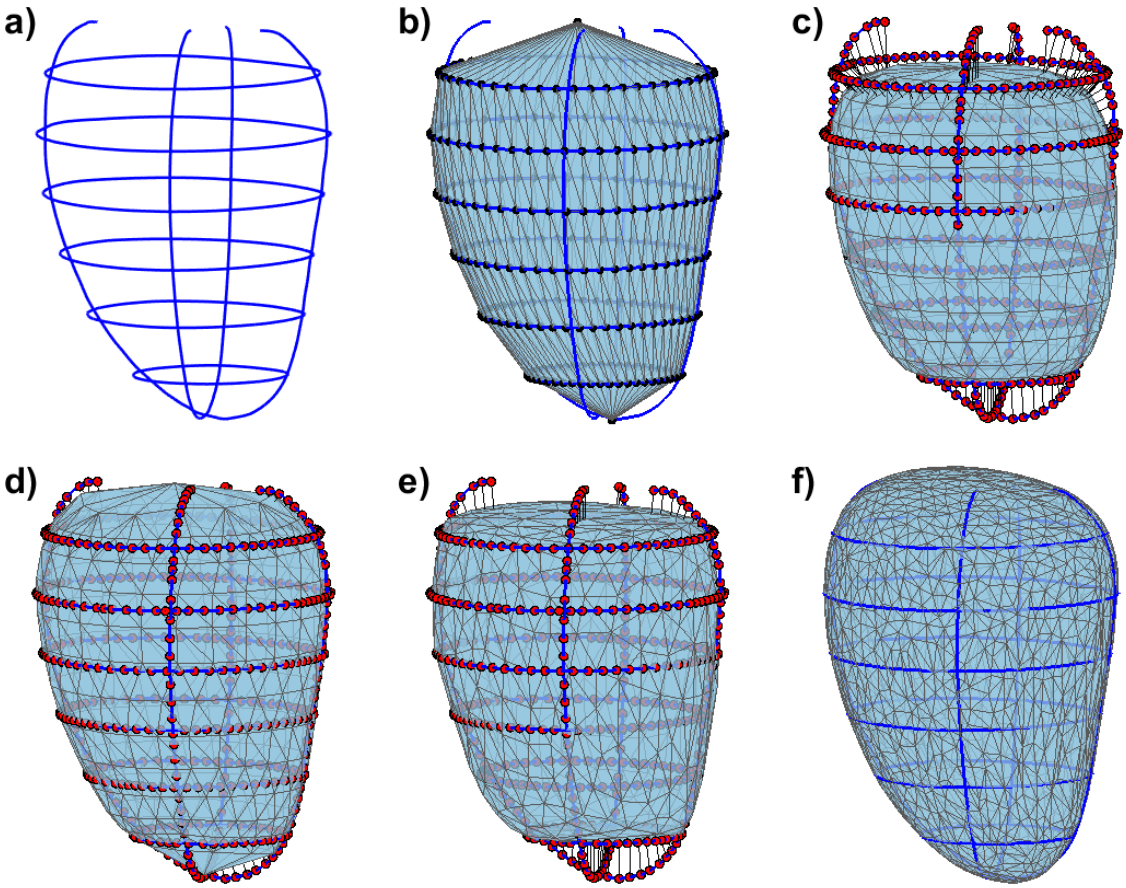

Figure 6.2: Overall pipeline for generating a surface mesh from contours illustrated on left epicardial contours. (**a**) the input contours; (**b**) *tubular* initial mesh with its vertices equally spaced along the contours; (**c**) the initial mesh after the Laplacian smoothing and showing the attractor points on the contours, with a visual illustration of their connections pulling the mesh; (**d**) the resulting mesh after several small deformations towards the attractor points; (**e**) the mesh after having undergone the process of subdivision, smoothing and decimation; (**f**) the resulting mesh, after several iteration of steps (**d**,**e**).

6.3 Results

In order to assess our algorithm, we performed a series of experiments on both synthetic and real data, allowing us to assess all the different aspects of our methodology in a pertinent manner.

Algorithm 2 Surface Mesh Reconstruction

procedure SURFACE MESH RECONSTRUCTION
Input: Contour points (PC), Initial Mesh ($\mathcal{M}_0$)
Output: Surface Mesh
Parameters:
- $FarthestPoint = 2$mm (Minimum distance between points)
- $SmoothDecIter = 20$ (Number of Smoothing and Decimation iterations)
- $TargetReduction = 0.72$ (Reduction at decimation step)
- $\beta = 0.1$ (Percentage of pushing force)
- $SmoothStrength = 200$ (Number of Smoothing iterations)
- $\lambda = 1e05$ (TPS trade-off parameter)

Get $\{Q\}$ by sampling contours with FarthestPointSampling
for $t = 0 : 250$ **do**
  for each q_i in $\{Q\}$ **do** find p_i in $\mathcal{M}_t$ such that $D(q_i, p_i)$ is minimal
  Deform mesh: $\mathcal{M}_{t+1} = F(\mathcal{M}_t; q_i, p_i, \beta, \lambda)$, via TPS, as per Eq. (6.2.3.1)
  if $!rem(t+1, SmoothDecIter)$ **then**
    Subdivide $\mathcal{M}_t$ using the interpolating butterfly subdivision algorithm
    Smooth the mesh using Laplacian Smoothing as per Eq. (6.2.3.3), having parameter $SmoothStrength$.
    Decimate the mesh using the Quadratic Decimation algorithm, having parameter $TargetReduction$.
  $SmoothDecIter = SmoothDecIter + 1$

6.3.1 Simulation of CMR Acquisitions

Our mesh construction methodology requires a number of parameters to be chosen. Some parameters, such as the TPS regularization term λ or the number of Laplacian smoothing iterations, have an explicit effect on the overall resulting shape, whilst for others such as the farthest points sample size or the decimation parameter, this effect is less obvious. In order to evaluate good candidate parameters, we performed a multi-parameter study using synthesized contours from several instances of the SSM fitted to real contours.

A parallel coordinate plot, as shown in Fig. 6.4, allowed us to select the default parameters to generate our meshes. Each parameter has a certain effect on the resulting surface mesh. It can be seen in Eq. (6.2.3.1) that the λ parameter controls the data matching strength. The higher λ is, the more reliant on the contours the method will be. In contrast, the lower the λ, the more spherical the resulting mesh will be. It can be noted, from the parameter exploration, that the Laplacian

smoothing term, which is represented by the amount of iterations of $\mathcal{L}$, and λ are linked with respect to data matching. Even with the lowest λ possible, if there is a low number of smoothing iterations, the resulting mesh will be close to the contours, but will be subject to the amount of decimations. As such, the mesh will tend to be quantitatively close to the contours but appear very wrinkled. It should be noted that, with a high number of decimations, the number of triangles increases, and can act as curvature control term. The higher the amount of decimations, the more curvature is possible, with regards to the overall mesh. Whilst this can result in a *smoother* looking mesh, the mesh has the possibility to be very convoluted. The farthest points parameter, controlled by the minimum inter-points distance allowable, does not appear to have much of an effect on the overall smoothness of the mesh, or the contour matching. However, as most of the normal case meshes, especially for the epicardium, can be geometrically simple (resembling a cylinder), the number of points needed to describe the geometry is low. With more convoluted geometries such as with some of the more severe abnormal cases, the lower the farthest points distance, the more likely the overall geometry will be captured. Figure 6.3 illustrates some resulting meshes for different extreme sets of parameters given the same input - left epicardial contours. Panel a) represents the mesh generated from the optimal parameters. (Smooth strength ($\mathcal{L}$) = 271, farthest points = -2 mm, target reduction (decimation) = 0.7233, $\lambda = 1e5$.) Panel b) represents the mesh generated with a higher target reduction parameter in the decimation step, resulting in meshes with smaller numbers of triangles and vertices. Although visually pleasing and accurate for the case showed in the figure, the mesh can be seen to be more decimated, reducing the degrees of freedom, and restricting the method to tailor to abnormal cases. (Smooth strength ($\mathcal{L}$) = 271, farthest points = -2 mm, target reduction (decimation) = 0.75, $\lambda = 1e5$.) Panel c) represents a mesh generated from a reduced amount of Laplacian smoothing iterations. This causes the geometrical structure of the mesh to be led more by the pulling process, creating more convoluted meshes, achieving meshes close to the contours, but with poor interpolation properties. (Smooth strength ($\mathcal{L}$) = 19, farthest points = -2 mm,

target reduction (decimation) = 0.7233, $\lambda = 1e5$.) Panel d) shows a mesh generated from a reduced amount of Laplacian smoothing iterations as well as lower sampling points. The mesh has therefore less attractor points on the contours, causing the mesh to become convoluted and have a higher contour to mesh distance discrepancy. (Smooth strength ($\mathcal{L}$) = 5, farthest points = -10 mm, target reduction (decimation) = 0.7233, $\lambda = 1e5$.) Panel e) represents a mesh generated with a low value for the λ parameter in the approximate TPS step (see Eq. (6.2.3.1)). This causes the mesh to become more elliptical and very smooth, as it has a high regularization weight with a low data matching weight. In addition, a large discrepancy with the contours can be observed. (Smooth strength ($\mathcal{L}$) = 1028, farthest points = -2 mm, target reduction (decimation) = 0.7233, $\lambda = 0.4642$.) Panel f) shows a mesh generated with a low value of λ and a lower target reduction parameter. This causes the mesh to be smooth and elliptical, but with a high number of vertices. Similarly to (e), the mesh has a high contour to mesh distance. (Smooth strength ($\mathcal{L}$) = 1028, farthest points = -10 mm, target reduction (decimation) = 0.6989, $\lambda = 0.4642$.) Panel g) represents a mesh similar to the one found in panel f), with a low Laplacian smoothing. This causes the high number of vertices to be pulled to the attractor points causing a convoluted appearance, although having a relatively low fitting error. (Smooth strength ($\mathcal{L}$) = 19, farthest points = -10 mm, target reduction (decimation) = 0.6989, $\lambda = 0.4642$.)

In our case, the best parameter sets were found with a very high λ, a target reduction of 72.3%, and a small farthest point distance (from 2 to 4 mm inter-point distance). The red bands in Fig. 6.4 show the cut off values for the parameters selection. Our method was initially tested in several scenarios and extreme values for its parameters were found empirically. Once extremal parameter values were found, each parameter was sampled accordingly. Some parameters, such as the target reduction, were sampled linearly, whilst others, such as λ, were sampled logarithmically. Using the cut off values on our requirements resulted in the most visually pleasing meshes with low mesh-to-mesh errors. These parameters were used for all of the resulting meshes in all of the experiments.

Table 6.1 represents the mean, median, min and maximum mesh-to-mesh distances for the epicardium, endocardium, and right ventricular for 40 cases, having used the simulated contours described in Section 6.2.1 as inputs. This experiment allowed us to simulate contours at the same spatial poses as CMR acquired images. Furthermore, having fitted an SSM to real data, allowed us to obtain various shape models based on patient population. This allows us to assess the robustness of our methodology to the variations of shapes seen in the population. By having a ground truth mesh, we were able to compare our results in a quantitative manner without the inherent contour discrepancies. The resulting meshes, showing the different shape variations, can be seen in Fig. A.1, Appendix A.1.

Measure	**Epicardium (mm)**	**Endocardium (mm)**	**Right Ventricle (mm)**
Mean	0.12	0.14	0.29
Median	0.09	0.09	0.22
Max	1.11	1.29	3.34
Min	1.35×10^{-5}	2.04×10^{-5}	7.49×10^{-5}

Table 6.1: Table depicting the mesh-to-mesh distances (in mm) between the ground truth surface meshes and our generated meshes for 40 different geometries.

6.3.2 Synthetic Data

6.3.2.1 Short Axis Contours

In a first instance, we applied our method to a set of parallel SAX contours and observed the ability of our technique to recover the ground truth mesh by evaluating the mesh-to-mesh distance error. Our validation consisted of calculating the distances between vertices of our generated mesh to the ground truth mesh (mean shape of the statistical model) as shown in Fig. 6.5. We also assessed the effect the number of SAX had on the ability of the mesh to recover the original mesh by increasing the amount of SAX contours provided initially from 2 to 15, between the base and the apex. The results can be seen in Fig. 6.6.

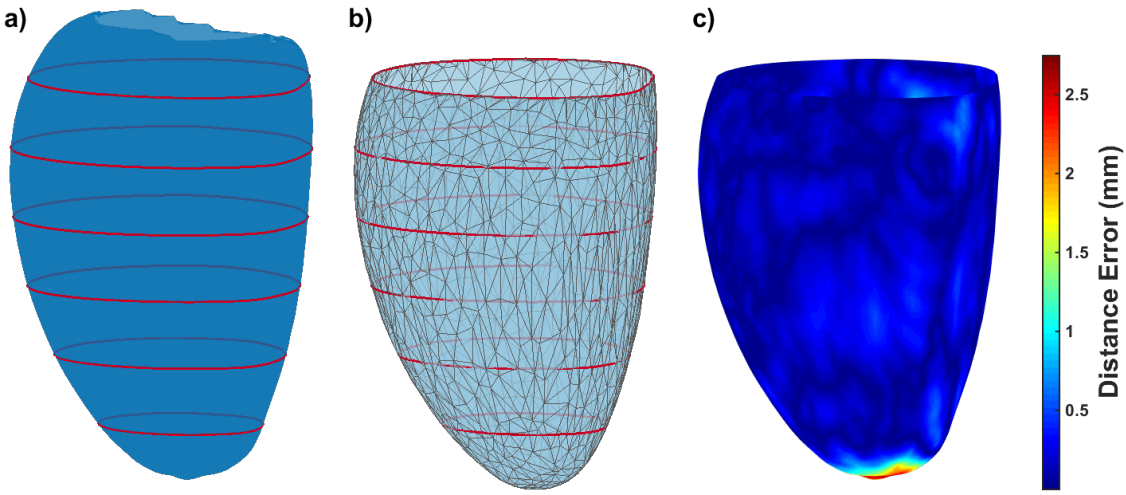

Figure 6.5: (**a**) the ground truth surface (left epicardium of the mean shape of the SSM) and the synthesized contours in red; (**b**) our resulting mesh given the red contours; (**c**) the surface-to-mesh distance error (in mm). The resulting mesh has been clipped at the level of the most basal contour.

6.3.2.2 Short Axis and Long Axis Contours

In a second experiment, we assessed the impact providing LAX contour information had. Similarly to the first experiment, we looked at varying the amount of LAX contours and the effect it had on the overall mesh result, as well as look at the influence LAX positions had. Four SAX contours were chosen, and an increasing amount of LAX contours were provided as the initial input by adding a rotated version of the initial LAX contour every iteration. From an initial LAX, rotated versions were iteratively added using the following rotations: 90, 45, 135, 22.5, 157.5, 112.5, 67.5 degrees, respectively. Similarly to Section 6.3.2.1, the same analysis was performed. The mesh-to-mesh contour distance was calculated and overlaid on the ground truth mesh, as shown in Fig. 6.7.

6.3.2.3 Non-Parallel, Non-Coincidental Contours

As our method is intended to be used with real data, we mimicked a standard acquisition. From the mean shape model, 10 SAX contours (resulting in 10 mm inter-slice distance) as well as two LAX contours (simulating two and four chambers'

views) were generated. The contours were then perturbed, mimicking breath-hold related misalignment, by applying small random translations and rotations (in plane as well as out of plane). The translation and rotation perturbations were performed within the ranges described in [20, 21]. The resulting non-parallel, non-coincidental contours were given as an input to our algorithm. The resulting mesh can be seen in Fig. 6.8a. The distance map overlaid on the contours is the distance from the resulting mesh to the input contours. The algorithm was also tested on a right ventricle, as seen in Fig. 6.8b.

6.3.3 CMR Acquired Data

Our reconstruction was applied to real data to show how it responds to both normal and abnormal cases. As we use expert delineated contours to build the mesh, there is no ground truth to assess our method on. However, we performed quantitative analysis on the real datasets by measuring the discrepancy between individual contours and the resulting mesh (contour-to-mesh distance). This measurement provides an evaluation of the "goodness of fit" of the mesh with regards to the input contours. An example of this can be seen in Fig. 6.11, as the blue box-plots (labelled OM). Each blue box-plot represents the contour-to-mesh distance distributions for each respective contour. It should be noted that there is an inherent contour discrepancy brought by the non-coincidental nature of the contours. It can be seen that the resulting mesh has a distance no bigger then 2 mm to each of the contours. The example showed was performed on the epicardium of an abnormal case with severe myocardial thickening, as seen in Fig. 6.9. This particular case was chosen as an example as we wanted a more complex geometry to demonstrate were the "fitting" effect would be noticeable. The resulting epicardial and endocardial surfaces can be observed in Fig. 6.10.

To further evaluate the individual impact each contour has on the resulting mesh geometry, we calculated the error distance from each of the contours to the output mesh, having removed that specific contour prior to building the mesh, in a *leave-one-out* fashion. In performing this experiment, we assess the ability of our

method to interpolate missing information from the dataset, as well as analyzing the structural and morphological information provided by each cardiac contour. The red box-plots (labelled CM) in Fig. 6.11 show the results applied to the same abnormal case (Fig. 6.9). It can be observed that the contours where the highest amount of error occurred were the LAX, and the basal and apical slices. It can also be seen that the second contour left of the apical slice has a significant error when having the mesh generated without it. This is the region where myocardial thickening is most prevalent, as seen in Fig. 6.9. We performed the same *leave-one-out* experiment on all 24 cases for the epi- and endocardium. The results can be seen in Fig. 6.12. In order not to bias the experiment, we chose similarly sized hearts, each containing 11 images. As we obtain an error distance distribution for each of the contours, since each of the distance is the closest point-to-mesh distance for a contour, we calculated the mean, median, minimum and maximum distance values of each of the distributions. This resulted in 24 values for each of the statistical measurements. The squares and error bars on the graphs represent their distribution. From these graphs, it can be seen that, for the various normal and abnormal cases, the LAX along with the basal and apical slices are the slices where the distance error was more significant due to their structural information. This effect is more pronounced in the endocardium, as the geometry is more complex. It appears that the epicardium meshes rely more on the LAX and apical contours, whilst the endocardial meshes rely more on the information provided by the LAX and basal slices.

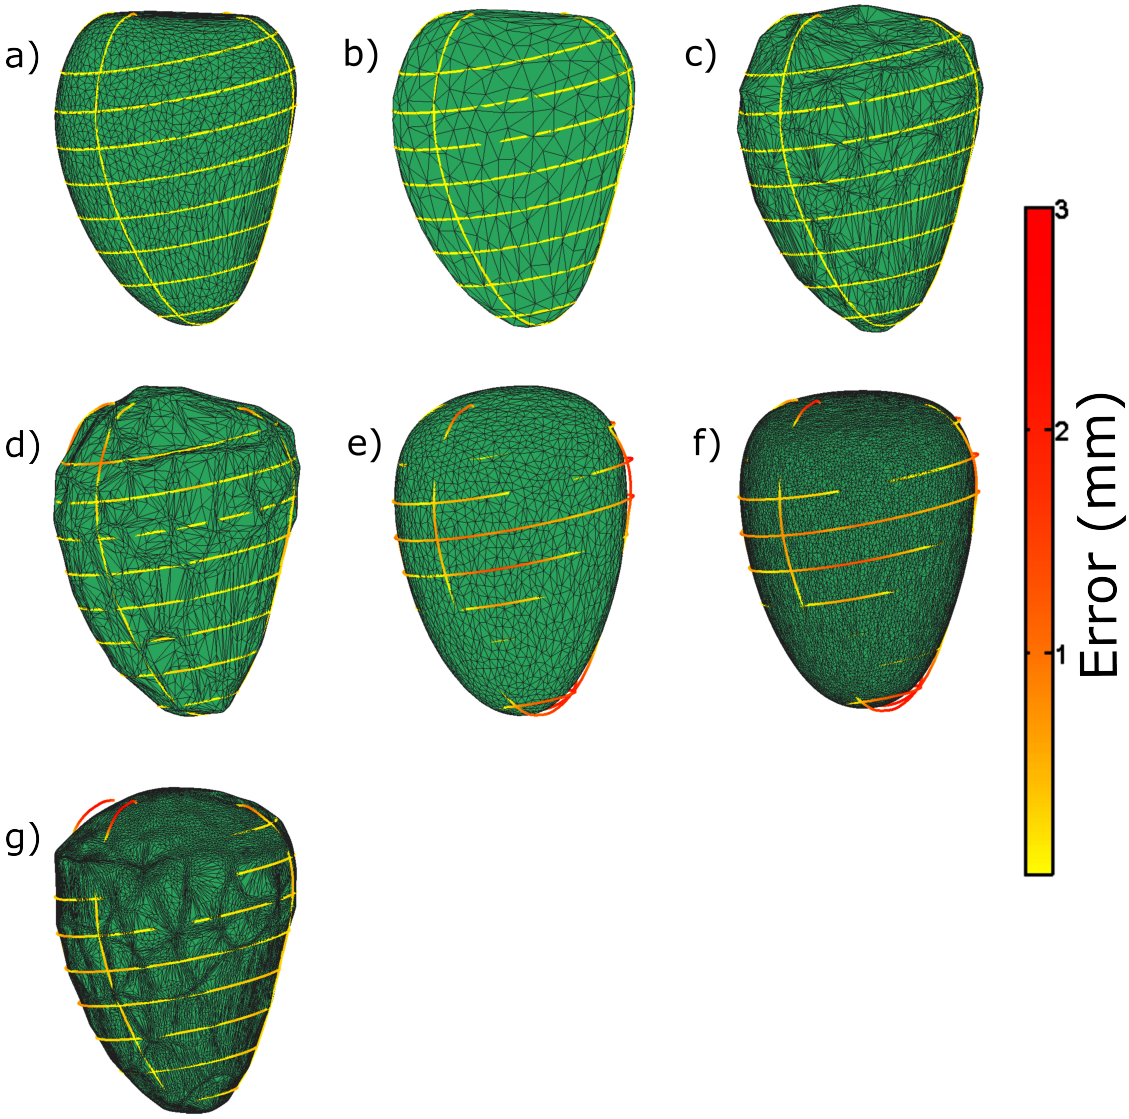

Figure 6.3: Illustration of resulting meshes for different extreme sets of parameters in the exploration and tuning of the parameters experiment. The color along the contours represents the distance to the resulting mesh. **a)** Mesh generated from the optimal parameters found in Section 6.3.1. (Smooth strength ($\mathcal{L}$) = 271, farthest points = -2 mm, target reduction (decimation) = 0.7233, $\lambda = 1e5$.) **b)** Mesh generated with a higher target reduction parameter in the decimation step. (Smooth strength ($\mathcal{L}$) = 271, farthest points = -2 mm, target reduction (decimation) = 0.75, $\lambda = 1e5$.) **c)** Mesh generated from a reduced amount of Laplacian smoothing iterations. (Smooth strength ($\mathcal{L}$) = 19, farthest points = -2 mm, target reduction (decimation) = 0.7233, $\lambda = 1e5$.) **d)** Mesh generated from a reduced amount of Laplacian smoothing iterations as well as lower sampling points. (Smooth strength ($\mathcal{L}$) = 5, farthest points = -10 mm, target reduction (decimation) = 0.7233, $\lambda = 1e5$.) **e)** mesh generated with a low value for the λ parameter in the TPS step. (Smooth strength ($\mathcal{L}$) = 1028, farthest points = -2 mm, target reduction (decimation) = 0.7233, $\lambda = 0.4642$.) **f)** Mesh generated with a low value of λ and a lower target reduction parameter. (Smooth strength ($\mathcal{L}$) = 1028, farthest points = -10 mm, target reduction (decimation) = 0.6989, $\lambda = 0.4642$.) **g)** Mesh with similar parameters to the mesh found in panel f), with a low Laplacian smoothing.(Smooth strength ($\mathcal{L}$) = 19, farthest points = -10 mm, target reduction (decimation) = 0.6989, $\lambda = 0.4642$.)

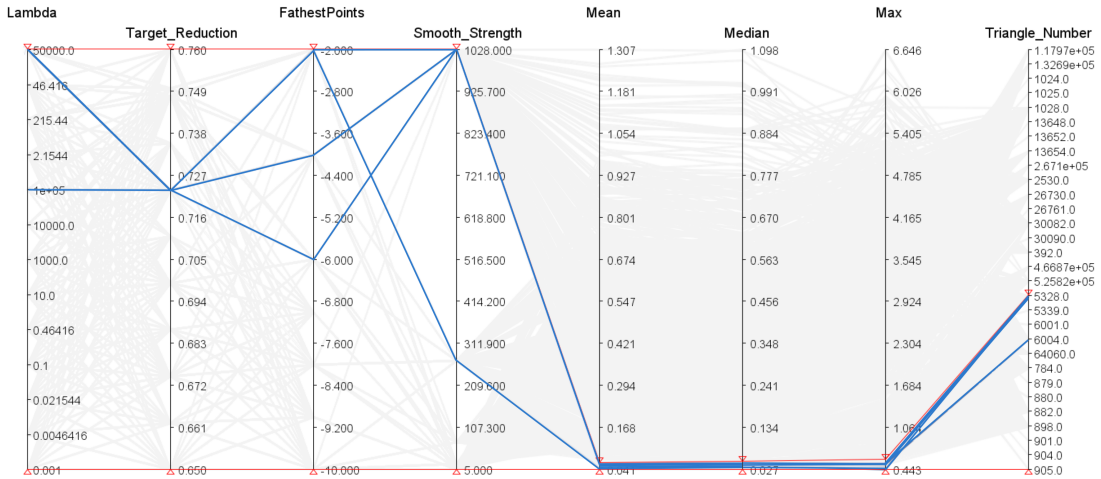

Figure 6.4: Parallel Coordinate graph showing the different combinations of parameters giving the minimum average, median, and min mesh to mesh distance, as well as the number of triangles generated. The red lines represent the selected bands for the y-axis. The highlighted blue lines represent the parameters that have generated meshes that satisfies the selected criteria. The faded lines represent every parameter permutation and their respective evaluation metrics. There are 3440 permutations. The fifth to seventh y-axis represent the mean, median and min mesh-to-mesh error distances in mm, between the ground truth meshes and the generated meshes given the respective parameters.

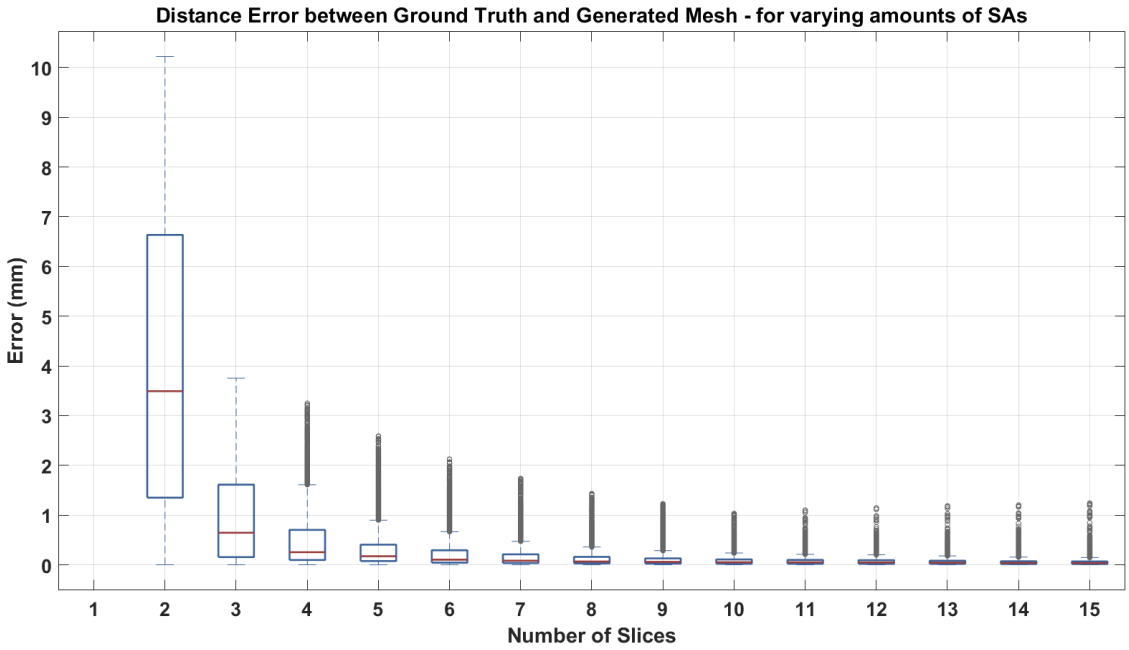

Figure 6.6: Impact of increasing the number of SAX contours on the error distance between the generated mesh and the ground truth surface.

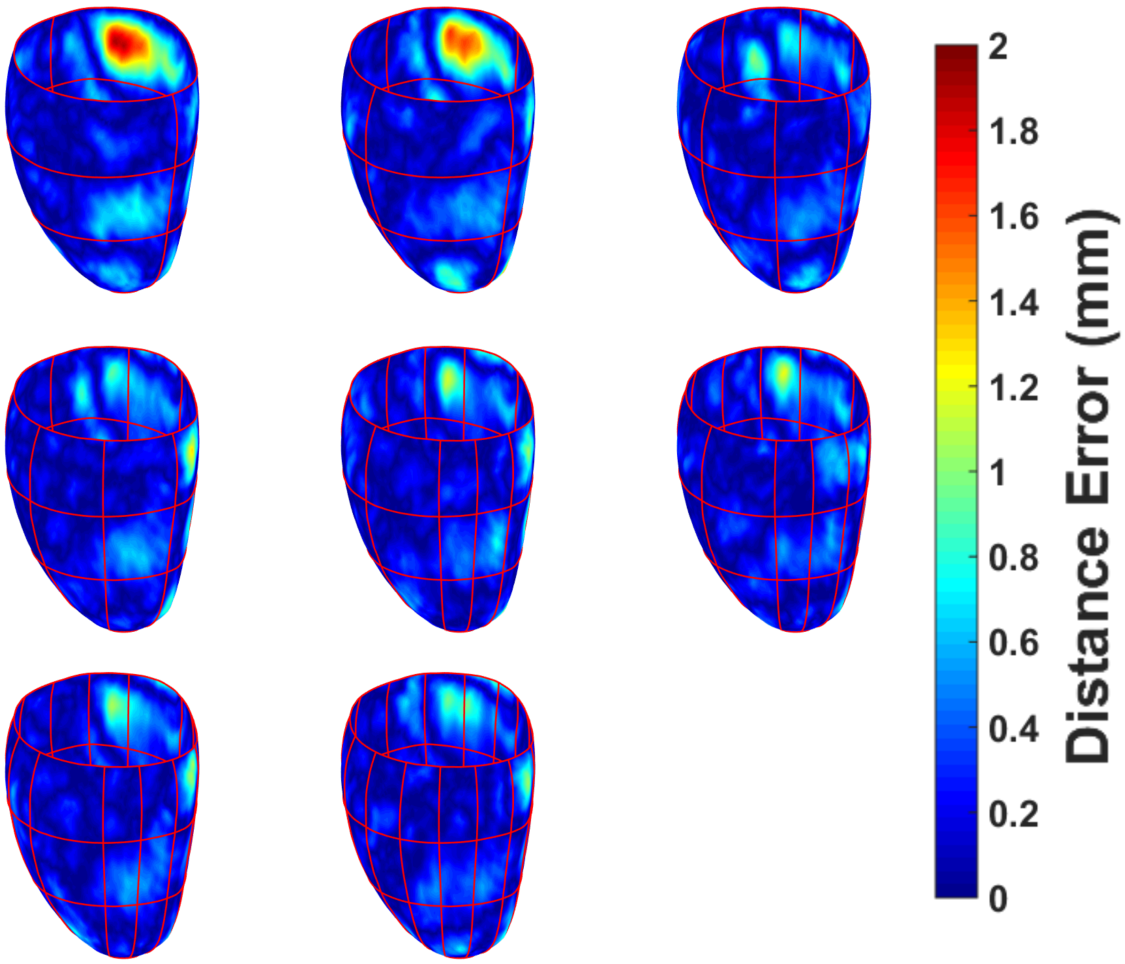

Figure 6.7: The effect of increasing the amount of left epicardial LAX contours have on the error distance between the generated mesh and the ground truth mesh. Top to bottom: 1–8 LAX slices added, successively. Each LAX added is a rotated version of the first LAX shown in the first row, first column. Rotations were generated as per the following: 90, 45, 135, 22.5, 157.5, 112.5, 67.5 degrees, respectively.

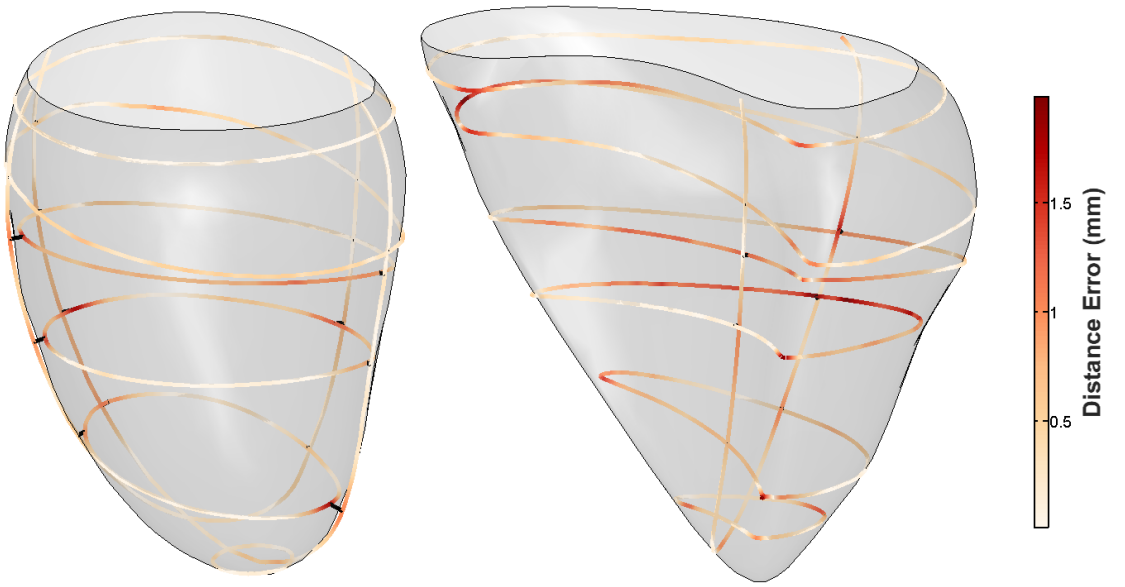

Figure 6.8: Distance map from the resulting geometrical mesh to the simulated cardiac input data. The contours were transformed by small out of plane rotations and translations. The left panel shows the resulting left ventricular epicardial mesh, and the right panel shows the right ventricular endocardial mesh. The distance errors represent the contour-to-mesh distances.

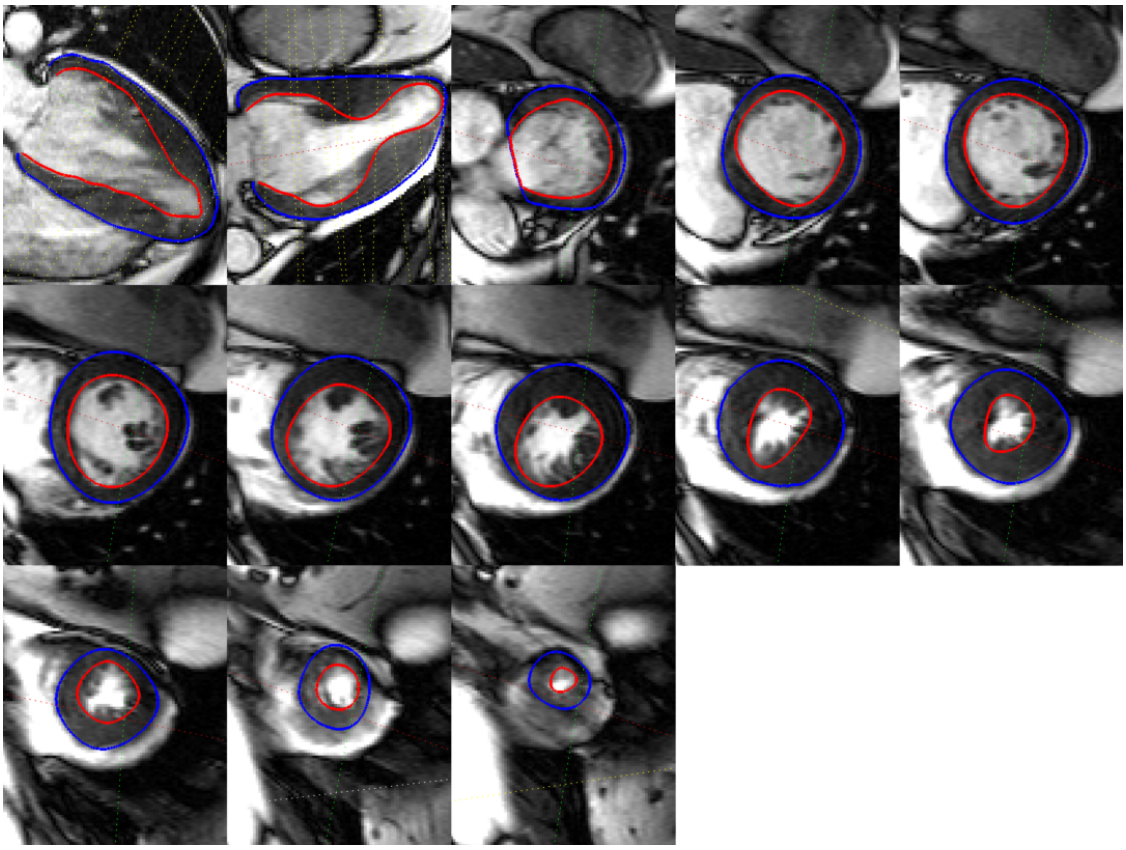

Figure 6.9: Different segmented CMR acquired contours on a severe abnormal anatomy.

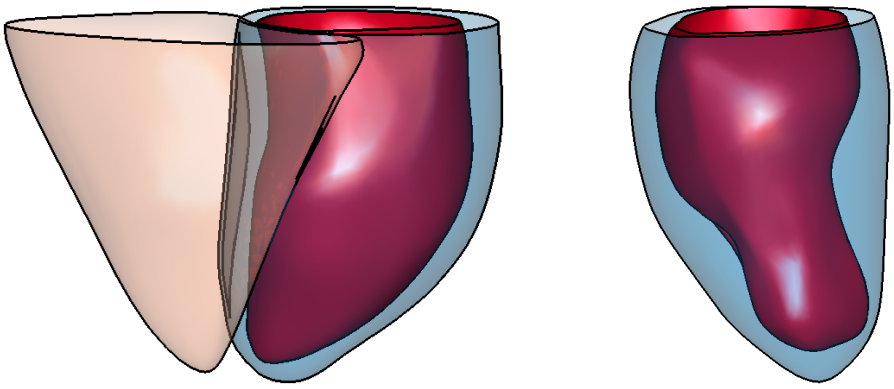

Figure 6.10: Example of resulting surface meshes generated from (**a**) contours for a normal case, and (**b**) contours belonging to a severe abnormal anatomy. The left panel shows surface meshes for the left endocardial (in red) and epicardial (in blue) ventricle as well as the right epicardial ventricle (in orange). The right panel shows surface meshes for the left endocardial (in red) and epicardial (in blue) ventricle.

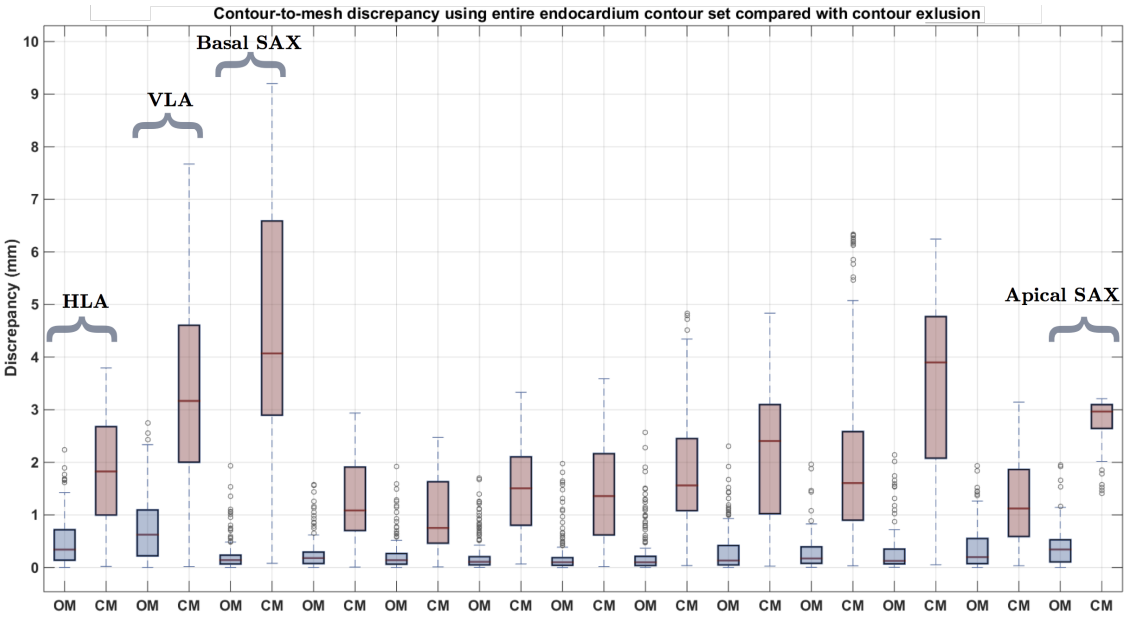

Figure 6.11: Endocardial contour-to-mesh distances having built the mesh with the entire set of contours (labelled OM, which stands for "Original Mesh", in blue) and excluded contour-to-mesh distance having built the mesh without the excluded contour (labelled CM, which stands for "Computed Mesh", in red).

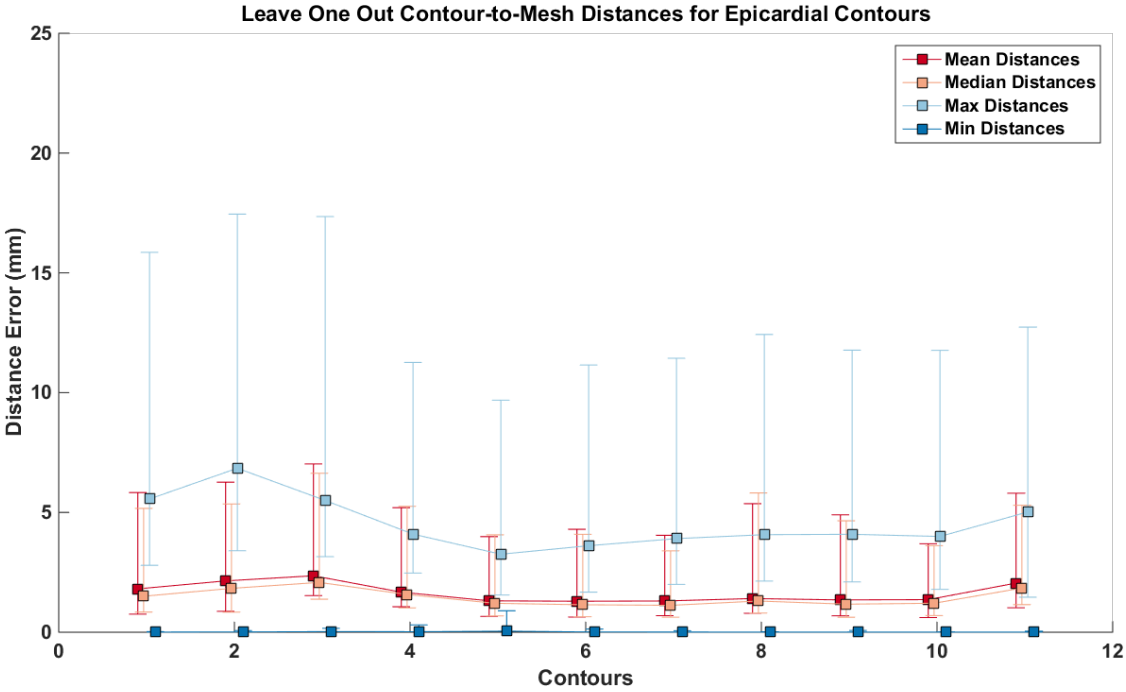

(a) Epicardium

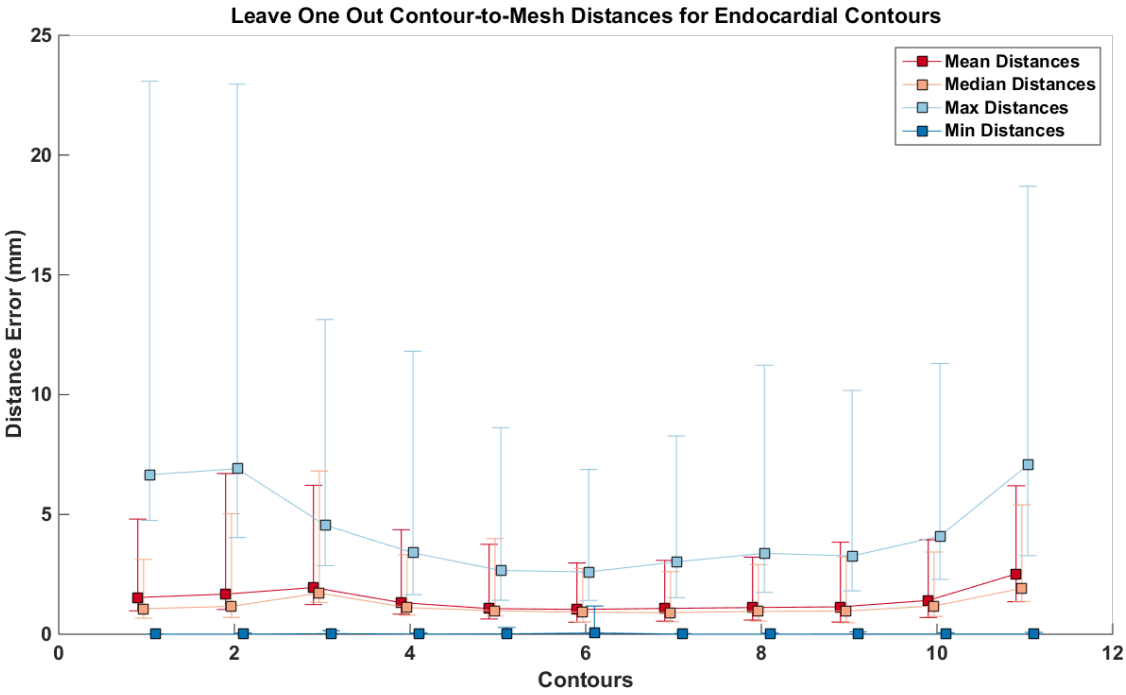

(b) Endocardium

Figure 6.12: Contour-to-mesh distances for all 24 cases, having built the meshes without the respectively excluded contour. Each error bar represents the distribution of distances for each of the respective statistical measurement. Each square represents the mean of each statistical measure, with the upper and lower notches representing the minimum and maximum values, respectively. **Top**: epicardial contour-to-mesh distances; **bottom**: endocardial contour-to-mesh distances. The y-axis represents the contour index. 1–2 represent the LAX, 3–11 represent the apex to the base, respectively.

6.4 Discussion

To be able to assess various aspects of our methodology, several experiments were carried out, using both synthetic and real data.

One of the biggest hurdles in validating methodologies on real data is the lack of ground truth. This can be due to the intra-inter user segmentation differences causing spatial inconsistencies, or slice misalignment due to breath hold related issues, or simply because the acquisition is sparse, providing partial views or information from the whole underlying anatomy. Due to the partial volume effect when acquiring CMR slices, the gap between contoured slices are significant, as CMR slices are typically 8 mm thick and separated by a 2 mm gap, creating around 10 mm spacing between each contour. This inherent gap between contour planes causes for existing anatomical features to be unobtainable or unrecoverable. Usually, the information between the gaps are interpolated, in-painted, or filled using prior information of what the geometry is believed to be. For this reason, we chose to use an SSM so that we had the global geometrical information, so that we may use it as ground truth.

As we wanted to quantify the ability of our method to work on standard CRM acquired data, without biasing the results to other indirect measurements, we simulated contours having the same spatial planes as real data, mimicking a real acquisition without inter-contour discrepancies. Furthermore, using the SSM allows us to generate legitimate, population-based, heart shapes. This is the main reason why we decided to use an SSM to generate ground truth with fully accessible anatomies, and to evaluate how well, as well as which parameters are required, for our method to mimic the completion of the gaps from only the provided partial contour information. By having possible entire heart-like shapes, we were able to validate and tune our method for the specific task of building cardiac anatomies. Using these synthesized contours, which include multiple SAX and LAX, we assessed the error between the generated mesh and the ground truth mesh given various parameter combinations. We assessed a total of 3440 parameter samplings over 40 fitted ground truth shapes generated from the SSM. Figure 6.4, in Section 6.3.1, shows a parallel coordinate plot of the different parameters evaluated, and the

resulting error metrics of the output results given those parameters. Assessing the combination of each parameter permutation allows for obtaining the most favorable set of parameters that result in the most similar shape given a set of contours. We define similar in terms of the shape that has the minimum mesh-to-mesh distance while maintaining a relatively low number of triangles. As such, our metrics of evaluation were the mean and medium mesh-to-mesh distances for each of the cases, as well as the number of triangles. Using the parallel coordinate plot allowed us to interactively select the parameter ranges of interest, which are highlighted in blue. These parameters were used as the default parameters for all of our resulting meshes.

We fitted an SSM to data acquired from a normal cohort, and generated contours from the resulting new shape at the same planes as the real contours. This experiment informs us on the ability to recover meshes from CMR acquired coincidental contours. Using these contours as inputs, we generated meshes, which can be seen in Fig. A.1 in Appendix A.1, and compared them to the fitted shape models. Table 6.1 shows the resulting statistical measures for all of the mesh-to-mesh distances for all of our output shapes. It can be seen that all of the mean and median values fall under sub-pixel accuracy. The maximum error distances show that the methodology performed worst on the right ventricles. As the right ventricles only had one horizontal long axis (which contains the four chambers of the heart), each of the right ventricles had one LAX contour. This long axis typically falls within one side of the right ventricle, and as such the most extreme distances were due to the geometrical information at the opposing areas only relying on the SAX information.

The method was applied to a stack of synthesized parallel contours and the result can be seen in Fig. 6.5. In order to capture the discrepancy between the ground truth and the resulting mesh, mesh-to-mesh distances were calculated. It can be seen in Fig. 6.5c that the maximum discrepancy is around 2.5 mm for a human heart size. This error is the distance between the ground truth shape and the reconstructed mesh's most apical areas. As we do not have information at the apex due to the absence of LAX, it is expected that this is where the discrepancy would be highest. Elsewhere, distances appear no greater than 2 mm. To be able to

assess the impact the amount of SAX contours have on the process, the distance between meshes was calculated by incrementing the number of equidistant SAX of the input data, in an iterative manner, starting with 2 SAX up to a stack of 15. The resulting box plot showing the distances can be seen in Fig. 6.6. It can be observed that there is a considerable reduction of distance discrepancies up until eight slices, after which the decrease in the distances is not as notable. In this experiment, 8–9 slices represent an average CMR acquired dataset as a SAX stack has generally an inter-slice distance of 10 mm. The outliers make up the error distances at the base of the mesh, where no information was provided.

The impact that long axes have on the provided data was then assessed. A stack of four short axis slices were selected at equidistant positions from base to apex, and 1–7 long axes were iteratively added simulating a radially oriented CMR acquisition. We evaluated the distance errors (see Fig. 6.7), which show that, as the amount of LAX increases, the error decreases. From the results, it can be seen that, after the addition of three LAX, there is negligible decrease in the error distances, as they all fall within 1 mm, which is of subpixel accuracy.

After assessing the SAX and LAX impact on the resulting meshes, we generated non-parallel, non-coincidental contours for both the left and right ventricle, as per Section 6.3.2.3. Due to breath-hold related issues, cardiac images are often misaligned, which must be corrected for, in order to minimize the coincidental discrepancy between contours. This step causes contours to lose their parallelism. Furthermore, as there also exist inherent contour discrepancies due to segmentation issues, there always remains slight contour-to-contour discrepancies in cross sectional pairs. Performing the experiment in Section 6.3.2.3 allowed us to assess our methodology's robustness to cope with this type of data. It can be seen in Fig. 6.8 that, at the contours' planes, the contour-to-mesh discrepancy falls under 1.7 mm for the left ventricle epicardium and under 2.5 mm for the right ventricle endocardium. Despite having non-parallel, non-coincidental contours, the resulting meshes are visually pleasing.

Finally, real data was used, for both a normal and abnormal case. The contours for the abnormal case can be seen in Fig. 6.9. The resulting meshes can be observed in Fig. 6.10. To assess our method's ability to deal with standard CMR acquired datasets, we measured the individual contour-to-mesh discrepancies, which can be seen in Fig. 6.11. It can be seen that, when all the contours are used (blue box-plots), the discrepancy is within 2 mm for the abnormal case (Fig. 6.11). As we proceed to iteratively remove a single contour, in a *leave-one-out* fashion, and rebuild the mesh, the distance for that excluded contour to the mesh increases. It can be seen that, when the most basal SAX is removed, there is a high discrepancy occurring when measuring the distance from that contour to the resulting mesh. As this is the most basal slice, this is expected as the basal slice contains important structural information about the global geometry of the heart. When the basal slice is removed, our process relies on the LAX to fill in the basal information, which can lead to an early closure of the basal part of the mesh, which might be at a lower geometrical position than the contour, leading to high discrepancy. The removal of contours allows us to assess the individual contributions of the contours to the geometrical result of the mesh. As such, it can be noted that, in the abnormal case (Fig. 6.10), at the positions where the thickening is greatest, when the contours containing valuable information are removed, the discrepancy will be high, which is also the case for the apex. Figure 6.12 shows the same experiment performed on all 24 cases. It can be seen that the highest discrepancies occur for the endocardial surfaces. It can be noted that the apical slices for the epicardial contours are significantly higher, and, when removed, the same error measurement occurs as the long axes. Both mesh results from normal and abnormal anatomies were pooled in the results. As the epicardial surfaces are typically comprised of simpler anatomy with respect to the endocardial surfaces, when the apical contour was removed, only a few attractor points remained on the long axes, causing the mesh to not fully recover the underlying anatomy of the epicardial surface. With regards to the endocardial surfaces, as they are more complex, and the long axes can be quite convoluted in an abnormal anatomy, more attractor points are selected

causing a lower contour-to-mesh error. These figures show that, even if the LAX contours are provided, the apical contour information should also be provided, in order to allow for a more accurate representation of the anatomy of the heart. Regardless, it can be seen that the main structural information is provided by the long axes contours. Furthermore, the information obtained from the most extremal slices is hard to interpolate, and causes the mesh to deviate from the underlying geometry, as opposed to the removal of a middle slice, which would not lead to a major loss of structural information. It should be noted that the methodology presented in this Chapter, works for different type of shapes, and not necessarily heart geometry. Appendix A.2 shows the application of SurFCo on other objects, such as the MIT bunny, and a pair of lungs.

Our method does come with some drawbacks, notably the ability to cope with branching only if branching information has been provided in the initial mesh. Once the 3D surface meshes have been obtained, several clinical anatomical parameters relevant in the clinic can be acquired such as the ventricular volumes, wall thicknesses, or sphericity indices. Furthermore, volumetric meshes can easily be produced for electrophysiology or mechanical assessments in a biophysical simulations framework.

6.5 Conclusions

We presented a tool for surface reconstruction applicable to sparse, non-parallel, non-coincidental, heterogeneous cardiac contours. Our technique works on the left ventricle as well as on the right ventricle. We have empirically shown by means of the exposition of a variety of results that our technique results in reconstructed interpolating surfaces with very good fitting to the input contours. In the areas to be filled by the interpolated surfaces, our technique presents very pleasant visual appearance. We also show the importance of the LAX information and its effect on the resulting geometries, as well as introducing a new validation protocol that allows the evaluation of the various aspects of our methodology.

7

Integrated Segmentation and Alignment Correction of Heart Images

Contents

7.1 Introduction 120
7.2 Proposed Framework 122
7.2.1 Automatic Segmentation 122
7.2.2 Methodology 126
7.3 Experiments 130
7.3.1 Data 132
7.3.2 Experiment 1: synthetic edgemaps 132
7.3.3 Experiment 2: Atlas template 134
7.3.4 Experiment 3: Real images 136
7.3.5 Experiment 4: Fully automated pipeline 137
7.4 Results 137
7.4.1 Automatic segmentation 137
7.4.2 Experiment 1 140
7.4.3 Experiment 2 141
7.4.4 Experiment 3 143
7.4.5 Experiment 4 144
7.5 Discussion 147

7.1 Introduction

Current methods for generating personalized CMR mesh models involve the steps of registration, segmentation and interpolation. Typically, slice misalignments occurring in 3D are first corrected for, as shown in Chapter 5, after which the segmented contours, obtained either manually or automatically, are used to create the mesh, such as was demonstrated in Chapter 6. As current contouring protocols extract objects of interest in 2D, when positioned in 3D, cross-sectional contours suffer from spatial inconsistency, as they are non-coincidental, as shown in Chapter 5 as well as in [87, 130]. These inconsistencies occur for several reasons. Any slice misalignments between the slices cause the images, and their respective contours to be misaligned. Moreover, uncertainties in the contouring because of slight differences in the image intensities at neighboring pixels, can cause the segmentations to over or under-estimate the region of interest in the respective intersecting slices. In cardiac imaging, false segmentation of papillary muscles in the LAX can occur, as it can be hard to distinguish from myocardium. There also exist an assumption that the underlying 3D shape from the slices is spatially consistent. This assumption can be lost as the slices are independently acquired, causing any breath-hold or triggering issues to acquire images having slightly different shapes. There is a clear need for an integrated method that combines the alignment of slices and a fine adjustment of the contours in order to arrange the segmented data consistently in a 3D environment, ultimately leading to the construction of geometrical meshes describing heart anatomy. It should be noted that throughout this Chapter we refer to the adjustment of the contours as correcting them. However, this does not necessarily imply that the contours were originally incorrect, but that the uncertainties can be improved towards a global description of the geometry.

A description of the prior work for slice misalignment correction can be found in Chapter 4. However, to the best of our knowledge, only one study has focused on an integration of slice alignment correction and mesh generation [129], and expanded in [128]. Paiement *et al* [128], demonstrate a technique that integrates these two steps by making use of level sets. They use level sets to automatically

segment and construct a 3D surface of the left ventricle on which they perform the alignment correction. By constraining the level sets to grow by using the combined information found in the intersecting slices, they use a registration methodology on the segmentations, integrating both methods. It should be noted that although they report registration results, no segmentation validation was performed, making the comparison of our proposed method unfeasible.

In this Chapter, a novel methodology is proposed, which we call ISACHI (Integrated Segmentation and Alignment of Heart Images) to correct for spatially inconsistent contours, allowing for a subsequent *patient-specific* surface mesh reconstruction, while aligning the image stacks. The approach is image driven, and does not necessitate any expert contouring, although an initial segmentation is required. By combining information from both the alignment as well as the segmentation steps in an integrated manner, an accurate shape that is spatially consistent with the information from the images can be obtained. By using an initial mesh, generated by using automatically segmented contours via a convolutional neural network (CNN), image derived forces can be used that rigidly transform individual image poses while deforming the contours from the mesh towards image edges. The resulting segmentation compiles both global and local information from the stack of short axis images as well as from the long axis views, being the most consistent possible at the intersection of slices.

To recapitulate, the initial segmentation allows for the creation of a surface mesh using the method described in Chapter 6 and in [87, 130]. The method provides a smooth interpolation fitting all the data points. At the points where the line intersections between intersecting slices (generally between SAX and LAX slices) intersect with the mesh, patches are computed, tangent to the mesh. The normal to the mesh triangle at that point is calculated and used as the normal to the patch. At the intersection of these patches and the slices, cross-sectional segments are generated, lying on the images, which guide the alignment and deformation parts of the algorithm. In the alignment step, each slice pose is optimized according to the image information lying at the patch segments, from all of the intersecting

slices. Once all planes have been transformed, the newly positioned patches undergo individual translations along the intersection lines between slices, following the images' information. These displacements are then extrapolated to their respective contours. The global process is performed in an iterative manner until convergence. The mesh is recomputed from the updated contours every iteration. The method incorporates a number of novel aspects, including the ability to correct for spatially inconsistent contours, and aligning the image stack together using both global and local information of intersecting slices, both resulting in a more correct representation of the anatomy and allowing the construction of accurate surface mesh.

7.2 Proposed Framework

Our proposed method aims to correct for contouring discrepancies in order to obtain 3D spatially consistent segmentations, resulting in a global description of the anatomy. This is achieved by relying on image information features. Any type of image information may be chosen by the user, as long as it is able to guide the segmentation towards similar areas reflected in the intersecting images, incorporating the 3D consistency. To that extent we calculate an *edgemap* in which pixels are assigned a certain edgeness, reflecting the likelihood of being part of an edge. The method to compute the edgemap is left to the user, and is not detrimental to our method as long as it follows the aforementioned premise. The method can work with fundamental techniques such as the magnitude of the gradient, although newer techniques can be used such as the one proposed by Xie *et al* [152]. In this work, the edgemap is computed using the local phase of the images, as introduced in Chapter 3 and Chapter 5.

7.2.1 Automatic Segmentation

The method requires an initial contouring of the slices, which are then combined to create an initial 3D surface mesh ($\mathcal{M}$). The contours can be obtained in any way; via human interaction, as well as automatically via convolutional neural networks (CNN), for example. In order to fully automate the method, a UNet

CNN based on [153] that have recently been shown to provide state-of-the-art segmentation results [154] is used.

7.2.1.1 Data

The U-Net CNN was trained on images from 100 cases obtained from a collection of both pathological and normal CMR cases from the John Radcliffe Hospital, Oxford, UK (IRAS project ID: 112814). The images were acquired on a 3.0 Tesla Trio MR system (Siemens Healthcare). The data consisted of 35 normal cases, and 65 cases that ranged from having slight to severe HCM. HCM was defined as having a LV hypertrophy $\geq$15 mm. Each of the cases consisted of end-systolic and end-diastolic phases, comprising of 2-3 LAX, and 7-12 SAX. All the cases had been manually segmented using the commercial software cmr42© (Circle Cardiovascular Imaging, Calgary, Canada), including the LV epicardium and endocardium contours, as well as the RV endocardium. Both normal and abnormal cases were used in order to increase the network's robustness in dealing with variable shapes. The segmentations were converted to labeled images containing three labels, as shown in Fig. 7.1. The labels represent LV myocardium, blood pool, and RV blood pool. The total dataset consisted of $1,827$ images. In order to train our CNN we performed data augmentation via random rotations as well as flipping, as this has been proven effective in improving segmentation results [155]. Finally a total of 3000 images were used. The data was split into training (2400 images) and validation (600 images) data, and optimized to minimizing categorical cross entropy.

7.2.1.2 Automatic cropping

In order to automatically crop the data, the 3D spatial information for each of the cases is used. All the images are positioned in their respective 3D spatial locations, as shown in Fig. 7.3a. Once in 3D, the line intersections between all planes are computed. As all the cases have two long axes slices (LAX), there exists an intersection line between the two LAX which crosses all the short axes slices (SAX). As the clinical field of view in cardiac MR images is most commonly focused on the left ventricle, the LAX line of intersection crosses all the SAX centered around

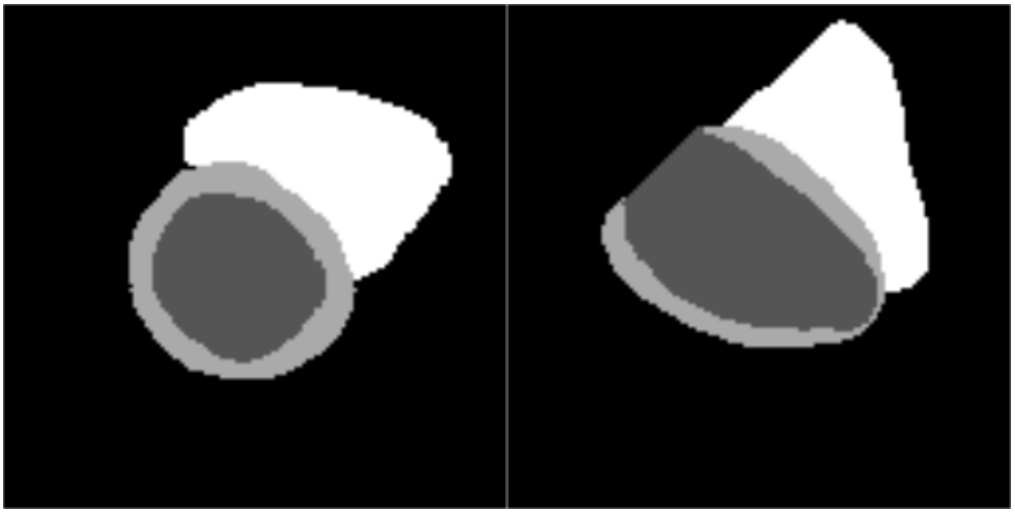

Figure 7.1: Training labels for the U-Net. There are 3 labels representing the blood-pool (dark gray), the myocardial wall (light gray), and the right ventricle (white). The left image represents a mid-ventricular SAX. The right image represents a LAX.

the left ventricle. The intersection line between all the SAX and the LAX are also computed. From these lines, a central point lying on the LAX-LAX intersection line and with the minimum distance to all of the other line segments can be computed, as shown in Fig. 7.3b. Finally, a cube is computed which encapsulates the heart, as shown in Fig. 7.3c. The slice information outside of this cube is masked out.

7.2.1.3 CNN architecture

A UNet CNN was chosen as they have been found to work particularly well on low volumes of training data as well as being able to segment touching objects of the same class [153, 154, 156]. There exist more complex approaches such as Res-net, SegNet or DenseNet [157–159], however, as only an initial segmentation is required, this particular choice of CNN performs to an appropriate standard. The network was trained following the architecture shown in Fig. 7.2. The images were cropped, following the method described in Section 7.2.1.2, to be 128x128 pixels, which incorporate both left and right ventricles (185 mm squared). The network consists of two parts; a down-sampling (right) and an up-sampling section (left). By using this architecture, both global and local information can be taken into account [154].

To be able to retrieve a segmentation map at the same spatial resolution as the

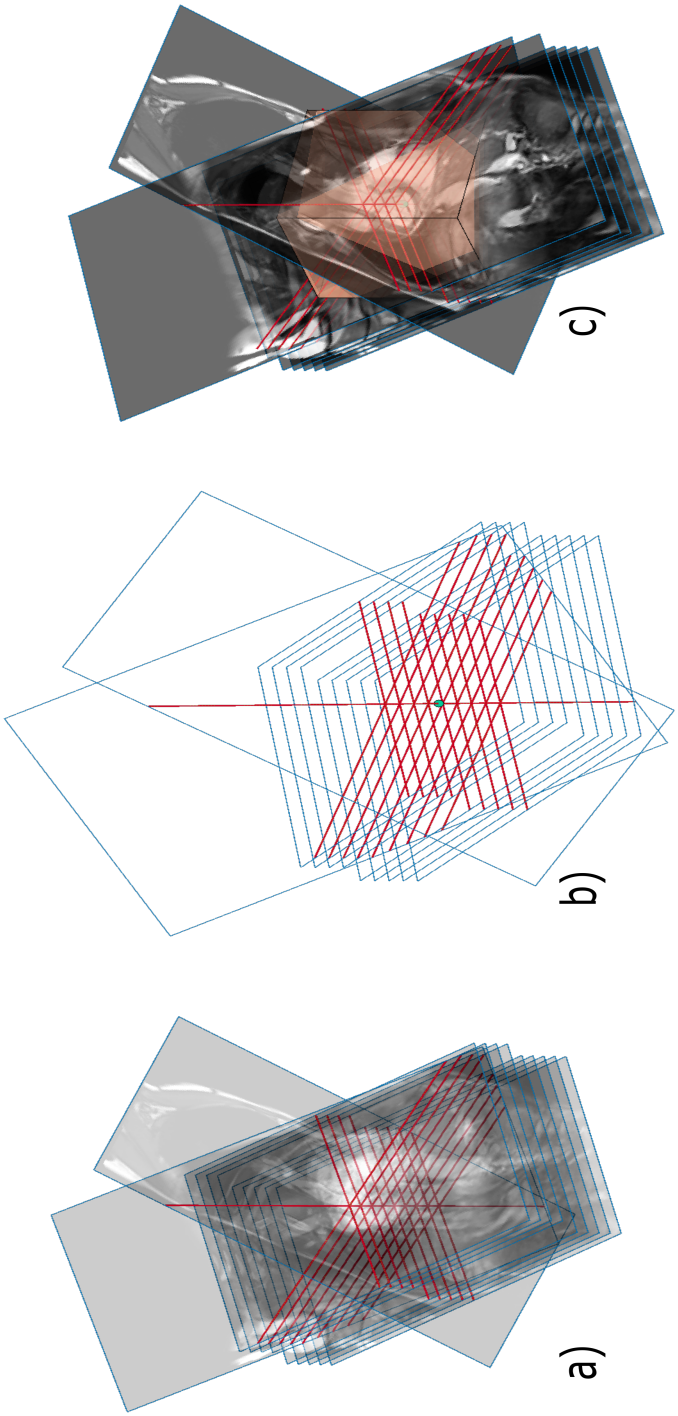

Figure 7.3: Automatic Cropping Pipeline. a) Image slices with line intersections between the various slices highlighted in red. b) Centre point (in green) computed as the minimum distance between all the line segments. c) Cube representing the cropping region encapsulating the heart.

input images, UNets use up-sampling operators on the right side of the network (shown as green arrows in Fig. 7.2).

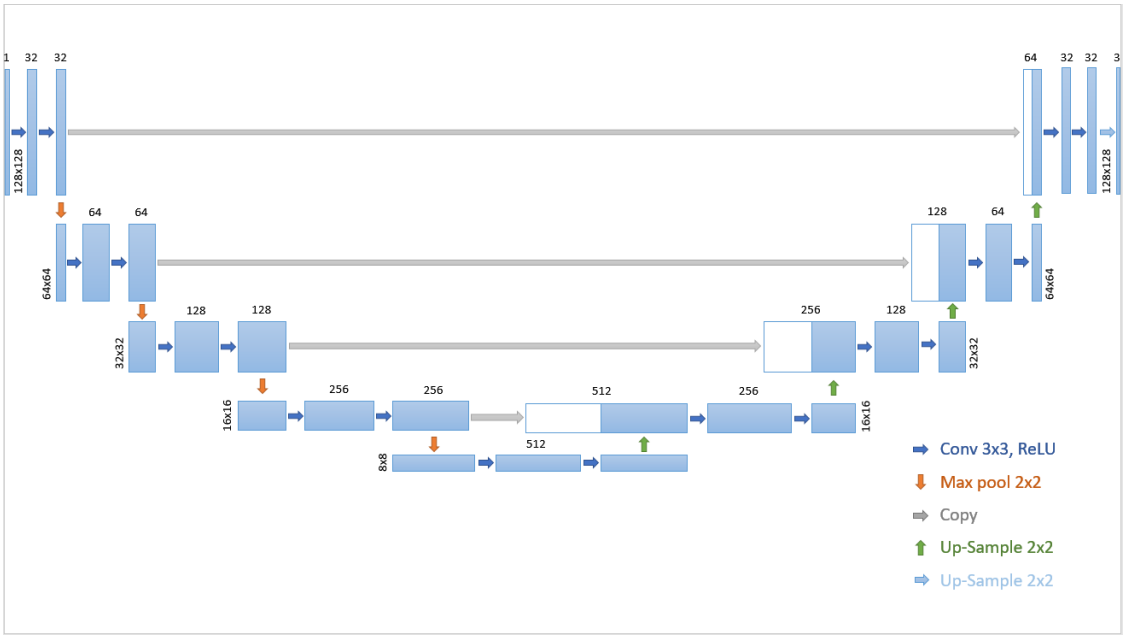

Figure 7.2: U-Net CNN architecture used for the automatic segmentation.

7.2.1.4 Evaluation of CNN Results

A validation dataset consisting of 10 normal cases and 5 abnormal cases were used to validate the network's performance on unseen data from the same cohort. The accuracy of the results was assessed against manual segmentations for the left ventricle epicardium and endocardium as well as the right ventricle endocardium, on a total of 81 segmentations for the normal cases, and 45 segmentations for the abnormal cases. In addition, we tested the network's performance on myocardial infarction cases, described in Chapter 5. The 15 Cine datasets at end-systole and end-diastole were used for validation, making up 30 datasets. In total, the left ventricular end-diastolic epicardium and endocardium segmentations accounted for 171 evaluation cases, while the right ventricular endocardium accounted for 84 evaluation cases. The left ventricular end-systolic epicardium and endocardium segmentations accounted for 159 evaluation cases, while the right ventricular endocardium accounted for 57 evaluation cases.

To assess the network's performance, the Dice coefficient (DC) between the network and expert segmentation was computed [160]. The Dice coefficient gives a measure of overlap between two objects and is given by

$$DC = \frac{2A_{X\cap Y}}{A_X + A_Y}\,, \tag{7.2.1.1}$$

were A_X, A_Y represent the individual areas of the objects X and Y, and $A_{X\cap Y}$ represents the area of the intersection between X and Y. The results will be shown in Section 7.4.1.

7.2.2 Methodology

Having an initial set of contours, an initial mesh can be generated following the methodology describe in Chapter 6. We compute the line intersections between all cross-sectional planes and obtain the points where the lines intersect with the generated mesh. Let S_i be the plane where contour C_i lies, with $i \in [1...N]$. Let l_{ij} be the line intersection between planes i and j, that is, $l_{ij} = S_i \bigcap S_j$; with $i < j$, or $l_{ij} \equiv l_{ji}$. Let $\{P^n\}$ be the set of points where the lines l_{ij}, intersect with the mesh $\mathcal{M}$. As such, $\{P^n\} = \{\mathcal{M} \bigcap l_{ij}, i < j\}$ with $n \in [1...T]$ and T being equal to the maximum amount of intersection points. At these points *patches* are created. Let D^n be a circular patch lying on the plane tangent to $\mathcal{M}$ with center P^n, that is, $D^n = (\top_{P^n}(\mathcal{M})) \bigcap (B_{P^n} R)$, where $\top_{P^n}(\mathcal{M})$ is the tangent to $\mathcal{M}$ at P^n and $B_{P^n} R$ is the ball centered at P^n with radius R . Let s_i be a segment of length $2R$, lying on S_i and D^n, that is, $s_i = D^n \bigcap S_i$ only if D^n has been built from S_i. It should be noted that each patch contains 2 segments, belonging to each of the respective intersecting planes, and that there are as many patches as there are intersections between the mesh and line intersection l_{ij}. See illustration in Fig. 7.4. The patches represent a first order (linear) description of the mesh in the surroundings of the intersection lines. These surroundings are expected to be the regions where the mesh consistently match the image boundaries.

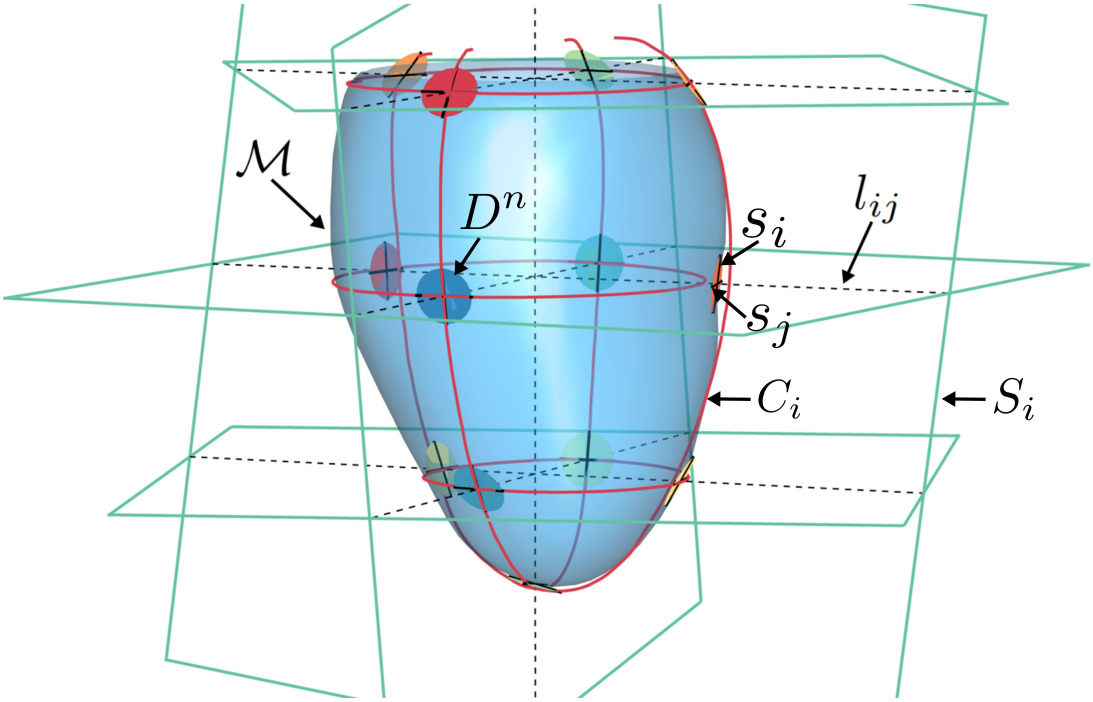

Figure 7.4: Representation of ISACHI definitions. The different colored circles represent the patches, tangent to the mesh. The black dotted lines represent the line intersection between planes, which are represented in green.

In our implementation, in order to get access to the patches, or slices during the optimization process, we define two functions: $g(n)$ and $p(i)$ that return slice indices when given a patch number, and all the patches having been built from a specific slice S_i^n, respectively. Let $g(n) = \{i, j\}$ and $p(i) = \{n : i \in g(n)\}$. For a slice i, we now define our energy function to be:

$$J_i = \sum_{n \in p(i)} \sum_{k \in g(n)} \xi(s_k^n, \Psi_k) \qquad (7.2.2.1) \qquad \xi(s, \Psi) = \frac{\int_{x \in s} \Psi_{(x)}}{2R} \qquad (7.2.2.2)$$

,

where Ψ is the edge map of an image, s_k^n is the n segment belonging to slice k , and ξ represents the energy function between a segment and an edge map. Ψ represents any edge map that the user chooses as long as the areas of high edgeness correlate with edges and boundaries. The energy J_i represents a measure of how well each of the segments agrees with the image boundary information. It relies on

all the segments lying on the slice, as well as the segments lying on the respective intersecting slice. A patch's global energy lies at the area of highest edgness between the two intersecting slices. At every step in the optimization, the intersection lines are recomputed, along with their respective points of intersection with the mesh and the normals at those points, allowing to obtain new patches, and segments.

There exists a global energy

$$\widetilde{J} = \sum_{i=1}^{T} J_i \,, \tag{7.2.2.3}$$

where T is the maximum number of patches, and inherently the number of intersection points.

One of our aims is to align slices together in order to obtain spatial consistency between the slices and inherently the contours. We assume that slices taken from a subject are acquired at the same cardiac phase and as such there is a global optimal alignment between the dataset where the 3D shape/anatomy of the heart remains fixed among slices. We thus consider rigid body transformations with 6 degrees of freedom, represented by: $\theta = \{t_x, t_y, t_z, \alpha, \beta, \gamma\}$. Our optimization can then be defined as: $\underset{\theta_i \in SE(3)}{\operatorname{argmin}}\ J_i(\theta_i)$, where SE(3) refers to the special Euclidean group of dimension 3 and

$$J_i(\theta_i) = \sum_{n \in p(i)} \sum_{k \in g(n)} \xi(s_k^n(\theta^1_{g(n)}, \theta^2_{g(n)}), \theta_k \star \Psi_k) \,, \tag{7.2.2.4}$$

where $\star$ denotes the application of θ on Ψ. Once we have iterated through all the slices until convergence, we consider the slices to have been aligned.

Practically, the optimization is constrained such that no slice is able to translate or rotate more than a certain amount. Furthermore, a transfer function is applied to allow for a smooth optimization space. Let θ^R be the 3D rotation parameters, θ^T be the 3D translation parameters, and $\mathcal{R}, \mathcal{T}$ be the rotation and translation bounds respectively. The minimization function becomes: $\underset{\theta_i \in SE(3)}{\operatorname{argmin}}\ J_i(\{\mathcal{R} \times \tanh\left(\frac{\theta_i^R}{\mathcal{R}}\right), \mathcal{T} \times \tanh\left(\frac{\theta_i^T}{\mathcal{T}}\right)\})$.

As the initial contours might not be a perfectly accurate representation of the segmentation, we proceed to transform individual patches $\{D^n\}$ spatially. As the

patches lie on the mesh, which is an approximation between the non intersecting contours, we consider them to be a good initial approximation of their spatial consistency. As we are not optimizing the entire plane, and just want to correct for the segmentation, we only allow the patches to move along the intersection lines. In this way, we search for the best spatial location of a patch taking into account the information from its respective intersecting slices. Every patch is optimized iteratively following the same method previously described, but taking into account both intersecting segments. We thus want $\mathrm{A}^* = \underset{\theta_\mathrm{i} \in \mathrm{SE(3)}}{\operatorname{argmin}} D^n(\theta_i)$, constrained to the line intersection l_{ij}. Once the patch has been optimized, the deformation A^* is extrapolated to the contours via thin plate splines (TPS) [149]. This is to ensure that the patches lie on their optimum calculated positions whilst approximating the mesh in a regularized manner

$$\xi_{TPS} = reg(u) + \lambda \sum_i ||u(m_i) - f_i||^2 , \tag{7.2.2.5}$$

where u represents the deformation of the space, m the moving points, f the target points. The parameter λ is a smoothing parameter controlling the trade off between the smoothness of the deformation and the data reliability. Our choice of λ was chosen empirically using a data driven approach. Once deformed, we proceed to recompute a mesh $\mathcal{M}$ to obtain a representation in accordance with the new spatial arrangement of the contours.

Practically, the edgemap Ψ is separated into its respective parts; left ventricular epicardium and endocardium, and right ventricular endocardium. This allows only corresponding edge maps to be matched. As the contours are obtained automatically via the CNN, or by manual segmentation, it is assumed that the *true* edge lies within a certain region from the contour. To that extent, a Gaussian mask modulated from each of the contours is applied to the edgemap in order to extract the respective parts. The mask represents a distance function from the contours, with respect to a Gaussian of a chosen width. This forces the patches to move towards the respective contour edge, as the energy is calculated with regards to the patch lying on its corresponding edgemap. An example can be seen in Fig. 7.5.

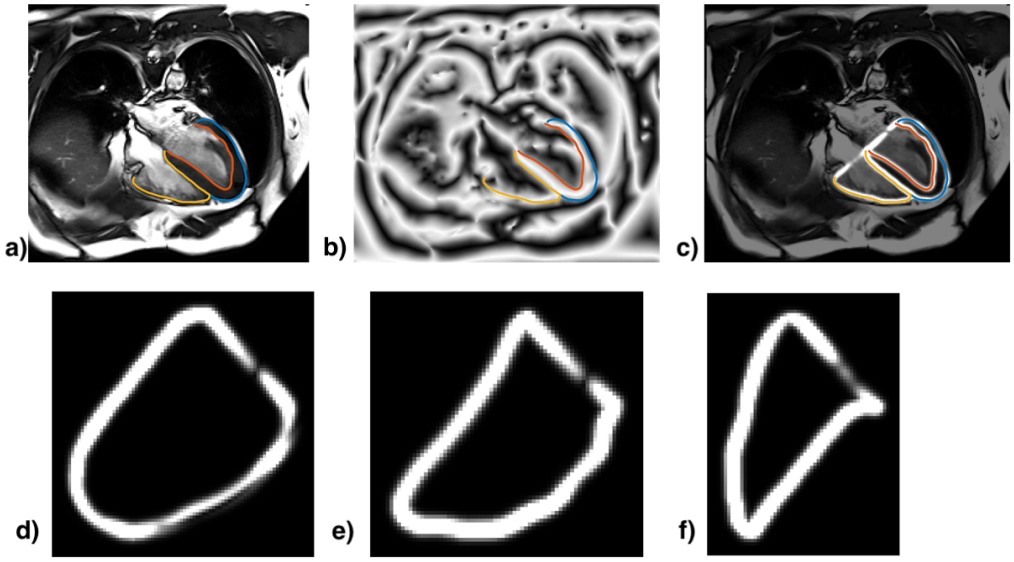

Figure 7.5: Example of ISACHI edgemaps. **a)** Original image with manual segmentations representing the left ventricular epicardium (blue), left ventricular endocardium (orange) and right ventricular endocardium (yellow). **b)** Edgemap computed from the local phase (LogGabor parameters: $\omega = 20$, $\sigma = 0.5$). **c)** Original image with the edgemap modulated by the Gaussian distance map with respect to the individual contours. **d-f)** The individual edgemaps representing the left ventricular epicardium, left ventricular endocardium and right ventricular endocardium respectively.

7.3 Experiments

To properly assess the methodology, both synthetic and clinical data were used. As the algorithm encompasses both alignment and segmentation correction, synthetic data allowed each component of the method to be evaluated, similar to Chapter 6, as we do not possess a *true* spatially consistent clinical ground truth. For the quantitative validation of clinical data, the normalized cross correlation of the line intensities between the images at the intersection lines (where the intensity profiles are expected to coincide) was used a an alternate measure to estimate the accuracy of the alignment process.

Several experiments were performed in order to assess the reliability and validity of the algorithm at every step. In a first experiment, synthetic images are generated from synthetic consistent contours obtained from a 3D mesh (as per Chapter 6) [146]. From the 3D mesh, the contours are obtained by slicing the mesh and random perturbations on the poses are applied. Images are synthesized, representing a

distance map from the contours. The methodology described in Section 7.2.2 is used, and validated assessing the method's ability to retrieve a mesh generated from spatially consistent contours. Our method aims at obtaining spatially consistent contours for the generation of a 3D mesh; however, due to the way the realignment is performed, the mesh can end up displaced from its original location. This does not have any negative effects on the final mesh, but can result in significant mesh-to-mesh distances in the validation. To compensate for this, we apply a final rigid registration between the output and the ground truth mesh, using the Iterative Closest Points (ICP) algorithm [161]. Point-to-point distances between the two meshes are computed, to quantify the algorithm's ability to produce an accurate mesh. This experiment allows to quantify ISACHI's ability to realign the slices and quantify whether the non-rigid deformation step does not stray from the correct shape.

In a second experiment, average images generated from an atlas were used to assess the algorithm's ability to work with images. Although generated from real images, these datasets are from here on referred to a synthetic. Slices are generated from the volumetric atlas, and contours synthesized by slicing the mesh at the slice locations. The slices are misaligned and the method is applied to evaluate the ability of ISACHI to realign the slices, based on edgemaps derived from synthetic images. A validation is performed by comparing the 3D mesh and the resulting ISACHI mesh output as in experiment one. The cross correlation between the line intensities for all slice intersections (as explained in Chapter 5) is also computed to validate the algorithm's ability to correct for slice misalignment.

A third experiment is conducted on clinical data, using expert delineated contours as the initial segmentation, to assess if the method is able to achieve spatially coherent contours whilst aligning the slices in a visually satisfying manner. As the output result contains spatially consistent contours, the contour-to-contour distance cannot be used as a validation measure. However, we compute the dice coefficients before and after the algorithm between the spatially consistent contours and the manual contours, to assess if the contours have deformed severely from

their initial ones. A qualitative result will be given by means of visual appearance. A quantitative validation is also performed using the alternate intensity measure.

A final experiment is performed using clinical data along with an initial segmentation obtained via a CNN. This represents a fully misalignment correction and segmentation pipeline on real-world data. Validation is performed in a similar manner to Experiment 3.

7.3.1 Data

Synthetic Data

An average MRI mesh obtained from [146] and available at `http://wp.doc.ic.ac.uk/wbai/data/` was used for our data. Biventricular atlas at end-systole and end-diastole are available with their respective surface mesh models.

Clinical Data

Clinical data was used consisting of DICOM files containing 2D CINE MRI and 2D LGE sequences of size 216 x 256 pixels (approx: 1.41 mm x 1.41 mm). 15 datasets (SIEMENS TrioTim 3 Tesla scanner at the John Radcliffe Hospital, Oxford, UK) from different subjects were used. Each dataset contains between 9 and 13 SA slices from apex to base separated by 8-10 mm and 2 LA slices (4 chamber view, and 2 chamber view).

7.3.2 Experiment 1: synthetic edgemaps

From a synthetic 3D mesh $\mathcal{M}_{\mathcal{GT}}$, synthetic and consisten contours $\mathcal{C}_{\mathcal{GT}}$ are generated, by slicing the mesh, at different intervals as to simulate a stack of SAX and a couple of LAX. Edgemaps are synthesized as a distance map from the contours modulated by a Gaussian radial basis function, where the width of the Gaussian (σ) controls the distance length. An example of edgemap and the choice of σ can be seen in Fig. 7.6 .

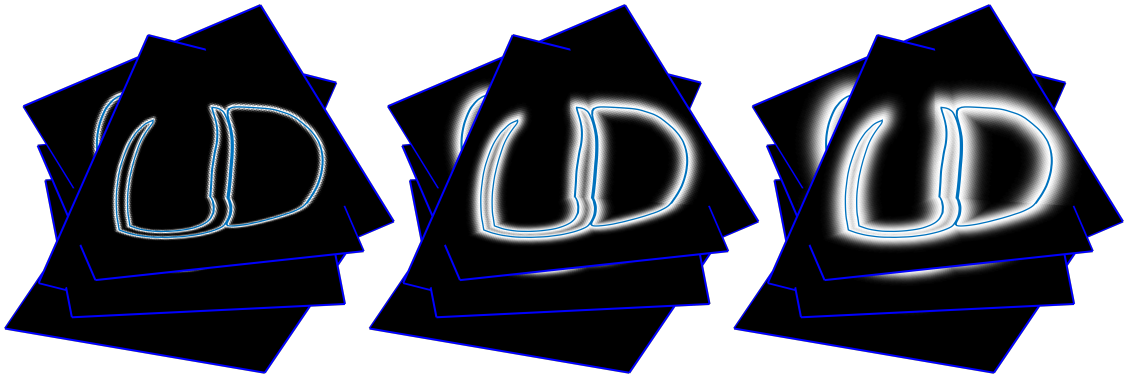

Figure 7.6: Synthetic edgemap obtained by a Gaussian radial basis function. Left: $\sigma = 1$. Center: $\sigma = 3$. Right: $\sigma = 5$. The contours (blue) have been obtained by slicing a 3D mesh.

In this initial experiment, ground truth contours, $\mathcal{C}_{\mathcal{GT}}$, were generated from the ground truth mesh $\mathcal{M}_{\mathcal{GT}}$, making up a total of six contours (four SAX ranging from the base to the apex, and two LAX - both including the right ventricle). The initial data can be seen in Fig. 7.7. Random rotations and translations (6 DoFs) were applied to three of the four SAX, to simulate slice misalignment, following the ranges found in [20]. The rotations applied ranged between -2 to 2 degrees in each of the Euler angles, and the translations applied between -5 to 5 mm. However, as the transformations are applied to each of the three slices independently of the other, the overall stack might be heavily misaligned. The cumulative range (the sum of all perturbation for all 3 slices) are $1 - 5$ degrees for the rotations, and 3 - 12 mm, for the translations. These misaligned contours $\mathcal{C}_{\mathcal{PT}}$ are considered to be the initial segmentation. A total of 10 different cases were generated with the aim to recover $\mathcal{M}_{\mathcal{GT}}$ having only been given $\mathcal{C}_{\mathcal{PT}}$. A table of the perturbations applied to the SAX slices can be seen in Table 7.1. The alignment method was applied to each of the cases, and validated in two ways; by using the normalized cross correlation of the line intensities as a surrogate measure as well as generating a new mesh from the output contours and comparing it with $\mathcal{M}_{\mathcal{GT}}$.

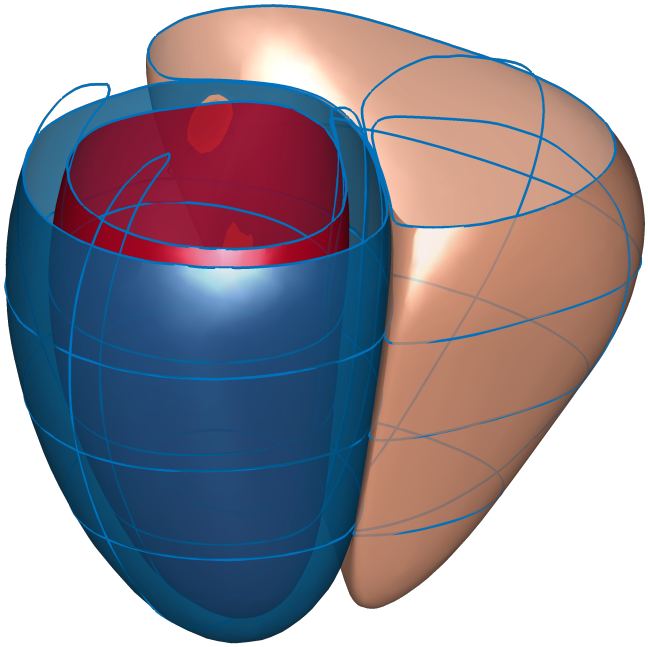

Figure 7.7: Initial data for ISACHI's first experiment. The ground truth mesh ($\mathcal{M}_{\mathcal{GT}}$ with its respective ground truth contours $\mathcal{C}_{\mathcal{GT}}$ can be seen.

7.3.3 Experiment 2: Atlas template

For this experiment, atlas data was used to assess the algorithm's ability to work with controlled images. A biventricular volumetric atlas at end-diastole is available from [146] with the respective surface mesh model ($\mathcal{M}_{\mathcal{GT}}$). The mesh along with the respective image template is used as the ground truth. Using the surface mesh, contours ($\mathcal{C}_{\mathcal{GT}}$) where synthesized by slicing it at various parallel and non parallel spatial positions. For this experiment, the dataset was also composed of four SAX and 2 LAX. The initial data can be seen in Fig. 7.7 (right). The same random translations and rotations as the first experiment were applied to the image planes. Furthermore, non-rigid deformations were applied to the contours ($\mathcal{C}_{\mathcal{PT}}$) to simulate uncertainties in the contouring resulting in non-spatial consistency between the contours. The contours were deformed in-plane, by adding random displacements in the x and y components of the contours, and applying a TPS deformation to the displaced points. Small deformations in the order of $< 2\,\text{mm}$ were chosen, in

		$\lvert\theta_\alpha\rvert$	$\lvert\theta_\beta\rvert$	$\lvert\theta_\gamma\rvert$	$\lvert\theta_x\rvert$	$\lvert\theta_y\rvert$	$\lvert\theta_z\rvert$
A	SAX_1	0.486	1.871	0.352	0.006	3.126	1.437
	SAX_2	0.308	1.637	0.721	1.582	2.715	2.508
	SAX_3	1.913	1.457	0.378	3.498	2.195	3.808
B	SAX_1	1.525	1.291	0.695	0.991	3.472	2.661
	SAX_2	1.152	0.246	0.184	3.361	1.284	1.397
	SAX_3	1.495	1.009	0.296	2.158	0.049	4.731
C	SAX_1	1.813	1.343	0.114	3.433	3.635	0.581
	SAX_2	0.785	1.674	0.901	3.597	1.869	0.288
	SAX_3	0.050	1.943	1.165	3.250	2.908	4.90
D	SAX_1	0.570	0.372	1.866	0.750	0.656	3.073
	SAX_2	1.190	0.386	0.781	1.985	2.175	0.055
	SAX_3	1.924	0.683	0.546	1.874	0.458	2.870
E	SAX_1	1.580	1.139	1.285	4.855	1.394	4.79
	SAX_2	0.471	0.123	0.443	4.232	3.733	3.101
	SAX_3	0.896	0.993	1.674	2.530	1.185	3.001
F	SAX_1	0.345	1.717	1.450	4.053	0.450	0.303
	SAX_2	0.181	1.822	0.460	2.019	1.605	3.628
	SAX_3	0.510	1.399	1.152	4.942	2.557	2.783
G	SAX_1	1.059	1.578	1.505	1.349	3.552	4.252
	SAX_2	1.660	0.636	0.220	2.623	1.559	4.558
	SAX_3	1.717	0.904	0.219	4.863	1.457	3.196
H	SAX_1	0.511	1.169	1.169	0.955	4.133	1.590
	SAX_2	0.177	1.896	0.570	2.213	3.384	0.669
	SAX_3	1.676	0.122	1.655	1.967	1.038	3.357
I	SAX_1	1.142	0.952	0.066	2.257	1.821	4.357
	SAX_2	0.339	1.816	0.108	1.913	2.662	1.643
	SAX_3	0.295	1.104	1.610	3.948	3.558	3.251
J	SAX_1	1.950	0.828	1.517	3.906	4.011	2.492
	SAX_2	0.152	0.618	1.990	0.979	2.121	4.045
	SAX_3	1.174	0.528	0.373	4.962	3.644	1.782

Table 7.1: Random transformations applied to ground truth contours. (10 cases)

order to simulate potential unintentional human error in contouring. An example of the perturbed dataset can be seen in Fig. 7.8.

In order to validate the method, a surface mesh $\mathcal{M}_{\mathcal{SF}}$ was reconstructed using the method found in Chapter 6 from $\mathcal{C}_{\mathcal{GT}}$, and the node-to-surface distance was computed, to assess the variability between the ground truth mesh and the mesh reconstructed from the contours. Another mesh was generated, from $\mathcal{C}_{\mathcal{PT}}$, and this mesh was compared to $\mathcal{M}_{\mathcal{SF}}$, to assess the initial error. This allowed a direct

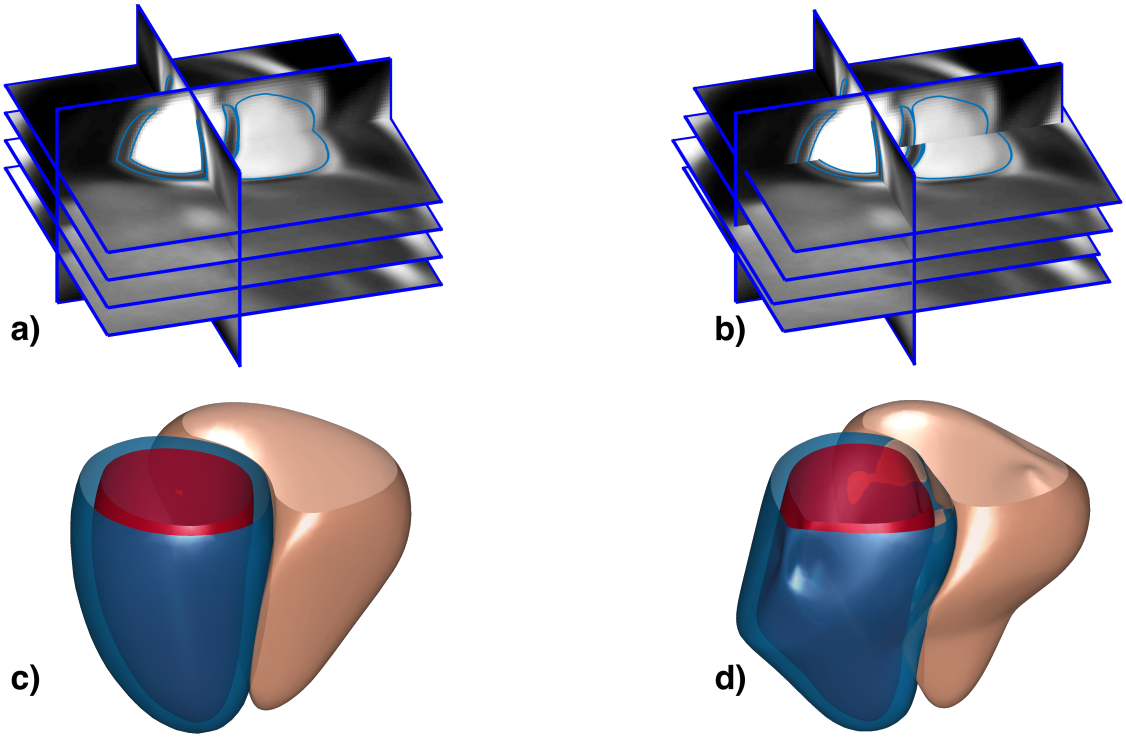

Figure 7.8: Overview of the second experiment using a biventricular atlas template. **a)** The original template along with its segmentation. **b)** Misaligned dataset, as well as the respective contours, which have also been deformed non-rigidly. This dataset represents the input to the algorithm. **c)** The mesh reconstructed from the contours in a). **d)** The mesh reconstructed from the contours in b).

quantification of the misalignment and contouring error, reflecting potential errors found in clinical data. Once the method was applied to the dataset, a final mesh was computed from the newly aligned and deformed contours $\mathcal{M}_{\mathcal{IA}}$. To allow a proper comparison between $\mathcal{M}_{\mathcal{IA}}$ and $\mathcal{M}_{\mathcal{SF}}$, the ICP algorithm was performed on the meshes, aligning them to the same space. The distances between the two meshes were computed by taking the closest point-to-mesh distances to assess the ability of the method to recover the ground truth shape of the heart.

7.3.4 Experiment 3: Real images

The previous experiments allowed to assess various parts of the algorithm, notably its ability to realign the slices based on the segments lying on appropriate edge information (Experiment 1), and to deform the contours towards the relevant edges (Experiment 2). However, the use of atlas-generated images limits the applicability

of the validation results to real-world scenarios. For this reason, in this experiment, real images are used which contain visible misalignment. From the data described in Section 7.3.1, six cases from the normal and abnormal cohorts were randomly selected. The manually segmented contours by experts were used as the initial segmentation. Visual results for the six cases are shown, and the surrogate measures of the line intensities are given.

7.3.5 Experiment 4: Fully automated pipeline

A final experiment was run in a fully automated way using the same data as Experiment 3. The images were segmented using the CNN, and ISACHI was run on the dataset. Similarly to Experiment 3, visual assessment of the 6 cases are given, along with the alternate measures of line intensities. A new Dice coefficient for the resulting segmentation is also given by comparing them to the manual segmentations.

7.4 Results

This section reports the results of the different experiments, starting with the evaluation of the CNN. All the reported distances are in mm.

7.4.1 Automatic segmentation

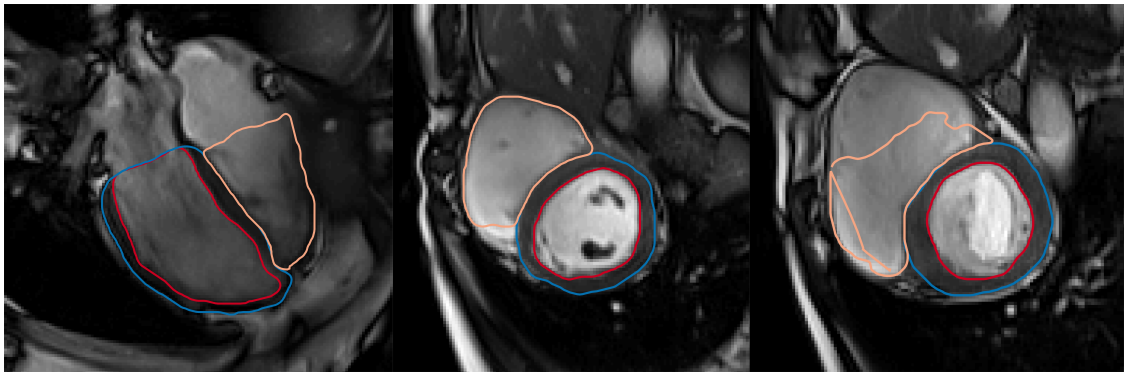

Figure 7.9: Example of CNN segmentation results on some HCM cases, showing a LAX, and two SAX. The rightmost image shows an example where the CNN performed poorly on the RV. The LV epicardium, endocardium and RV endocardium segmentations are shown in blue, red, and pink respectively.

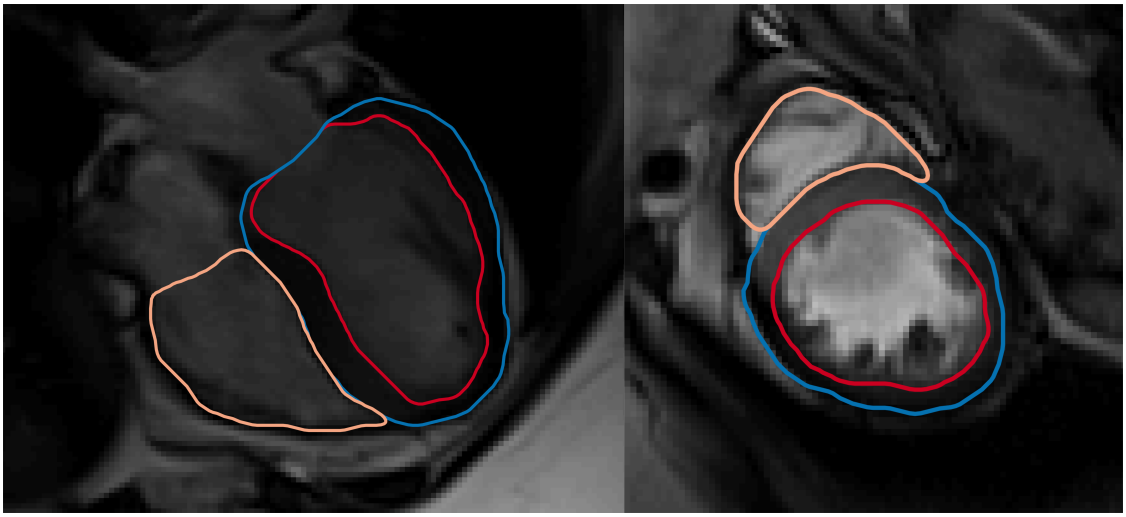

Figure 7.10: Example CNN segmentation result on the myocardial infarction dataset, for an LAX, and a SAX. The LV epicardium, endocardium and RV endocardium segmentations are shown in blue, red, and pink respectively.

Evaluation for the segmentation accuracy of the CNN was performed on a subset of the data described in Section 7.2.1, which we had excluded from the training dataset for validation purposes. Figure 7.9 shows an example of the CNN segmentation result on a LAX and two SAX. The leftmost figure represents an example where the CNN segmented the RV poorly. Figure 7.11 shows the resulting Dice coefficients for the 15 cases belonging to the same cohort as the training data. The figure represents the mean, median and standard deviation of the Dice coefficients for the 81 segmentations belonging to the normal cases, and 45 segmentations belonging to the abnormal cases. For the normal cases, the mean and standard deviations, as well as the median Dice coefficient for the LV epicardium were 0.87 ± 0.16 and 0.94. The LV endocardium Dice coefficient values for the normal cases were 0.87 ± 0.20 and 0.95, respectively. The mean and standard deviation, as well as the median Dice coefficient values for the RV endocardium were 0.73 ± 0.21 and 0.77, respectively. As for the abnormal cases, the mean and standard deviation along with the median values for the LV epicardium and LV endocardium were, 0.91 ± 0.05 and 0.92, and 0.92 ± 0.06 and 0.94, respectively. The values for the RV endocardium for the abnormal cases were 0.65 ± 0.27 and 0.77, respectively.

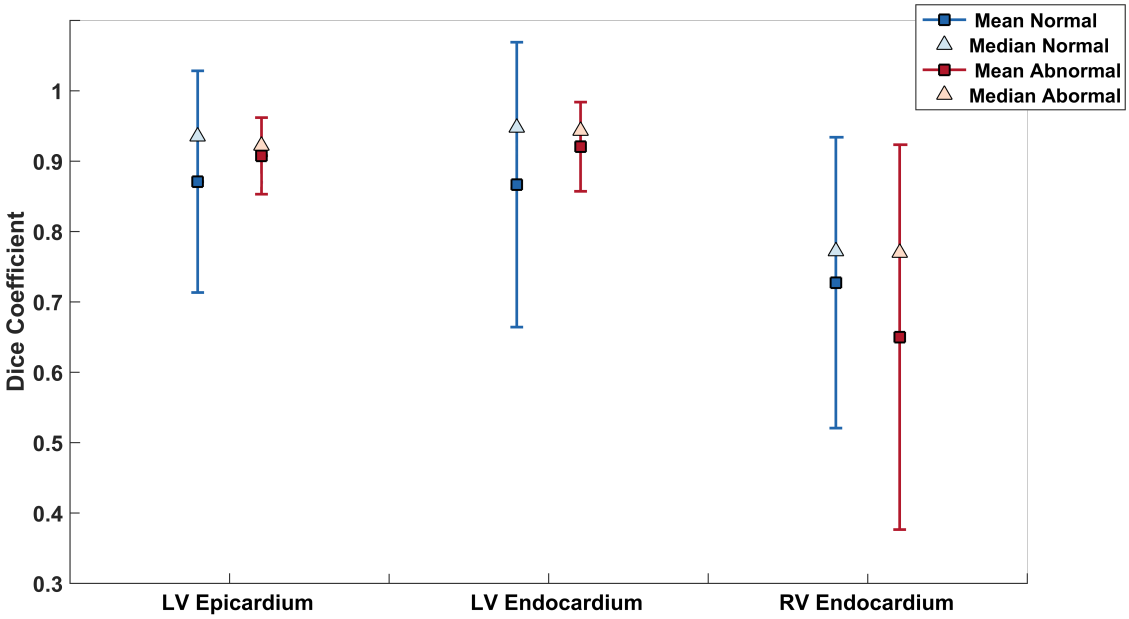

Figure 7.11: CNN performance validation on the data presented in Section 7.2.1.1. 15 cases were used for validation (10 normal, and 5 abnormal). In total 81 segmentations were compared for the normal cases and 45 for the abnormal ones. The blue and red lines represents the resulting mean Dice values of the left ventricle epicardium and endocardium as well as the right ventricle endocardium for the normal and abnormal data, respectively. The triangles represent the median values.

Another validation was performed on the myocardial infarction dataset to assess the CNN's ability to work on different images acquired from a different study. The Cine acquisitions at end-diastole and end-systole were used for the 15 datasets, making up 30 datasets. Figure 7.10 shows an example of the CNN segmentation on the myocardial infarction data for a LAX and a SAX. The results can be seen in Fig. 7.12. The mean and standard deviation and median values at end-systole for the LV epicardium and endocardium were 0.90 $\pm$ 0.11 and 0.93, and 0.85 $\pm$ 0.13 and 0.90, respectively. The RV endocardium values were 0.65 $\pm$ 0.22 and 0.73. The mean and standard deviation along with the median values at end diastole for the LV epicardium and endocardium were 0.91 $\pm$ 0.10 and 0.94, and 0.89 $\pm$ 0.13 and 0.93, respectively. The RV endocardium values were 0.76 $\pm$ 0.22 and 0.85.

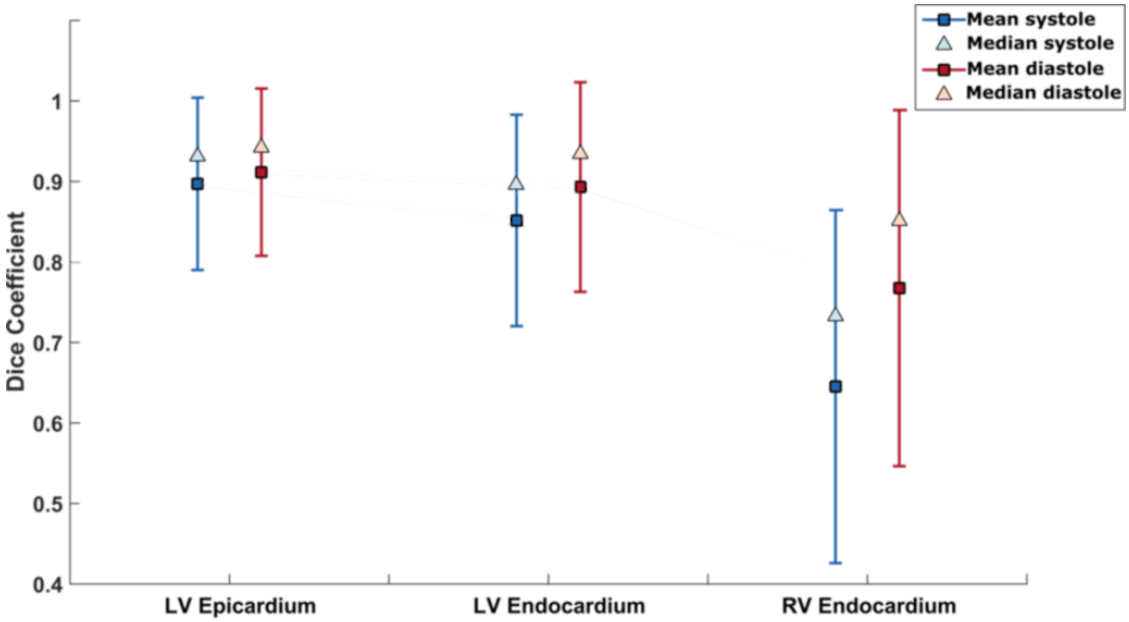

Figure 7.12: CNN performance validation on the myocardial infarction dataset presented in Chapter 5. The 15 datasets at end-systole and end-diastole were used for validation, making up 30 datasets. In total, the left ventricular end-systolic epicardium and endocardium segmentations accounted for 159 evaluation, while the right ventricular endocardium accounted for 57 evaluations. The left ventricular end-diastolic epicardium and endocardium segmentations accounted for 171 evaluation, while the right ventricular endocardium accounted for 84 evaluations. The blue and red lines represents the resulting mean Dice values of the left ventricle epicardium and endocardium as well as the right ventricle endocardium at end-systole and end-diastole, respectively. The triangles represent the median values.

7.4.2 Experiment 1

Figure 7.13 presents the input and output of a sample dataset used in Experiment 1. The resulting output can be observed as having recovered the initial alignment, and producing a similar mesh to $\mathcal{M}_{\mathcal{GT}}$. For this experiment, the contours were not deformed as the main objective was to assess the algorithm's capability to align misaligned data. For each of the 10 cases (labeled A-J from here on), the resulting mesh was registered to the initial mesh, as the slices might have shifted spatial position during the alignment process. The resulting average node-to-mesh distances were taken and can be seen in Table 7.2. An evaluation of the NCC for the line intensities can be seen in Table 7.3.

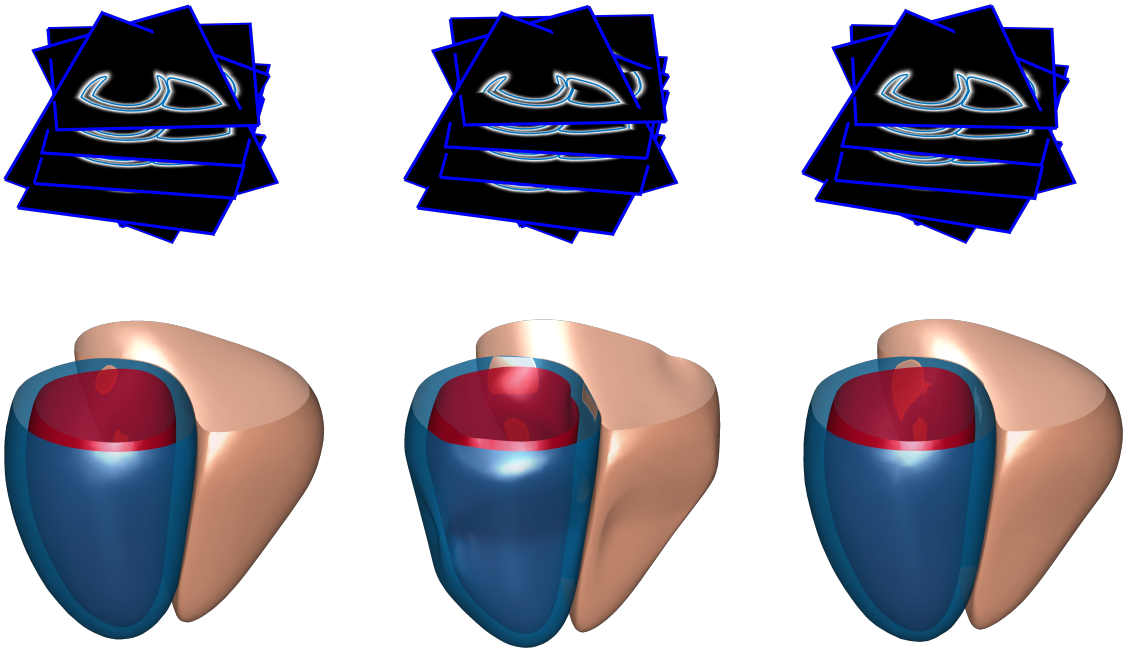

Figure 7.13: Visual Result for ISACHI Experiment 2. The left column represents the initial ground truth data. The center column represents the misalignment applied to the initial data (case **H**) and is the input to the algorithm. The right most column represents ISACHI's output, and the respective mesh derived from the aligned and spatially consistent contours.

7.4.3 Experiment 2

Figure 7.14 shows an example of the evolution of the algorithm after having applied in-plane misalignments (ranging between 3-5 mm) to the end-diastolic atlas template. Figure 7.14 shows and mean (red squares), and standard deviations for the NCC of the cross-sectional intensities between slices, as a function of ISACHI iterations. At the left of the vertical red line is the mean and standard deviation for the NCC values at the initial state. The right shows the subsequent evaluations for each slice update. The gray vertical lines represent the iterations. The bottom blue line represents the energy function $\widetilde{J}$ as described in Eq. (7.2.2).

After having run the method on the 10 cases (shown in Table 7.1), the node-to-mesh distance was taken between the resulting mesh, and the mesh generated from the ground truth contours. The meshes were first registered via ICP, to eliminate contour drift that might have occurred during slice realignment. The mean of the mean distance errors before the misalignments were: 2.03 mm, 1.94 mm, 2.21 mm

	Before			After		
Case	LV epi	LV endo	RV endo	LV epi	LV endo	RV endo
A	1.20	1.17	1.11	0.42	0.18	0.30
B	1.06	0.99	1.26	0.54	0.31	0.33
C	1.50	1.47	1.79	0.50	0.37	0.28
D	0.76	0.69	0.96	0.63	0.43	0.34
E	1.48	1.46	1.85	0.57	0.60	0.39
F	1.44	1.39	1.60	0.55	0.27	0.27
G	1.27	1.20	1.33	0.39	0.35	0.44
H	1.18	1.14	1.41	0.27	0.29	0.23
I	1.32	1.26	1.41	0.36	0.31	0.30
J	1.31	1.20	1.30	0.31	0.25	0.41
Mean	1.25	1.20	1.40	0.45	0.34	0.33

Table 7.2: Isachi results for Experiment 1. The mean left ventricular epicardium, left ventricular endocardium and right ventricular endocardium node-to-mesh distances are shown before and after the algorithm. Each before case represents the synthetic misalignment seen in Table 7.1 applied to $\mathcal{C}_{\mathcal{GT}}$.

	A	B	C	D	E	F	J	H	I	J
Mean (Before)	0.871	0.823	0.747	0.876	0.726	0.741	0.777	0.856	0.784	0.798
Mean (After)	0.994	0.994	0.994	0.994	0.993	0.995	0.993	0.995	0.994	0.994
STD (Before)	0.122	0.153	0.227	0.101	0.247	0.252	0.201	0.122	0.185	0.255
STD (After)	0.006	0.003	0.003	0.004	0.007	0.003	0.004	0.003	0.004	0.005
Median (Before)	0.896	0.801	0.704	0.849	0.752	0.798	0.762	0.836	0.726	0.926
Median (After)	0.995	0.995	0.995	0.995	0.995	0.995	0.995	0.995	0.996	0.996

Table 7.3: ISACHI line intensity NCC results for Experiment 1. The table shows the mean, standard deviation and median for the NCC of the inter slice line intensities before and after ISACHI.

for the left ventricular epicardium, endocardium and right ventricular endocardium, respectively. After ISACHI, the respective errors were: 1.12 mm, 0.83 mm, 0.83 mm. All results showing the mean of the distances can be found in Table 7.5. The normalized cross correlation of the intensities between the slices was also taken as a surrogate measure. The results can be seen in Table 7.4. An example of the visual result for case **H** can be seen in Fig. 7.15.

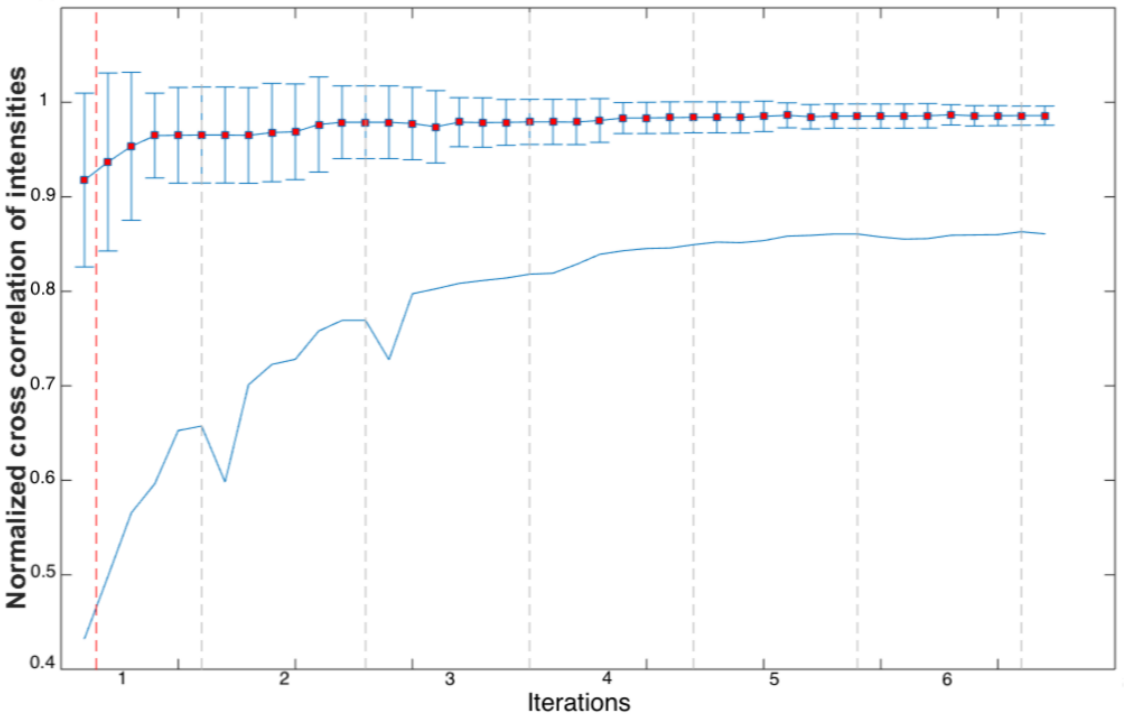

Figure 7.14: Alignment correction validation. The top part of the graph represents the mean and standard deviation values for the NCC of the cross-sectional slice intensities, at every step of ISACHI. The left of the vertical red line represents the synthetically misaligned dataset. The bottom line represents the energy function $\widetilde{J}$.

	A	B	C	D	E	F	J	H	I	J
Mean (Before)	0.954	0.974	0.954	0.984	0.955	0.959	0.967	0.971	0.970	0.945
Mean(After)	0.950	0.980	0.984	0.995	0.980	0.990	0.992	0.974	0.992	0.984
STD (Before)	0.037	0.032	0.054	0.024	0.055	0.064	0.055	0.036	0.039	0.096
STD (After)	0.064	0.037	0.021	0.005	0.031	0.009	0.008	0.036	0.007	0.023
Median (Before)	0.948	0.990	0.962	0.992	0.968	0.983	0.988	0.978	0.979	0.989
Median (After)	0.971	0.995	0.996	0.996	0.993	0.993	0.995	0.984	0.996	0.994

Table 7.4: ISACHI line intensity NCC results for Experiment 2. The table shows the mean, standard deviation and median for the NCC of the inter slice line intensities before and after ISACHI.

7.4.4 Experiment 3

The method was applied on real data from a normal and abnormal cohort. Six cases (3 normal and 3 HCM cases) were selected for visual assessment. Figure 7.16 shows the result of an abnormal case before alignment, and Fig. 7.17 shows the result for the same case, after the application of ISACHI. The resulting meshes

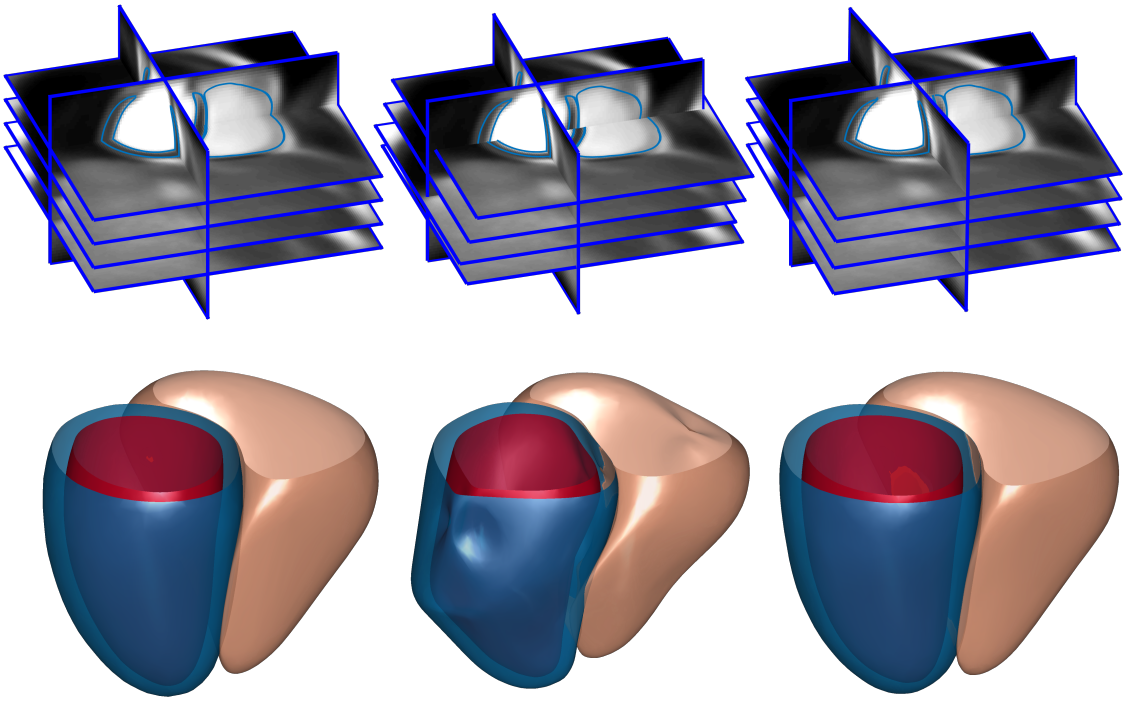

Figure 7.15: Visual Result for ISACHI Experiment 2. The left column represents the initial ground truth data. The center column represents the misalignment applied to the initial data (case **H**) and is the input to the algorithm. The left most column represents ISACHI's output, and the respective mesh derived from the aligned and spatially consistent contours.

for all 6 cases can be seen in Appendix B. The NCC of the line intensities at the cross intersections of the image planes was used as an alternate measure for alignment validation, to allow for some quantitative evaluation of the alignment process, and can be seen in Table 7.6.

7.4.5 Experiment 4

In a final experiment, the full automatic pipeline was applied to the data. The images were initially segmented via the CNN, and ISACHI applied on the datasets. Figure 7.18 shows an example of the segmentation result from the CNN on the same abnormal case as seen in Experiment 3 (Fig. 7.16). The resulting alignment and segmentation correction can be observed in Fig. 7.19. The contours can be seen matching the edges, whilst being spatially consistent.

Figure 7.20 shows an example of the initial and resulting output meshes from the expert segmented dataset (Experiment 3, Fig. 7.16), as well as from the CNN

	Before			After		
Case	LV epi	LV endo	RV endo	LV epi	LV endo	RV endo
A	1.51	1.52	1.79	0.77	1.05	1.25
B	1.89	1.78	1.97	0.74	0.64	0.77
C	2.37	2.25	3.05	2.42	0.72	1.07
D	1.30	1.29	1.23	0.58	0.46	0.49
E	2.26	2.10	2.25	2.23	0.93	1.00
F	2.49	2.36	2.45	0.51	0.52	0.57
J	2.10	2.05	2.14	0.48	0.49	0.67
H	1.96	1.85	2.29	0.85	2.38	1.03
I	1.99	1.90	2.06	0.53	0.55	0.58
J	2.43	2.29	2.79	2.07	0.58	0.86
Mean	2.03	1.94	2.20	1.12	0.83	0.83

Table 7.5: Isachi results for Experiment 2. The mean left ventricular epicardium, left ventricular endocardium and right ventricular endocardium node-to-mesh distances are shown before and after the algorithm. Each before case represents the synthetic misalignment seen in Table. 7.1 applied to $\mathcal{C}_{\mathcal{GT}}$.

	Normal_1	Normal_2	Normal_3	Abnormal_1	Abnormal_2	Abnormal_3
Mean (Before)	0.705	0.862	0.817	0.571	0.573	0.895
Mean (After)	0.754	0.644	0.819	0.560	0.603	0.847
STD (Before)	0.229	0.075	0.124	0.278	0.243	0.056
STD (After)	0.165	0.253	0.112	0.303	0.199	0.100
Median (Before)	0.849	0.887	0.860	0.631	0.633	0.924
Median (After)	0.804	0.714	0.869	0.721	0.670	0.881

Table 7.6: Experiment 3 NCC of the line intensities between image planes before and after applying ISACHI on 3 normal and 3 abnormal cases for the LV epicardium, endocardium and RV endocardium.

(Experiment 4, Fig. 7.18). The two resulting meshes were registered using ICP, and the node-to-mesh distances were calculated to obtain a measure of difference between the two meshes. For this particular example, the LV epicardial and endocardial mesh, had 1.7 mm and 1.78 mm average discrepancy between each respective mesh. The resulting LV endocardial meshes had a 2.1 mm average discrepancy. The Dice coefficient was computed between the manual segmentation and the CNN segmentation before and after ISACHI. For this particular example, the mean and standard deviation and median values before the algorithm were 0.95 ± 0.02 and 0.95 for the LV epicardium, 0.92 ± 0.08 and 0.95 for the LV endocardium, and 0.89 ± 0.04 and 0.89 for the RV epicardium, respectively. After the alignment

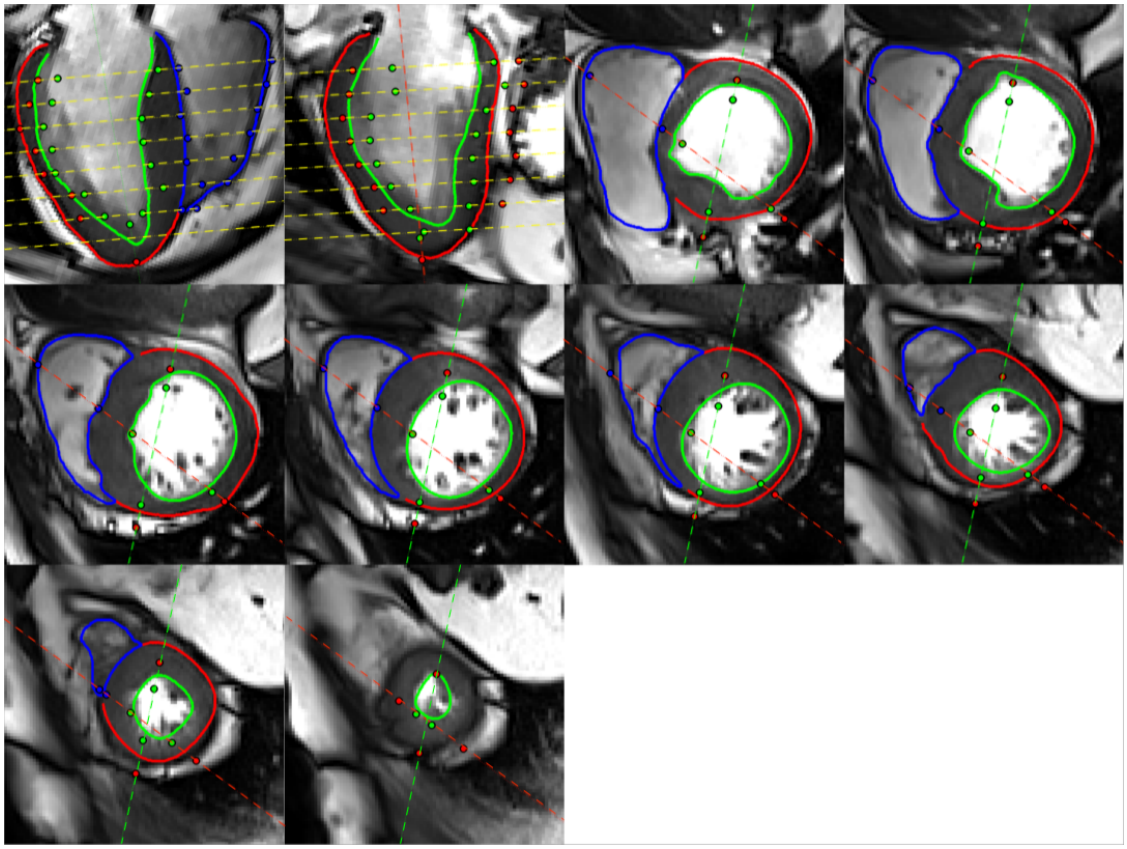

Figure 7.16: Example of an abnormal case for Experiment 3 corresponding to the abnormal case 1. The red, green and blue segmentations represent the LV epicardium, LV endocardium and RV endocardium contours, respectively. The yellow dashed lines represent the SAX plane intersections with the LAX. The red dashed lines represent the 4CH LAX, and the green dashed lines represent the 2CH LAX. The respective plane-to-contour intersections between each of the slices and the contours are represented as the respective colored dots. The contour points in the second image (2CH LAX) can be seen lying far away from the respective intersecting contours.

correction and contour deformations towards spatial consistency, the mean and standard deviation and median values were 0.92 ± 0.02 and 0.92 for the LV epicardium, 0.89 ± 0.07 and 0.92 for the LV endocardium, and 0.86 ± 0.05 and 0.85 for the RV epicardium, respectively.

The Dice coefficients before and after the alignment process were computed and can be seen in Table 7.7.

Similarly to the previous experiments, the NCC of the line intensities was used as a surrogate measure, the results of which can be seen in Table 7.8.

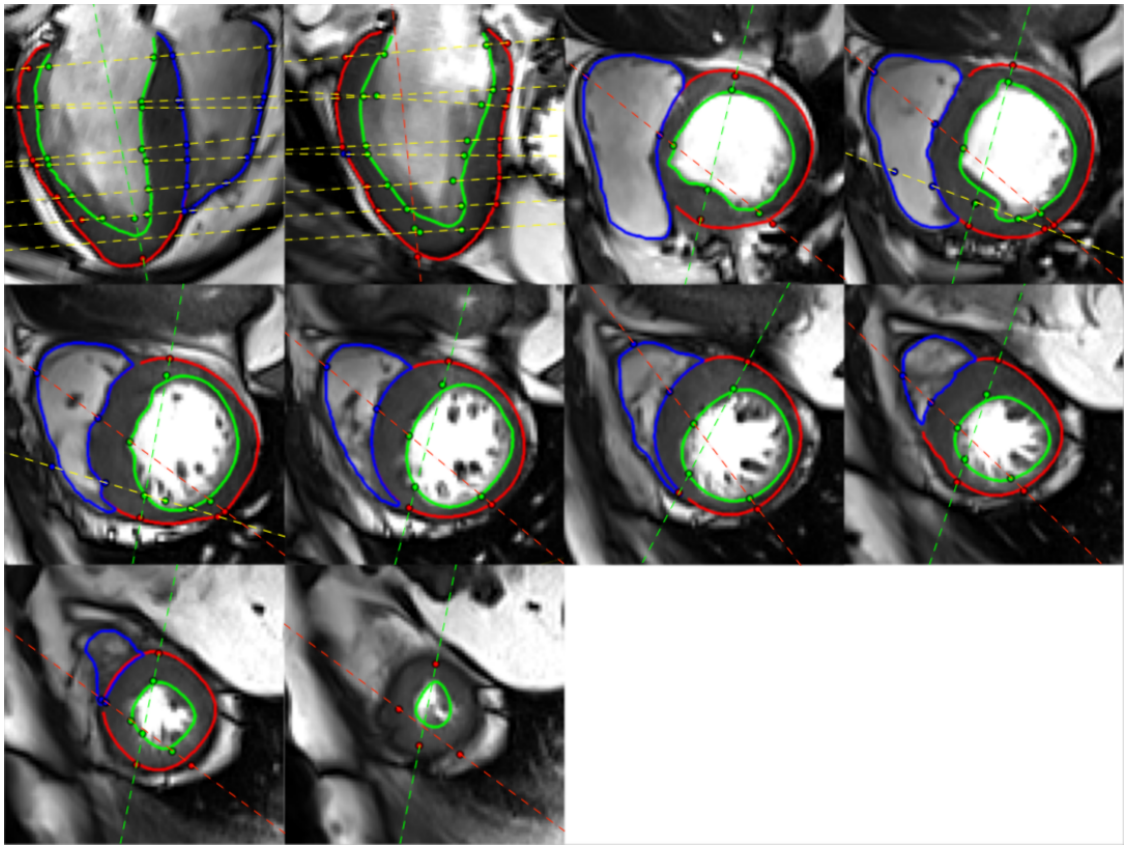

Figure 7.17: Visual result for Experiment 3 after applying ISACHI on the abnormal case show in Fig. 7.16. The red, green and blue segmentations represent the LV epicardium, LV endocardium and RV endocardium contours, respectively. The yellow dashed lines represent the SAX plane intersections with the LAX. The red dashed lines represent the 4CH LAX, and the green dashed lines represent the 2CH LAX. The respective plane-to-contour intersections between each of the slices and the contours are represented as the respective colored dots. The contour points can be seen matching with respective contours.

7.5 Discussion

We presented a novel methodology to perform simultaneous misalignment and segmentation adjustments for cardiac images. Our methodology combines both information from image data as well as the 3D arrangement of the contours. The premise of our method lies in that for slices to be aligned, the image features in a slice must match respective features in cross-sectional slices. Furthermore, structural information must also lie on those feature of interest allowing for the smooth aspect of the mesh. By having an initial segmentation, which can be obtained in any way, a surface mesh can be generated and used to combine both

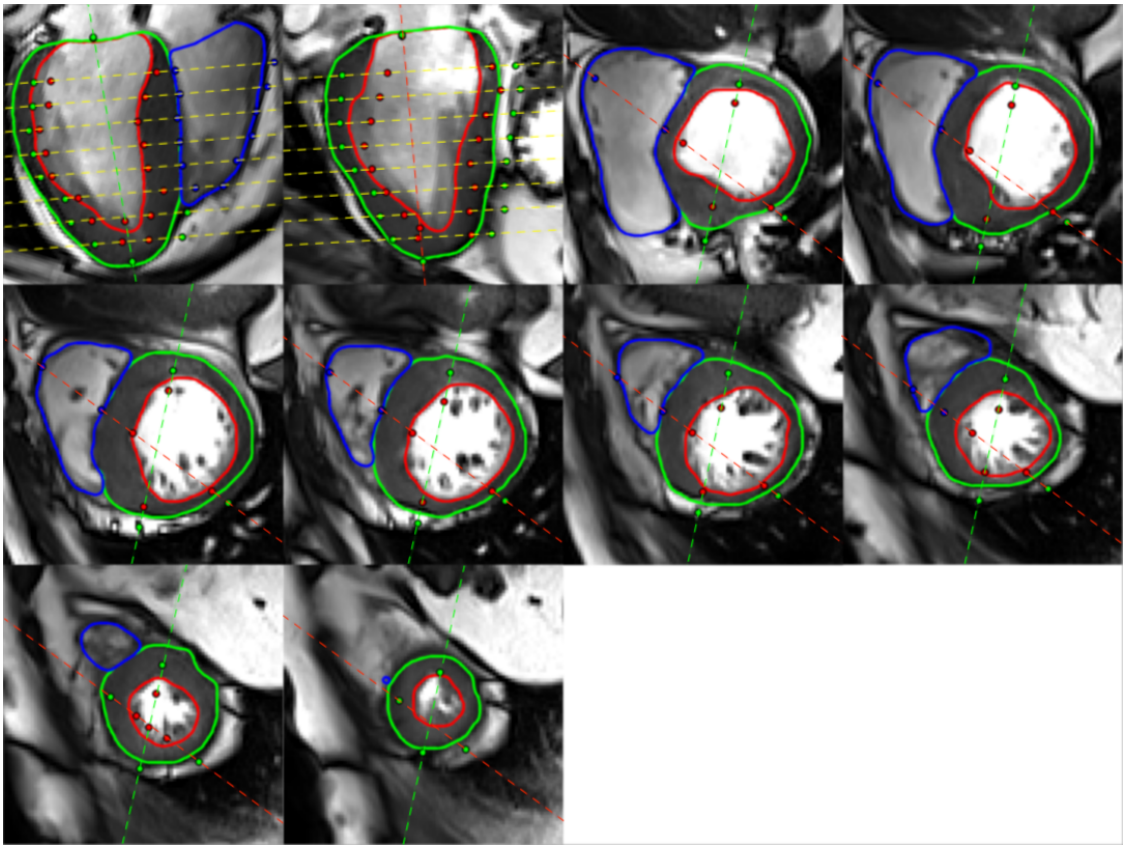

Figure 7.18: Example of the CNN segmentation result on an abnormal case for Experiment 4. The red, green and blue segmentations represent the LV epicardium, LV endocardium and RV endocardium contours, respectively. The yellow dash lines represent the SAX plane intersections with the LAX. The red dash lines represent the 4CH LAX, and the green dash lines represent the 2CH LAX. The respective plane-to-contour intersections between each of the slices and the contours are represented as the respective colored dots. The contour points can be seen lying far away from the respective intersecting contours.

structural and feature information allowing for the optimization to incorporate both global and local data. Via a set of experiments, we show that the method is able to correct for misalignment issues, as well as recover a realistic mesh, close to the ground truth. In order to validate the initial experiments, the node-to-mesh distances between the resulting meshes and the ground truth meshes were taken. This measurement was taken instead of node-to-node measurements, as the meshes do not have point-to-point correspondences. Instead the closest distance between each node of the resulting mesh and the surface plane of the output mesh was taken.

To illustrate a fully automatic integration of the methodology pipeline, a CNN

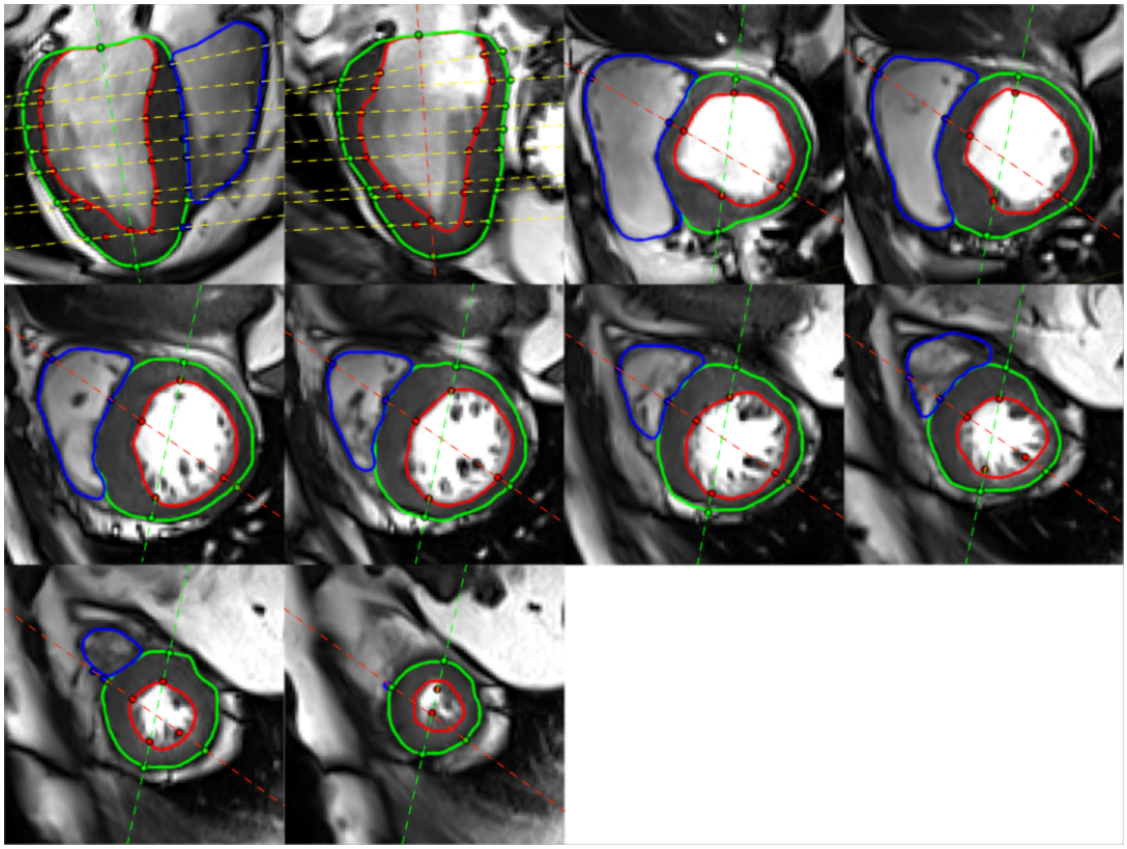

Figure 7.19: Visual result for Experiment 4. ISACHI was applied on the abnormal case show in Fig. 7.18. The segmentations were automatically obtained via the CNN. The red, green and blue segmentations represent the LV epicardium, LV endocardium and RV endocardium contours, respectively. The yellow dash lines represent the SAX plane intersections with the LAX. The red dash lines represent the 4CH LAX, and the green dash lines represent the 2CH LAX. The respective plane-to-contour intersections between each of the slices and the contours are represented as the respective colored dots.

was trained and validated against two different datasets. One dataset coming from an HCM cohort, on which the network was trained, and an alternative dataset, coming from a myocardial infarction cohort. Regarding the HCM cohort, the Dice coefficients were calculated from 81 segmentation comparisons for the normal cases, and 45 for the abnormal cases. It can be seen in Fig. 7.11 that the LV segmentations perform better than the RV, which is to be expected as the training data for the network included more LV labeled images, since the 2CH LAX did not have RV segmentations. Another validation was performed on the myocardial infarction dataset, described in Chapter 5, to assess the networks accuracy on data from another clinical study. The assessment was separated into end-systolic and

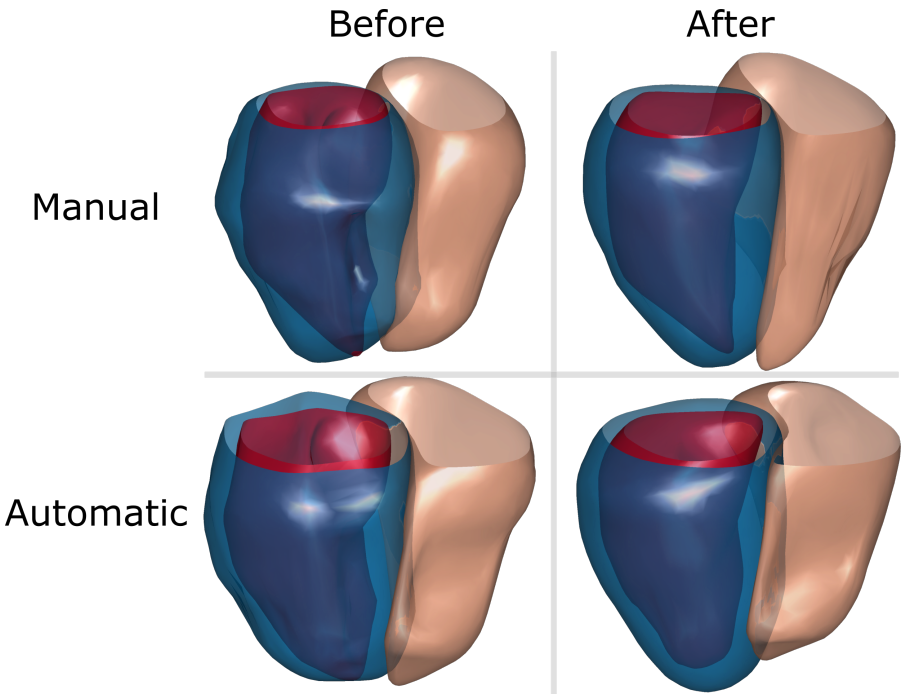

Figure 7.20: Resulting surface meshes before and after ISACHI using manual (Experiment 3) and automated (Experiment 4) contouring.

end-diastolic data which accounted for 159 segmentation comparisons for the LV, and 57 for the RV. The resulting Dice measures are similar to the training data cohort, demonstrating the networks ability to work with data acquired differently.

In a first experiment, the ability of the ISACHI algorithm to realign images based on their edgemaps is assessed. Synthetic edgemaps are generated by using a Gaussian modulation of the distance to the contours, as shown in Fig. 7.6. The contours were generated from a synthetic mesh, resulting in spatially consistent contours. In this initial experiment, the slices were misaligned using the random perturbations found in Table 7.1. The node-to-mesh distances between the resulting mesh and the initial, ground truth, mesh $\mathcal{M}_{\mathcal{GT}}$, after registering them via ICP, was taken, and the results can be seen in Table 7.2. The discrepancy before the alignment ranged between 0.7 mm to 1.8 mm. After applying ISACHI, the range decreased between 0.18 mm to 0.63 mm. The NCC of the line intensities was used as a surrogate measure to assess the alignment before and after using ISACHI, and can be seen in Table 7.3. As the data used in this experiment is synthetic, with distinct matching areas, the surrogate measure shows an extent of the simulated

	$Normal_1$	$Normal_2$	$Normal_3$	$Abnormal_1$	$Abnormal_2$	$Abnormal_3$
Mean (Before - LV Epi)	0.942	0.925	0.935	0.871	0.941	0.917
Mean (After - LV Epi)	0.942	0.948	0.943	0.923	0.931	0.948
STD (Before - LV Epi)	0.030	0.043	0.040	0.057	0.024	0.041
STD (After - LV Epi)	0.024	0.018	0.018	0.023	0.025	0.028
Median (Before - LV Epi)	0.947	0.945	0.942	0.870	0.935	0.922
Median (After - LV Epi)	0.923	0.931	0.948	0.924	0.933	0.959
Mean (Before - LV Endo)	0.938	0.952	0.950	0.909	0.918	0.935
Mean (After - LV Endo)	0.943	0.954	0.935	0.889	0.910	0.956
STD (Before - LV Endo)	0.032	0.020	0.02	0.054	0.076	0.035
STD (After - LV Endo)	0.022	0.027	0.017	0.073	0.043	0.027
Median (Before - LV Endo)	0.951	0.950	0.960	0.919	0.941	0.941
Median (After - LV Endo)	0.889	0.910	0.956	0.915	0.919	0.966
Mean (Before - RV Endo)	0.694	0.718	0.747	0.687	0.657	0.591
Mean (After - RV Endo)	0.869	0.912	0.893	0.855	0.543	0.864
STD (Before - RV Endo)	0.296	0.193	0.164	0.276	0.284	0.342
STD (After - RV Endo)	0.120	0.049	0.024	0.050	0.311	0.071
Median (Before - RV Endo)	0.808	0.767	0.819	0.742	0.768	0.754
Median (After - RV Endo)	0.855	0.543	0.864	0.853	0.622	0.886

Table 7.7: Dice coefficient results before and after applying ISACHI on 3 normal and 3 abnormal cases for the LV epicardium, endocardium and RV endocarrdium.

	$Normal_1$	$Normal_2$	$Normal_3$	$Abnormal_1$	$Abnormal_2$	$Abnormal_3$
Mean (Before)	0.705	0.862	0.817	0.571	0.573	0.895
Mean (After)	0.762	0.779	0.741	0.640	0.666	0.857
STD (Before)	0.229	0.075	0.124	0.278	0.243	0.056
STD (After)	0.157	0.160	0.226	0.284	0.189	0.099
Median (Before)	0.849	0.887	0.860	0.631	0.633	0.924
Median (After)	0.826	0.833	0.826	0.678	0.710	0.882

Table 7.8: Experiment 4 NCC of the line intensities between image planes before and after applying ISACHI on 3 normal and 3 abnormal cases for the LV epicardium, endocardium and RV endocardium.

misalignment. Case **E** has the lowest mean measure bound (0.73), and case **D** the highest mean measure bound (0.88). It can be seen that after applying the algorithm, all cases result in a mean NCC value of 0.99 . The algorithm's ability to realign the slices was assessed in this experiment. The values in Table 7.3 showed that the algorithm was able to realign the slices as the surrogate measure reported near identical values for all the cases.

To assess the algorithm's ability to work with images, Experiment 2 was performed, using an atlas template. By using this synthetic dataset, a controlled edgemap was obtained. Non-rigid deformations were applied to the contours as well

as applying random transformation to the SAX slices. It can be seen in Table 7.5 that, after ISACHI, the distances between the meshes have been reduced to sub pixel accuracy. However in some cases, there still exists a high discrepancy, such as the LV epicardium in case **C**, **E**, and **J**, as well as LV endocardium in case **H**. This is because the transformations applied to the ground truth contours were quite severe for each of these cases, as can be observed in Table 7.1. For example, case **E** had a *cumulative* θ_x displacement of 11.6 mm. The algorithm's main aim is to correct for slice misalignment as well as any contour misalignment that might occur. To that extent, the end result might be achieved having a different shape from the ground truth, but having a plausible shape fitting those criteria. As the slices in the aforementioned cases are severely misaligned, the overall resulting shape might differ from the ground truth, although they are well aligned. It can be seen that in the three cases (**C**, **E** and **H**), only one of their respective three meshes (LV epicardium, LV endocardium, and RV endocardium) is badly misaligned, while the others have low error distances. Furthermore when looking at the NCC of the line intensities they all have increased, signifying a better alignment. The mean inter-slice NCC for the misaligned line intensities for the cases **C**, **E**, **J**, **H** increased from 0.954, 0.955, 0.967, 0.971 to 0.9840, 0.980, 0.992, 0.974, respectively. A potential reason for the resulting misalignment the smoothness of the atlas template. As the atlas was acquired as an average from numerous datasets, it is very smooth, and lacks distinct boundaries. As such, there can exist multiple spatial locations for the segments to lie on, as the edgemap can span several pixels.

It can be seen from Table 7.2 and Table 7.5 that the node-to-mesh discrepancy between the misaligned and ground truth meshes differs for both experiments. This is due to the input contour data being different for both cases, which is shown in Fig. 7.7. The contours in Experiment 2 are generated by slicing the mesh at different spatial positions than Experiment 1. The LAX contours in Experiment 1 both contain RV information, as opposed to Experiment 2. Furthermore, the contours in Experiment 2 were deformed non-rigidly in plane, as to simulate contouring error. These differences, increased the distance errors slightly for the second experiment.

In a third experiment, the algorithm was run on real data from a normal (3 cases) and abnormal (3 cases) cohort. As no ground truth data on the alignment was available, a visual assessment for one of the abnormal cases was given before (Fig. 7.16) and after the use of ISACHI (Fig. 7.17). It can be seen in Fig. 7.16 that the second LAX (2CH) suffers from high misalignment. The resulting dataset shown in Fig. 7.17 has corrected for this misalignment as the corresponding segmentations from the other planes lie on respective features. Furthermore, the contours have been deformed to give spatial consistency to the dataset. One particular area where the algorithm fails, in this particular case, is at the apical slice. The apical contours in the second LAX view can be seen below the apex given by the LAX slice. However, if the apex SAX image is observed, the epicardium contours (given in red), can be seen to lie on the LV epicardium. A comparison of the spatial consistency cannot be made as no epicardium contour exists for that slice. It can be observed that this results in a better position than the initial pose, as the apical SAX epicardial contours are severely misaligned in Fig. 7.16. Meshes generated before and after applying the algorithm, for all 6 cases, can be seen in Appendix B. Panels *a* and *b* show the meshes before and after, respectively. The NCC of the line intensity was also used as an alternate measure of validation, the results of which can be seen in Table 7.6. Most of the cases can be seen having similar values, at the exception of the second normal case. However, the resulting mesh (Fig. B.4, Appendix B) appears visually correct. NCC of the line intensity profiles works appropriately with synthetic images because the profiles are extracted from the 3D mesh. However, this is not the case with real data, where many factors can influence the result, such as the different contrast in the blood pool, as well as the different organs outside of the heart.

A final experiment was performed using the entire automated pipeline. Given the images from the 6 cases, a CNN was used to automatically segment the slices, and ISACHI was then applied to that data. An initial accuracy measurement for the CNN's performance on the cases was computed using the Dice Coefficient with respect to the manual annotations from an expert. An example of the segmentation

result can be seen in Fig. 7.18 applied to the same example case (abnormal) used in Experiment 3. The segmentations from Fig. 7.18 and Fig. 7.16 can be directly compared, and appear very similar. One of the positive aspects of the network, which can be observed in Fig. 7.18 is that all of the slices were segmented and the segmentation of the papillary muscles was avoided. The same alignment issues with the 2CH LAX can be seen in Fig. 7.18 as with Fig. 7.16. The result of applying ISACHI to the CNN segmented example data can be seen in Fig. 7.19. It can be observed on the first LAX (4CH) that the resulting alignment caused the slices to merge towards a similar position than in Fig. 7.17. Furthermore the contours are spatially coincidental and appear to be located at the correct cardiac borders. The fully automatic pipeline results for the 6 cases can be seen in Appendix B. A comparison of the resulting meshes for the abnormal example can be seen in Fig. 7.20. The initial meshes using the expert and automatic segmentations can be seen in the left column, and their respective results on the right. Both resulting meshes appear to have similar shape, with the exception of the apical area. As the CNN was able to segment the apical slices, more information was could be used to generate the mesh.

It should be re-iterated that, while some particular results in Table. 7.7, and Table. 7.8 do not show significant improvement in the method, the dice coefficient and line intensity were used as surrogate measures of validation, as no ground truth was available for such clinical cases, and thus do not adequately reflect an accurate validation measure of the method.

The aim of the network was to fully automize the pipeline. Although acceptable for a preliminary study, the Dice values show that the networks performance can still be improved, notably for the right ventricle. However, any segmentation tool can be used as the correction technique just requires an initial adequate segmentation.

One of the benefits of ISACHI, is the ability to rely on the input data. By using the edgeness profiles, derived from the data, the slices can be realigned by taking the information evaluated at the patches. Similar to the method in Chapter 5, ISACHI takes into account both global and local information. The method utilizes the information obtained at all the patch locations. Compared to using the information

at the line intersections of slices, this approach is more focalized on the important features, surrounding the mesh, to match. As the segments are positioned sparsely on the image plane, they capture a greater global description of the image intensities, as well as providing local information. The method combines both geometrical and image information, although it relies on a very mild smoothing assumption such that the mesh can be locally represented by a patch. If the mesh was highly convoluted, the patched would not be an appropriate representation. The geometrical and image information play a symbiotic role which allow to align the slices and deform the contours by relying solely on the image information allowing the generation of a mesh that follows data information. The method defines the ability to obtain spatial consistency such that all the concurrent information obtained at every pixel from the mesh, contours, segment intensities, and edges should hold the same information. Furthermore, with the advent of machine learning techniques, allowing for reliable segmentations, ISACHI has the ability to be fully automatic.

Our method has certain limitations. Firstly, it relies on the ability to derive reliable edgemaps form the images, where areas of high edgeness describe an edge or boundary. Secondly, the method relies on first order patches, *i.e.* linear patches tangent to the current estimation of the surface mesh. As these patches are tangent to the mesh, they do not accurately convey the boundary informations. Having higher order patches would be more suitable, as they would allow the method to deform the contours following the curvature of the edgness profile.

8

Conclusion and Future Work

The overall aim of this Thesis was to develop new tools and methodologies to correct for slice misalignment in cardiac datasets for the reconstruction of patient-specific geometry. The resulting models can be used in the area of precision medicine, or for population studies. The area of medical imaging suffers from a wide range of challenges which can cause issues in the images, and subsequent analysis. Motion related artifacts are present in several medical research areas and applications, such as in the brain, liver, fetus, or in the heart. Different types of motion artifacts exist such as inherent organ motion, breathing related motion, as well as patient motion inside the scanners.

In cardiac imaging, any 3D analysis derived from the medical images necessitates a pre-processing step correcting for any motion misalignment between slices, due to their independent acquisitions, as well as being subject to different breath holds. Current techniques for solving this issue commonly involve using segmentations to align data. However, these techniques do not take into account any image information and can stray away from the underlying patient geometry. Furthermore, as the segmentations are generally performed in 2D, they are not 3D spatially consistent. These discrepancies can be due to slight differences in the image intensities at neighboring pixels, causing the segmentations to over or under-estimate the region of interest and lie on different pixels in the respective intersecting slices.

Large discrepancies between the intersecting contours can also be found due to the incorrect segmentation of papillary muscles in the LAX, which can be hard to distinguish from myocardium. Furthermore, by using contours to correct for motion misalignment between the slices necessitates expert segmentations which can be costly and timely to obtain.

The first novel aspect of this thesis consisted of developing a methodology to correct for the misalignments between slices using only the information from the images, allowing the method to be used directly in the clinics. The method aligns the respective slices together by using the line intensities at their intersections. By optimizing each slice's spatial position in 3D with regards to all its intersecting slices, the method tries to obtain the best alignment for that slice with respect to all the others, in an iterative manner. Instead of using the line intensities, the line profiles of the local phase images were used. This allowed to match features of interest between slices instead of matching their line intensities, which can be different due to field inhomogeneities, or LGE contrast. The method was applied to LGE datasets and validated against expert segmentations. A validation methodology was also introduced to assess the possible quantitative lower bound values. The results showed that the methodology had similar results than if it was used on labeled data generated from the expert segmentations. The minimum quantitative values for the expert segmented contours were calculated to show the maximum possible quantitative results for the datasets, as the contours suffer from spatial inconsistency.

Once the misalignment between slices is corrected, a 3D geometrical mesh can be obtained using the aligned contours. However, because of the alignment process, the contours lose their parallelism. Furthermore, they are heterogeneous as the individual points forming the contours are not equidistantly spaced. Typical methods to build surface meshes necessitate dense oriented point clouds (having normals) or require for the data to be spatially consistent, or parallel. As such, to allow for the generation of a 3D geometrical surface mesh from cardiac segmentations, a method was developed to deal with sparse, non-parallel, non-coincidental, heterogenous contours. In order to avoid a trade-off between data fitting or the smoothness of

the mesh, the method maximizes the data fitting aspect by composing smooth deformations, in an iterative manner. By using TPS as an interpolation function, smooth deformations can be obtained, passing through the data points. A validation methodology is presented for the method, assessing its ability to recover a template mesh generated from cardiac data, given sparse sets of contours. Furthermore, the method is applicable to other types of data, such as other organs, although a multi-parameter study was performed to tune the method specifically for cardiac data, the parameters of which are presented.

The aim of obtaining personalized 3D meshes is typically achieved in an iterative manner. The images need to be initially segmented, then any misalignment corrected for, in order for the surface meshes to be finally generated. Each of these steps rely on several assumptions, and do not make use of information from each step. A final method was presented which integrates all these steps, in a fully automatic manner. Using the images, segmentations are obtained automatically using a UNet CNN. From these segmentations, an initial mesh is built, and segments, tangent to the mesh and lying on the images are generated. These segments guide the alignment process by using information from respective intersecting slices, both locally and globally. Using edgemaps generated from the images, the patches are guided towards areas of high edgeness, having a trade-off between its respective slices. In such a way, the segments find areas of high edgeness that lie at the intersection of both cross-sectional images. A rigid transformation is applied to the images following the optimization of the segments, which corrects for the alignment. In order to correct for the spatial inconsistency issues, the position of the segments are then optimized individually, and the remaining transformations are extrapolated to deform the contours non-rigidly. The entire process is performed iteratively. Similarly to the previous works, the method is validated at different levels. The preliminary results show that the novel idea of combining both geometric and image information to guide the alignment works well.

The work presented in this Thesis does have some limitations. The generated surface meshes are assumed to represent accurate geometries as they rely on a

minimum amounts of regularization while relying on the image data as much as possible. However, no ground truth comparison has been given on real data, as such ground truth was not available, and instead non real data was used to obtain some accuracy measurements. Regarding the last research work, the impact of the choice of edgemap was not evaluated. Better edgemaps calculations are available, and should be compared against. Although the use of the CNN in the Thesis allowed to demonstrate an initial application, better CNN models should be used.

Future Work

The same methodology as the last research work (ISACHI) could be fully integrated using neural networks. Given 3D volumes, such as SSFP acquisitions, a 3D CNN could be trained to segment datasets having spatial consistency. Once trained, the network could be applied to the misaligned data, and the network's feature maps could be used as the edgemaps. Another possibility would be to incorporate the segmentation process to use an alternative CNN architecture such as the one in [152]. Furthermore, the calculation of the segments could be solely obtained using the contours, and the mesh built at the final stage. Higher order patches could be used so that the information guiding the registration relies on segments fitted to the contours.

Finally, the entire pipeline presented in the Thesis could be applied to a large cohort such as the UK Biobank or Cardiac Atlas project to perform global population studies on shape analysis. Furthermore, by using Cine acquisitions, the pipeline can be applied on all cardiac frames, allowing for motion studies.

The tools found in this Thesis have been developed to be easily implementable in the clinics as well as for any research purposes. They have been made publicly available, and can be found on the following GitHub pages: `http://github.com/BenVillard/surfco`, and `BenVillard/repositories`.

Appendices

A

Cardiac Surface Meshes

A.1 Simulation of CMR Acquisitions

Figure A.1 shows the resulting meshes for the 40 simulated cases described in Section 6.3.1.



Figure A.1: Resulting output mesh using our methodology on simulated CMR acquired contours (see Section 6.3.1). The figure shows the various shapes of the 40 generated meshes.

A.2 Application of SurFCo on alternate objects

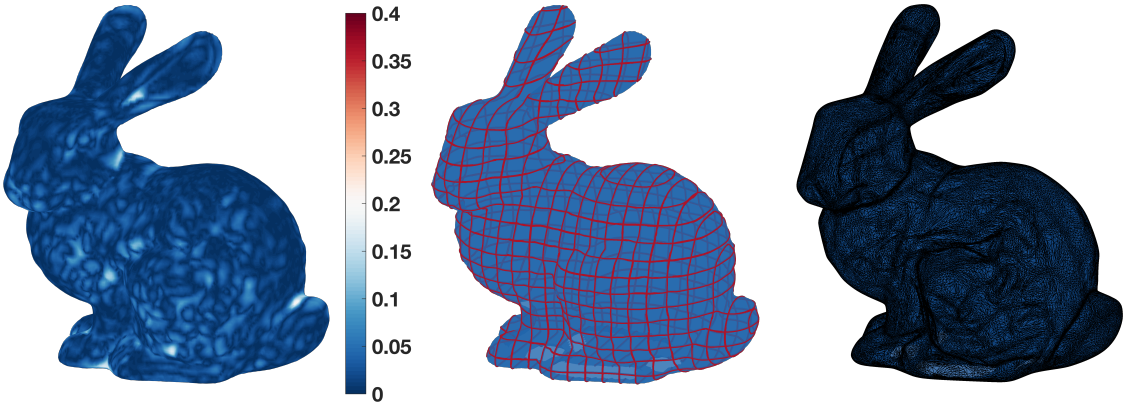

Figure A.2: Application of SurFCo on the MIT Bunny. The original mesh was sliced 20 times equidistantly from the top most to the bottom most plane, as well from the tail to the nose. SurFCo was then applied to the contours, and the mesh-to-mesh distances was taken (in mm). The **left** image represents the original mesh. The colormap represents the mesh-to-mesh distance error between the SurFCo output mesh and the original mesh. The **center** image represents the original mesh sliced at various positions, synthesizing contours that will be used as the input to SurFCo. The **right** image shows SurFCo's resulting mesh.

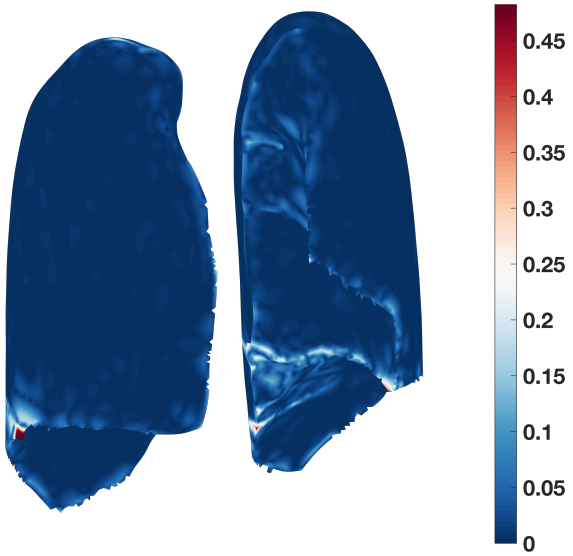

Figure A.3: Application of SurFCo on a pair of lungs. Similarly to Fig. A.3, the original lung meshes were sliced 10 times equidistantly from the top most to the bottom most plane, as well as sliced 5 times from the right to the left. SurFCo was then applied to the contours, and the mesh-to-mesh distances was taken (in mm).

B
ISACHI Visual Results

B.1 Experiment 4

This Appendix shows the visual results for Experiment 4 in Chapter 7, Section 7.4.5. The following images represent the result of automatically segmenting the images via a U-Net CNN, and applying the ISACHI algorithm. The green, red and blue segmentations represent the LV epicardium, LV endocardium and RV endocardium contours, respectively. The yellow dash lines represent the SAX plane intersections with the LAX. The red dash lines represent the 4CH LAX, and the green dash lines represent the 2CH LAX. The respective plane-to-contour intersections between each of the slices and the contours are represented as the respective colored dots.

The first three images represent data from a normal cohort, and the last three images represent data from a pathological cohort.

Afterwards, the images represent the respective meshes built from the segmentation results. The meshes in panels *a*) represent the meshes generated from the non corrected expert segmented contours. The meshes in panels *b*) represent the meshes generated after applying ISACHI using the expert segmented contours. The meshes in panel *c*) represent the meshes generated from the non corrected, misaligned, contours obtained from the CNN. The meshes in panel *d*) represent the meshes generated from the CNN segmentations after applying ISACHI. The resulting

output meshes were registered via ICP, and the node-to-mesh distance between them computed to allow for some quantitative evaluation of the differences between the fully automatic method and after applying ISACHI using expert segmented contours. For the 3 normal cases, the LV epicardial meshes differed by 1.25 mm, 2.24 mm, and 3.59 mm. The LV endocardial meshes differed by 1.26 mm, 1.56 mm, and 3.91 mm and the RV endocardial meshes had a discrepancy of 3.59 mm, 3.28 mm, and 4.24 mm. Regarding the data from the abnormal cohort, the LV epicardial meshes differed by 1.7 mm, 3.66 mm, and 5.62 mm, the LV endocardial meshes by 1.77 mm, 3.86 mm, and 4.43 mm, and the RV endocardial meshes by 2.01 mm, 7.35 mm, and 5.64 mm.

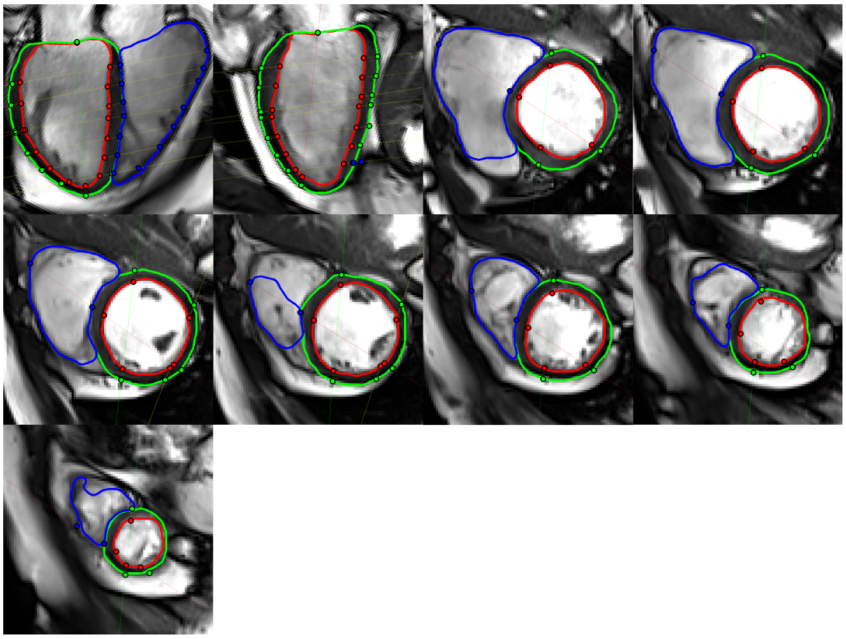

Figure B.1: ISACHI result for Experiment 4, using the using the fully automated approach, on the first normal case. The red, green and blue contours represent the LV epicardium, endocardium and RV endocardium, respectively.

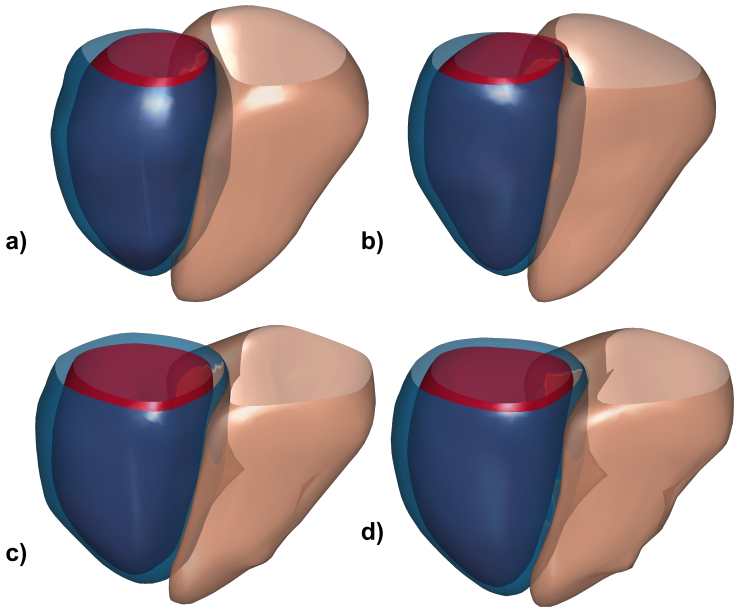

Figure B.2: Meshes generated before and after applying ISACHI on the expert segmented and CNN contours. The meshes in this figure represent the first normal case. a) Mesh generated from the manual contours without any alignment or contour adjustment. b) Mesh generated from the manual contours after applying ISACHI (Experiment 3). c) Mesh generated from the CNN contours, without any alignment or contour adjustment. d) Mesh generated from the CNN contours after applying ISACHI (Experiment 4).

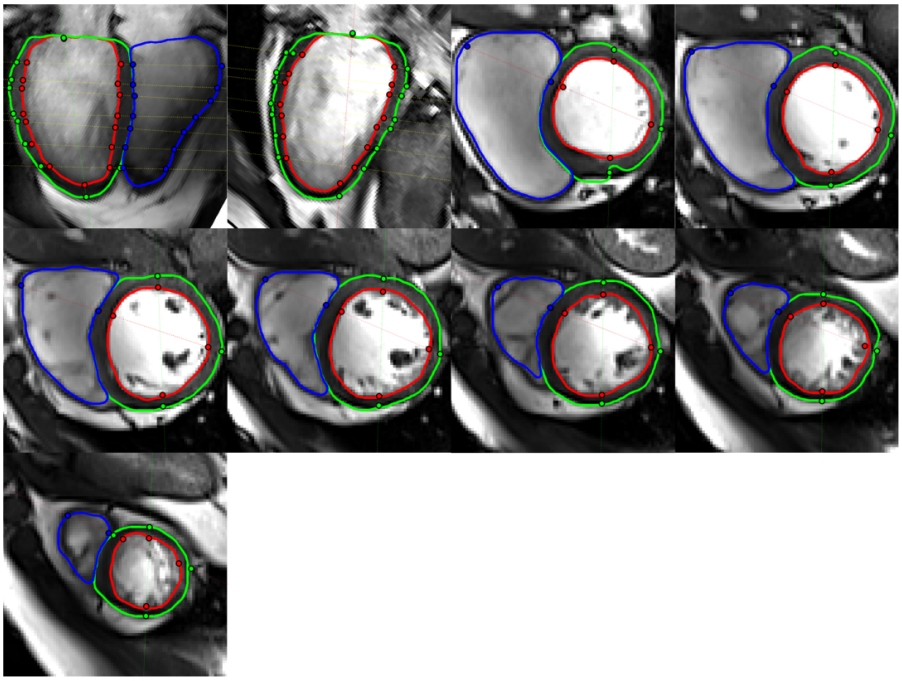

Figure B.3: ISACHI result for Experiment 4, using the using the fully automated approach, on the second normal case. The red, green and blue contours represent the LV epicardium, endocardium and RV endocardium, respectively.

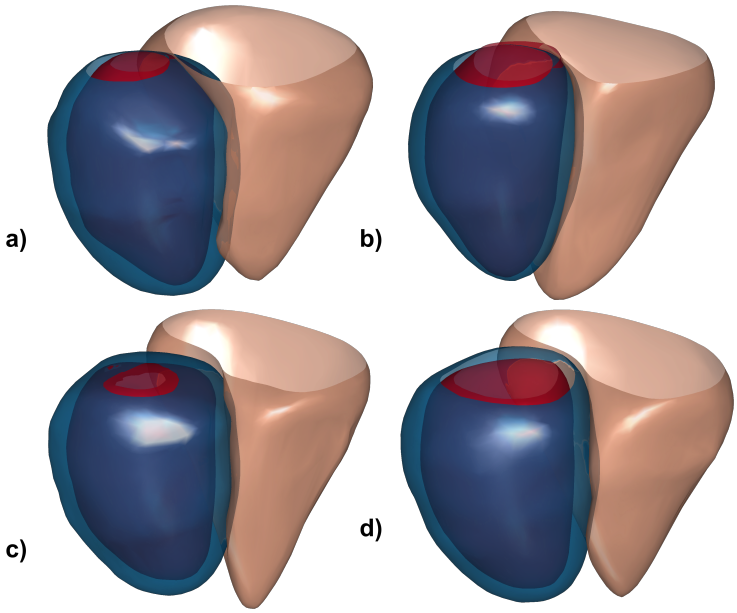

Figure B.4: Meshes generated before and after applying ISACHI on the expert segmented and CNN contours. The meshes in this figure represent the second normal case. a) Mesh generated from the manual contours without any alignment or contour adjustment. b) Mesh generated from the manual contours after applying ISACHI (Experiment 3). c) Mesh generated from the CNN contours, without any alignment or contour adjustment. d) Mesh generated from the CNN contours after applying ISACHI (Experiment 4).

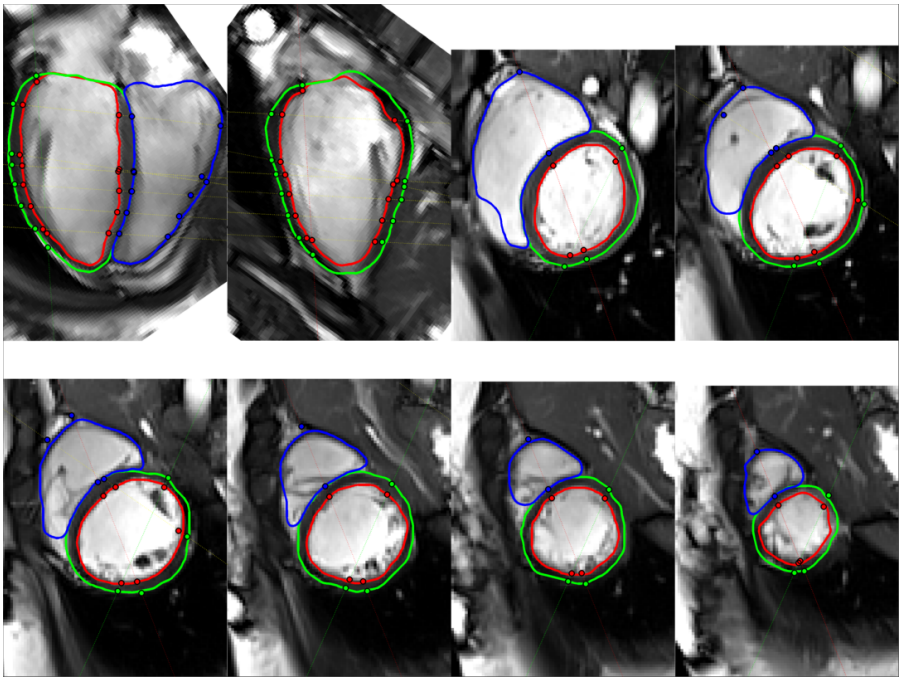

Figure B.5: ISACHI result for Experiment 4, using the using the fully automated approach, on the third normal case. The red, green and blue contours represent the LV epicardium, endocardium and RV endocardium, respectively.

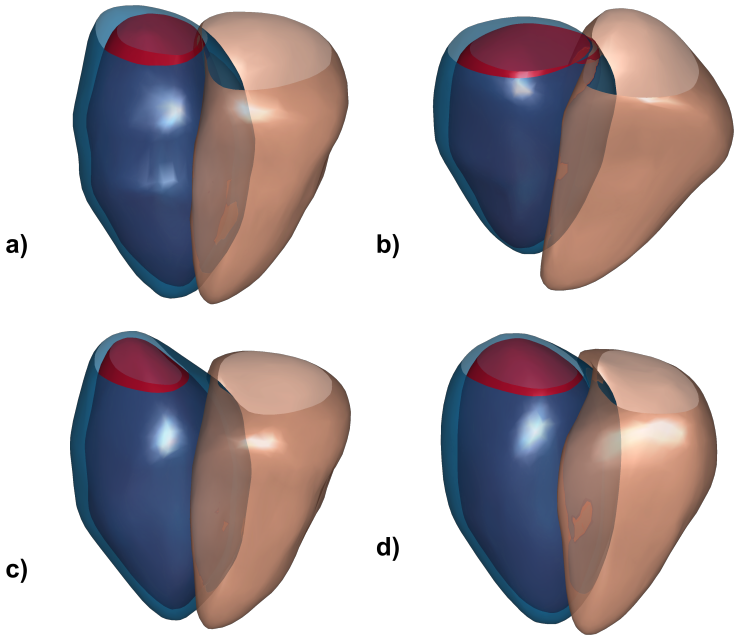

Figure B.6: Meshes generated before and after applying ISACHI on the expert segmented and CNN contours. The meshes in this figure represent the third normal case. a) Mesh generated from the manual contours without any alignment or contour adjustment. b) Mesh generated from the manual contours after applying ISACHI (Experiment 3). c) Mesh generated from the CNN contours, without any alignment or contour adjustment. d) Mesh generated from the CNN contours after applying ISACHI (Experiment 4).

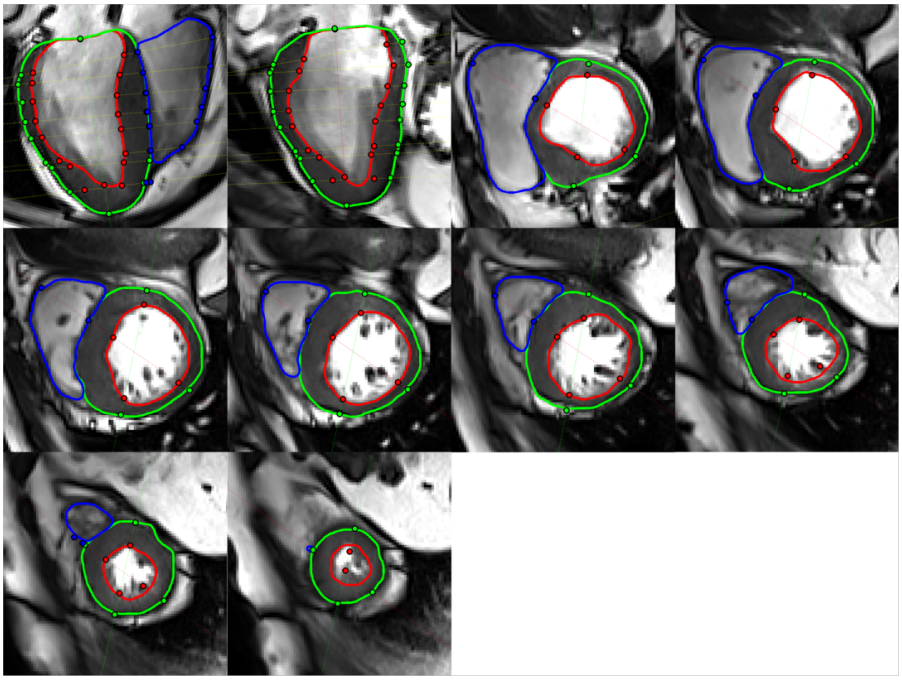

Figure B.7: ISACHI result for Experiment 4, using the using the fully automated approach, on the first abnormal case. The red, green and blue contours represent the LV epicardium, endocardium and RV endocardium, respectively.

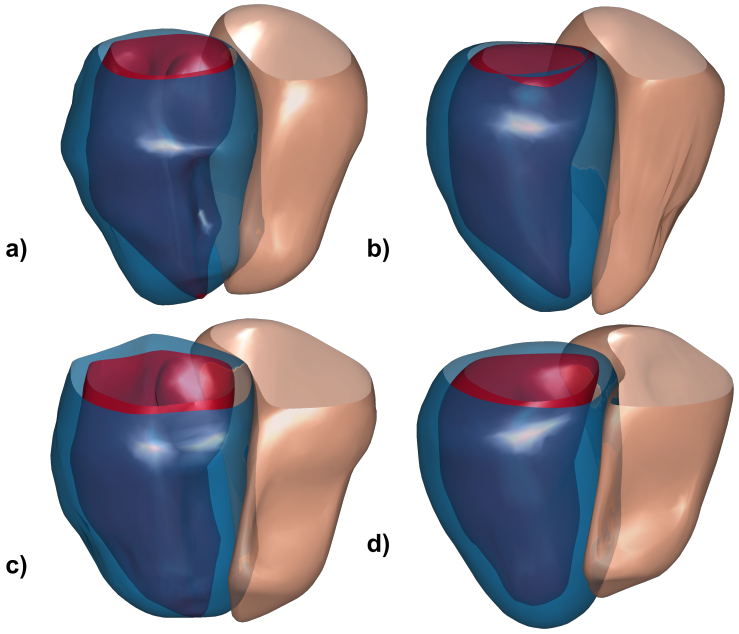

Figure B.8: Meshes generated before and after applying ISACHI on the expert segmented and CNN contours. The meshes in this figure represent the first abnormal case. a) Mesh generated from the manual contours without any alignment or contour adjustment. b) Mesh generated from the manual contours after applying ISACHI (Experiment 3). c) Mesh generated from the CNN contours, without any alignment or contour adjustment. d) Mesh generated from the CNN contours after applying ISACHI (Experiment 4).

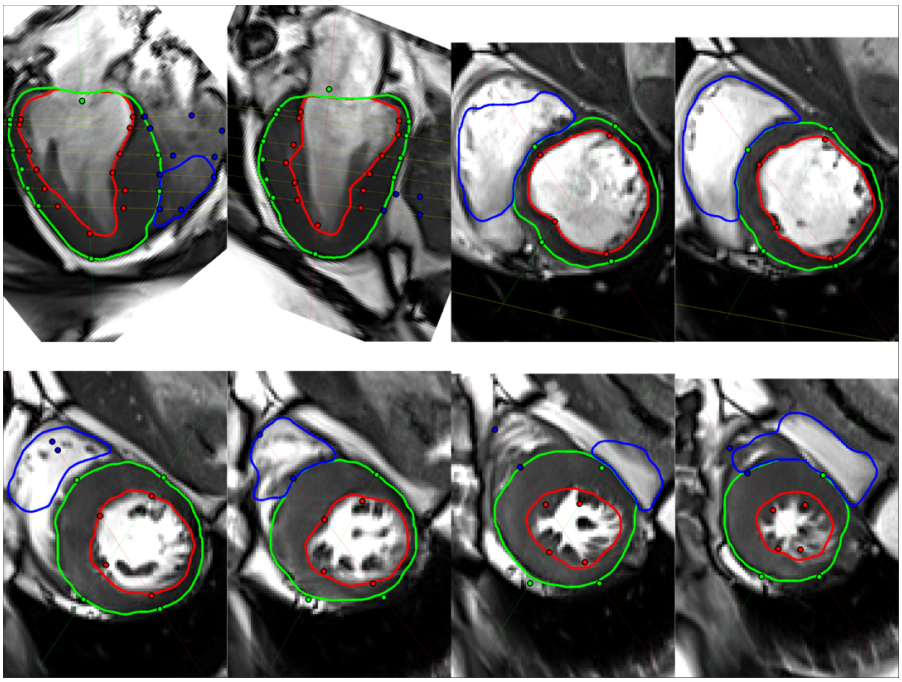

Figure B.9: ISACHI result for Experiment 4, using the using the fully automated approach, on the second abnormal case. The red, green and blue contours represent the LV epicardium, endocardium and RV endocardium, respectively.

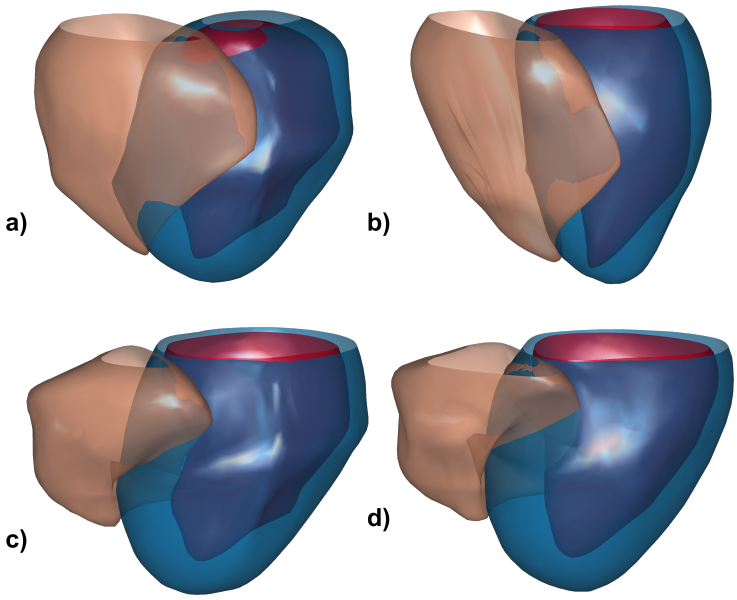

Figure B.10: Meshes generated before and after applying ISACHI on the expert segmented and CNN contours. The meshes in this figure represent the second abnormal case. a) Mesh generated from the manual contours without any alignment or contour adjustment. b) Mesh generated from the manual contours after applying ISACHI (Experiment 3). c) Mesh generated from the CNN contours, without any alignment or contour adjustment. d) Mesh generated from the CNN contours after applying ISACHI (Experiment 4).

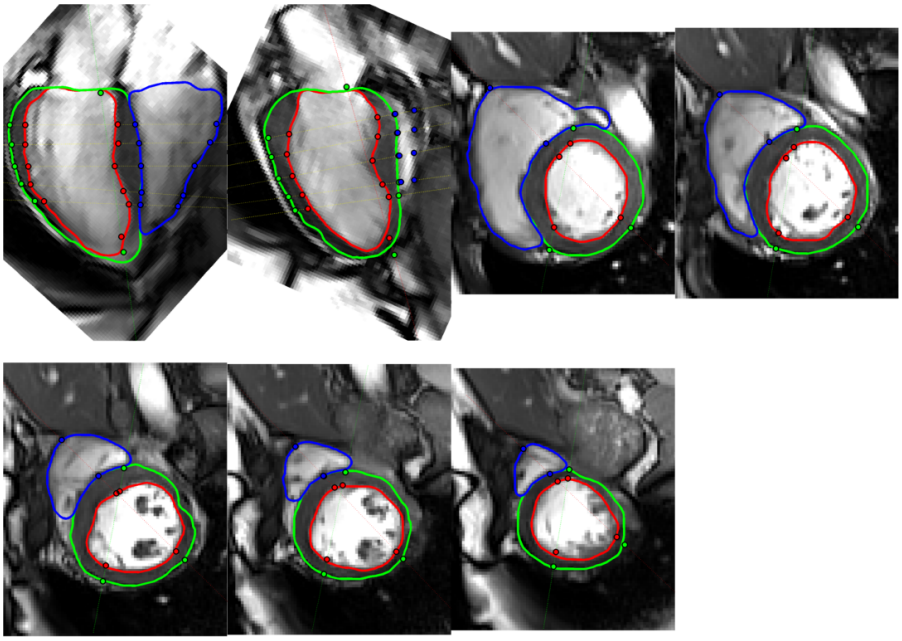

Figure B.11: ISACHI result for Experiment 4, using the using the fully automated approach, on the third abnormal case. The red, green and blue contours represent the LV epicardium, endocardium and RV endocardium, respectively.

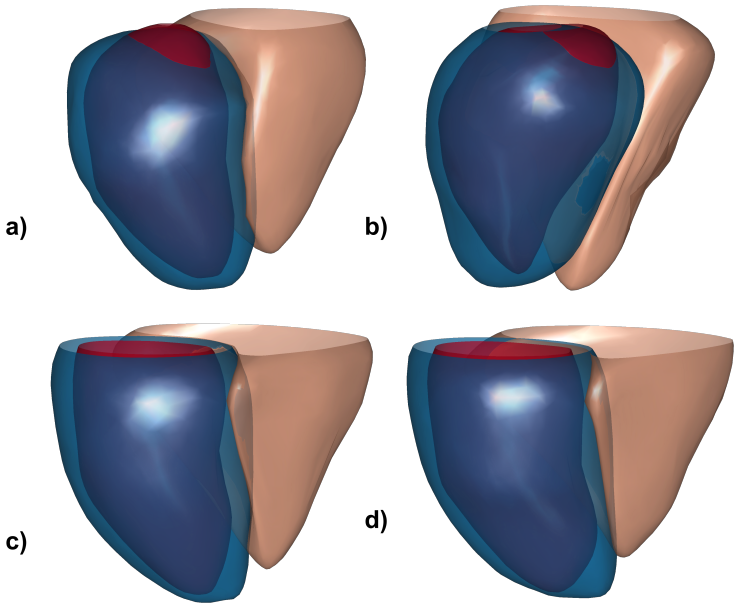

Figure B.12: Meshes generated before and after applying ISACHI on the expert segmented and CNN contours. The meshes in this figure represent the third abnormal case. a) Mesh generated from the manual contours without any alignment or contour adjustment. b) Mesh generated from the manual contours after applying ISACHI (Experiment 3). c) Mesh generated from the CNN contours, without any alignment or contour adjustment. d) Mesh generated from the CNN contours after applying ISACHI (Experiment 4).

Bibliography

[1] C. Y. Liaw and M. Guvendiren. Current and emerging applications of 3D printing in medicine. *Biofabrication*, 9(2):024102, 2017. (Cited on page 1.)

[2] P. Shah and B. S. Chong. 3D imaging, 3D printing and 3D virtual planning in endodontics. *Clinical Oral Investigations*, 22(2):641–654, 2018. (Cited on page 1.)

[3] G. Sakas. Trends in medical imaging: from 2D to 3D. *Computers & Graphics*, 26(4):577–587, 2002. (Cited on page 1.)

[4] C. P. Botha, B. Preim, A. E. Kaufman, S. Takahashi, and A. Ynnerman. *From individual to population: challenges in medical visualization*, pages 265–282. Springer London, London, 2014. (Cited on page 1.)

[5] EchoPixel. EchoPixel: A revolution in 3D medical imaging. (Cited on page 1.)

[6] F. S. Collins and H. Varmus. A new initiative on precision medicine. *New England Journal of Medicine*, 372(9):793–795, 2015. (Cited on page 1.)

[7] S. Anwar, T. Rockefeller, D. A. Raptis, P. K. Woodard, and P. Eghtesady. 3D printing provides a precise approach in the treatment of tetralogy of fallot, pulmonary atresia with major aortopulmonary collateral arteries. *Current Treatment Options in Cardiovascular Medicine*, 20(1), 2018. (Cited on page 1.)

[8] NIH Precision Medicine Initiative. http://allofus.nih.gov. (Cited on page 1.)

[9] G. A. Roth, M. H. Forouzanfar, A. E. Moran, R. Barber, G. Nguyen, V. L. Feigin, M. Naghavi, G. A. Mensah, and C. J. L. Murray. Demographic and

epidemiologic drivers of global cardiovascular mortality. *New England Journal of Medicine*, 372(14):1333–1341, 2015. (Cited on page 2.)

[10] European Cardiovascular Disease Statistics 2012. http://www.escardio.org. (Cited on pages 2 and 16.)

[11] Uk biobank imaging study. `https://imaging.ukbiobank.ac.uk`. Accessed: 25-08-2018. (Cited on page 2.)

[12] C.G. Fonseca, M. Backhaus, D. A. Bluemke, R. D. Britten, J. D. Chung, B. R. Cowan, I. D. Dinov, J. P. Finn, P. J. Hunter, A. H. Kadish, D. C. Lee, J. A. C. Lima, P. Medrano-Gracia, K. Shivkumar, A. Suinesiaputra, W. Tao, and A. A. Young. The cardiac atlas project – an imaging database for computational modeling and statistical atlases of the heart. *Bioinformatics*, 27(16):2288–2295, 2011. (Cited on page 2.)

[13] J. Bogaert, S. Dymarkowski, A. M. Taylor, and V. Muthurangu. *Clinical cardiac MRI*. Springer, 2012. (Cited on pages 9, 12, 15, 16, 26, 29, 31, and 61.)

[14] OpenStax CNX. Anatomy and physiology, 2017. (Cited on pages 9, 10, and 14.)

[15] J. R. Levick. *An introduction to cardiovascular physiology*. Hodder A., 4th edition, 2003. (Cited on page 9.)

[16] N. Dagres and G. Hindricks. Risk stratification after myocardial infarction: is left ventricular ejection fraction enough to prevent sudden cardiac death? *European Heart Journal*, 34(26):1964–1971, 2013. (Cited on page 12.)

[17] A. F. Parisi, P. F. Moynihan, C. L. Feldman, and E. D. Folland. Approaches to determination of left ventricular volume and ejection fraction by real-time two-dimensional echocardiography. *Clinical Cardiology*, 2(4):257–263, 1979. (Cited on page 14.)

[18] C. M. Haggerty, S. P. Kramer, O. Skrinjar, C. M. Binkley, D. K. Powell, A. C. Mattingly, F. H. Epstein, and B. K. Fornwalt. Quantification of left ventricular volumes, mass, and ejection fraction using cine displacement encoding with stimulated echoes (DENSE) MRI. *Journal of Magnetic Resonance Imaging*, 40(2):398–406, 2013. (Cited on page 15.)

[19] S. Natori, S. Lai, J. P. Finn, A. S. Gomes, W. G. Hundley, M. Jerosch-Herold, G. Pearson, S. Sinha, A. Arai, J. A. C. Lima, and D. A. Bluemke. Cardiovascular function in Multi-Ethnic Study of Atherosclerosis: normal values by age, sex, and ethnicity. *American Journal of Roentgenology*, 186, 2006. (Cited on pages 15 and 16.)

[20] K. McLeish, D. L. G. Hill, D. Atkinson, J. M. Blackall, and R. Razavi. A study of the motion and deformation of the heart due to respiration. *IEEE Transactions on Medical Imaging*, 21(9):1142–1150, 2002. (Cited on pages 15, 16, 78, 106, and 133.)

[21] G. Shechter, J. R. Resar, and E. R. McVeigh. Displacement and velocity of the coronary arteries: cardiac and respiratory motion. *IEEE Transactions on Medical Imaging*, 25(3):369–375, 2006. (Cited on pages 16 and 106.)

[22] J. N. Cohn, R. Ferrari, and N. Sharpe. Cardiac remodeling—concepts and clinical implications: a consensus paper from an international forum on cardiac remodeling. *Journal of the American College of Cardiology*, 35(3):569–582, 2000. (Cited on page 19.)

[23] T. Mansi, I. Voigt, B. Leonardi, X. Pennec, S. Durrleman, M. Sermesant, H. Delingette, A. M. Taylor, Y. Boudjemline, G. Pongiglione, and N. Ayache. A statistical model for quantification and prediction of cardiac remodelling: application to tetralogy of fallot. *IEEE Transactions on Medical Imaging*, 30(9):1605–1616, 2011. (Cited on page 19.)

[24] T. M. Deserno. *Fundamentals of biomedical image processing*. Springer Berlin Heidelberg, 2011. (Cited on pages 21, 22, 38, and 43.)

[25] W. Birkfellner. *Applied Medical Image Processing, Second Edition: A Basic Course*. CRC Press, 2016. (Cited on pages 25, 37, and 52.)

[26] C. Santelli, R. Nezafat, B. Goddu, W. J. Manning, J. Smink, S. Kozerke, and D. C. Peters. Respiratory bellows revisited for motion compensation: preliminary experience for cardiovascular MR. *Magnetic Resonance in Medicine*, 65(4):1097–1102, 2011. (Cited on page 29.)

[27] P. C. Yang, A. B. Kerr, A. C. Liu, D. H. Liang, C. Hardy, C. H. Meyer, A. Macovski, J. M. Pauly, and B. S. Hu. New real-time interactive cardiac magnetic resonance imaging system complements echocardiography. *Journal of the American College of Cardiology*, 32(7):2049–2056, 1998. (Cited on page 29.)

[28] T. Chen, S. Chung, and L. Axel. 2D motion analysis of long axis cardiac tagged MRI. In Nicholas Ayache, Sébastien Ourselin, and Anthony Maeder, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2007*, pages 469–476. Springer Berlin Heidelberg, 2007. (Cited on page 32.)

[29] S. Kuruvilla, N. Adenaw, A B. Katwal, M. J. Lipinski, C. M. Kramer, and M. Salerno. Late gadolinium enhancement on cardiac magnetic resonance predicts adverse cardiovascular outcomes in nonischemic cardiomyopathy: a systematic review and meta-analysis. *Circulation: Cardiovascular Imaging*, 7(2):250–258, 2014. (Cited on page 32.)

[30] A. Sotiras, C. Davatzikos, and N. Paragios. Deformable medical image registration: A survey. *IEEE Transactions on Medical Imaging*, 32(7):1153–1190, 2013. (Cited on page 35.)

[31] J. A. Schnabel, M. P. Heinrich, Papież B. W., and Sir Brady J. M. Advances and challenges in deformable image registration: From image fusion to complex motion modelling. *Medical Image Analysis*, 33:145 – 148, 2016. (Cited on page 35.)

[32] W. Hu, Y. Xie, L. Li, and W. Zhang. A total variation based nonrigid image registration by combining parametric and non-parametric transformation models. *Neurocomputing*, 144(Supplement C):222–237, 2014. (Cited on page 35.)

[33] J. P. Thirion. Image matching as a diffusion process: an analogy with Maxwell's demons. *Medical Image Analysis*, 2(3):243–260, 1998. (Cited on page 37.)

[34] B. K. P. Horn and B. G. Schunck. Determining optical flow. *Artificial Intelligence*, 17(1):185–203, 1981. (Cited on page 37.)

[35] J. Mitra, R. Marti, A. Oliver, X. Llado, J. C. Vilanova, and F. Meriaudeau. A comparison of thin-plate splines with automatic correspondences and B-splines with uniform grids for multimodal prostate registration. In *Medical Imaging 2011: Visualization, Image-Guided Procedures, and Modeling*, volume 7964, pages 7964–1–7964–8, 2011. (Cited on page 38.)

[36] J. Duchon. Interpolation des fonctions de deux variables suivant le principe de la flexion des plaques minces. *R.A.I.R.O. Analyse Numérique*, 10(R3):5–12, 1976. (Cited on page 38.)

[37] J. Meinguet. Multivariate interpolation at arbitrary points made simple. *Zeitschrift für angewandte Mathematik und Physik ZAMP*, 30(2):292–304, 1979. (Cited on page 38.)

[38] G. Donato and S. Belongie. Approximate thin plate spline mappings. In *Computer Vision — ECCV 2002*, pages 21–31. Springer Berlin Heidelberg, 2002. (Cited on page 39.)

[39] D. Perperidis. *Spatio-temporal registration and modelling of the heart using cardiovascular MR imaging*. PhD thesis, Imperial College London, 2006. (Cited on page 41.)

[40] M. Chiara Carminati, F. Maffessanti, and E. Gianluca Caiani. Nearly automated motion artifacts correction between multi breath-hold short-axis and long-axis cine CMR images. *Computers in Biology and Medicine*, 46:42–50, 2014. (Cited on pages 42, 61, and 67.)

[41] C. Studholme, D. L. G. Hill, and D. J. Hawkes. An overlap invariant entropy measure of 3D medical image alignment. *Pattern Recognition*, 32:71–86, 1999. (Cited on page 42.)

[42] N. D. C. Cahill. Normalized measures of mutual information with general definitions of entropy for multimodal image registration. In *Biomedical Image Registration*, volume 6204, pages 258–268. Springer Berlin Heidelberg, 2010. (Cited on page 42.)

[43] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens. Multimodality image registration by maximization of mutual information. *IEEE Transactions on Medical Imaging*, 16(2):187–198, 1997. (Cited on page 42.)

[44] C. P. Bridge and J. A. Noble. Object localisation in fetal ultrasound images using invariant features. In *2015 IEEE 12th International Symposium on Biomedical Imaging (ISBI)*, pages 156–159. IEEE, 2015. (Cited on page 43.)

[45] V. Grau, H. Becher, and J.A. Noble. Phase-based registration of multi-view real-time three-dimensional echocardiographic sequences. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2006*, pages 612–619. Springer Berlin Heidelberg, 2006. (Cited on page 43.)

[46] K. Rajpoot, A. Noble, V. Grau, and N. M. Rajpoot. Feature detection from echocardiography images using local phase information. In *12th Medical Image Understanding and Analysis (MIUA)*, 2008. (Cited on page 43.)

[47] W. Zhang, J.A. Noble, and J.M. Brady. RT3D US to MR cardiovascular image registration using a phase-based approach. In *IEEE International Symposium on Biomedical Imaging (ISBI)*, 2006. (Cited on page 43.)

[48] P. Kovesi. Image features from phase congruency. *Videre: A Journal of Computer Vision Research*, 1(3):1–26, 1999. (Cited on pages 43 and 81.)

[49] M. Felsberg and G. Sommer. The monogenic signal. *IEEE Transactions on Signal Processing*, 49(49), 2001. (Cited on page 44.)

[50] A. V. Oppenheim and J. S. Lim. The importance of phase in signals. *Proceedings of the IEEE*, 69(5):529–541, 1981. (Cited on pages 46, 47, and 75.)

[51] M. Mellor and M. Brady. Phase mutual information as a similarity measure for registration. *Medical Image Analysis*, 9(4):330–343, 2005. (Cited on pages 48 and 82.)

[52] M. Alessandrini, H. Liebgott, D. Barbosa, and O. Bernard. Monogenic phase based optical flow computation for myocardial motion analysis in 3D echocardiography. In O. Camara, T. Mansi, Mihaela Pop, Kawal Rhode, Maxime Sermesant, and Alistair Young, editors, *Statistical Atlases and Computational Models of the Heart. Imaging and Modelling Challenges*, pages 159–168. Springer Berlin Heidelberg, 2013. (Cited on page 48.)

[53] K. Kelly, S. Neubauer, R. Choudhury, E. Dall'Armellina, and V. Grau. A local phase-based algorithm for registration of CMR scans from multiple visits. pages 937–940. IEEE, 2014. (Cited on page 48.)

[54] J. S. Greenfield. A proof for a QuickHull algorithm. Technical report, 1990. (Cited on page 51.)

[55] T. M. Chan. Optimal output-sensitive convex hull algorithms in two and three dimensions. *Discrete & Computational Geometry*, 16(4):361–368, 1996. (Cited on page 51.)

[56] K. Sugihara. Robust gift wrapping for the three-dimensional convex hull. *Journal of Computer and System Sciences*, 49(2):391–407, 1994. (Cited on page 51.)

[57] P. T. An. A modification of graham's algorithm for determining the convex hull of a finite planar set. In *Annales Mathematicae et Informaticae*, volume 34, pages 3–8. Eszterházy Károly College, Institute of Mathematics and Computer Science, 2007. (Cited on page 51.)

[58] W. J. Schroeder, J. A. Zarge, and W. E. Lorensen. Decimation of triangle meshes. *SIGGRAPH Comput. Graph.*, 26(2):65–70, 1992. (Cited on pages 52 and 57.)

[59] D. Zorin, P. Schröder, and W. Sweldens. Interpolating subdivision for meshes with arbitrary topology. In *Proceedings of SIGGRAPH 96, Computer Graphics Proceedings, Annual Conference Series*, volume 30, 1998. (Cited on pages 53, 54, and 55.)

[60] L. Kobbelt, S. Campagna, J. Vorsatz, and H. P. Seidel. Interactive multi-resolution modeling on arbitrary meshes. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques - SIGGRAPH '98*, pages 105–114. ACM Press, 1998. (Cited on page 53.)

[61] L. Kobbelt. $\sqrt{3}$-subdivision. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques - SIGGRAPH '00*, pages 103–112. ACM Press, 2000. (Cited on page 53.)

[62] C. T. Loop. Smooth subdivision surfaces based on triangles. Master's thesis, University of Utah, 1987. (Cited on pages 53 and 54.)

[63] C. Loop. Smooth spline surfaces over irregular meshes. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques - SIGGRAPH '94*, pages 303–310. ACM Press, 1994. (Cited on page 53.)

[64] D. Doo and M. Sabin. Behaviour of recursive division surfaces near extraordinary points. *Computer-Aided Design*, 10(6):356–360, 1978. (Cited on page 53.)

[65] N. Dyn, J. A. Gregory, and D. Levin. A Butterfly subdivision scheme for surface interpolation with tension control. *ACM Transactions on Graphics*, 9(2):160–169, 1990. (Cited on pages 53, 54, and 99.)

[66] E. Catmull and J. Clark. Recursively generated B-spline surfaces on arbitrary topological meshes. *Computer-Aided Design*, 10(6):350–355, 1978. (Cited on page 53.)

[67] H. R. Pakdel and F. F. Samavati. Incremental subdivision for triangle meshes. *International Journal of Computational Science and Engineering*, pages 80–92, 2007. (Cited on page 54.)

[68] J. Warren. Subdivision methods for geometric design. Technical report, 1995. (Cited on page 54.)

[69] J. O. Talton. A short survey of mesh simplification algorithms. Technical report, University of Illinois at Urbana-Champaign, 2004. (Cited on pages 57 and 58.)

[70] M. Garland and P. S. Heckbert. Surface simplification using quadric error metrics. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. ACM Press, 1997. (Cited on pages 58 and 99.)

[71] T. Gabriel. A signal processing approach to fair surface design. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques - SIGGRAPH '95*. ACM Press, 1995. (Cited on pages 59 and 99.)

[72] D.A. Field. Laplacian smoothing and Delaunay triangulations. *Communications in Applied Numerical Methods*, 4(6):709–712, 1988. (Cited on page 59.)

[73] W. Jun and Y. Zeyun. A novel method for surface mesh smoothing: Applications in biomedical modeling. In *Proceedings of the 18th International Meshing Roundtable*, pages 195–210. Springer Berlin Heidelberg, 2009. (Cited on page 59.)

[74] T. R. Jones, F. Durand, and M. Desbrun. Non-iterative, feature-preserving mesh smoothing. *ACM Transactions on Graphics*, 22(3):943–949, 2003. (Cited on page 59.)

[75] H. Zhang and E. Fiume. Mesh smoothing with shape or feature preservation. *Advances in Modelling, Animation and Rendering*, pages 167–181, 2002. (Cited on page 59.)

[76] D. Le Bihan, C. Poupon, A. Amadon, and F. Lethimonnier. Artifacts and pitfalls in diffusion mri. *Journal of Magnetic Resonance Imaging*, 24(3):478–488, 2006. (Cited on page 60.)

[77] B. W. Papież, J. Franklin, M. P. Henrich, F. V. Gleeson, and J. A. Schnabel. Liver motion estimation via locally adaptive over-segmentation. In N. Navab, J. Hornegger, William M. Wells, and A. F. Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, volume 9351, pages 427–434. Springer Berlin Heidelberg, 2015. (Cited on page 60.)

[78] F. Odille, A. Bustin, B. Chen, Vuissoz P. A., and J. Felbinger. Motion-corrected, super-resolution reconstruction for high-resolution 3d cardiac cine mri. In N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, volume 9351, pages 435–442. Springer Berlin Heidelberg, 2015. (Cited on page 60.)

[79] F. Godenschweger, U. Kägebein, D. Stucht, U. Yarach, A. Sciarra, R. Yakupov, F. Lüsebrink, P. Schulze, and O. Speck. Motion correction in mri of the brain. *Physics in Medicine and Biology*, 61(5):R32–R36, 2016. (Cited on page 60.)

[80] M.R. Makowski, A. J. Wiethoff, C. H. P. Jansen, S. Uribe, V. Parish, A. Schuster, R. M. Botnar, A. Bell, C. Kiesewetter, R. Razavi, T. Schaeffter, and G. F. Greil. Single breath-hold assessment of cardiac function using an accelerated 3D single breath-hold acquisition technique–comparison of an intravascular and extravascular contrast agent. *Journal of Cardiovascular Magnetic Resonance*, 14:53, 2012. (Cited on page 61.)

[81] J. Jaroni, R. Meier, A. Beer, K. Herrmann, M. Settles, E. J. Rummeny, and A. Huber. Three-dimensional magnetic resonance imaging using single breath-hold k-t BLAST for assessment of global left ventricular functional parameters. *Academic Radiology*, 20(8):987–994, 2013. (Cited on page 61.)

[82] W. Shi, X. Zhuang, H. Wang, S. Duckett, D. V. N. Luong, C. Tobon-Gomez, K. Tung, P. J. Edwards, K. S. Rhode, R. S. Razavi, S. Ourselin, and D. Rueckert. A comprehensive cardiac motion estimation framework using both untagged and 3-D tagged MR images based on nonrigid registration. *IEEE Transactions on Medical Imaging*, 31(6):1263–1275, 2012. (Cited on page 61.)

[83] C. Swingen, R. T. Seethamraju, and M. Jerosch-Herold. An approach to the three-dimensional display of left ventricular function and viability using MRI. *The International Journal of Cardiovascular Imaging*, 19(4):325–336, 2003. (Cited on pages 61 and 62.)

[84] P. Peng, K. Lekadir, A. Gooya, L. Shao, S. E. Petersen, and A. F. Frangi. A review of heart chamber segmentation for structural and functional analysis using cardiac magnetic resonance imaging. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 29(2):155–195, 2016. (Cited on page 61.)

[85] H.C. van Assen, M. G. Danilouchkine, M. S. Dirksen, J. H. C. Reiber, and B. P. F. Lelieveld. A 3-D active shape model driven by fuzzy inference: Application to cardiac CT and MR. *IEEE Transactions on Information Technology in Biomedicine*, 12(5):595–605, 2008. (Cited on page 61.)

[86] A. A. Young, B. R. Cowan, S. F. Thrupp, W. J. Hedley, and L. J. Dell'Italia. Left ventricular mass and volume: Fast calculation with guide-point modelling on MR images. *Radiology*, 2(216):597–602, 2000. (Cited on pages 61 and 97.)

[87] B. Villard, V. Grau, and E. Zacur. Surface mesh reconstruction from cardiac MRI contours. *Journal of Imaging*, 2018. (Cited on pages 61, 72, 92, 94, 120, and 121.)

[88] J. Lötjönen, J. Koikkalainen, D. Smutek, S. Kivistö, and K. Lauerma. Four-chamber 3-D statistical shape model from cardiac short-axis and long-axis MR images. In Randy E. Ellis and Terry M. Peters, editors, *Lecture Notes in Computer Science*, pages 459–466. Springer Berlin Heidelberg, 2003. (Cited on page 62.)

[89] H. C. van Assen, M. G. Danilouchkine, A. F. Frangi, S. Ordás, J. J. M. Westenberg, J. H. C. Reiber, and B. P. F. Lelieveldt. SPASM: A 3D-ASM for segmentation of sparse and arbitrarily oriented cardiac MRI data. *Medical Image Analysis*, 10(2):286–303, 2006. (Cited on page 62.)

[90] R. Casero, U. Siedlecka, E. S. Jones, L. Gruscheski, M. Gibb, J. E. Schneider, P. Kohl, and V. Grau. Transformation diffusion reconstruction of three-dimensional histology volumes from two-dimensional image stacks. *Medical Image Analysis*, 38:184–204, 2017. (Cited on pages 62 and 74.)

[91] G. Malandain, É. Bardinet, K. Nelissen, and W. Vanduffel. Fusion of autoradiographs with an MR volume using 2-D and 3-D linear transformations. *NeuroImage*, 23(1):111–127, 2004. (Cited on page 62.)

[92] Y. Su, M. L. Tan, C. W. Lim, S. K. Teo, K. S. Senthil, L. T. Zhong, and R. S. Tan. Automatic correction of motion artifacts in 4D left ventricle model reconstructed from MRI. volume 41, pages 705–708, 2014. (Cited on pages 62 and 74.)

[93] M. L. Tan, Y. Su, C.W. Lim, S. K. Selvaraj, L. Zhong, and R. S. Tan. A geometrical approach for automatic shape restoration of the left ventricle. *PLoS ONE*, 8(7):1–11, 2013. (Cited on page 62.)

[94] M. Wan, W. Huang, J. M. Zhang, X. Zhao, J. C. Allen, R. S. Tan, X. Wan, and L. Zhong. Correcting motion in multiplanar cardiac magnetic resonance images. *BioMedical Engineering OnLine*, 15(1), 2016. (Cited on page 63.)

[95] P. Medrano-Gracia, B. R. Cowan, B. Ambale-Venkatesh, D. A. Bluemke, J. Eng, J. P. Finn, C. G. Fonseca, J. A. C. Lima, A. Suinesiaputra, and A. A. Young. Left ventricular shape variation in asymptomatic populations: the multi-ethnic study of atherosclerosis. *Journal of Cardiovascular Magnetic Resonance*, 16(1):56, 2014. (Cited on pages 63, 93, and 97.)

[96] J. Ector, S. De Buck, J. Adams, S. Dymarkowski, J. Bogaert, F. Maes, and H. Heidbüchel. Cardiac three-dimensional magnetic resonance imaging and fluoroscopy merging. a new approach for electroanatomic mapping to assist catheter ablation. *Circulation*, 2005. (Cited on page 63.)

[97] A. G. Chandler, R. J. Pinder, T. Netsch, J. A. Schnabel, D. J. Hawkes, D. L. G. Hill, and R. Razavi. Correction of misaligned slices in multi-slice cardiovascular magnetic resonance using slice-to-volume registration. *Journal of Cardiovascular Magnetic Resonance*, 10(1):1–9, 2008. (Cited on page 63.)

[98] C. Zakkaroff, A. Radjenovic, J. Greenwood, and D. Magee. Recovery of slice rotations with the stack alignment transform in cardiac MR series. In *Procedings of the British Machine Vision Conference 2012*. British Machine Vision Association, 2012. (Cited on pages 64 and 67.)

[99] C. Zakkaroff, A. Radjenovic, J. Greenwood, and D. Magee. Stack alignment transform for misalignment correction in cardiac MR cine series. Technical report, University of Leeds, 2012. (Cited on page 64.)

[100] J. Lötjönen, Mika Pollari, S. Kivistö, and K. Lauerma. Correction of movement artifacts from 4-D cardiac short- and long-axis MR data. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2004*, pages 405–412. Springer Berlin Heidelberg, 2004. (Cited on page 64.)

[101] J. Lötjönen, M. Pollari, S. Kivistö, and K. Lauerma. Correction of motion artifacts from cardiac cine magnetic resonance images. *Academic Radiology*, 12(10), 2005. (Cited on pages 64 and 67.)

[102] J. Moore, M. Drangova, M. Wierzbicki, J. Barron, and T. Peters. A high resolution dynamic heart model based on averaged MRI data. In *Lecture Notes in Computer Science*, pages 549–555. Springer Berlin Heidelberg, 2003. (Cited on page 64.)

[103] B. Kim, J. L. Boes, P. H. Bland, T. L. Chenevert, and C. R. Meyer. Motion correction in fMRI via registration of individual slices into an anatomical volume. *Magnetic Resonance in Medicine*, 41(5):964–972, 1999. (Cited on page 64.)

[104] J. Barajas, K. L. Caballero, J. Garcia Barnés, F. Carreras, S. Pujadas, and P. Radeva. Correction of misalignment artifacts among 2-D cardiac MR images in 3-D space. In *First International Workshop on Computer Vision for Intravascular and Intracardiac Imaging, MICCAI 2006*, volume 1, pages 114–121, 2006. (Cited on page 65.)

[105] K. Kim, P. A. Habas, F. Rousseau, O. A. Glenn, A. J. Barkovich, and C. Studholme. Intersection based motion correction of multislice MRI for 3-D in utero fetal brain image formation. *IEEE Transactions on Medical Imaging*, 29(1):146–158, 2010. (Cited on page 65.)

[106] P. J. Slomka, D. Fieno, A. Ramesh, V. Goyal, H. Nishina, L. E. J. Thompson, R. Saouaf, D. S. Berman, and G. Germano. Patient motion correction for multiplanar, multi-breath-hold cardiac cine MR imaging. *Journal of Magnetic Resonance Imaging*, 25(5):965–973, 2007. (Cited on pages 65, 66, and 76.)

[107] A. Elen, J. Hermans, J. Ganame, D. Loeckx, J. Bogaert, F. Maes, and P. Suetens. Automatic 3-D breath-hold related motion correction of dynamic multislice MRI. *IEEE Transactions in Medical Imaging*, 29(3):868–878, 2010. (Cited on pages 65, 66, 67, and 76.)

[108] Y. M. Liew, R. A. McLaughlin, B. T. Chan, Y. F. Abdul Aziz, K. H. Chee, N. M. Ung, L. K. Tan, K. W. Lai, S. Ng, and E. Lim. Motion corrected LV quantification based on 3D modelling for improved functional assessment in cardiac MRI. *Physics in Medicine and Biology*, 60(7):2715–2733, 2015. (Cited on pages 66, 67, 75, 76, and 78.)

[109] K. K. Bhatia, J. V. Hajnal, B. K. Puri, A. D. Edwards, and D. Rueckert. Consistent groupwise non-rigid registration for atlas construction. In *2004 2nd IEEE International Symposium on Biomedical Imaging: Nano to Macro (IEEE Cat No. 04EX821)*, volume 1, pages 908–911, 2004. (Cited on page 67.)

[110] S. K. Balci, P. Golland, and W. M. Wells III. Non-rigid groupwise registration using B-spline deformation model. volume 10, pages 105–21. Med Image Comput Comput Assist Interv. (MICCAI), 2007. (Cited on page 67.)

[111] J. C. Bezdek and R. J. Hathaway. Convergence of alternating optimization. *Neural, Parallel, and Scientific Computations*, 11(4):351–368, 2003. (Cited on pages 67 and 76.)

[112] M. Sinclair, W. Bai, E. Puyol-Antón, O. Oktay, D. Rueckert, and A. P. King. Fully automated segmentation-based respiratory motion correction of multiplanar cardiac magnetic resonance images for large-scale datasets. In *Lecture Notes in Computer Science*, volume 10434, pages 332–340. Springer International Publishing, 2017. (Cited on page 68.)

[113] B. Villard, E. Zacur, E. Dall'Armellina, and V. Grau. Correction of slice misalignment in multi-breath-hold cardiac MRI scans. *Statistical Atlases and Computational Models of the Heart. Imaging and Modelling Challenges.*

STACOM 2016. Lecture Notes in Computer Science (LNCS), 10124, 2017. (Cited on pages 68, 70, 71, 73, and 94.)

[114] A. Lopez-Perez, Rafael S., and J. M. Ferrero. Three-dimensional cardiac computational modelling: methods, features and applications. *BioMedical Engineering OnLine*, 14(1):35, 2015. (Cited on page 68.)

[115] M. Zou, M. Holloway, Carr N., and T. Ju. Topology-constrained surface reconstruction from cross-sections. *ACM Transactions on Graphics*, 34(4):128:1–128:10, 2015. (Cited on pages 68, 69, and 93.)

[116] O. Schall and M. Samozino. Surface from scattered points. In *A Brief Survey of Recent Developments. 1st International Workshop on Semantic Virtual Environments, Page (S*, pages 138–147, 2005. (Cited on page 68.)

[117] M. Berger, A. Tagliasacchi, L. M. Seversky, P. Alliez, G. Guennebaud, J. A. Levine, A. Sharf, and C. T. Silva. A survey of surface reconstruction from point clouds. *Computer Graphics Forum*, 36(1):301–329, 2016. (Cited on page 68.)

[118] A. Khatamian and H. R. Arabnia. Survey on 3D surface reconstruction. *Journal of Information Processing Systems*, pages 338–357, 2016. (Cited on pages 68 and 69.)

[119] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Surface reconstruction from unorganized points. *ACM SIGGRAPH Computer Graphics*, 26(2):71–78, 1992. (Cited on page 68.)

[120] F. Bernardini, J. Mittleman, H. Rushmeier, C. Silva, and G. Taubin. The ball-pivoting algorithm for surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics*, 5:349–359, 1999. (Cited on page 68.)

[121] C.W. Lim, Y. Su, S. Yong Yeo, G. M. Ng, V. T. Nguyen, L. Zhong, R. S. Tan, K. K. Poh, and P. Chai. Automatic 4D reconstruction of patient-specific

cardiac mesh with 1-to-1 vertex correspondence from segmented contours lines. *PLoS ONE*, 9(4):e93747, 2014. (Cited on pages 69 and 97.)

[122] K. Sunderland, B. Woo, C. Pinter, and G. Fichtinger. Reconstruction of surfaces from planar contours through contour interpolation. In R. J. Webster and Ziv R. Yaniv, editors, *Medical Imaging 2015: Image-Guided Procedures, Robotic Interventions, and Modeling*. SPIE, 2015. (Cited on pages 69 and 93.)

[123] L. Liu, C. Bajaj, J. O. Deasy, D. A. Low, and T. Ju. Surface reconstruction from non-parallel curve networks. *Computer Graphics Forum*, 27(2), 2008. (Cited on page 69.)

[124] Z. Zhang, K. Konno, and k. Tokuyama. 3D terrain reconstruction based on contours. *Ninth International Conference on Computer Aided Design and Computer Graphics*, pages 325–330, 2005. (Cited on page 69.)

[125] W. Zhenhua, G. Nan, and Z. Zhiyi. Surface mesh reconstruction based on contours. In *2009 International Conference on Computational Intelligence and Software Engineering*. IEEE, 2009. (Cited on page 69.)

[126] P. Lamata, S. Niederer, D. Nordsletten, D. C. Barber, I. Roy, D. R. Hose, and N. Smith. An accurate, fast and robust method to generate patient-specific cubic Hermite meshes. *Medical Image Analysis*, 15(6):801–813, 2011. (Cited on pages 69, 93, and 97.)

[127] M. Kazhdan, M. Bolitho, and H. Hoppe. Poisson surface reconstruction. In *Eurographics Symposium on Geometry Processing*, 2006. (Cited on page 70.)

[128] A. Paiement, M. Mirmehdi, X. Xie, and M. C. K. Hamilton. Registration and modeling from spaced and misaligned image volumes. *IEEE Transactions on Image Processing*, 25(9):4379–4393, 2016. (Cited on pages 70 and 120.)

[129] A. Paiement, M. Mirmehdi, X. Xie, and M. C. K. Hamilton. Integrated segmentation and interpolation of sparse data. *IEEE Transactions on Image Processing*, 23(1):110–125, 2014. (Cited on pages 70 and 120.)

[130] B. Villard, V. Carapella, R. Ariga, V. Grau, and E. Zacur. Cardiac mesh reconstruction from sparse, heterogeneous contours. In *Communications in Computer and Information Science*, pages 169–181. Springer International Publishing, 2017. (Cited on pages 72, 92, 94, 120, and 121.)

[131] B. Villard, V. Grau, and E. Zacur. SurFCo: Surface From Contours. In International Symposium on Biomedical Imaging (ISBI), 2018. (Cited on pages 72 and 92.)

[132] A. A. Goshtasby and D. A. Turner. Fusion of short-axis and long-axis cardiac MR images. *Computerized Medical Imaging and Graphics*, 20:202–211, 1996. (Cited on page 75.)

[133] Jorge Nocedal and Stephen J Wright. *Numerical Optimization*. Springer New York, 2006. (Cited on page 79.)

[134] X. Li, X. Wang, Z. Sheng, and X. Duan. A modified conjugate gradient algorithm with backtracking line search technique for large-scale nonlinear equations. *International Journal of Computer Mathematics*, 95(2):382–395, 2018. (Cited on page 80.)

[135] Neculai A. A modified polak–ribiére–polyak conjugate gradient algorithm for unconstrained optimization. *Optimization*, 60(12):1457–1471, 2011. (Cited on page 80.)

[136] S. Bernstein, J.-L. Bouchot, M. Reinhardt, and B. Heise. Generalized analytic signals in image processing: comparison, theory and applications. *Quaternion and Clifford Fourier Transforms and Wavelets*, pages 221–246, 2013. (Cited on pages 80 and 81.)

[137] K. G. Derpanis. Quadrature filters. Technical report, Department of Computer Science and Engineering, York University, 2005. (Cited on page 81.)

[138] U. Jentsch and G. Ivanova. Circular correlation methods for the analysis of oscillations in dependent time series. *Proceedings of the 29th Annual International Conference of the IEEE EMBS*, pages 1981–1984, 2007. (Cited on page 83.)

[139] H. J. Arevalo, F. Vadakkumpadan, E. Guallar, A. Jebb, P. Malamas, K. C. Wu, and N. A. Trayanova. Arrhythmia risk stratification of patients after myocardial infarction using personalized heart models. *Nature Communications*, 7, 2016. (Cited on page 93.)

[140] N. Smith, A. de Vecchi, M. McCormick, D. Nordsletten, O. Camara, A. F. Frangi, H. Delingette, M. Sermesant, J. Relan, N. Ayache, M. W. Krueger, W. H. W. Schulze, R. Hose, I. Valverde, P. Beerbaum, C. Staicu, M. Siebes, J. Spaan, P. Hunter, J. Weese, H. Lehmann, D. Chapelle, and R. Rezavi. euHeart: Personalized and integrated cardiac care using patient-specific cardiovascular modelling. *Interface Focus*, 1(1):349–364, 2011. (Cited on page 93.)

[141] D. Deng, J. Zhang, and L. Xia. Three-dimensional mesh generation for human heart model. In *Communications in Computer and Information Science*, pages 157–162. Springer Berlin Heidelberg, 2010. (Cited on page 93.)

[142] M. Vukicevic, B. Mosadegh, J. K. Min, and S. H. Little. Cardiac 3D printing and its future directions. *JACC: Cardiovascular Imaging*, 10(2):171–184, 2017. (Cited on page 93.)

[143] H. Lehmann, R. Kneser, M. Neizel, J. Peters, O. Ecabert, H. Kühl, M. Kelm, and J. Weese. Integrating viability information into a cardiac model for interventional guidance. In *Functional Imaging and Modeling of the Heart*, pages 312–320. Springer Berlin Heidelberg, 2009. (Cited on page 93.)

[144] X. Albà, M. Pereañez, C. Hoogendoorn, A. J. Swift, Jim M. Wild, A. F. Frangi, and K. Lekadir. An algorithm for the segmentation of highly abnormal hearts

using a generic statistical shape model. *IEEE Transactions on Medical Imaging*, 35(3):845–859, 2016. (Cited on page 93.)

[145] A. de Marvao, T. J. W. Dawes, W. Shi, C. Minas, N. G. Keenan, T. Diamond, G. Durighel, G. Montana, D. Rueckert, S. A. Cook, and D. P. O'Regan. Population-based studies of myocardial hypertrophy: high resolution cardiovascular magnetic resonance atlases improve statistical power. *Journal of Cardiovascular Magnetic Resonance*, 16(1), 2014. (Cited on page 95.)

[146] W. Bai, W. Shi, A. de Marvao, T. J. W. Dawes, D. P. O'Regan, S. A. Cook, and D. Rueckert. A bi-ventricular cardiac atlas built from 1000+ high resolution MR images of healthy subjects and an analysis of shape and motion. *Medical Image Analysis*, 26(1):133–145, 2015. (Cited on pages 95, 130, 132, and 134.)

[147] J. G. Gonzalez and G. R. Arce. Optimality of the myriad filter in practical impulsive-noise environments. *IEEE Transactions on Signal Processing*, 49:438–441, 2001. (Cited on page 97.)

[148] C. Moenning and A. D. Neil. Fast marching farthest point sampling for implicit surfaces and point clouds. Technical report, University of Cambridge, 2003. (Cited on page 98.)

[149] K. Rohr, H. S. Stiehl, R. Sprengel, T. M. Buzug, J. Weese, and M. H. Kuhn. Landmark-based elastic registration using approximating thin-plate splines. *IEEE Transactions on Medical Imaging*, 20(6):526–534, 2001. (Cited on pages 98 and 129.)

[150] R. Sprengel, K. Rohr, and H. S. Stiehl. Thin-plate spline approximation for image registration. In *Proceedings of 18th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, volume 3, pages 1190–1191, 1996. (Cited on page 98.)

[151] A. Nealen, T. Igarashi, O. Sorkine, and M. Alexa. Laplacian mesh optimization. In *Proceedings of the 4th international conference on Computer graphics and interactive techniques in Australasia and Southeast Asia - GRAPHITE '06*. ACM Press, 2006. (Cited on page 99.)

[152] S. Xie and Z. Tu. Holistically-nested edge detection. In *Proceedings of IEEE International Conference on Computer Vision*, 2015. (Cited on pages 122 and 159.)

[153] D. M. Vigneault, W. Xie, D. A. Bluemke, and J. A. Noble. Feature tracking cardiac magnetic resonance via deep learning and spline optimization. In *Functional Imaging and Modelling of the Heart*, pages 183–194. Springer International Publishing, 2017. (Cited on pages 123 and 124.)

[154] O. Ronneberger, P. Fischer, and T. Brox. U-Net: Convolutional networks for biomedical image segmentation. In *Lecture Notes in Computer Science*, pages 234–241. Springer International Publishing, 2015. (Cited on pages 123 and 124.)

[155] J. Wang and L. Perez. The effectiveness of data augmentation in image classification using deep learning. Technical report, Standford University, 2017. (Cited on page 123.)

[156] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2015. (Cited on page 124.)

[157] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2016. (Cited on page 124.)

[158] G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. Technical report, 2018. arXiv:1608.06993. (Cited on page 124.)

[159] V. Badrinarayanan, A. Kendall, and R. Cipolla. SegNet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017. (Cited on page 124.)

[160] K. H. Zou, S. K. Warfield, A. Bharatha, C. M. C. Tempany, M. R. Kaus, S. J. Haker, W. M. Wells, F. A. Jolesz, and R. Kikinis. Statistical validation of image segmentation quality based on a spatial overlap index: scientific reports. *Academic Radiology*, 11(2):178–189, 2004. (Cited on page 126.)

[161] P. J. Besl and N. D. McKay. A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992. (Cited on page 131.)