Evaluation of Tissue Evolution after Myocardial Infarction: Correspondence of CMR Scans from Multiple Visits



Christopher Kelly

Balliol College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy in Engineering Science

Trinity 2018

Acknowledgements

I would firstly like to thank my supervisor, Professor Vicente Grau, for his constant support during my D.Phil. studies - for the insightful discussions during our weekly meetings, and his guidance and encouragement throughout the thesis write-up period.

My thanks also to my research group, and colleagues in the wider Biomedical Image Analysis Laboratory for their company and support that has made for such an enjoyable working environment.

I am grateful to our collaborators at the John Radcliffe Hospital for the provision of data and their helpful clinical consultation. I would also like to thank the Engineering and Physical Sciences Research Council for providing the funding that has allowed me to carry out my D.Phil. research.

Away from my studies, I am indebted to the welcoming and vibrant environment of Balliol's graduate community at Holywell Manor, where I have had the opportunity to form so many important friendships. My thanks to all my Oxford friends, and those further afield, who have been a constant source of support and encouragement.

Finally, I am forever grateful to my parents and family for their continued backing, unwavering belief, and inexhaustible patience throughout my undergraduate and graduate academic life.

Abstract

In response to myocardial infarction injury, the heart undergoes a remodelling process that alters chamber structure and affects cardiac function. Due to its excellent anatomical and functional tissue characterisation, cardiovascular magnetic resonance (CMR) imaging has become a well-established diagnostic tool in such cases of ischemic heart disease; in particular, CMR imaging is increasingly used to assess cardiac state in the acute phase of the disease. The longitudinal analysis of these CMR acquisitions can help to elucidate the remodelling process and evaluate the efficacy of clinical interventions.

A limiting factor for all post-processing operations is the constrained resolution of clinically-acquired CMR data. The standard planar images obtained are highly anisotropic, leading to poor quality volume reconstructions when image data are integrated. This can have detrimental effects on operations such as segmentation and registration, reducing confidence in the clinical conclusions resulting from image analysis.

In this thesis, we look to develop registration methods that enable the comparative analysis of longitudinally acquired CMR data, in the form of complementary cine and late gadolinium-enhanced (LGE) studies. Phase-based registration approaches are proposed to address the challenges of data misalignment observed in individual imaging sessions, as well as to compensate for the differences between longitudinal follow-up data.

Poor quality volumetric reconstructions can have a detrimental effect on the 3D registration accuracy. Therefore, in the final part of this thesis, we present an anisotropic reaction-diffusion-based reconstruction process, which reduces some of the undesirable structure-blurring outcomes from linear interpolation. In addition, we explore a number of adaptations to conventional interpolation of short-axis slice data, in which we aim to maximise the utilisation of structural information contained in the imaging studies acquired.

Contents

1 Introduction **1**
- 1.1 Clinical Motivation . . . 1
- 1.2 Thesis Outline . . . 3

2 Background - Cardiac Anatomy and Imaging **7**
- 2.1 Cardiac Physiology . . . 7
- 2.2 Cardiac Pathophysiology . . . 13
- 2.3 Cardiac Imaging . . . 16
  - 2.3.1 Magnetic resonance imaging (MRI) . . . 18
  - 2.3.2 Cardiac magnetic resonance imaging (CMRI) . . . 25

3 Clinical Data and Associated Challenges **32**
- 3.1 Clinical Data . . . 33
  - 3.1.1 Data delineations . . . 34
  - 3.1.2 Cardiac mesh generation . . . 37
  - 3.1.3 Scar calculation from LGE image data . . . 40
- 3.2 Challenges and Limitations of MRI Acquisition . . . 43
  - 3.2.1 Image resolution and partial-volume effect . . . 43
  - 3.2.2 Cardiac and respiratory motion - Breath-hold misalignments 45
  - 3.2.3 Intensity inconsistencies . . . 47
  - 3.2.4 Image acquisition reproducibility - Misalignment of slices from different studies . . . 47

4 Background - Image Analysis **54**
- 4.1 Registration . . . 55
  - 4.1.1 Basics . . . 56
  - 4.1.2 Optical flow . . . 61
  - 4.1.3 Applications of image registration in the analysis of CMR data 66
- 4.2 Reconstruction . . . 80
  - 4.2.1 Volumetric reconstruction methodologies applied to CMR data 83
- 4.3 Local Phase . . . 85
  - 4.3.1 Local phase in nD . . . 92
  - 4.3.2 Further phase-based properties . . . 99
  - 4.3.3 Applications in image registration . . . 101

5 Image Registration of Cardiac Datasets **105**
5.1 Introduction . . . 105
5.1.1 Materials . . . 107
5.2 Slice Realignment . . . 107
5.2.1 Methods . . . 107
5.2.2 Experiments and results . . . 114
5.2.3 Discussion . . . 126
5.3 LGE to Cine Registration . . . 129
5.3.1 Methods . . . 131
5.3.2 Experiments and results . . . 133
5.3.3 Discussion . . . 139
5.4 3D CMR Registration - Longitudinal Analysis . . . 142
5.4.1 Methods . . . 142
5.4.2 Experiments and results . . . 150
5.4.3 Discussion . . . 169

6 Volume Reconstruction of Cardiac Datasets from Multiple Slices 175
6.1 Introduction . . . 175
6.2 Diffusion-based Processing . . . 178
6.3 Methods . . . 188
6.3.1 Reaction-diffusion interpolation . . . 188
6.3.2 Multi-frame reconstruction - Combining information from all cardiac phases . . . 194
6.3.3 Registration-based reconstruction . . . 196
6.3.4 Integration of LA and SA data . . . 199
6.4 Experiments and Results . . . 204
6.5 Discussion . . . 213

7 Conclusions and Future Work **222**
7.1 Conclusions . . . 222
7.2 Future Work . . . 224

References **227**

1 Introduction

Contents

1.1 Clinical Motivation 1
1.2 Thesis Outline 3

1.1 Clinical Motivation

Ischaemic heart disease, alternatively referred to as coronary heart disease (CHD), is the leading cause of death worldwide, responsible for an estimated 8.75 million deaths in 2015 [1]. Despite dramatic improvements made during the past half a century, it remains the most common cause of premature death in the UK [175]. In 2014, 15% of male deaths and 10% of deaths in women were attributed to CHD, making it the biggest single cause of death [201]. While improvements in medical treatments and drug therapies have reduced the rates of mortality, they have simultaneously shifted the burden of the disease on society; far greater numbers now live with the disease, with over one and a half million people living in the UK having had a heart attack, and an estimated 2 million suffering from angina and/or heart failure [175]. As a result, improvements in the understanding of how the pathological heart changes in response to injury are becoming increasingly important, in order that the optimal treatment outcomes can be achieved to minimise complications in later life.

Cardiovascular magnetic resonance (CMR) imaging has a key role to play in the diagnosis and care of patients with ischaemic myocardial injury. Due to its provision, by non-invasive means, of complementary structural and functional information, CMR is considered by some a "one-stop-shop" for cardiac assessment [160]. Its all-encompassing evaluation makes it an invaluable tool in the modern clinical setting, with the potential to greatly increase understanding of the mechanisms of heart failure and improve patients' outcomes.

In cases of myocardial ischaemic injury, the longitudinal changes in tissue composition lead to global left ventricle remodelling, which constitutes the alteration of chamber structure and configuration [95]. This left ventricle remodelling is an important stage in the progression of cardiac disease and can be used as an early indicator of patient prognosis [118]. In the short term, remodelling is due to fibrotic repair of the infarcted area to form scar tissue, whilst in the longer term remodelling occurs predominantly in the non-ischaemic regions as it responds to the increased load put upon it [192]. Due to the anticipated dilation of the left ventricle that occurs during remodelling, progression into heart failure is often assessed by means of the cardiac chamber volumes. However, such approaches do not provide insight into localised tissue changes and their mechanics. An improved understanding of how local tissue alterations affect the final ventricular remodelling process could prove beneficial [95], leading to a greater appreciation of the merits and limitations of different interventions.

The first contribution of this thesis, therefore, is a framework for the 3D analysis of longitudinal changes in tissue, based on the acquisition of CMR images at multiple longitudinal time-points. Specifically, the data acquired from patients in the first 24 hours after cardiac intervention and their 6 month follow-up assessment is analysed. We identified a number of challenges associated with such an analysis, and look to address many of these in chapter 5 of the thesis.

While working on the 3D registration of cardiac cine data it became apparent that the linear interpolation of the SA data can result in qualitatively unsatisfactory volumetric data representation, and that the poor definition of important tissue

structures could have an adverse effect on the registration results. Therefore, we look to address this issue in the second part of the thesis; we propose a novel anisotropic reaction-diffusion-based interpolation method that generates a final isotropic volumetric tissue representation which has improved image quality compared to the more common linear approaches.

1.2 Thesis Outline

A brief description is given below of the work presented in each of the thesis chapters.

Chapter 2 provides an introduction to both cardiac physiology and pathophysiology. In particular, we discuss pathologies that lead to myocardial infarction (MI) and consider the subsequent response of the heart to this injury. An introduction to magnetic resonance imaging (MRI) follows this, where the physics of image acquisition and the applications in cardiac analysis are discussed.

Chapter 3 introduces the data used in this thesis. In addition to the acquired image data, we are provided with clinical delineations of the important cardiac structures, which are to be used for the validation of the image processing tasks. Specifically, for this purpose an integration of the data to form 3D surfaces is required; a method is proposed here that accomplished this using the open-source mesh manipulation tool-kit, MeshLab [207]. In addition to validation, the delineations of the late gadolinium enhanced (LGE) data are further used to estimate the scar tissue present in the images. In the final part of this chapter, we discuss some of the main challenges that affect our MR datasets which will be addressed later in the thesis. In particular, we present results from a study where we quantified the differences in SA slice locations between two follow-up MR studies. These results, which were published as part of the abstracts of the 17th annual SCMR scientific sessions [87], illustrate the problematic nature of slice-to-slice assessment for the evaluation longitudinal tissue changes.

Chapter 4 discusses the background to the key image analysis approaches used in this thesis. Registration methodology is introduced, before a discussion of the state-of-the-art literature relating to the three key registration tasks addressed is given. Following this, a brief discussion of approaches for improving image resolution is presented, in addition to a review of state-of-the-art methods for 3D volume reconstruction from planar cardiac MR acquisitions. Finally, we present the theory of local phase-based analysis, which is used extensively in the remainder of the thesis.

Chapter 5 is comprised of three main sections, each of which presents an image registration challenge related to cardiac MR data acquisitions. Firstly, a local phase-based method is proposed for correcting the positions of 2D slices that are misaligned as a result of breath-hold variation. Having corrected for the relative alignment of short-axis (SA) and long-axis (LA) data in both cine and LGE studies, we next consider the alignment of the LGE images to the cine data. The relative positions of these slices will have been similarly corrupted as a result of breath-hold imaging and patient movements. Therefore, a further registration is required to bring these complementary studies into alignment. Once again we employ a local phase-based approach with which we obtain a 3D rigid transformation that compensates for misalignments. Finally, a 3D non-rigid registration of the longitudinal acquired data is presented. Preliminary results using a phase-based similarity metric were presented at the 2014 Computing in Cardiology conference [88]. In this chapter a combined local-global variational optical flow framework is used, with a local phase vector similarity metric. The required deformation is calculated according to the cine data correspondences, before the same calculated deformation field is applied to the previously-aligned LGE data. This brings all cine and LGE studies into the same coordinate frame, enabling the assessment of tissue changes as a result of cardiac remodelling, such as the modifications observed in the myocardial scar region. A journal paper discussing the results using the proposed local phase registration is currently under preparation. An

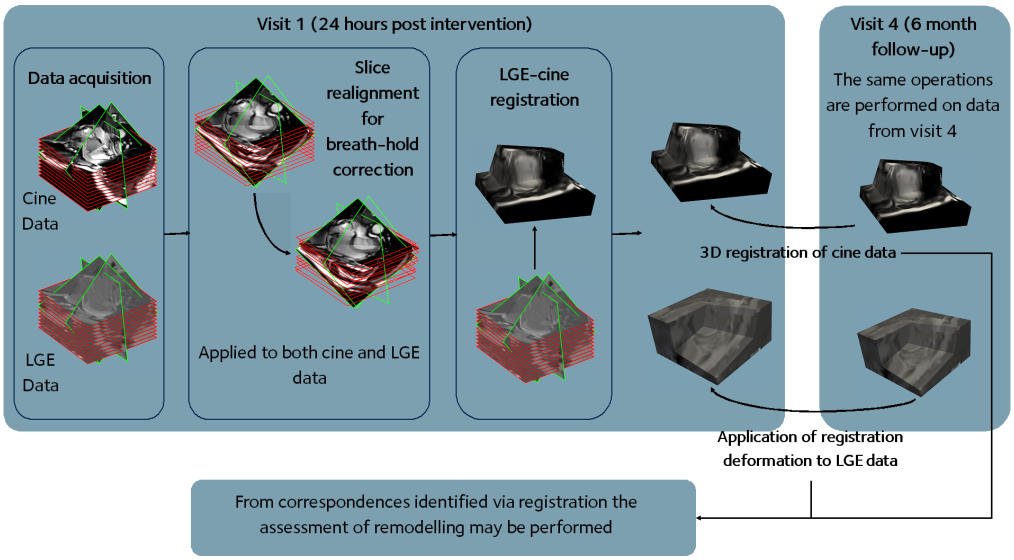


Figure 1.1: Outline of the 3D registration framework used for longitudinal multi-CMR-image analysis.

illustration of the steps required for the assessment discussed in this chapter is given in figure 1.1.

Chapter 6 describes a novel method for generating isotropic volumetric reconstructions of cardiac tissue, imaged by means of multiple thick planar slices at varying positions and orientations. The proposed method uses an anisotropic diffusion system, coupled with a reaction component that maintains the high in-plane detail available in the original anisotropically acquired slices. The first results presented for the reaction-diffusion-based reconstruction are based on an initial linear interpolation of the SA slice data. These results illustrate the dependency of the method on the initialisation of the volume reconstruction; therefore, we explore a number of alternative initialisation strategies that look to utilise the acquired slice data in different ways. These are: the combined utilisation of information from the entire cardiac cycle, captured in the multiple frames of the cine sequence; the registration of adjacent SA slices, for a volumetric linear interpolation that is guided by image content; and the combined utilisation of SA and LA data. The merits and limitations of each approach are discussed in relation to the basic linear method. The methods

explored in this chapter are to be presented in a journal paper currently under preparation.

Chapter 7 provides a final discussion of the work presented in the thesis, in addition to an overview of potential directions for future work based on the results shown. As well as the possibility for further interesting post-processing research, we discuss the feedback that can be provided to clinicians to improve the acquisition of data such that post-processing studies can be optimised.

2

Background - Cardiac Anatomy and Imaging

Contents

2.1 Cardiac Physiology 7
2.2 Cardiac Pathophysiology 13
2.3 Cardiac Imaging 16
2.3.1 Magnetic resonance imaging (MRI) 18
2.3.2 Cardiac magnetic resonance imaging (CMRI) 25

In this chapter, a description is given of cardiac physiology and the role played by the heart as part of the cardiovascular system. We examine the effect of pathophysiological states on the heart's ability to function, and consider how it remodels itself in response. Finally, we introduce cardiac imaging as a tool for the diagnosis and study of disease.

2.1 Cardiac Physiology

[106] and [107] were used as primary references for the following introduction to cardiac anatomy and pathophysiology.

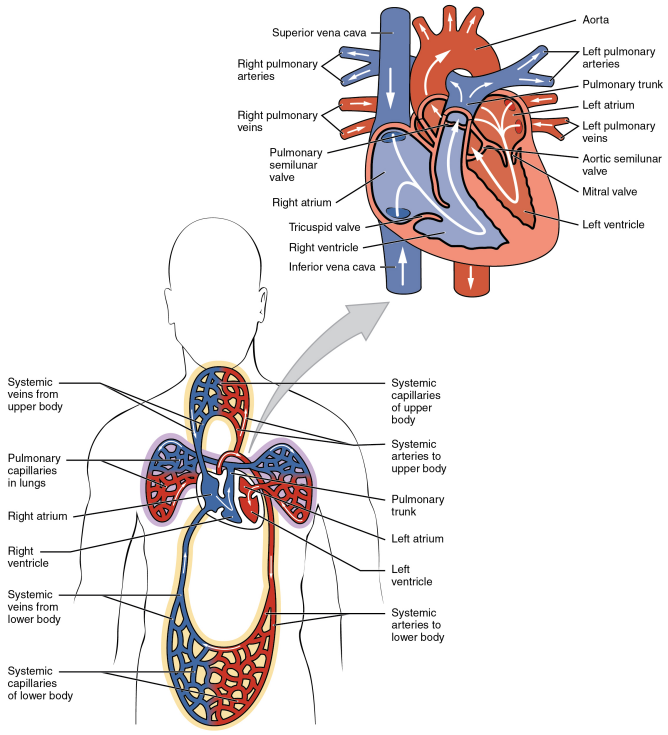


Figure 2.1: An illustration of the cardiovascular system is given, composed of both the systemic and pulmonary parts. In addition, the internal structure of the heart is depicted in the upper right-hand portion of the diagram. Figure reproduced from [19]. Reprinted under the terms of the Creative Commons Attribution License v4.0 (https://creativecommons.org/licenses/by/4.0/).

The cardiovascular system

The principal responsibility of the cardiovascular system is to facilitate the efficient transport of oxygen and other nutrients to the body, while simultaneously removing metabolic waste products such as carbon dioxide. Crucially, the cardiovascular system plays a major role in homeostasis, the preservation of the body's internal environment.

Enabling this functionality is a complex network of vessels with the heart acting as a pump at its core. The heart is often regarded as a two pump system; working in series, these pumps drive the pulmonary and systemic cardiovascular systems which are responsible for distributing blood through the lungs and through the rest of the body respectively, as shown in figure 2.1.

Each vessel's structure is dependent on its position within the circulatory network, and are resultantly named as they appear in series within the system as: arteries,

arterioles, capillaries, venules and finally veins. Except for the capillaries, the vessels have the same core component layers. The innermost surface, referred to as the endothelium, is composed of a thin layer of cells that form the barrier between the blood and the vessel walls. These endothelial cells allow the selective diffusion of materials between the tissue and blood, while also playing an essential role in reducing the occurrence of blood clotting in the vessels. Blood clots that form in the vessels are referred to as thrombi and can lead to numerous cardiovascular complications. Smooth muscle cells, elastin, and collagen are the key constituents of the central layer, providing much of the structural support for the vessel and affording contractile capabilities. The primary role of the final outermost layer is to tether vessels to their surrounding tissues.

The heart

The human heart is a muscular four chamber organ that is typically 12cm long, 8cm wide, and 6cm thick, weighing between 250g and 350g [19]. It is located medially within the thoracic cavity, lying between the left and right lungs. The heart's principal axis lies approximately along a line connecting the right shoulder to the left hip, with a slight angular rotation of the apical end towards the body's anterior surface. The heart's location within the chest is maintained by its containment within a double-walled sac, called the pericardium, which is attached to nearby structures such as the sternum and diaphragm. While anchoring its global position, the pericardium plays a secondary role providing a low friction environment for the cardiac surfaces during contraction.

The heart wall itself is formed of three layers: the epicardium, myocardium, and endocardium. The epicardium lies just beneath the pericardium and provides a further thin layer of protection for the heart. The myocardium is the thickest of the layers; it is principally responsible for the contractile action of the heart and is therefore primarily composed of cardiac muscle cells, referred to as cardiomyocytes. The endocardium is the innermost layer of the heart wall; as in the cardiac vessels, this innermost layer is composed of endothelial cells that provide a smooth, non-adherent surface between the blood and tissue.

The internal structure of the heart is indicated in the upper right-hand side of figure 2.1. The heart is divided into four hollow chambers; two upper chambers called atria and two lower chambers called ventricles. Functionally, the heart can be partitioned into two sides, each composed of a single atrium and ventricle, that pump blood through the pulmonary and systemic systems respectively. The septal wall separates the two sides preventing the mixing of oxygenated and de-oxygenated blood.

The upper and lower chambers work sequentially during the contraction period. The heart receives blood into the atria via the pulmonary vein and the vena cava. As the atria begin to contract, blood is driven into the ventricles through either the tricuspid valve or the mitral valve; collectively these are referred to as the atrioventricular valves. The ventricles subsequently pump the blood out into the two arterial systems through the pulmonary and aortic, semilunar valves.

Atrial walls are thinner than ventricle walls as they are not required to provide as much contractile force. The left and right ventricle walls also differ; the left ventricle wall is approximately three times thicker than that of the right due to the higher pressure required to pump blood throughout the entire body.

The coronary circulation

As the most critical muscle in the human body, it is essential that the myocardium has a reliable supply of blood. To ensure this provision, the heart has its own dedicated system of vessels as shown in figure 2.2 - this system is referred to as the coronary circulation system. Two main coronary arteries branch away from the systemic system immediately after the aortic valve; they supply the left and right ventricles respectively. The coronary system is responsible for ensuring that the cardiac muscle cells always have a sufficient supply of oxygen; during increased excursion, the coronary arteries are able to match the rate of blood flow in response to increased oxygen demand via the regulation of the arterial lumen diameter.

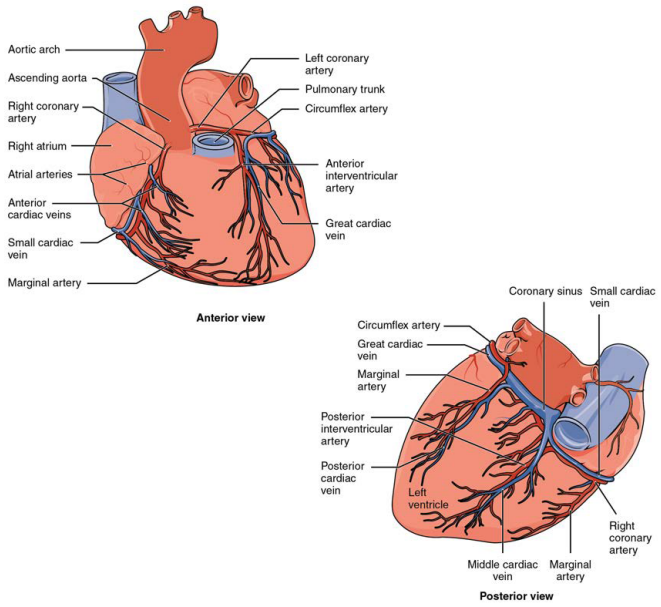


Figure 2.2: The heart's coronary circulation. Figure reproduced from [19]. Reprinted under the terms of the Creative Commons Attribution License v4.0 (https://creativecommons.org/licenses/by/4.0/).

The cardiac cycle

In the following subsection, the electro-mechanical operation of the heart during a typical cardiac cycle is summarised. Figure 2.3 depicts the key pressure changes that occur as well as the timings of valvular transitions.

The cardiac cycle refers to the sequence of events that occur during a single contraction and relaxation of the heart. As implied in the previous sentence, the cycle is conventionally divided into two phases referred to as diastole and systole. Under normal physical conditions the diastolic phase, which corresponds to the period of ventricular relaxation and filling, occupies up to two-thirds of the duration of the entire cardiac cycle; systole, which comprises ventricular contraction, occupies the final third.

At the early diastolic stage of the cycle, while the atrioventricular valves are open, blood is passively moving through the atria into the ventricles. Myocardial contraction occurs in response to an electrical signal that stimulates the cardiac cells. This electrical impulse is provided by the sinoatrial node, a collection of specialised cells located just below the epicardial layer at the top of the right atrium, that acts as the heart's natural pacemaker. As this spontaneously generated

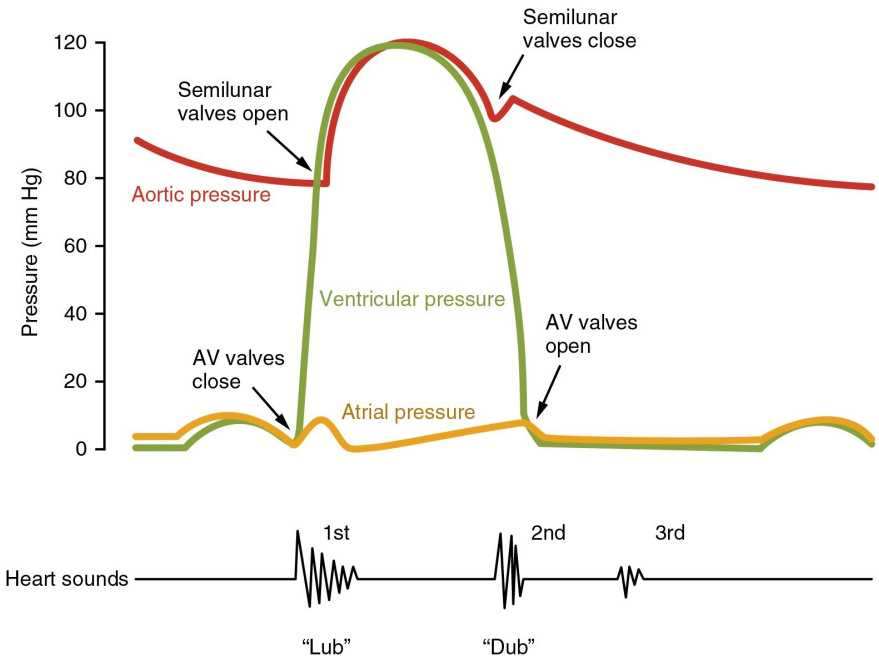


Figure 2.3: The pressure variation in the chambers of the heart during the cardiac cycle and the resultant operations of the valves. Figure reproduced from [19]. Reprinted under the terms of the Creative Commons Attribution License v4.0 (https://creativecommons.org/licenses/by/4.0/).

signal propagates through the atria towards the atrioventricular node it causes the serial contraction of the cardiomyocytes in the atrial wall. This forces a final volume of blood into the ventricles.

Upon reaching the atrioventricular node, the electrical impulse is briefly delayed. This delay is significant as it allows for the total contraction of the atria before ventricular response. From the atrioventricular node, the signal is sent down the bundle of His, a fast-conducting muscle fibre that lies in the septal wall. This bundle branches into the Purkinje fibers that wrap around the apex of the ventricles, ensuring that contraction occurs upwards from the apex to the base of the heart. At the outset of the systolic phase, when contraction begins, the semilunar valves are closed and the ventricular pressure starts to rise. As the pressure rises above that in the atria, the atrioventricular valves are forced shut, and the ventricles become isolated. This initiates a period of isovolumetric contraction that is accompanied by a rapid increase in ventricular pressure. The ventricular pressure eventually exceeds arterial pressure causing the semilunar valves to open. An initial rapid ejection of

blood occurs, during which time the pressures in both the ventricles and proximal arteries continue to rise. Eventually, after ventricular contraction has ended, the intraventricular pressure falls below the arterial pressure and the semilunar valves close. The ventricles are once again isolated, and muscular relaxation results in a further, rapid drop in pressure. The pressure drops below that of the atria and the atrioventricular valves re-open. Ventricular filling restarts, signalling the beginning of a new cardiac cycle.

Indices of cardiac function

The efficacy of the described cardiac function can be characterised by a series of indices, for example the stroke volume (SV), the ejection fraction (EF), and the cardiac output (CO). These quantities can be estimated based on measurements derived from cardiac imaging modalities.

The SV is the difference calculated between the end-diastolic volume (EDV) and the end-systolic volume (ESV), which correspond to the maximal and minimal ventricular volumes, $SV = EDV - ESV$. The SV therefore quantifies the total amount of blood expelled from a ventricle during a single contractile action.

A second, related index is the EF defined by the equation, $EF = \frac{EDV-ESV}{EDV} = \frac{SV}{EDV}$. This measures the fraction of the originally contained blood volume expelled. The ejection fraction is used as a primary indicator of systolic heart failure.

As the primary function of the heart is to adequately supply blood to the rest of the body according to demand, the third index, the CO, is a measure of the total volume of blood expelled as a function of time. It is calculated according to the equation, $CO = SV \cdot HR$, where HR is the heart rate, the number of heartbeats per unit time.

2.2 Cardiac Pathophysiology

The brief introduction to the basics of cardiac anatomy and function has already conveyed the complex interdependent systems that are relied upon for efficient cardiac function. In this section, we will look at some of the common pathologies of

the cardiovascular system that result in myocardial infarction, and we will consider how the heart attempts to adapt in response.

Atherosclerosis

While atherosclerosis is not a disease of the heart itself, rather being an affliction of the blood vessels, it can have important implications for the condition of the heart. Atherosclerosis involves the narrowing of arteries due to the abnormal accumulation of fatty substances on the arteries' inner walls. Commonly referred to as plaques, their presence reduces blood flow to distal tissues due to the narrowing of the arterial lumen. In addition to this, plaque sites are areas that have a heightened risk of rupture; rupture induces a clotting response, referred to as thrombosis, which can cause further narrowing at the plaque location or can become dislodged, moving through the arterial system and causing a blockage at a more distal part of the network.

Ischemic heart disease (Coronary heart disease)

Ischemia refers to the insufficient supply of oxygen to tissue. In other words, it is an imbalance between the provision of oxygen by the blood and the demand for oxygen by the perfused tissue. The most common cause of ischemia is reduced blood flow due to narrowing of the arteries (atherosclerosis). In ischemic heart disease, it is the coronary arteries, the vessels supplying the myocardium with blood, that are affected by atherosclerotic narrowing.

During an ischemic myocardial incident, the deficiency of oxygen supply to tissue - hypoxia - forces myocytes to convert from aerobic to anaerobic respiration. This affects myocardial function, results in labored breathing, and is the cause of mild or severe pain. However, the long-term effect on myocardial tissue depends on the severity of the ischemic incident and its duration.

The least severe presentation of ischemia is angina pectoris, which occurs in response to a short period of reduced oxygen supply. Manifested as a sharp tightening of the chest, the return of sufficient oxygen supply alleviates the condition and no long term damage to the myocardium occurs.

A second potential outcome is referred to as myocardial stunning and presents as a prolonged period of myocardial dysfunction: inhibition of the myocardium's ability to contract or relax lasts beyond the point when normal blood flow has returned. Despite the extended period of reduced functionality, there is no associated long-term damage as the oxygen supply has been restored.

The most severe manifestation is a result of prolonged oxygen starvation. In this case, myocardial tissue necrosis, called myocardial infarction, is expected. The duration of hypoxia is sufficient for irreversible damage and scar formation, regardless of the potential restoration of oxygenated blood flow to the affected area.

Acute coronary syndromes

Acute coronary syndromes are life-threatening, and can occur in coronary heart disease (CHD) patients at any time. While the ischemic events referred to above are a result of a gradual build-up of plaque in the arteries, it is also possible for the occlusion of an artery to occur suddenly. Such incidents can be more severe, often causing complete occlusion of the arterial vessel. These sudden events are commonly the result of a rupture of a forming atherosclerotic plaque. Again, the precise outcome of this imbalance is dependent on the degree of occlusion and, therefore, level of ischemia.

The least severe manifestation is again a form of angina. In this case, it is referred to as unstable angina, so called as it does not occur predictably as a result of increased physical excursion. This form of angina comes on quickly and unexpectedly and is often more painful than its stable counterpart. Unlike angina pectoris - stable angina - rest is unlikely to alleviate the pain experienced during an episode of unstable angina.

In cases where the blockages in the coronary arteries are not cleared quickly, by intervention or otherwise, cell death is inevitable. The severity of the infarction may be recognised via the profile of the electrocardiogram (ECG), which is a record of the electrical activity of the heart; the most severe cases, where the complete occlusion of the supply artery occurs, are referred to as ST segment elevation MI (STEMI) due to the elevated state of the ST portion of the ECG.

Cardiac remodelling

Cardiac remodelling refers to gradual changes in the shape and size of the heart over time. In the case of pathological remodelling, this occurs in response to an injury to the heart muscle. Remodelling occurs initially as a means of maintaining necessary cardiac function; however, the long term effects tend to be detrimental to the condition of the heart.

The remodelling process begins immediately in the first few days post-infarct. Initial remodelling occurs predominantly in the infarct area. Myocytes undergo necrosis, and there is an associated reduction in collagen fiber quality - there is an expansion of the infarcted region along with a corresponding thinning of the wall in this region. As inflammatory cells infiltrate the necrotic tissue and collagen is gathered in the infarct region, a scar is formed. The appearance of scar tissue now resists further expansion of the region.

The remodelling process continues in the remaining healthy tissue, over the course of many weeks. The non-infarcted myocardial tissue dilates initially to compensate for the loss of function in the infarcted tissue; however, in the long term the detrimental effects dominate, resulting in compromised function that may be recognised by a reduced ejection fraction.

2.3 Cardiac Imaging

The ability to provide accurate information regarding the state of cardiac tissue means that imaging has become an integral tool for cardiovascular clinicians. The combined information gathered from modalities such as echocardiography, computed tomography (CT) and magnetic resonance imaging (MRI) aid in the diagnosis of pathology and the long-term management of patients. In this thesis we work exclusively with MRI data provided by clinical collaborators, therefore here we give only brief descriptions of both echocardiography and CT, before presenting a more detailed introduction to fundamental MRI physics and its application to clinical cardiology and research.

Echocardiography

Echocardiography is the most widely used imaging technique for the assessment of cardiac structure and function [179], largely as a result of its non-invasiveness, wide availability, comparatively low cost, and ease-of-portability [122]. During an echocardiographic exam, a transducer is placed on the chest through which sound waves are delivered. An image is formed according to the interactions of these ultrasound waves with the tissue and allows for describing cardiac structure. Echocardiography has a variety of applications including the characterisation of morphology, the assessment of valvular performance, and the description of cardiac haemodynamics [179].

Classically, echocardiographic data are 2D images captured across the cardiac cycle at specifically chosen image planes, however, more recent technological developments are enabling the acquisition of full 3D temporal datasets [90].

The primary drawbacks of echocardiography include operator variability, limited acoustic windows, and the poor tissue contrast afforded [72]. As a result, initial diagnosis via echocardiography may be complemented by assessment using other modalities such as CT and MR.

Cardiac CT

Cardiac CT utilises fast x-ray imaging and specialised reconstruction techniques [64] to visualise cardiac structure. Multiple x-ray projections are acquired as an x-ray tube and detector are rotated around the patient and these are subsequently combined to produce the desired cross-sectional images.

In addition to the anatomy of the heart, a primary use of cardiac CT is the visualisation of cardiac vessels [217]. Cardiac vessel imaging, CT angiography, is aided by the use of radio-contrast agents that improve the contrast of the vessel structures provided in the reconstructed images.

CT has superior image resolution compared to the other modalities described here, however, this comes at the undesirable cost of x-ray radiation exposure.

2.3.1 Magnetic resonance imaging (MRI)

MRI uses the nuclear magnetic resonance property of atoms in the body to visualise its internal structures. By combining the effects of static and dynamic magnetic fields, the natural nuclear spin states of the body's atoms are perturbed such that their relative changes can be measured and analysed to form a visual representation of tissue.

All nuclei in the body have a characteristic quantum property called spin. Spin, which is based on the number of protons and neutrons in the nucleus, is an intrinsic form of angular momentum that induces a magnetic moment in all non-zero-spin nuclei. To adhere to the theory of classical physics, these nuclei are often represented by spheres that rotate on their own axis, generating magnetic dipole moments that point from south to north magnetic poles, as indicated in figure 2.4a. While this model does not correlate with quantum physical explanations of spin, it is a commonly utilised illustration that we will employ during our own discussion of MR imaging.

In their natural state, nuclear dipoles point in all orientations (figure 2.4b) such that the net magnetisation vector, the sum effect of all dipoles, is zero. During an MRI acquisition, a patient lies inside a large, powerful magnet, typically 1.5 or 3 Tesla, and the strong magnetic field causes the nuclei of atoms in the body to align along the field direction, as shown in figure 2.4c.

A second consequence of turning on the external magnetic field is that the nuclei begin to precess, which is the rotation of a nuclear axis around the axis of the external magnetic field. The frequency of precession is given by the Larmor equation, $\omega = \gamma B_0$, and is an important property for MR image generation. The net magnetisation vector is no longer zero but oriented in the direction of the external magnetic field. While the individual nuclei precess, the net magnetisation vector itself does not, as all the individual precessions are out of phase and their effects sum to give a zero transverse component.

The diagram originally presented here cannot be made freely available via ORA because of copyright.

The diagram was sourced at:
http://199.116.233.101/index.php/File:Proton_Figure.png#filelinks
D. Adair, M. Walton, A. Kroker, K. Rivard, S. Sudhindra, and S. Pandey. Physics of MRI. MediaWiki

Figure 2.4: (a) indicates the dipole nature of nuclear magnetisation. In the resting state depicted in (b), the nuclei are randomly oriented, while in the presence of a stationary magnetic field illustrated in (c), the nuclei align with the direction of the applied field. The resultant net magnetisation vector, M_0 is aligned in the direction of the external magnetic field, B_0. Figure reproduced from [2].

Generating the MR signal

In the steady state, with the net magnetisation vector aligned with the external magnetic field, there is no measurable signal. A detectable signal results from the application of a radio-frequency (RF) pulse at a right angle to the stationary magnetic field. If the frequency of the RF pulse matches that of the nuclei's Larmor frequencies, then the RF pulse adds energy to the system - a resonance between the applied signal and the natural rotation is required. If resonance occurs, and energy is added to the system, the alignment of the net magnetisation vector is moved away from the direction of the strong magnetic field towards the transverse, XY, plane.

The response of nuclei to an RF pulse is depicted in figure 2.5. Two processes occur; firstly, the addition of energy to the system causes some nuclei to jump to a higher level state which is oppositely oriented to their resting state - this reduces the magnitude of the Z-axis magnetisation. Simultaneously, the application

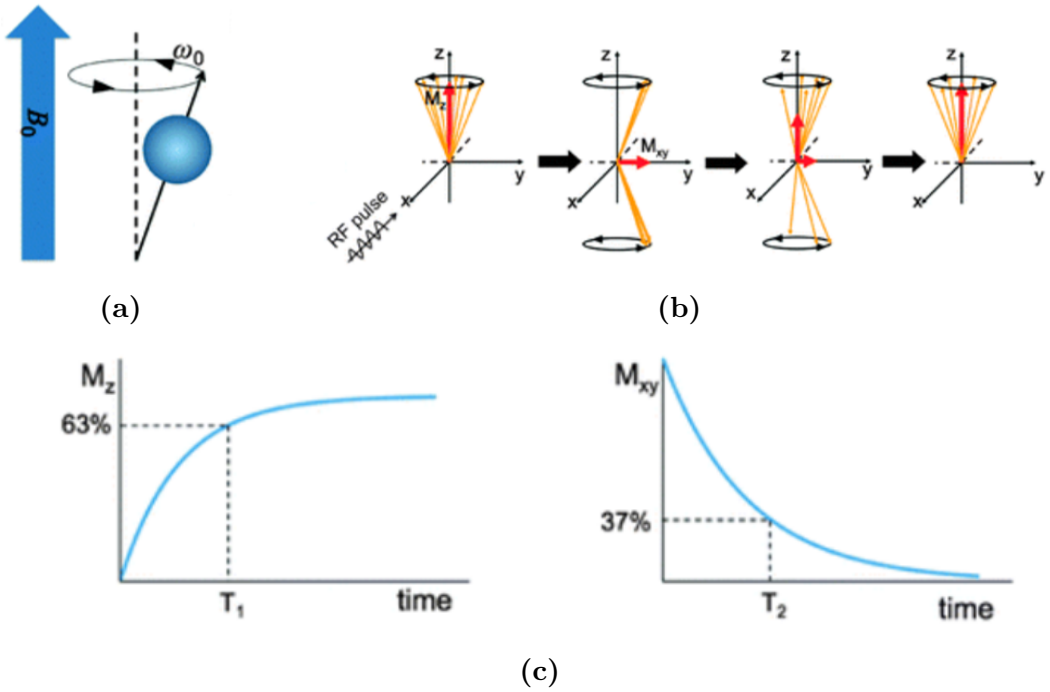


Figure 2.5: (a) indicates the alignment of a nuclear spin with the stationary magnetic field applied. (b) illustrates how the motion of the nuclei contribute towards the net magnetisation vector after the application of an RF pulse. The plots in (c) indicate the respective changes in the components of the net magnetisation vector after the RF pulse application. Adapted from [103] with permission of The Royal Society of Chemistry.

of the RF pulse in the transverse plane brings the precession of individual nuclei around the stationary field into phase, thus creating a transverse magnetisation in the XY plane. For the duration of the RF pulse, the Z-axis magnetisation decreases while the XY-plane magnetisation increases. The effect of these changes on the net magnetic vector is that it rotates towards the XY-plane in a process referred to as nutation. The total extent to which the net magnetisation vector is perturbed is often described by the flip-angle, which measures the angle between the net vector and the static magnetic field; it is a function of both the RF pulse magnitude and the duration of application.

Immediately following the RF pulse, the nuclei begin returning the gained energy to the environment. This is governed by two processes referred to as spin-lattice and spin-spin relaxation.

The spin-lattice relaxation process determines the recovery of the net magnetisation vector's Z-axis-alignment. It is alternatively referred to as either the

longitudinal- or T1-relaxation, and can be mathematically described by $M_z(t) = M_0(1 - e^{-\frac{t}{T1}})$, where $T1$ is the time constant - this function is plotted on the left-hand side of figure 2.5c. The rate at which energy can be transferred back to the surrounding environment is dependent on the difference between the Larmor frequency and the natural motional frequencies of the environment's composition - as a result, the governing time constant, $T1$, is a property of the underlying tissue. For example, lipid molecules have a rate of molecular motion that is conducive to energy transfer, meaning fat has a particularly short T1 relaxation time. Conversely, the higher molecular motion of water molecules result in a long T1 relaxtion time.

The second relaxation process, spin-spin relaxation, is associated with a loss of precessional phase coherence between nuclei and is characterised by the decay of magnetisation in the XY-plane. Alternatively referred to as T2 relaxation, the process can be mathematically described by the equation $M_{xy}(t) = M_0 e^{-\frac{t}{T2}}$. This relationship is plotted on the right-hand side of figure 2.5c, where $T2$ indicates the decay time constant. The rate of decay is again dependent on the chemical make-up of the tissue - tissues with higher water content have longer T2 values due to greater distances between molecules that reduce the occurrences of spin-spin interaction. However, the decay is also influenced by local non-uniformities in the magnetic field. A second decay constant is often used therefore, defined as $\frac{1}{T2^*} = \frac{1}{T2} + \frac{1}{2\gamma\Delta B}$. The observed transverse decay is consequently faster than that associated with the tissue spin-spin relaxation alone.

The generated MR signal is a result of the transverse component of the net magnetisation vector. The precession of the net magnetisation vector places the RF coil (that which previously applied the pulse) in the presence of a changing magnetic field, which consequently induces a current in the coil. This induced current is the source signal for MR image generation and is known as the free induction decay (FID) signal. The oscillating, decaying nature of the transverse magnetisation is illustrated in the FID signal indicated in figure 2.6, where the difference between theoretical $T2$ and observed $T2^*$ decay rates is also illustrated.

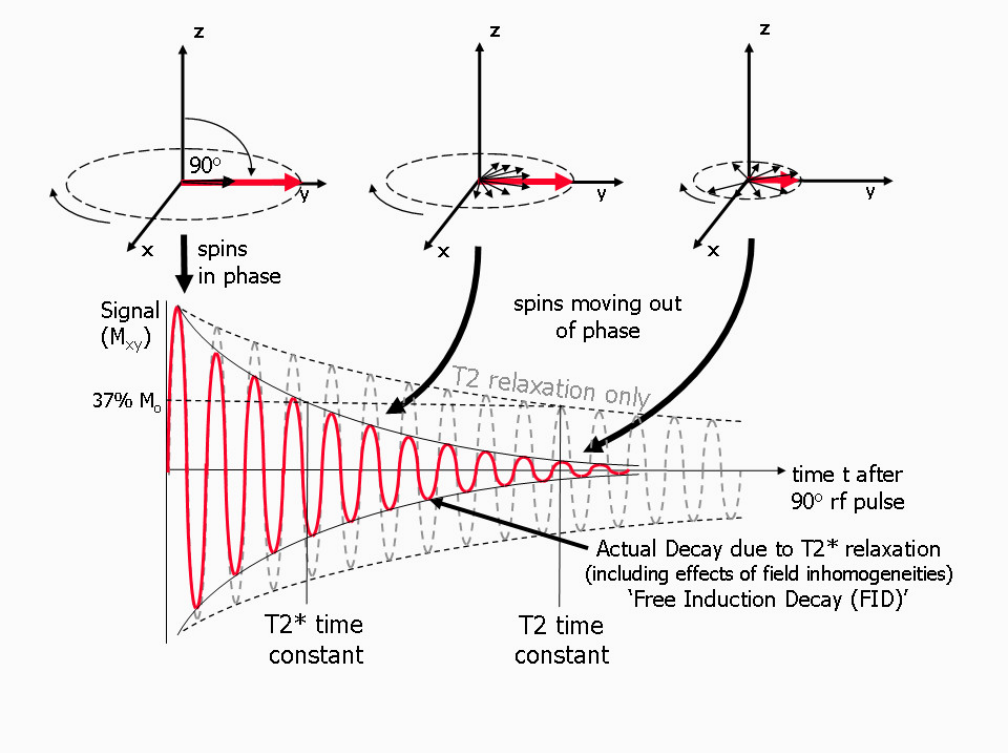


Figure 2.6: The free induction decay (FID) signal is generated by the transverse magnetisation of the net vector. The figure illustrates the decay of this component as the nuclei de-phase. Figure reproduced from [172]. Reprinted under the terms of the Creative Commons Attribution License v2.0 (http://creativecommons.org/licenses/by/2.0) via BioMed Central Ltd.

Spatial localisation

The above theory describes how nuclei in a sample can be excited using an RF signal and the subsequent FID signal observed to help characterise the local composition of tissue. However, the RF coil receives a composite signal, formed of FID signals from the entire sample with no information regarding their spatial origins. In order to generate a meaningful MR image, a means of localising the component signals is required. Such spatial localisation is achieved by utilisation of magnetic field gradients, which are applied in addition to the original magnetic field, thus generating a spatially varying magnetic field strength. By manipulating the local magnetic field strength it is possible to alter Larmor frequencies in such a way that the origin of FID signals become differentiable.

Three gradients are used to define the 3D position. The first gradient is referred to as the slice selection gradient and is used to isolate the plane of interest in the

imaged object. A linear gradient is applied in the perpendicular direction to the desired plane. This causes the Larmor frequency to vary in the same linear fashion along this gradient direction. Under these conditions, when the RF excitation pulse is applied, only nuclei whose Larmor frequency matches that of the RF signal resonate. Different planes of the volume can be excited therefore by changing the central frequency of the radio frequency pulse.

Having defined a plane, the second and third gradients are used to localise the signals emitted from that plane. The gradients are applied in orthogonal axis and are referred to as frequency- and phase-encoding gradients.

The operation of the frequency-encoding gradient is similar to slice selection. A linear gradient applied along one of the plane axes generates a spatially variable Larmor frequency. This gradient is applied during the read-out period, and is therefore alternatively referred to as the read-out gradient. Where previously a single frequency signal would have been detected, the received signal is now composed of multiple frequency components; the relative contributing amplitudes of these frequency components represent the signal amplitudes according to their position along this axis. The partitioning of the received signal into constituent frequency components can be performed using the Fourier transform.

This still leaves an ambiguity in the third dimension, as each frequency component signal is formed by the combined effect of all protons that lie at a certain gradient strength. A third gradient is required in order to differentiate the contributions along this axis, and is referred to as the phase-encoding gradient. This gradient is applied after the slice-selection gradient but prior to the read-out (or frequency-encoding) gradient. Before its application, all nuclei in the plane precess at the same frequency and are in phase. The phase-encoding gradient is turned on for a short time period during which the Larmor frequency now varies linearly along this axis - the nuclei rotate at different frequencies according to their position. This variation in rotational frequency generates precessional phase differences along this axis. Upon removal of the phase-encoding gradient, the Larmor frequencies of all nuclei return to the original value, but the phase offsets are retained.

In this state, the application of the read-out gradient still generates a signal that is composed of frequency components defined according to the frequency-encoding gradient, however, as a result of the phase-encoding gradient the signal also has embedded in it spatial information along the final axis by means of phase-offsetting. In order to use this embedded information, the repeated application of the three encoding gradients is necessary, with each signal acquisition using a different phase-encoding gradient magnitude. It is the comparison between the final incrementally changed phase-encoded signals that enables the identification of spatial localisation according to relative precessional phase offsets. The final in-plane localisation, using both the frequency- and phase-encoding schemes, can be assessed using the 2D Fourier transform.

The MR signal echo

The application of spatial encoding gradients disrupts the generated FID by causing de-phasing of the nuclei. By carefully-timed application of further gradients, the spins may be re-phased and an MR signal can be detected that is an echo of the originally desired FID. The important timing parameter used here is commonly referred to as the echo time, TE.

Generating image contrast

Another important timing parameter is the repetition time, TR, which determines the time between successive pulse applications. Clearly this affects the total imaging time, but more significantly, in combination with TE, it provides the ability to manipulate the observed soft tissue image contrast. Specifically, the timings involved in image acquisition sequences are designed such that the final contrasts provided are dependent on the amount of the different relaxation processes allowed to occur.

T1-weighted imaging: image contrast is primarily dependent on the T1 relaxation process, with tissues characterised by short T1 values having a higher signal intensity than those with long T1.

T2-weighted imaging: image contrast is primarily dependent on the T2 relaxation process, with tissues characterised by long T2 values having a higher signal intensity than those with short T2.

Proton density imaging: neither T1 nor T2 relaxation plays a significant role in the contrast observed in these generated images, rather the signal intensity simply reflects the amount of signal originating from each location

2.3.2 Cardiac magnetic resonance imaging (CMRI)

The importance of MRI in cardiovascular medicine is ever increasing due to the accuracy and reliability with which the assessment of cardiac anatomy and function can be performed [108]. MRI has a number of advantages that make it preferable to other imaging modalities for the assessment of ischaemic heart disease - it has excellent spatial resolution, yields good soft tissue contrast, can provide functional information, and is safer than many other modalities as it does not require any ionising radiation. However, there are some noteworthy drawbacks compared to alternative modalities such as echocardiography and CT: the potentially extended examination time required; the high costs of the equipment; and the suitability of the patient for examination, due to issues such as claustrophobia, acute stability, and MR contraindications.

A further strength of MR imaging is the slice selection process, which allows the acquisition of image planes in any orientation. This is particularly important in cardiac MR, where the orientation of the heart does not correspond to the standard sagittal, coronal, and axial body planes, and is also highly variable between patients. The flexibility of MR slice selection permits the acquisition of specific planes that relate most appropriately to cardiac tissue. Standardised imaging planes were set out in [21] to facilitate consistent cardiac analysis; those being the short-axis (SA) view, lying perpendicular to the long-axis of the left ventricle, and the vertical and horizontal long-axis (LA) which slice through the SA views, as illustrated in figure 2.7.

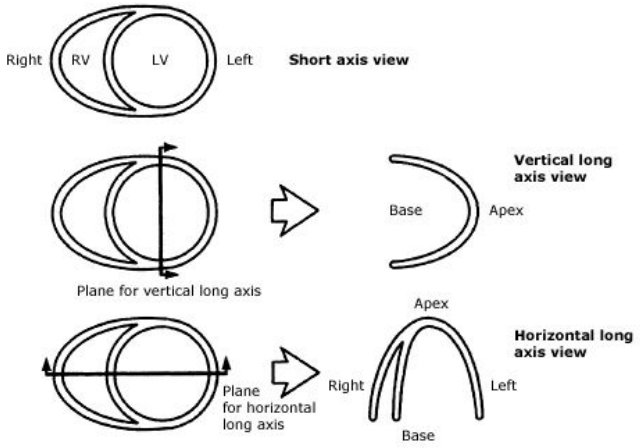


Figure 2.7: Illustration of the standard cardiac acquisition planes, referred to as the short axis (SA), the vertical long axis (VLA) and the horizontal long axis (HLA). Reproduced from [21] (Bonow RO, Gibbons RJ, Berman DS, Johnson LL, Rumberger JA, Schwaiger M, Taubert KA, Wackers FJ, Ritchie JL, Brundage BH, Dreifus LS. Standardization of cardiac tomographic imaging. Circulation. 1992;86(1):338-9. https://www.ahajournals.org/journal/circ) with permission of Wolters Kluwer Health, Inc.

A number of MR acquisition sequences exist and the appearance of the resulting images depends heavily on the technique used [158]. In this next section, a summary is given of the main techniques used in cardiac applications, with particular attention given to their use in the assessment of ischaemic heart disease.

Cine MRI

Cardiac cine imaging is the cornerstone of functional assessment of the heart [58]. Images from each stage of the cardiac cycle are acquired over multiple heart beats and combined to capture the motion of the heart. A typical cine acquisition consists of an SA stack and a number of LA images, examples of which are shown in figure 2.8. The MRI protocol most commonly used in current clinical practice is the steady-state free precession (SSFP) gradient echo sequence, due to its speed of acquisition, high contrast to noise ratio and its excellent blood-myocardial contrast [58].

The ability to see the motion of the heart over a cardiac cycle is very useful, providing a means of identifying wall motion abnormalities, which can be indicative of damaged tissue. The excellent blood-myocardial contrast provided allows for

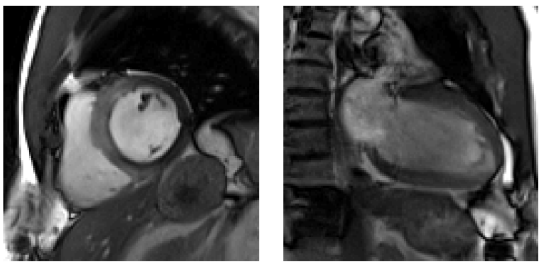


Figure 2.8: An example SA and LA cine MR image acquisition.

accurate and potentially automated delineation of the cardiac chambers. This makes measurement and assessment of important cardiological indices easier.

Tagged MRI

The images depicted in figure 2.9 illustrate the appearance of tagged cardiac MR data. The observed grid pattern is a result of alternating saturated and equilibrium magnetisation spin states at these locations, resulting from the application of a specific combination of radio frequency pulses and magnetic gradients. As the applied magnetisation is now a property of the underlying tissue, the grid pattern deforms in accordance with the contraction of the muscle. This deformation can be captured using a cine MR acquisition protocol.

The benefit of the additional tagging pattern is that it allows the analysis of motion within the myocardium, a difficult task using conventional cine acquisitions due to the homogeneity of signal intensity in the cardiac wall. However, there are challenges associated with this analysis. Firstly, changes in the magnetisation states with time mean that the grid pattern gradually fades, and secondly, the acquisition of 2D planar data prohibits the direct assessment of out-of-plane motion. Specific reviews of cardiac MR tagging procedures and subsequent post-processing analyses can be found in [8, 83].

Perfusion MRI

Seminal papers proposing the use of contrast agents to visualise myocardial perfusion were published by Brown et al in [24] and by Miller et al in [138]. The most commonly used contrast agents are gadolinium-based, where the interaction of gadolinium with the tissue nuclei reduce both the T1 and T2 relaxation times. In perfusion

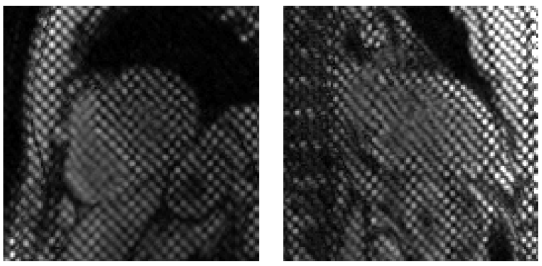


Figure 2.9: An example SA and LA tagged MR image acquisition.

MRI, T1-weighted acquisitions are used meaning that gadolinium affected regions see an increased signal intensity due to the reduced T1 relaxation time.

For the assessment of perfusion, a bolus of contrast agent is intravenously injected and its passage through the heart analysed using a series of images. Signal enhancement is sequentially observed in the right ventricle, left ventricle, and finally the myocardium. Myocardial regions with reduced perfusion are recognisable by their limited relative signal enhancement.

The procedure is often performed while under vasodilator stress. The use of a stressor emphasises the difference between healthy and diseased tissue. The arterial beds fed by a diseased artery will already be maximally dilated and, therefore, will see no increase in blood supply when stressed [142] - the resulting intensities will be slow to enhance and appear dimmer than healthy tissue. In order to enable the analysis, an initial nulled myocardial image is acquired with subsequent images, taken with every heart beat, showing myocardial signal enhancement due to contrast agent uptake. Areas of enhancement can be qualitatively recognised or alternatively a quantitative analysis can be performed by plotting time-intensity curves. From these curves, the myocardial tissues can be rated according to perfusion indices. Perfusion imaging is a very important tool in the diagnosis of ischaemia: it achieves 80% - 90% sensitivity as well as 70% - 80% specificity and is thus just as successful a tool as nuclear perfusion imaging [142].

Late enhancement MRI

It was first recognised that the intensity of infarcted tissue becomes enhanced with extracellular contrast agents in the late 1980s [203], and delayed enhancement

imaging is now an accepted tool for the detection of both acute and chronic cases [89].

While imaging soon after the administration of a bolus of contrast agent enables the visualisation of the perfusion of blood, imaging after a delay is a means of identifying regions of scar tissue. The initial distribution of contrast agent after bolus administration is dependent on the myocardial blood flow, but long-term distribution is dependent on extracellular fluid volume [58]. Myocardial scar tissue is associated with a larger extracellular fluid volume than non-infarcted tissue. As a result, the concentration of contrast agent in the scar tissue is much higher, leading to greater enhancement. Significantly, scarred tissue also experiences slower rates of contrast wash-out, making delayed image acquisition highly effective. As discussed in the previous section, the gadolinium present reduces the T1, longitudinal relaxation of local tissue meaning that the regions of scar tissue in the final T1-weighted images are clearly indicated by higher signal intensities, as shown in figure 2.10.

A number of factors can affect the diagnostic benefits of late enhancement imaging. Kellman et al [86] provide a detailed discussion of many of the issues faced as well as the state-of-the-art clinical acquisition protocols used to address them.

Identification of infarcted cardiac tissue is possible due to the intensity contrast provided by the nulling of the MR signal from remote myocardium versus the hyper-enhancement of damaged tissue as a result of contrast accumulation. While the optimality of this contrast depends on factors including acquisition timings and imaging parameters, of greater concern are scenarios such as sub-endocardial infarcts where there is poor contrast between the hyper-enhancement and blood pool. This lack of contrast can result in missed or inaccurate identification of scar, and while solutions have been proposed to improve this contrast these generally come at the expense of other factors such as prolonged acquisition times or reduced SNR [86].

Further difficulties are presented by the presence of fat; similar to the contrast-enhanced myocardium, fat is characterised by low T1 values resulting in high signal intensities in LGE images. As a result, it can be difficult to distinguish fat from fibrosis if the myocardium is damaged sub-epicardially or if fat is present intramyocardially. Fat causes further issues in the form of chemical shift artefacts

which are characterised by signal voids at fat-water tissue intersections. Kellman et al [86] again present state-of-the-art acquisition methodologies to tackle these issues including fat-water-separated imaging.

A key observation of Kellman's discussion is an increasing demand for quantitative characterisation of LGE data in terms of scar shape, size, transmurality, and border-zone. To meet these clinical demands, continued advancements in acquisition technologies are required to improve image SNR and reduce the influence of image artefacts on analysis.

Finally, regardless of the discussed imaging limitations, the nephrotoxic effect of gadolinium-based contrast agents is a known risk [94] and can, in the case of patients with reduced renal function in particular, act as a contraindication for LGE data acquisition.

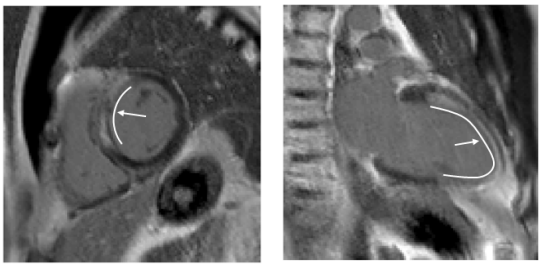


Figure 2.10: An example SA and LA LGE image acquisition where the arrows and white lines indicate the areas of infarction.

Oedema imaging

In the acute stage of ischemia, potentially before the injury becomes irreversible, regional increases in myocardial water content can be seen - this accumulation of excess water is known as oedema. Oedema occurs due to the disruption of energy-dependent ionic transport mechanisms across cell membranes [84, 216]. It was previously noted that water has a long T2 relaxation time. Therefore, a local increase in tissue water content is correspondingly characterised by an increase in T2 relaxation time. A T2-weighted acquisition captures these changes, with regions of oedema associated with higher signal intensities. The comparison between oedema images and delayed enhancement images is useful for the identification

of salvageable myocardium, as they respectively indicate the area at risk and the regions of irreversible scar.

Phase contrast imaging

Phase contrast imaging enables the measurement of tissue velocities [142]. The phase contrast method relies on the fact that moving protons in a magnetic field gradient experience a phase shift. The magnitude of this phase shift is proportional to their position along the gradient, so faster moving protons experience a greater shift. Subsequent application of an opposite gradient returns the phase values of all stationary protons to their original value, while moving protons are left with a phase value proportional to their position and, therefore, velocity. Phase contrast imaging is most commonly used to quantify blood flow, and such measures can be indicators of overall cardiac performance [188].

MR coronary angiography

In addition to the anatomical and functional assessment of the myocardium, MR imaging can be used to evaluate the state of the coronary circulation. However, using MR coronary angiography as a diagnosis tool in cases of ischaemic heart disease has many associated difficulties. The small scale and large movements that are inherent properties of the arteries make their imaging very difficult. As a result, current clinical utilisation is limited to the assessment of anomalous coronary arteries, aneurysms, and coronary bypass grafts [38].

3

Clinical Data and Associated Challenges

Contents

3.1 Clinical Data 33
3.1.1 Data delineations 34
3.1.2 Cardiac mesh generation 37
3.1.3 Scar calculation from LGE image data 40
3.2 Challenges and Limitations of MRI Acquisition 43
3.2.1 Image resolution and partial-volume effect 43
3.2.2 Cardiac and respiratory motion - Breath-hold misalignments 45
3.2.3 Intensity inconsistencies 47
3.2.4 Image acquisition reproducibility - Misalignment of slices from different studies 47

While fully isotropic, volumetric cardiac images may be desirable, such 3D acquisition methods have a number of associated limitations: they require long breath-hold durations, compromises in spatial and temporal resolution, and additionally experience reduced contrast-to-noise ratio compared to 2D techniques [126]. As a result, it remains the clinical standard for cardiac datasets to be acquired in the form of 2D planar images that span the tissue volume.

A number of imaging sequences are used during a cardiac study; depending on the selected parameters, data may be acquired at a range of temporal states across the cardiac cycle, and with varying image contrast. Therefore, by careful selection of sequence parameters image acquisition may be optimised according to

the specific analysis of cardiac condition desired. However, despite providing the clinician with high-quality functional and anatomical information, these cardiac MR datasets do come with notable limitations.

In this chapter, we will describe the clinical data used in this thesis before subsequently discussing some of the common challenges faced when dealing with such cardiac MR images.

3.1 Clinical Data

The data used in this project is part of the Oxford study in Acute Myocardial Infarction (Ox-AMI). The aim of the study is to improve the understanding and treatment of acute coronary syndromes. The study is distinctive in its approach, combining state-of-the-art clinical research facilities with the standard emergency clinical response pathway.

As part of the study, ST-elevation myocardial infarction patients underwent 3T CMR on four separate time points: 12 to 48 hours (24H); 5 to 7 days (1W); 14 to 17 days (2W); and 6 months (6M) after percutaneous coronary intervention (PCI).

In the work presented in this thesis, we specifically consider the changes in anatomy observed in the cine and LGE data from the first and fourth visit only. Functional data was acquired in the form of cine cardiac MR images, using an SSFP sequence with typical protocol parameters: TE = 1.5ms, TR = 50.0ms, flip angle = 50°, acqusition matrix = 256x208, reconstructed image size = 256x208, and voxel size 1.54 x 1.54 x 8mm. The cine sequence is generated by retrospectively gating the MR data with each slice acquired during breath-hold. Scar is analysed by means of LGE images; 10 to 15 minutes after the administration of a gadolinium-based contrast agent (0.1mmol/kg (gadoterate meglumine, Dotarem, Guerbet, Villepinte, France)), a set of LGE images was acquired, triggered to the diastolic phase, using a T1-weighted segmented inversion recovery gradient echo-phase sensitive-inversion recovery (GRE-PSIR) sequence with typical protocol parameters: TE = 2.5ms, TR = 850ms, flip angle = 50°, acqusition matrix = 256x184, reconstructed image size = 256x230 and a voxel size of 1.41 x 1.41 x 8mm; the inversion time was

adjusted during acquisitions to achieve optimal nulling of remote myocardium. In both cases the gap between slice acquisitions is 2mm, meaning the average distance between acquired slice centres is 10mm.

Equivalent sets of images were acquired using both the cine and LGE sequences. A single study is generally composed of three LA slices - the horizontal-long-axis (HLA), the vertical-long-axis (VLA), and the left-ventricle-outflow-tract (LVOT) plane - in addition to contiguous SA slices, parallel to the AV groove, with full left ventricle coverage - an example dataset is illustrated in figure 3.1. The SA stacks are composed of between 8 and 13 slices. While the slice locations were set during the acquisition of the cine data, the same imaging locations were used to acquire LGE data. The LGE data are acquired at a later point during the imaging session as indicated in figure 3.2.

Data from 19 patients were processed. For each patient, cine and LGE studies from 2 visits were available, those being from the first and fourth visits. Approval was obtained from the OxAMI project board for work to be undertaken on these initial 19 patients. In the future, access to the remaining large cohort will be requested on the basis of the results presented in this thesis.

For the work presented in this thesis, the data was manually cropped thus utilising a reduced field-of-view containing only the cardiac structures. This cropping of the datasets helps by, firstly, reducing the required computation times, but more importantly reducing the effect of other anatomical structures on the results obtained. In particular, by excluding large proportions of the lungs we expect that their potentially detrimental effect on registration will be reduced. Figure 3.1 illustrates the extent of the original image data acquired in comparison to the typical cropping performed - after cropping, all images are 128 x 128 pixels, therefore having an approximate physical extent of 20x20cm^2

3.1.1 Data delineations

Post-processing analysis of the image data was performed by expert clinical collaborators using cmr[42] (Circle Cardiovascular Imaging, Calgary, Alberta, Canada), a

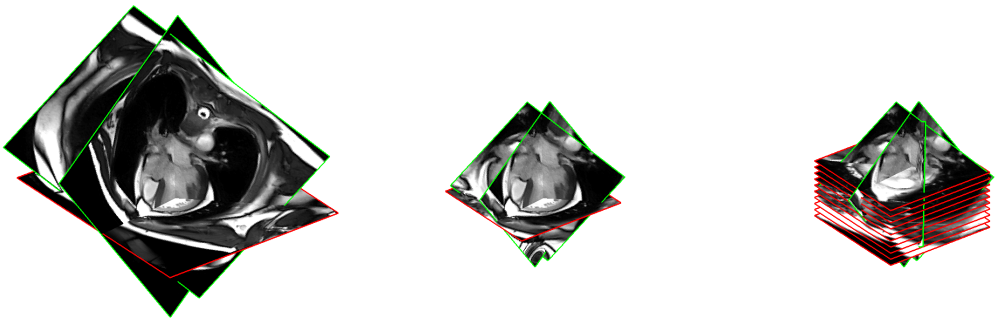


Figure 3.1: Illustration of the original data acquisition field-of-view and then the cropped data used in the remainder of this thesis. Finally, an illustration of the slices that constitute a standard cardiac cine study.

certified clinical software suite. For the cine data, this required the delineation of left ventricle endo- and epi-cardial borders in SA and LA data, in addition to SA segmentation of the right ventricle. An example set of images and corresponding delineations is indicated in figure 3.3.

In the work presented in this thesis the endo- and epi-cardial contours are used for validation of image alignment processes by calculating minimum distances between delineations. In cases where the 2D delineations are not capable of describing the quality of data alignment due to out-of-plane differences in contour location, the delineations are combined to form mesh-based representations of cardiac surfaces -

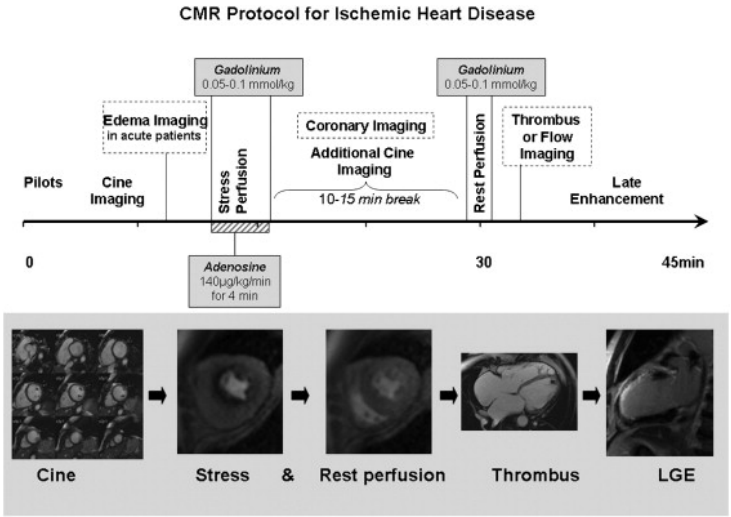


Figure 3.2: Approximate timings of images that can be acquired during a CMR study in ischemic heart disease patients. Reprinted from [84], American Heart Journal, 162/1, Karamitsos TD, Dall'Armellina E, Choudhury RP, Neubauer S. Ischemic heart disease: Comprehensive evaluation by cardiovascular magnetic resonance, 16-30, ©2011, with permission from Elsevier.

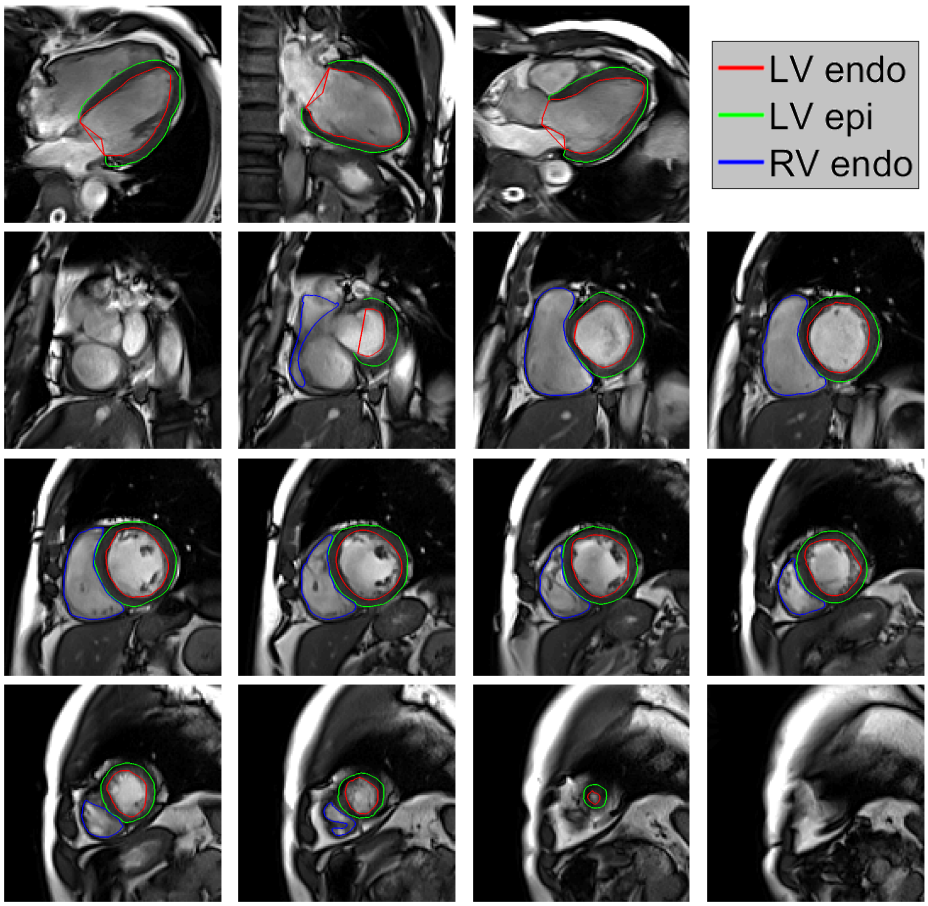


Figure 3.3: Delineations of cine cardiac data, as provided by clinical collaborators.

the mesh construction method used is described later in section 3.1.2 - in particular the endo- and epi-cardial surfaces are used for quantitative validation. While LA endo- and epi-cardial contours are used for validation of slice realignment they are not utilised in the mesh-surface reconstruction due to potential inconsistencies in the delineations of SA and LA data respectively. Constructing RV meshes using the approach presented in this chapter did not prove consistently successful due to the non-convexity of the RV shape and therefore it is not used for validation. We show one partial example of an RV mesh later during discussion.

In figure 3.4, the corresponding LGE image data are shown; the endo- and epi-cardial borders are again segmented in LA and SA data. We note that further outlines are provided in the SA LGE data; these indicate specifically selected reference areas used for the identification of regions of scar tissue.

As for the cine data, the endo- and epi-cardial delineations are used for validation

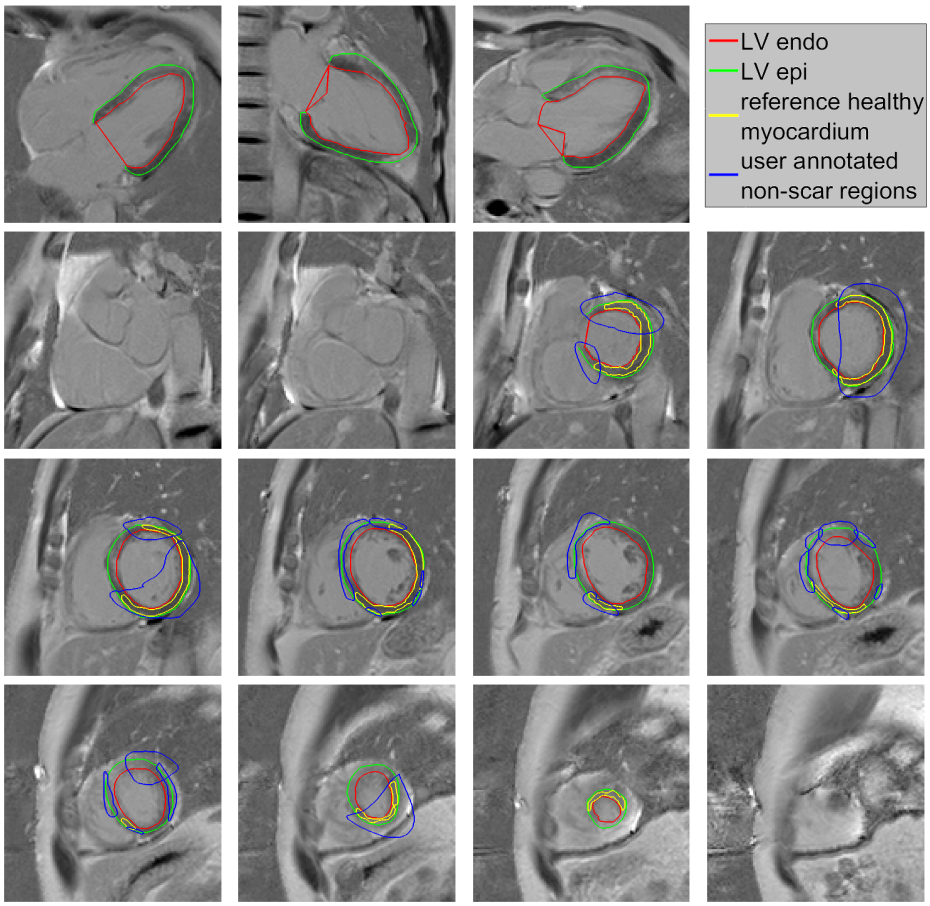


Figure 3.4: Delineations of LGE cardiac data, as provided by clinical collaborators.

purposes. The additional delineations provided with the LGE data are used to identify regions of myocardial damage; the specific method for this is described in section 3.1.3.

3.1.2 Cardiac mesh generation

The delineations provided are in the form of 2D contours located on sparse image slices. However, much of the work presented in this thesis considers the clinical data in the 3D coordinate domain. For the purpose of validating the 3D post-processing operations applied to the data, these 2D contours in 3D space are insufficient - in particular, we note that the validation of the 3D registration that is presented in chapter 5 would not be possible using the delineated contours as there is no guarantee that contours from each input dataset lie in equivalent planes before or after registration. Therefore, to appropriately utilise the contour data available, we integrate the in-plane delineations to form 3D surface meshes.

Due to our primary interest in the left ventricle, a meshing of the planar left-ventricle endo- and epi-cardial contours in the SA stack slices is performed to generate corresponding surface representations. The meshing was accomplished using the open-source software package MeshLab [207], which provides a wide variety of in-built tools for manipulating triangular meshes. The primary steps in the proposed mesh generation framework presented here are:

Point sets normal estimation: the normals to the input points must first be estimated. By considering a local neighbourhood of points, a tangent plane is approximated from which a normal vector is calculated. A number of smoothing steps are subsequently applied to the calculated normals for a more robust estimation. The MeshLab filter *'Compute normals for point sets'* is used with the following parameters: number of neighbours = 150, number of smoothing iterations = 5. To ensure the consistency of the estimated normal orientations, a reference point is also provided located at the centre of the left ventricle blood pool.

Surface mesh generation: the Poisson surface reconstruction approach [85] is used to generate the endo- and epi-cardial surface meshes. A number of algorithm parameters may be specified: tree-depth = 6, solver-divide = 6, samples-per-node = 1, surface-offsetting = 1. The tree-depth and samples-per-node parameters, in particular, influence the detail in the final reconstructed mesh, weighting the relative importance of adherence to input data and smoothness of the final solution.

Surface mesh sub-division: the generated Poisson mesh is subdivided using Loop's approximating subdivision surface scheme [111], in which each triangular element is recursively partitioned to give four component elements.

The meshing steps are indicated in figure 3.5. We see in figure 3.5c that the generated meshes extend beyond the limits of the available delineated contours - these extensions are a result of the necessary smoothness of the generated mesh.

(a) Input data points and calculated normals.

(b) Poisson generated endocardial surface and subsequent subdivision.

(c) Poisson generated epicardial surface and subsequent subdivision.

(d) Final meshes after cropping to limits of input contour data.

Figure 3.5: Illustration of the steps taken in MeshLab to generate endo- and epi-cardial surface meshes.

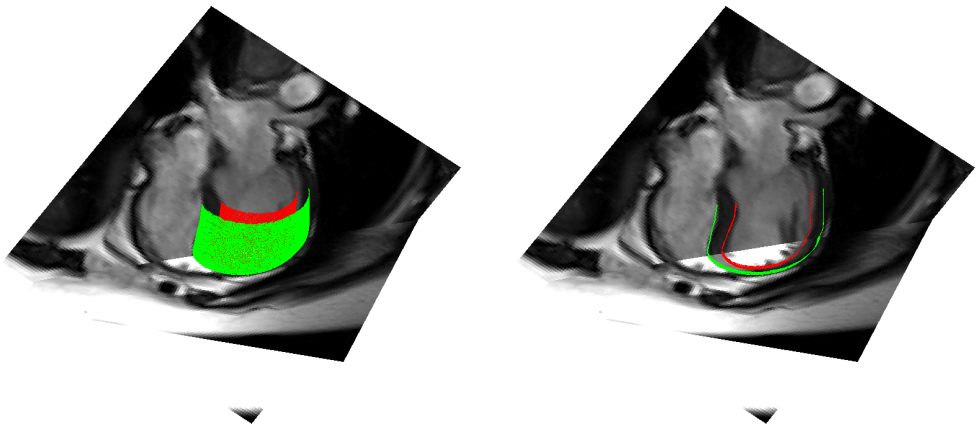


Figure 3.6: Illustration of generated meshes in relation to image data. A section of each mesh is excluded in the right-hand image for further clarity.

The meshes are cropped of these extensions as indicated in figure 3.5d, such that our final surface meshes are limited to the regions where we have available input data. Finally, in figure 3.6, the generated meshes are illustrated as they relate to the SA and LA image data.

It was indicated in the previous section that delineations were available for both SA and LA data; however, we have outlined a mesh generation that used the SA data only. The primary reason for this is that the contours drawn on SA and LA data respectively do not consistently intersect, even after the slice realignment process that is described in chapter 5. While the meshing process employed would be capable of dealing with this - generating the best fitting surface from all the data points - the confidence in the accurate representation of SA contours is reduced. Given that the meshes will be used to assess the registration accuracy of volumes generated from SA data alone, it was chosen to use meshes that accurately represent the original SA contour data.

3.1.3 Scar calculation from LGE image data

In order to analyse the LGE images, in addition to the left ventricle borders, clinicians require a number of further data annotations. These extra annotations provide the necessary data for the automated characterisation of myocardial scar tissue based on tissue intensities. Here, a simple explanation of the common automated scar characterisation method employed by the cmr[42] software will be



(a) LGE annotations **(b)** Myocardial region **(c)** Myocardial reference zone **(d)** No enhancement zones

Figure 3.7: The clinical annotations used for estimating regions of scarred tissue.

given; using the delineations provided by our collaborators, we mimic this process, forming our own estimate of damaged cardiac tissue. We then generalise this approach for an analysis of scar tissue in our generated 3D volume data.

In figure 3.7a, an illustration is given of the annotations provided by the clinicians. As with all cine and LGE data, the endo- and epi-cardial contours are provided, and these are used to define the myocardial region indicated in figure 3.7b. The yellow contour in 3.7a represents a region of healthy myocardial tissue; this reference region (figure 3.7c) is used to define the mean and standard deviation (SD) that characterise the intensities of healthy tissue within the myocardium. Scar tissue may subsequently be approximated by identifying those pixels, within the myocardial region, that are a prescribed intensity distance away from the mean value. There is little consensus in the literature regarding the exact SD value to be used for LGE assessment, with values ranging from 2 SD [44] to 6 SD [18] previously used. Based on a qualitative assessment of the scar tissues identified it was concluded that the use of 3 SD gave the best correspondence to that observed in the images. Due to delineation errors and noise, false positive scar identification can occur. The blue contours are used to designate regions, such as those in figure 3.7d, that do not contain scar and can therefore be used to discount pixels identified as scar within them.

The results of the simple scar calculation outlined are given in figure 3.8 for values of SD ranging from 1 to 5. Based on these images, we can see that a value of approximately 3 SD gives a good indication of the scar region present.

The above analysis is applied to the acquired 2D slices. We desire an equivalent analysis that we can apply to the reconstructed LGE volumes to define 3D renderings

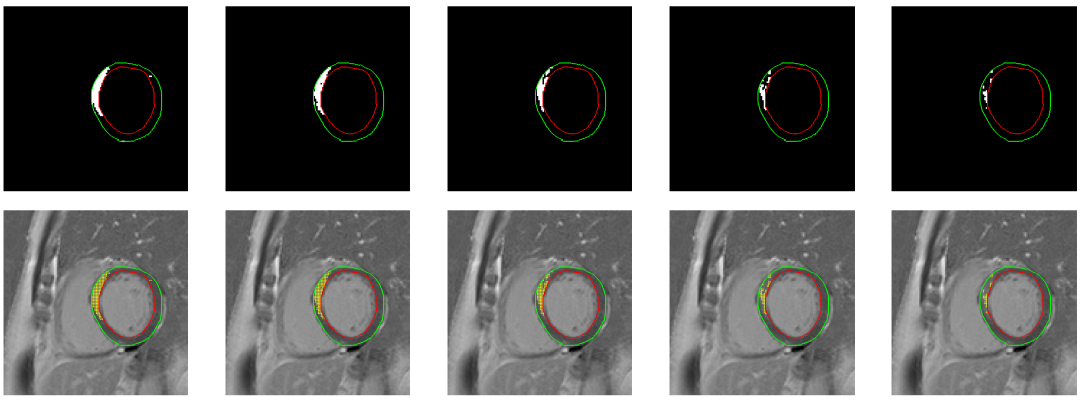


Figure 3.8: Regions of scar tissue calculated with different standard deviation offset thresholds.

of scar. Volumetric myocardial regions can be straight-forwardly defined using the surface meshes, generated using the methodology described in the previous section. The definition of reference myocardial regions and non-enhancement regions is less straightforward. These selected regions do not necessarily show correlation between slices, therefore an approach such as shape-based interpolation is not appropriate.

Rather, we choose to use the reference myocardial regions defined at the slice positions to provide mean and standard deviation measures at their respective locations. From these values, defined at the slice locations, a linear interpolation is performed that determines the individual mean and standard deviation values that each voxel may be compared to.

The no-enhancement regions are interpolated using a nearest neighbour approach. This choice is based on the assumption that any erroneous enhancement that was present in the original slices - necessitating the definition of a no-enhancement zone - may be locally propagated in the image volume by means of the linear interpolation performed. Therefore, the volumetric no-enhancement regions must also be locally similar to those defined in the slices.

In figure 3.9 the scar regions identified within the myocardial borders are depicted over the same range of standard deviation ranges as for the 2D case.

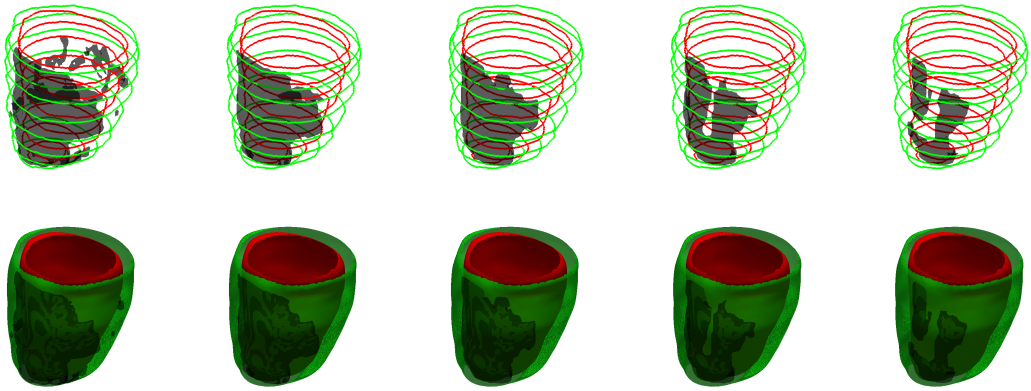


Figure 3.9: 3D renderings of scar identified from volumetric data over a range of standard deviation offset thresholds.

3.2 Challenges and Limitations of MRI Acquisition

As was indicated at the beginning of the chapter, there are limitations associated with data in the form of 2D planar acquisitions such as those used in the work presented in this thesis. In the next few sections we discuss some of the primary issues that are faced, all of which we look to address in subsequent chapters.

3.2.1 Image resolution and partial-volume effect

The first limitation we discuss is that of the resolution of the images. The 2D planar datasets acquired are highly anisotropic. Each pixel in the output images displayed is in reality a representation of the signal received from a voxel (volume-element) of the imaged tissue. The out-of-plane dimension of these voxels is generally considerably larger than the in-plane dimensions. Referred to as the slice thickness, this dimension is controlled by two factors: the magnitude (and therefore slope) of the slice selection gradient, and the bandwidth of the RF excitation pulse. To excite a single thin slice in the tissue, an RF pulse constrained to a single frequency would be required. The true RF excitation pulse applied contains a spread of frequencies, described as the bandwidth of the pulse. As a result, a finite slab of tissue is excited, composed of nuclei whose Larmor frequencies fall within the bandwidth of the RF pulse. The thickness can be adjusted, either by changing the bandwidth of the

excitation signal, or by changing the gradient field strength, which would result in a different rate of change of Larmor frequency along the slice selection axis.

In practice, the choice of slice thickness is a trade-off between two factors; increasing the slice thickness improves the SNR due to acquisition of signal from a greater number of nuclei, however, it additionally makes it increasingly likely that the voxel contains multiple different tissue structures. There is no way of distinguishing these different tissue types and the received signal is a mix of those emitted - this phenomenon is referred to as the partial volume effect. In the output MR image the partial volume effect can be observed as a blurring of structure edges.

The in-plane dimensions of the voxels are much smaller than the out-of-plane dimension. The in-plane resolution is controlled by two factors: the field-of-view and the matrix size. These factors are dependent on the frequency- and phase-encoding steps during image acquisition; improvements in resolution in the respective dimensions can be achieved by increasing the number of phase-encoding steps or the sampling rate of the frequency-encoded read-out signal.

In addition to the issue of data anisotropy, many acquisitions do not capture data from the entire 3D volume, rather leaving gaps between the slabs of tissue from which images are generated. There are multiple reasons for this: firstly, the total acquisition time for a certain volume extent can be reduced by introducing these small gaps between the image data acquired; secondly, and more significantly, is its ability to reduce what is referred to as cross-talk. The shape of an RF pulse's frequency distribution is not perfectly rectangular; as a result, adjacent slices may have overlapping excitation areas. Due to the rapid application of RF pulses, this may result in multiple consecutive excitations of the same area of tissue, which can lead to saturation of the magnetisation vector. Introducing gaps between slices is one method used to compensate for this, as it decreases the area of overlap between adjacent pulses.

Here we have discussed the limitations of the attainable data resolution based on the restrictions of MR acquisition. The primary concern is the anisotropy of

the 2D planar data acquired. Later in the thesis we will discuss proposed post-processing solutions for improving the resolution in this out-of-plane direction such that isotropic data representations are generated.

3.2.2 Cardiac and respiratory motion - Breath-hold misalignments

Organ motion during MR acquisition causes a number of issues such as image blurring, and image ghosting [178]. Such issues are of particular note when imaging the heart due to its dynamic nature. It undergoes deformations due to contraction as well as positional changes such as those due to respiration. A number of methods have been developed to minimise their effect on image quality.

For example, gated acquisitions are used to deal with the contractile component of heart motion. Images are acquired at a number of time points during the cardiac cycle with the information contained in the electrocardiogram (ECG) signal used to either prospectively control the image acquisition timings [101], or retrospectively reorder and reconstruct the appropriate cardiac phases [67, 104]. Either way, each frame of the final cine sequence is generated by combining data captured of the same cardiac phase over a number of successive heartbeats.

Of similar importance are the strategies proposed for dealing with respiratory induced motions, that occur as a result of changes in lung volume and the related movement of the diaphragm during the breathing cycle. In [128], McLeish et al quantified the potential scale of these displacements by registering MR images acquired at a series of respiratory stages, finding that the predominant motion is observed in the superior-inferior direction.

As was the case for dealing with contractile motion, respiratory compensation schemes look to exploit the recurrent nature of the breathing action. However, without a clear equivalent to the ECG signal used for cardiac gating, a variety of approaches for monitoring respiration have been proposed, such as: belts applied to the upper abdomen for estimating changes in volume [52]; analysis of the alteration in measured ECG signal due to the changing impedance with thoracic volume variation

[55]; or the application of additional RF pulses, referred to as navigators, that can be used to assess the relative position of each repeated image acquisition [51].

Based on the respiratory signals measured, the triggering system may be synchronised to capture images at the optimal phases of the respiratory cycle, or continuous acquisitions may be coupled with image reconstruction methods that accept data only if captured within a pre-defined window of the respiratory cycle. Due to the inefficiencies associated with these gated approaches, there is significant interest in the utilisation of respiratory models [127]. By modelling the relationship between the measured respiratory signal and the observed image motion, these methods can be used to compensate for breathing induced motion, potentially improving image quality and scan efficiency.

A simpler means of mitigating the effects of respiratory motion is the breath-hold approach proposed in [150]. This involves the patient suspending breathing and maintaining a constant lung volume for the duration of an acquisition. The quality of the results in this case is limited by the duration for which the patient can comfortably maintain breath-hold. Despite this, it is still a common approach in clinical practice, and was employed for the acquisition of data used in this thesis. The time required to obtain each MR image necessitates that this process be repeated a number of times in order to fulfil an entire cardiac study; a study that comprises a number of 2D planar slices at a variety of angles and locations in relation to the cardiac anatomy. The necessity for multiple breath-holds presents the challenge of reproducibility. Differences in breath-hold volumes between distinct 2D acquisitions means that the image content in each slice has the potential to be displaced in relation to each other. This makes the subsequent integration of 2D data into a 3D reconstruction difficult.

Adding to the distortion due to breath-hold reproducibility, larger patient movements within the scanner during the imaging session result in further misalignments between corresponding tissues in intersecting images.

3.2.3 Intensity inconsistencies

There exist a number of intensity-based MR image acquisition issues. The more widely featured in the literature is intensity inhomogeneity, alternatively referred to as the bias field. This artefact is characterised by a slowly varying intensity across the image domain that has no relation to tissue properties - it is caused primarily by imperfections in the imaging device. The result is that homogeneous tissues may be characterised by variable signal intensities in the generated MR image. The degree of variation can cause problems when looking to utilise automated image processing tools based on intensity. A review of proposed methods for correcting these intensity variations is given in [208].

The second potential intensity-based issue is caused by the lack of a fixed intensity scale: intensity values between different MR images have no fixed meaning. As a result corresponding tissues displayed in two MR images may be characterised by different intensity values. Moreover, the intensity differences are variable across tissues, meaning a simple scaling of the image values is not possible. As was the case for bias field inhomogeneity, the intensity differences encountered cause problems when using automated image processing tools. In addition to making operations such as segmentation and registration difficult, these intensity inconsistencies can make the integration of complementary information from multiple image orientations challenging.

The data dealt with in this thesis is subject to potentially further distortion as a result of the contrast agent used during the imaging session. As was indicated in figure 3.2, the acquisition of cine data can be dispersed between the administration of gadolinium contrast and therefore any agent that persists in the tissue for longer periods can have an effect on the observed intensities.

3.2.4 Image acquisition reproducibility - Misalignment of slices from different studies

The difficulties of dealing with cardiac and respiratory motion seem apparent, however perhaps less acknowledged is the fundamental challenge presented to image

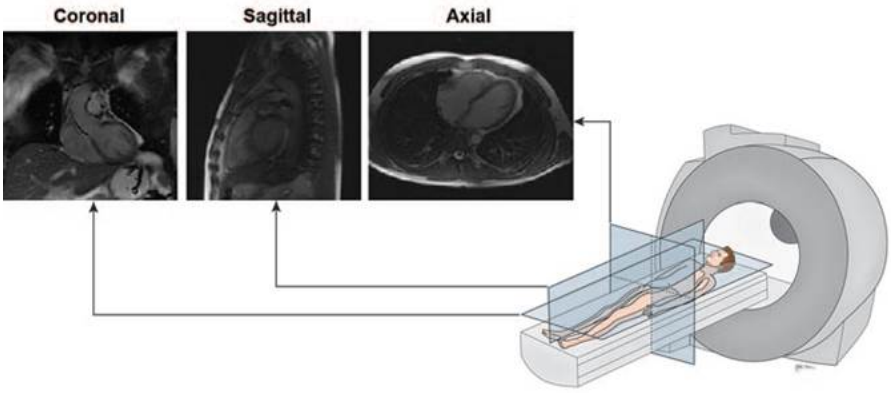


Figure 3.10: The standard body imaging planes. Figure reproduced from [66]. Reprinted with permission from the American Journal of Roentgenology.

acquisition by the position and orientation of the heart within the body. The standard imaging planes of the body are the axial, sagittal and coronal planes, and they are used for imaging various structures in the body. However, in figure 3.10 we see how these planes relate to the heart, and while they provide an adequate overall illustration of the heart's morphology, they do not provide the desirable anatomical and functional characterisation of the muscle [66]. For this, images are required that are aligned to the primary orientations of the heart itself as indicated in figure 3.11

The process of acquiring the desired images uses the initial body planes as a starting point for locating the appropriate cardiac planes. The successive images acquired contain new tissue features which can guide the localisation of the appropriate planes, until a final study comprising SA and multiple LA views, as indicated in figure 3.11, is obtained.

Despite the standardised procedure used and the care taken in acquiring the correct cardiac planes, the process is clearly complex, and much of the plane selection is dependent on the accurate identification of key landmarks. As a result, it is unlikely that images taken on two independent visits to the hospital are of the exact same anatomical area - rather differences in both location and orientation are likely to exist.

Current longitudinal studies analyse remodelling under the assumption that a pair of images acquired during distinct imaging sessions are of the same tissue. In this section, we will present the results of an initial study carried out to ascertain

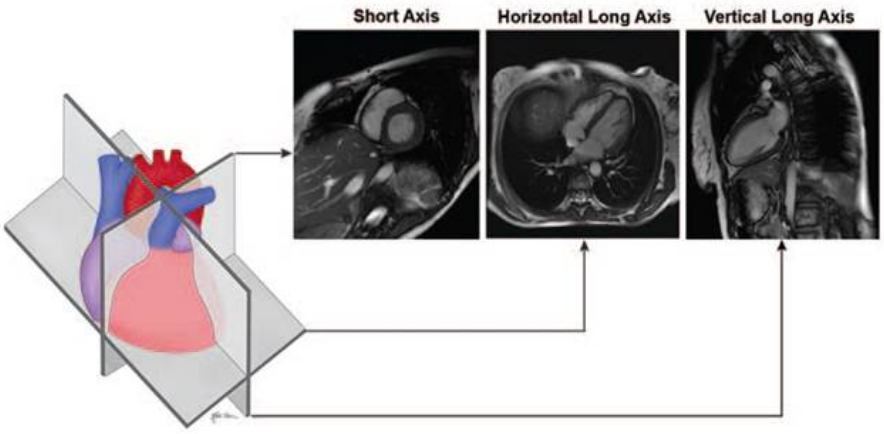


Figure 3.11: The standard cardiac imaging planes. Figure reproduced from [66]. Reprinted with permission from the American Journal of Roentgenology.

the potential errors introduced in current analysis due to the limitation of data alignment to two dimensions.

In previous studies of CMR reproducibility [140, 74, 139] results have been presented that consider only the relative values of cardiac metrics such as left ventricular volumes and ejection fraction. These studies indicate good interstudy reproducibility, however, they do not report results that quantify specific slice misalignments. Here, we propose to analyse the anatomical misalignment between images acquired of the same patient during distinct imaging sessions, reporting quantitative results in the form of orientation and out-of-plane distances. Therefore, the goal of the experiment is to quantify the location error between slices that, in clinical routine, are identified visually as "corresponding".

The results from this study were previously presented at the 17th annual Society of Cardiovascular Magnetic Resonance (SCMR) scientific sessions in New Orleans [87].

Methods

To ascertain the differences in location of slices acquired during two longitudinal imaging sessions of the same patient, a 2D in-plane registration between a single LA slice and its corresponding slice from the fourth visit data is performed and the resultant positions of the SA slices assessed.

There are a few key assumptions that we have made regarding the data that allows this estimation of typical slice positional discrepancies to be made. Firstly, we assume that the shape of the heart is not significantly changed between the two scans; i.e. we assume no major remodelling during this six month period and that variations in breath-hold positions do not have significant implications on the geometry of the heart. We have additionally assumed that breath-hold positions within a single scan are consistent enough not to cause major misalignment between LA and SA slices. Finally, we make the assumption that the LA slice selected for 2D in-plane registration has been accurately acquired and depicts the same plane of tissue.

The results for this study were obtained from a small subset of datasets that were initially provided for our project and for which we were able to qualitatively verify the validity of our assumptions.

The positions of all images may be ascertained from the contents of the DICOM file headers, and are given relative to a specific fixed point in the MRI machine. Before carrying out an alignment based on the LA image content, we first move the entire dataset such that the LA image pair selected occupy the same spatial location. The transformation required is computed according to the respective DICOM header content of the two LA slices.

For the subsequent alignment of the chosen LA images, a point-based registration approach was used. Using the MATLAB interface, at least three corresponding points were manually selected in both images - in particular points near the left ventricle valves and an apical point were chosen. Based on the selected points a 2D affine transform was calculated via a least-squared estimation. The corresponding 3D transformation was used to move all other slices in the dataset. Clearly this registration is highly dependent on control point selection with any discrepancies likely to negatively impact the quality of the registration result. We chose to use a small number of points as this allowed us to select clear salient points in the image for which it was straightforward to identify its correspondence in the second image. The validity of all registration results were qualitatively evaluated before further analysis.

Quantification of the SA slice positional differences is performed between a mid-cavity slice from the first visit and the fourth visit slice which, qualitatively, is considered to be taken from the "corresponding" anatomical location, as done in clinical practice. From their respective positions, the angular differences and pixel offset values were calculated. Angular differences were estimated according to the respective plane normals. The offset differences were first calculated between all points in the images. However, due to the potential for misleadingly large offset values at the image edges resulting from small angular differences, we also calculated the offset values for only a small region of interest around the cardiac tissue. The selection of a region of interest was again performed manually.

Results

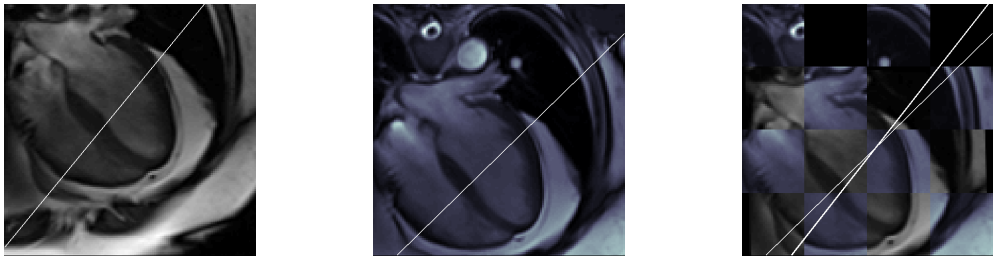
For each dataset a single LA view was selected for alignment. The initial angular difference and offset difference between the two selected planes was corrected before point-based registration was used for 2D alignment. Figures 3.12a and 3.12b show the respective LA slices and the position of one of their intersecting SA slices before registration. In figure 3.12c a checkerboard display is used to illustrate the relative alignment of the images after registration.

The error in SA alignment was quantified via the angular difference between plane normals and the mean z-offset of the entire plane and a manually selected region of interest (ROI) containing only the heart and close surrounding tissues.

The mean angular difference for all patients analysed was $8.34^o \pm 7.47^o$, whilst the mean offsets were 8.16mm $\pm$ 8.54mm and 6.50mm $\pm$ 6.54mm for the entire plane and the ROI respectively. Figure 3.13 shows the relative positions of the SA slices for a single patient from the two hospital visits. The arrows drawn connect corresponding pixels - the calculated z-offsets were taken as the z-components of these connecting vectors.

Discussion

The results presented indicate the potential error that exists in the position of two SA slices taken from two different hospital visits. Results showed important



(a) LA slice - visit 1 **(b)** LA slice - visit 4 **(c)** Data alignment post-registration

Figure 3.12: (a) and (b) show corresponding LA slices from two visits. White lines represent the intersection of the selected SA slices. In (c), the result of the co-registration of (a) and (b) is shown, with the two images occupying alternate squares in a checkerboard pattern (note, the different colour maps have been used to facilitate differentiation).

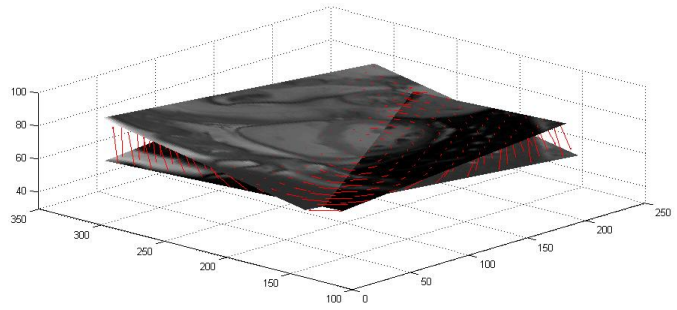


Figure 3.13: The position of SA slices after applying LA co-registration, where the misalignment is represented by the red vertical arrows.

errors in all tested datasets. There was a fairly large variation in the magnitude of these errors - this in fact emphasises the difficulty and lack of consistency involved in the cardiac image acquisition process. It is important that these misalignment errors are considered - subsequent analysis, after alignment, such as the calculation of myocardial wall thickness in the two images could lead to inaccurate conclusions, if the two tissue areas being compared do not in fact correspond to the same anatomical tissue region.

There are limitations to this analysis. As previously mentioned, it is likely that there will have been error in the positioning of slices of all orientations between the two datasets. However, for the purpose of our analysis it was necessary to assume that a single image orientation was positioned in the same plane as its

corresponding image in the second dataset, enabling 2D registration between these two images. In our analysis we choose a LA slice for this purpose to enable the comparison of SA slices. The analysis could be performed similarly by assuming SA slices were in the same plane and comparing the resulting position of the LA slices. In reality both these investigations involve inaccuracies, though we hypothesise that despite the initial assumption taken, the values of error quantified are indicative of the true errors seen between the anatomical positions assessed.

4

Background - Image Analysis

Contents

4.1 Registration 55
4.1.1 Basics 56
4.1.2 Optical flow 61
4.1.3 Applications of image registration in the analysis of CMR data 66
4.2 Reconstruction 80
4.2.1 Volumetric reconstruction methodologies applied to CMR data 83
4.3 Local Phase 85
4.3.1 Local phase in nD 92
4.3.2 Further phase-based properties 99
4.3.3 Applications in image registration 101

In this chapter, a number of image analysis topics that are important to the remainder of this thesis will be presented. Firstly, the theoretical foundation of image registration is discussed, before a specific review of the literature relevant to the registration problems tackled in subsequent chapters is given. This will be followed by a discussion of methods for the interpolation of data in the context of improving the resolution of CMR data. Finally, the mathematical basis of local phase-based properties will be reviewed as this signal representation is essential for the work presented in later chapters.

4.1 Registration

Image registration is the process by which a mapping is found that brings multiple data sources into a common coordinate system. It enables the comparison of data acquired at different times, from a variety of sources, from different viewpoints as well as of similar, but distinct objects. It is a highly active research area that has found application in a range of fields including computer vision, remote sensing, and medical image processing. In medical image analysis, the objective is to bring together complementary information that can aid in a number of clinically important tasks such as diagnosis and therapeutic planning.

The modern clinical setting is equipped with an assortment of imaging devices, and a primary benefit of image registration is its ability to bring together the complementary information contained in the various modalities. For example, in studies of the brain, positron emission tomography (PET) is used to identify indications of function, but is itself limited in its provision of anatomical information: MR, on the other hand, provides high-quality anatomical data. Through the utilisation of registration methods [218], the two data sources can be combined, giving a more complete physiological account. Alternatively, registration may be used simply to improve anatomical description. A typical example is the combination of CT and MR data [190], which are respectively proficient in identifying bony structures and soft tissues.

A further clinical role is found for image registration in the long-term observation of patient condition. Follow-up image acquisitions enable the observation of important clinical determinants, such as disease progression [123] and the body's response to treatment [78].

An application frequently seen in medical research is the alignment of data acquired from different patients. While this initially enables the direct comparison of subjects, larger scale studies can be used to generate atlases of healthy and pathological tissue [199], which help in identifying common trends of variability. Such registrations also serve a purpose in enabling the automation of other image

processing operations, such as segmentation, through subject to labelled-atlas correspondences [29].

While there exists a wide breadth of applications, each approach is fundamentally composed of the same constituent parts:

- a transformation model - which defines how an image may be deformed to match the other
- a similarity metric - which is the means of quantifying the data correspondence
- an optimisation methodology - which is the framework used to traverse the search space in order to identify the transformation parameters corresponding to the optimal alignment.

The constituent components referred to above are briefly discussed in the next sections, with specific reference to medical application where appropriate. For further reference, introductions to image registration methodology are available in [25] and [226], while methods that deal specifically with medical imaging data are discussed in [120], [81] and [77]. The application of image registration to cardiac data is reviewed in [121].

4.1.1 Basics

The fundamental aim during registration - that of finding the optimal image transformation, T_{opt} - can be written formulaically as in equation 4.1.

$$T_{opt} = \arg \min_{T} / \max S(I_F, T(I_M)) \tag{4.1}$$

where we have introduced notation for the moving image, I_M, the fixed image, I_F, the transformation, T, and the (dis)similarity criterion denoted by S.

Similarity metrics

The quantification of current image alignment quality is done by means of a similarity (or dissimilarity) criterion. The two primary classifications of criterion are feature-based and voxel-based. Feature-based approaches utilise correspondences between either extrinsic or intrinsic marker locations; extrinsic markers, invasive or non-invasive, are physical objects artificially placed on the body for easy

identification; intrinsic markers are salient locations such as points, lines, or surfaces that are identified in the images during post-processing. Feature identification can be performed manually or by automated processes specifically designed for landmark identification.

Both approaches have associated drawbacks; while extrinsic markers may be undesirably invasive or suffer from marker displacement during acquisition, intrinsic markers require significant time and expertise, making them susceptible to inter-user discrepancies, or require highly robust automated identification algorithms.

Using the identified features, alignment is achieved by the minimisation of distances between corresponding landmarks. For example in point-based registration, assuming the use of a fairly simple transformation model (i.e. rigid or affine), this problem is often solved given the solution to the well-known orthogonal Procrustes problem [59].

Voxel-based similarity metrics utilise the image voxel values to quantify alignment quality. Metric selection is based on the images to be aligned, the key factor being the expected relationship between values at corresponding image locations.

The sum-of-squared-differences can be used when the intensities of two images are expected to match when correctly aligned; therefore a minimisation of the squared intensity difference, expressed in equation 4.2, provides an optimal mapping, where x indicates the position in the images and Ω defines the image domain. The equal-intensity assumption can be somewhat relaxed by using the cross-correlation coefficient, or more appropriately its normalised form, expressed in equation 4.3, where the $\bar{I}$ indicates the mean image intensity. This provides a measure of the degree to which two signals are linearly related.

$$\textbf{Sum of Squared Differences, } SSD = \frac{1}{N}\sum_{x\in\Omega}\left(I_1(x) - I_2(x)\right)^2 \tag{4.2}$$

$$\textbf{Normalised Cross Correlation, } NCC = \frac{\sum_{x\in\Omega}(I_1(x) - \bar{I_1}) \cdot (I_2(x) - \bar{I_2})}{\sqrt{\sum_{x\in\Omega}(I_1(x) - \bar{I_1})^2 \cdot (I_2(x) - \bar{I_2})^2}} \tag{4.3}$$

Given their basic intensity assumptions, the above metrics are most commonly restricted to the registration of mono-modal data. The more complex intensity

relationships encountered in multi-modality registration require greater flexibility in the similarity metric assumptions.

For multi-modal applications, information theoretic measures such as mutual information(MI) [41, 206] are popular. Expressed in equation 4.4, MI is calculated based on entropy values, H, and is related to the distribution of pixel grayscale values, specifically the co-occurance of intensity values, $H(I_1, I_2)$, in two images. For quantification that is invariant of image overlap, normalised mutual information (NMI), expressed in 4.5, was presented in [191]. A survey of the application of mutual information in medical image registration literature is given in [159].

$$\textbf{Mutual Information, } MI = H(I_1) + H(I_2) - H(I_1, I_2) \quad (4.4)$$

$$\textbf{Normalised Mutual Information, } NMI = \frac{H(I_1) + H(I_2)}{H(I_1, I_2)} \quad (4.5)$$

The metrics described provide only a brief and limited insight into the utilisation of similarity metrics. The specific metrics used are often designed for the application in mind.

Transformation models

To align the moving image, I_M, to the fixed image, I_F, it is necessary first to define the type of spatial mapping that I_M can undergo. The selection of a transformation model provides a description of the deformation permitted.

Transformation models can be categorised broadly as being global, where the transformation is applied homogeneously to the entire image space, or local, where the defined deformation is dependent on position in the image space. The choice of transformation is important and depends on the expected nature of the underlying misalignment, in addition to some consideration of computational complexity introduced by increasing parameter space.

For global transformations the new image coordinates can be related to the original coordinates via simple equations, such as those given in equations 4.6 and 4.7 for rigid and affine models respectively. Rigid transformations, defined in 2D by three parameters (six in 3D), allow for rotation (θ) and translation (t_X, t_Y), and are appropriate in cases where the retention of shape and size is sought-after. The

affine model incorporates further degrees of motion in the form of shear and scaling and can be defined in 2D by the six parameters (twelve in 3D) of equation 4.7.

$$\begin{bmatrix} x_{new} \\ y_{new} \\ 1 \end{bmatrix} = \begin{bmatrix} cos(\theta) & -sin(\theta) & t_X \\ sin(\theta) & cos(\theta) & t_Y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \tag{4.6}$$

$$\begin{bmatrix} x_{new} \\ y_{new} \\ 1 \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \tag{4.7}$$

While the simplicity of global model definitions may be advantageous, such transformations are rarely appropriate for medical image registration. The deformable nature of tissue requires more complex models that have the flexibility to capture local image differences. These models, where deformations are described according to position in the image space, may be referred to in the literature as non-rigid, non-linear or deformable transformations.

While global transformations can be compactly described by small sets of parameters, non-rigid models can be either parametric or non-parametric. Spline-based techniques are a commonly utilised parametric non-rigid registration; deformation is orchestrated by the movement of control points that span the image domain, where the splines provide a smoothly varying deformation field between the control points. The objective is to find the set of transformation parameters for each control point that optimises the similarity criterion over the entire image.

The cubic B-spline is the fundamental component of the free-form deformation model and has been applied in numerous registration tasks [174]. A primary advantage of B-spline-basis functions is their compact support, meaning a change in parameters for one control point does not affect the deformation of the entire image, only those within a certain neighbourhood. This aids in the correction of local deformation in addition to reducing the computational cost. Local registration accuracy can be improved by defining a finer control point lattice, however, the increased number of parameters required similarly increases computational cost.

Non-parametric models define a different deformation vector for each pixel in an image with no specific functional relationship. The lack of any functional constraint leaves an ill-posed problem that necessitates the introduction of an additional regularisation term as part of the cost function definition. Optical flow is one of the most common non-parametric approaches, a specific introduction to which will be given later in this chapter.

Equation 4.8 is of a similar form to that which we gave at the beginning of the section, but now includes the regularisation term, R, therefore providing a more general definition encompassing non-rigid formulations.

$$E(T) = \arg\min_{T} \Big(S(I_F, T(I_M)) + \alpha R(T) \Big) \tag{4.8}$$

The regularisation serves a general purpose in controlling deformation behaviour, ensuring the plausibility of the final result. Regulariser selection is motivated by the particular characteristics of the expected deformation - some common regularisers are:

- diffusion - ensures the overall smoothness of the registration result by penalising large gradients in the deformation field, $\mathbf{u}$.

$$R(\mathbf{u}) = \frac{1}{2} \sum_{i,j=1}^{d} u_{ij}^2 \quad \text{where} \quad u_{ij} = \frac{\delta u_i}{\delta x_j}$$

- elastic [23] - the smoothness of the deformation field is governed by the parameters λ and μ, set to reflect the underlying elasticity characteristics of the deformed object. The regularisation acts as the internal force that responds to the application of an external force - here quantified by the similarity criterion - according to the elastic behaviour of the object. The final deformation reflects the equilibrium state.

$$R(\mathbf{u}) = \frac{1}{2} \sum_{i,j=1}^{d} \left(\lambda u_{ii} u_{jj} + \mu \left(\frac{u_{ij} + u_{ji}}{2} \right)^2 \right)$$

- fluid [39] - the regularisation here considers the smoothness of the instantaneous velocity field rather than that of the displacement. As a result, fluid regularisation allows for greater deformations than those possible when using elastic regularisation.

Further details regarding regularisers and non-rigid registration can be found in a number of surveys [105, 42, 187].

Optimisation

Without a closed-form solution, image registration is framed as an optimisation problem, whereby the image transformation is iteratively modified such that the calculated similarity metric is improved. Optimisation methods such as gradient-descent, Gauss-Newton, and the simplex method, which are utilised in voxel-based registrations, update parameters by traversing the search space according to specified directions and step size until a convergence of the cost function is reached. For non-parametrised non-rigid registration methods, an analytic optimisation approach may be used that leads to the Euler-Lagrange equations; these can be solved using numerical solvers such as the successive-over-relaxation approach.

To avoid falling into local minima of the optimisation search-space, multi-resolution or multi-scale strategies are employed during registration. These approaches register the data using coarse-to-fine-resolutions or coarse-to-fine-features respectively. The deformations calculated at each level of the scheme may be used as an initial estimate of the registration result at the next level.

4.1.2 Optical flow

In this subsection, we pay particular attention to optical flow, as it forms the basis of the registration approach used later in this thesis.

Optical flow is the perceived 2D motion field observed as a result of the relative 3D motion between an object and viewer. Many techniques exist that strive to calculate this 2D image motion; classically, methods are based on the assumption that the intensity of moving objects in the scene is retained. Based on this premise, any change in observed intensity is associated directly with object motion. The intensity constancy assumption may therefore be expressed as in equation 4.9 where, $\mathbf{x} = [x, y]$ is the image spatial coordinate vector, $\mathbf{dx} = [dx, dy]$ is the motion vector, and t is used to denote time.

Amongst the variety of methods for estimating optical flow [17, 60], approaches based on the calculation of spatio-temporal gradients are perhaps the most well known and extensively used. These methods may be formulated by considering the first order Taylor expansion of the right hand side of the intensity constancy formulation. The resultant expression is shown in equation 4.10. A series of simplifying steps leads to equation 4.11, where $\nabla I(\mathbf{x},t) = [\frac{dI(\mathbf{x},t)}{dx}, \frac{dI(\mathbf{x},t)}{dy}]$ is the image spatial gradient, and $\mathbf{u}(\mathbf{x}) = [u(\mathbf{x}), v(\mathbf{x})] = [\frac{dx(\mathbf{x})}{dt}, \frac{dy(\mathbf{x})}{dt}]$ represents the displacement vectors of the 2D motion field. This final formulation is referred to as the optical flow (or gradient) constraint equation, and relates image motion to the spatial and temporal gradients of the sequence.

$$I(\mathbf{x},t) = I(\mathbf{x}+\mathbf{dx}, t+dt) \tag{4.9}$$

$$I(\mathbf{x},t) = I(\mathbf{x},t) + \frac{dI(\mathbf{x},t)}{d\mathbf{x}}\delta\mathbf{x} + \frac{dI(\mathbf{x},t)}{dt}\delta t \tag{4.10}$$

$$\nabla I(\mathbf{x},t)\cdot\mathbf{u} + I_t(\mathbf{x}) = 0 \tag{4.11}$$

This expression is under-constrained, having multiple unknown variables for a single equation. No unique solution can be calculated from this solitary constraint; rather, it enables the estimation of the motion's projection along the orientation of the intensity gradient only - this issue is commonly referred to as the aperture problem.

To allow for the calculation of an unique motion field solution, further constraints are required. There are two classical choices. The first, proposed by Lucas and Kanade [117], forms the cornerstone to what are referred to as local methods - these methods assume that the combined information in a region local to a given pixel provide sufficient information to calculate the pixel's optical flow. The second choice is that of Horn and Schunk [82], a global method that enforces constraints on the nature of the optical flow field over the entire image domain. For these global methods, the optical flow vector computed at any pixel has an effect, and is affected by, every other pixel in the image.

Local methods - Lucas and Kanade

Local methods, such as the Lucas-Kanade approach, are based on the assumption of spatial coherence, meaning neighbouring pixels undergo similar motions. Based on this assumption, more equations are made available from which the motion vector at each pixel can be estimated. Mathematically we say that the motion vector $\mathbf{u}$ must satisfy each of the constraints in equation 4.12, where there are N pixels in the defined neighbourhood.

$$\begin{aligned} \nabla I(\mathbf{x}_1, t) \cdot \mathbf{u} + I_t(\mathbf{x}_1) &= 0 \\ \nabla I(\mathbf{x}_2, t) \cdot \mathbf{u} + I_t(\mathbf{x}_2) &= 0 \\ &\vdots \\ \nabla I(\mathbf{x}_N, t) \cdot \mathbf{u} + I_t(\mathbf{x}_N) &= 0 \end{aligned} \tag{4.12}$$

The equations can be reformulated to give the linear system, $\mathbf{Au} = \mathbf{b}$, shown in equation 4.13.

$$\begin{bmatrix} I_x(\mathbf{x}_1) & I_y(\mathbf{x}_1) \\ I_x(\mathbf{x}_2) & I_y(\mathbf{x}_2) \\ \vdots & \vdots \\ I_x(\mathbf{x}_n) & I_y(\mathbf{x}_n) \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} -I_t(\mathbf{x}_1) \\ -I_t(\mathbf{x}_2) \\ \vdots \\ -I_t(\mathbf{x}_n) \end{bmatrix} \tag{4.13}$$
$$\mathbf{Au} = \mathbf{b}$$

This system is over-determined, but can be solved in the least-squared sense, if we minimise $(\mathbf{Au} - \mathbf{b})^2$. In equation 4.14, we see the result of the minimisation procedure, where the matrix product $\mathbf{A}^T\mathbf{A}$ is referred to as the structure tensor whose components are expressed in 4.15.

$$\begin{aligned} \mathbf{A}^T\mathbf{Au} - \mathbf{A}^T\mathbf{b} &= 0 \\ \mathbf{u} &= \left(\mathbf{A}^T\mathbf{A}\right)^{-1}\mathbf{A}^T\mathbf{b} \end{aligned} \tag{4.14}$$

$$\mathbf{A}^T\mathbf{A} = \begin{bmatrix} \sum_{i \in N} I_x^2(\mathbf{x}_i, t) & \sum_{i \in N} I_x(\mathbf{x}_i, t) I_y(\mathbf{x}_i, t) \\ \sum_{i \in N} I_x(\mathbf{x}_i, t) I_y(\mathbf{x}_i, t) & \sum_{i \in N} I_y^2(\mathbf{x}_i, t) \end{bmatrix} \tag{4.15}$$

In the above definition, equal influence has been given to each of the neighbourhood pixels. More commonly, a proximity-based weighting is used. The sum of contributions from weighted neighbourhood pixels can be compactly expressed by

the convolution with a Gaussian-shaped kernel, denoted by K_ρ in equation 4.16.

$$\begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} K_\rho * I_x^2 & K_\rho * I_x I_y \\ K_\rho * I_x I_y & K_\rho * I_y^2 \end{bmatrix}^{-1} \begin{bmatrix} -K_\rho * I_x I_t \\ -K_\rho * I_y I_t \end{bmatrix} \tag{4.16}$$

The solution of the above expression is dependent on the invertibility of the matrix product $A^T A$; if the matrix product becomes singular, such as in locations where there are no local intensity gradients, no solution is possible; this is the primary drawback of local solution methods - that we are not guaranteed a dense optical flow field solution.

Global methods - Horn and Schunk

Global methods are a means of guaranteeing dense motion fields. Rather than making an assumption regarding the local coherence of motion, global approaches place constraints on the characterisation of the flow field across the entire image domain. Horn and Schunk's proposed method is the foundation of such global approaches - it is formulated as the minimisation of an energy function of the form shown in equation 4.17, where Ω defines the image domain, α the relative weighting of the energy's terms, and other symbols take on their previously noted definitions.

$$\arg\min_{\mathbf{u}} E(\mathbf{u}) = \int_\Omega \left(\nabla I(\mathbf{x},t) \cdot \mathbf{u}(\mathbf{x},t) + I_t(\mathbf{x},t)\right)^2 + \alpha \|\nabla \mathbf{u}(\mathbf{x},t)\|^2 d\mathbf{x} \tag{4.17}$$

The energy is composed of two terms. The first term is referred to as the data term, and is defined based on the familiar optical flow constraint equation in 4.11. The second term is the regularisation term, which in the Horn-Schunk approach is defined as the squared magnitude of the velocity gradients. During minimisation of the total energy, this term will penalise spatially abrupt changes in the motion field, characterised by high gradients. The degree of penalisation is dependent on the α coefficient, which controls the relative weighting and, therefore, importance of the data and regularisation terms respectively.

This energy function is often reformulated by considering the optical flow constraint equation in its homogeneous form ($Ax = 0$). This alternative notation

requires the definition of the spatio-temporal gradient ($\nabla_{\mathbf{x}t} I$), spatio-temporal structure tensor ($\mathbf{J}(\nabla_{\mathbf{x}t} I)$), and the homogeneous vector, $\mathbf{u}_{\mathbf{x}t}$.

$$\begin{aligned} \nabla_{\mathbf{x}t} I &= \begin{bmatrix} I_x & I_y & I_t \end{bmatrix}^T \\ \mathbf{J}(\nabla_{\mathbf{x}t} I) &= \left(\nabla_{\mathbf{x}t} I \nabla_{\mathbf{x}t}^T I \right) \\ &= \begin{bmatrix} I_x^2 & I_x I_y & I_x I_t \\ I_y I_x & I_y^2 & I_y I_t \\ I_t I_x & I_t I_y & I_t^2 \end{bmatrix} \\ \mathbf{u}_{\mathbf{x}t} &= \begin{bmatrix} \mathbf{u} & 1 \end{bmatrix} \\ &= \begin{bmatrix} u & v & 1 \end{bmatrix}^T \end{aligned} \tag{4.18}$$

With this notation the Horn-Schunk energy minimisation can be expressed according to equation 4.19.

$$\arg\min_{\mathbf{u}} E(\mathbf{u}) = \int_{\Omega} \mathbf{u}_{\mathbf{x}t}^T \mathbf{J}(\nabla_{\mathbf{x}t} I) \mathbf{u}_{\mathbf{x}t} + \alpha \|\nabla \mathbf{u}\|^2 d\mathbf{x} \tag{4.19}$$

It can be shown that the solution $\mathbf{u}$ that satisfies the minimisation of the Horn-Schunk energy function must correspondingly satisfy the pair of Euler-Lagrange equations shown in equation 4.20, similarly expressed using the components of the spatio-temporal structure tensor in equation 4.21. Via the discrete approximation of the Laplace operator and utilisation of an iterative solver, such as the successive over-relaxation (SOR) approach, a solution, $[u, v]$, can be found to these equations.

$$\begin{aligned} \Delta u - \frac{1}{\alpha} I_x \left(I_x u + I_y v + I_t \right) = 0 \\ \Delta v - \frac{1}{\alpha} I_y \left(I_x u + I_y v + I_t \right) = 0 \end{aligned} \tag{4.20}$$

$$\begin{aligned} \Delta u - \frac{1}{\alpha} \left(J_{11} u + J_{12} v + J_{13} \right) = 0 \\ \Delta v - \frac{1}{\alpha} \left(J_{21} u + J_{22} v + J_{23} \right) = 0 \end{aligned} \tag{4.21}$$

As initially alluded to, the key advantage of global methods is the provision of a dense motion field solution. In highly uniform areas of the image, where no gradients are available for motion estimation, the solution is provided, based on the necessitated smoothness of the motion field across the entire image domain. Therefore, global methods fill in holes that would exist if a local approximation was used. However, global methods have their own drawbacks, primarily their susceptibility to noise and the higher computational cost of finding a solution.

Combined local and global optical flow

A desirable algorithm is one that combines the advantages of both local and global methods. Such an approach was proposed in [26], in which the formulations of Lucas-Kanade and Horn-Schunk were directly combined. The formulation straightforwardly integrates both methods by adapting the data term in the Horn-Schunk approach such that it resembles the Lucas-Kanade minimisation cost. The data term now includes a convolution by a Gaussian smoothing kernel term, utilising the available local information in the image. The new energy function in equation 4.22, closely resembles the Horn-Schunk formulation, however now the spatio-temporal structure tensor, defined in 4.18, has been filtered according to a Gaussian of variance ρ, as indicated by the subscript notation $\mathbf{J}_\rho$ (equation 4.23). The solution $\mathbf{u}$ can be calculated using the same Euler-Lagrange equation with the new values for the structure tensor components.

$$E(\mathbf{u}) = \int \mathbf{u}_{\mathbf{x}t}^T \mathbf{J}_\rho(\nabla_{\mathbf{x}t} I)\mathbf{u}_{\mathbf{x}t} + \lambda \left(\|\nabla \mathbf{u}\|^2 \right) d\mathbf{x} \tag{4.22}$$

$$\mathbf{J}_\rho(\nabla_{\mathbf{x}t} I) = \begin{bmatrix} K_\rho * I_x^2 & K_\rho * (I_x I_y) & K_\rho * (I_x I_t) \\ K_\rho * (I_y I_x) & K_\rho * I_y^2 & K_\rho * (I_y I_t) \\ K_\rho * (I_t I_x) & K_\rho * (I_t I_y) & K_\rho * I_t^2 \end{bmatrix} \tag{4.23}$$

While local-global optical flow is the focus of work presented later in this thesis, a discussion of some more recent developments in the field of optical flow can by found in the review by Fortun et al [62].

4.1.3 Applications of image registration in the analysis of CMR data

In chapter 5 we consider three registration-based challenges relating to CMR datasets. Specifically, we propose methods for:

- the realignment of intersecting CMR images acquired during a single imaging session
- the registration of LGE slices to volumetric cine reconstructions generated from 2D planar data acquired in the same imaging session

- the 3D registration of cine volumes generated from longitudinal studies of the same patient

Therefore, in the following sections we will look at the current state of the art registration methodology relating to each of these applications.

Slice misalignment

Here, we discuss some of the methods proposed in the literature for correcting slice misalignments that occur as a result of multiple breath-hold acquisitions, an issue that was previously discussed in section 3.2.2.

Feature-based methods of slice realignment commonly involve the delineation of the left ventricle endo- and epi-cardial surfaces and subsequent alignment of these extracted features.

In [202], slice misalignment correction is required to ensure the accurate utilisation of active shape models for left ventricle segmentation. The approach employed first considers the positions of the manual SA endocardial delineations centres in relation to the intersection positions of LA views, before refining the SA slice positions until the best match between the SA endocardial contours and one of the LA endocardial contours is obtained.

A similar assumption, based fundamentally on the symmetric nature of left ventricle shape, is made in [193] where a polynomial is fitted to the centroids of the SA endo-/epi-cardial borders, and the difference between this polynomial and the true contour centre used as a correction vector.

Another recent publication by Wan et al [209], uses a point-based approach; slices are registered using the generalised iterative closest point (ICP) algorithm to match SA contours to LA slice contours - it is suggested that by using the generalised formulation of ICP, the probabilistic nature of the metric accounts for the anticipated delineation variability.

Tan et al [195], take a different approach, whereby the image delineations are used to generate a left ventricle surface mesh - correction of slice positions is then

based on the geometric assumption that, when accurately aligned, the generated mesh should fulfil certain criteria in terms of principal curvatures.

In [134], a left ventricle surface model is fitted to the SA delineations. Due to an implicit stiffness in the formulation of the model, misaligned contours do not intersect precisely with the surfaces, and in-plane shift vectors can be defined to correct for these mis-registrations.

All of the above techniques are reliant on the accurate segmentation of the original cardiac MR data. This requires either robust automated segmentation methods, or significant time in order to undertake manual delineations. Therefore, voxel-based approaches to registration are often preferred in order to avoid such tasks. These approaches utilise the image data directly to find the optimal positions of the acquired cine data slices. Clinical datasets are commonly composed of multiple LA images and an SA stack. Thus, a significant number of methods utilise their intersection profiles.

In [13], Barajas et al undertake a single simultaneous optimisation of all slice positions, using the normalised mutual information criterion as an intersection similarity metric. Whereas Barajas et al use only a single cardiac frame, additional complementary information is available in all other image frames (all slice frames being acquired in a single breath-hold). Therefore, in [186], Slomka et al utilise a cost function based on the combination of correlation coefficient and the sum of absolute voxel differences over all these time frames to find the optimal alignment. Similarly, Elen et al, in [54], utilised all cine frames and presented results for a number of different metrics. In addition to the intersection-based cost function, a second term was introduced by Elen et al, which limits the allowed out-of-plane motion for SA slices.

In [113], a different dataset composition is available - both an SA and LA stack were acquired. Such a dataset is accompanied by significantly more intersections on which to base realignment; unfortunately such datasets are rarely acquired clinically. More recently this method was applied to a study where a more standard dataset consisting of only a few LA images was available [114]. A different form

of dataset is also utilised by Chandler et al [32]. In addition to the SA stack, a full 3D acquisition of the heart is available. While such acquisitions have been developed, they too are not considered a standard acquisition in clinical practice; a significant drawback is the associated reduced temporal resolution, which reduces their usefulness for functional analysis. The proposed realignment in this case is by means of a slice-to-volume registration, using NMI as the correspondence metric between the slice and intersection plane.

In [185], Sinclair et al propose to correct for breath-hold related image misalignment using segmentation-based representations of the image data. Using a fully convolutional neural network (FCNN) segmentation approach, labelled data distinguishing LV myocardium, LV blood pool and background pixels in both SA and LA data are available. An iterative registration approach was then used which sequentially translates each image according to the minimisation of the SSD metric with stationary intersecting images. Using their iterative slice translation framework, they provide quantitative results that compare their proposed FCNN algorithm to uncorrected data, an image-based alignment approach using NMI, and a manual segmentation-based approach. Based on the reported distances between SA and LA contours, their method achieves the best results.

Another recent approach presented by Villard et al [205] used a similar iterative slice realignment framework, however, they proposed the utilisation of a local phase-based image representation. Their algorithm was applied to the correction of slice misalignment in LGE cardiac MR data, with their results showing particular improvements in the most extremely misaligned data reducing the mean distance between SA and LA contours from 7.73mm before alignment to 2.72mm after. They draw attention in their paper to the fact that achieving a distance error of zero between contours is highly unlikely due to potential inconsistencies in the manual contouring of data - in fact they found that using an approach that aimed instead to minimise the distances between contours directly, they were limited to a final mean distance of 0.43mm in the most severely misaligned cases.

It is worth noting a key difference between the above methods, namely the degrees of freedom afforded by the applied transformation models. These differences were noted by Carminati et al in [30], who themselves restrict their transformations to 2D in-plane correction. While they presented valuable improvements in a high proportion of cases, they also observed that a more complex model may be needed to deal with some of the observed misalignments.

If we consider the other publications mentioned above: [195], [134], and [185] are also restricted to in-plane translations; [202] and [13] extend this, allowing also for out-of-plane translation; [186] increases this further, with some rotational differences considered; however, only in [209], [54], [32] and [205] are all 6 degrees of freedom allowed.

LGE to cine registration

The analysis of disease progression is a key determinant in clinical treatment selection. A wide range of data are acquired in this endeavour, providing complementary anatomical and functional information. An example is the anatomical and functional data available in cine MRI, and the complementary scar data available from LGE images. For maximal utilisation of this data, it is of benefit if both MR modalities are registered allowing direct integration of findings.

Despite both data sources being acquired in a single imaging session, and protocols being in place to acquire the same imaging planes in both cases, the images acquired are not always well registered. LGE images are acquired some time after the acquisition of cine slices. This, as the name suggests, is because time is required for the contrast agent to enter the tissue, and for this same contrast agent to wash out of the healthy tissue. This delayed imaging results in contrast agent being present only in tissue areas that are damaged, as has previously been discussed. Unfortunately, one issue that results from the need to acquire LGE images a significant period later than the corresponding cine data, is that there is an increased likelihood that the patient will have moved. In many cases these movements will be relatively small - a result of minor positional adjustments or

differences in breath-hold. However, on occasion these may be more significant, and in both cases a correction for these misalignments must be performed, if subsequent co-analysis is to be performed.

Much of the literature relating to the registration of cine and LGE data is focused on the propagation of cine segmentations to LGE images, as discussed in the next paragraphs. This is due to the difficulty in automated segmentation of LGE data, a result primarily of intensity inhomogeneity associated with scar presence within the myocardium [196]. Due to this, a common approach for automation is to propagate delineations from cine data (for which automated segmentation is well researched, see [157]) via 2D slice registration, as is evident from the literature discussed in the next section.

The requirement for a registration step between cine and LGE images is demonstrated in [116]. Lu et al present a study where cine contours were utilised directly for LGE image analysis; it was found that up to 49% of contours required significant adjustments due to mis-registrations. In a later publication, [115], they introduced an affine alignment step to accommodate these issues; however, the contours still required local refinement to accurately segment the LGE images. In [211], Wei et al follow an affine initialisation with a B-spline-based local registration using the pattern intensity metric [210]. Despite the introduction of this non-rigid registration step, a final contour adjustment stage was still required. In a subsequent paper, [212], they drop the non-rigid registration in favour of the direct modification of cine surface meshes to the LGE images. This is after a translation initialisation, which they claim is sufficient for capturing a large proportion of the misalignment. In [47], they perform no initialisation, rather opting immediately to non-rigidly register the cine images to match the LGE-MRI.

Given the sources of mis-registration - patient movements and breath-hold variance - methods have been proposed that allow 3D data realignment. In [220], a rigid registration is performed between 2D LGE slices and a 3D cine volume using mutual information in a multiscale TV-flow framework, as proposed in [194, 6]. This multiscale decomposition allows for initially aligning larger scale structure before

refining using finer structural detail. Hennemuth et al, who in [80] are looking to analyse a range of contrast-enhanced MRI data, also undertake a 2D rigid alignment of LGE to a volumetric representation of the heart. In [40], Ciofolo et al perform a 3D rigid alignment between the estimated LA axis in both datasets; this is then used to transform a mesh, which they subsequently refine to the LGE data. Finally, in [37], they registered the 3D cine volume to the 3D LGE volume using a rigid transformation and the NMI metric; this was undertaken on a sub-volume that contained only the heart and the surrounding structures.

In [196] an initial rigid transformation between a 3D LGE acquisition and a cine SA stack is applied. Using the registration result, the cine SA contours give an initial delineation of the LGE myocardium; to accommodate the higher resolution of the LGE data, extra contours are generated between those defined in the cine stack using linear interpolation. The initial contour positions are refined to better match the LGE myocardium using a multi-stage process that considers endo- and epi-contour edge directions, the 3D continuity of the contours, and the specific locations of scarred tissue in the images.

3D cardiac MR registration

In this section, we consider registration methodologies that lead to the volumetric alignment of cardiac MR data. There are two primary application areas in which we encounter such methods - cardiac motion estimation and cardiac atlas analyses.

Due to the correlation between cardiac disease and myocardial motion, there is an extensive literature related to the automated analysis of tissue deformation. Tagged MRI is considered the gold-standard for cardiac motion analysis [7, 222], generally preferred to cine-based analysis. This is due to tagged MRI's ability to describe the deformations within the cardiac muscle; such an analysis is significantly more challenging using cine MRI, due to the lack of contrast available within the myocardium.

Many registration methods have been designed specifically for the analysis of tagged data. Feature-based approaches, for example, involve the identification of

grid lines and intersection points for tracking frame-to-frame motion. Classical voxel-based methods, such as optical-flow, require modifications that enable them to deal with the application's characteristics, such as the tag-fading property of the data [75, 163, 49]. Of greater interest during this review are papers in which either motion analysis is performed using cine data, or those in which methods are applied to tagged data, but are not primarily characterised by their ability to deal with tag patterns.

Free-form-deformation(FFD)-based registration has been used extensively in the literature for both cine and tagged data analysis. A series of papers was published on the topic from Daniel Rueckert's group, who first introduced the FFD framework to medical registrations [174]. Chandrashekara initially proposed a multi-level FFD [177] approach to approximate the cardiac deformation fields in tagged data [34, 35]. The registration was optimised based on the weighted sum of the NMI for both SA and LA data - the inclusion of both orientations enables the capture of out-of-plane motions, which is essential for an accurate 3D derivation of motion. Later publications, [33, 36], used an adapted method where the multi-level FFD was omitted, preferring instead to embed the single-level-FFD in a 4D framework - a reference 4D volume is composed of duplicates of a single cardiac phase, while the second contains all the originally acquired frames. In [36], they show the flexibility of the approach by registering cine data in addition to tagged data. The objective of the study is to investigate the differing abilities of tagged and cine data to capture myocardial deformation. As is generally accepted in the field, the results showed that cine data was less capable of retrieving the twisting motion of the muscle during contraction, but could still adequately estimate the radial and longitudinal displacements. In [183], Shi et al look to combine the benefits of both modalities by performing a registration using the cine and tagged data within an FFD framework. The contributions of both modalities to the computed deformation is weighted based on image position, with the tagged data having a higher weighting within the myocardium, while cine data are highly weighted at the endo- and epi-cardial borders.

The multi-level FFD approach was used in [168, 169] to derive patient deformation fields. These papers then also used the single-level-FFD to define inter-patient registration, thus enabling the comparison of cardiac deformations. This inter-subject registration was performed on the untagged cine data, which were acquired at the same location as the tagged data. In addition to enabling the direct comparison between two subjects, the second paper [169] presented a broader study, which enabled the generation of a cardiac motion atlas.

Perperidis proposed more complex inter-subject registrations based on FFDs [155, 156, 154] to enable the correction of both spatial and temporal misalignments between cardiac cine sequences. Different transformation models were utilised in each: an affine model in both space and time [155]; a non-rigid spatial model with an affine temporal model [156]; and non-rigid spatial and temporal models [154]. This final framework was designed to deal with both differences in cardiac cycle length and cardiac dynamics. In [153], much of the above work was brought together; the spatio-temporal alignment of untagged data was performed to generate a cardiac atlas, while deformation analysis using methods from [34, 168] can be associated to the generated atlas.

More recently, De Craene et al [45] introduced an extended FFD model in which they sum spatio-temporal B-spline kernels forming a velocity field which provides temporal consistency to the registration result; their method is further regularised by a myocardial incompressibility term. The approach, which they call temporal diffeomorphic FFD, is driven by the assessment of image similarity in terms of mean squared intensity differences, and was originally proposed for application to 3D ultrasound data [45]. However, it has subsequently been used for estimating strain curves from tagged MR data [46] and was further applied to cine MR data as part of the MICCAI cardiac motion challenge [200].

Shi et al [182] developed a sparse multi-level FFD approach for application to cardiac motion analysis - the method is motivated by the difficulty in selecting appropriate resolution scales for the B-spline grids during registration. In their approach an extra sparsity term is added to regularise the registration in such a

way that the optimal level of deformation detail is provided by the result. They compared the results obtained using their sparse FFD approach with two previously described methods; firstly the classic FFD scheme [174], and secondly with the approach presented by Metz et al in [137]. The method presented by Metz et al is another B-spline-based transformation scheme for registering temporal data; in their paper the optimal transformations were found by minimising the variance in intensity across the temporal dimension in a groupwise registration scheme. While application to cine-MR is not exemplified in their paper, they do provide results obtained on cardiac CT data and MR data of the lungs and discuss a cyclical implementation ideal for cardiac applications. In Shi et al [182], they showed that the additional sparsity constraint resulted in improvements in accuracy compared to both the classic FFD scheme [174] and the temporal FFD approach of Metz et al [137] when registering cine-MR data. This method has since been applied in [9] to study the potential for defining cardiac regions based on estimated tissue motion.

In [164], B-spline transformations are used again to model cardiac motion, here from tagged MR data. The registration method composes B-spline transformation fields calculated based on adjacent frame similarities to define deformations from all frames to the end-diastolic frame used as reference; temporal smoothness of the transformation is subsequently ensured by fitting a single 3D+t B-spline model to the composed transformation. The transformations estimated are used to help define a multi-modal motion atlas from both cardiac MR and electrocardiography data.

Examples of the application of optical flow techniques to cine MR registration are also to be found in the literature. Early examples, such as [5], are restricted to 2D analysis. However, with the increased availability of data, improved computational power, and - as noted above - the more functionally accurate analysis possible, subsequent efforts have been concentrated on 3D analysis. In [15], Baron et al assessed the feasibility of utilising optical flow for motion extraction from cine MR data by applying the 3D extensions of the optical flow methods of Horn and Schunk [82], and of Lucas and Kanade [117]. Qualitatively, it was observed that the Horn-Schunk results were generally more aesthetically pleasing, but could perhaps benefit

from a regulariser that would allow for some discontinuities, as they are inherent at certain locations and times in the imaged tissue. Another 3D implementation of Horn-Schunk's optical flow method is utilised in [110, 119]. However, in these papers, the classical problem definition is accompanied by further constraints, based on already known vectors at a subset of boundary points.

Queiros et al [166] chose to recover observed cardiac deformations using an optical flow scheme and a locally affine motion model. Motion is estimated based on local image information - in particular their method relies on the definition in a single phase of the myocardial region - this way, an anatomically constrained motion field is estimated based on image information within a myocardial ROI defined in spherical coordinates. They couple the proposed tracking algorithm with previously described segmentation approaches [14, 165] to define delineations for every phase of the cardiac cycle, applied to both echocardiographic and cine CMR data. While a 3D optical flow formulation is described, for the analysis of cine MR slices they choose instead to track the motion in each 2D cine slice separately using a cylindrical coordinate space due to the anisotropy of the data.

A similar approach to optical flow is the Demons algorithm - it is formulated in the same way, but is solved via a different optimisation procedure. Recently, an implementation of the Demons approach discussed in [98] was used by Popescu et al in [162] to derive cardiac motion; based on these results, slic-supervoxels were used to define regions of myocardial tissue showing similar motions, and these were then correlated to the presence of scar tissue identified from LGE data.

In [124] a diffeomorphic Demons approach is used, with specific constraints related to the physical properties of cardiac tissue. Myocardial incompressibility is enforced by constraining the divergence of the calculated deformation field, and tissue elasticity is modelled by appropriate selection of regularisation term. The proposed method is validated by computing the myocardial strain, and comparing the obtained results with estimates acquired from tagged-MRI and 2D strain estimates.

The performance of this method presented by Mansi et al was further demonstrated in [129], where the authors used the iLogDemons approach to analyse

echocardiography, cine-MR, and tagged-MR data - this is another of the methods compared in the benchmarking publications of Tobon-Gomez [200].

A potential disadvantage of non-rigid methods such as the iLogDemons approach is the number of parameters associated, as the deformation is defined uniquely for each voxel - whilst this provides flexibility to capture localised motion, it makes the comparison of motion between patients or longitudinal time points difficult to interpret. McLeod et al therefore proposed another algorithm in [130] for cardiac motion analysis whereby a poly-affine transformation is estimated from the deformation field calculated via the logDemons approach; this allows cardiac motions to be described in local regions via a small number of parameters. The poly-affine model developed has since enabled advanced statistical analysis of cardiac motion between subjects and patients cohorts [132, 131, 133]

In [102], Folgoc et al propose a probabilistic registration approach that is able to automatically determine certain registration parameters such as the weightings for image similarity and deformation regularisation terms. Additionally, the Bayesian approach described is spatially adaptive, capable of modifying the local transformation parametrisation thus achieving optimal accuracy and regularity; this characteristic is similar to the sparse FFD algorithm described by Shi et al [182] and ensures that the transformation is described by the minimal required degrees of freedom. The probabilistic approach described is applied to the temporal registration of cardiac ultrasound, cine-MR and tagged-MR data. This sparse registration framework has since been applied as part of a combined segmentation registration approach for short-axis cine-MR deformation analysis with improved robustness [125].

We saw earlier, during our discussion of LGE registration, the potential for using alignment results to propagate segmentations. This is one means by which researchers have attempted to alleviate the time-consuming nature, and undesirable inter-/intra-observer variability of cardiac cine data segmentation. Proposals have been made to either propagate segmentations from a single frame to all other frames [143], or - more generally - to propagate segmentation atlases to subject

cases [173]. Specific examples are presented below of the use of segmentation atlases to cine MR data.

As in the case of cardiac motion analysis, FFD-based approaches have found application in the area of segmentation-propagation. Lorenzo-Valdes et al [112] use a propagation approach to achieve the full segmentation of the 3D+t representation of cardiac cycle. They first utilise the multi-level FFDs to generate a so-called subject-specific atlas (a cleaner volumetric representation that results from calculating the mean image of all registered frames). This gray-scale subject specific atlas is then registered, via the combined global and local components of the method previously proposed in [174], to the general multi-subject-atlas with a label segmentation. The end-diastolic atlas that has been acquired from the multi-subject-atlas is then propagated back through the frames of the subject's 3D+t data for full cardiac cycle segmentation.

Bai et al have proposed multiple atlas-based segmentation approaches that use FFD-based image registration for propagating labelled data to new cine-MR acquisitions. Whilst in [11] the similarity metric used for estimating the registration between the multi-atlas data and the new acquisition is the commonly utilised NMI metric, in [12] this NMI term is combined with a second metric that considers the consistency of the transformed labelled atlases, a strategy which they refer to as registration refinement. Bai et al have also used an FFD-based approach to create a bi-ventricular cardiac atlas from a large set of 3D cardiac MR datasets [10].

In [225], the transformation between atlas and subject is restricted to a locally affine approach - coined LARM (locally affine registration method). They choose this approach, due to the belief that this provides a degree of accuracy not available by utilisation of a solely global initial transformation, but avoids unrealistic topological changes that may result by direct application of a non-rigid approach. After the initial local transformations, which are specific to the different sub-structures of the heart, they refine using a non-rigid approach. The alignment between the subject cine-data and the atlas intensity image is quantified via NMI, and the final result is used to propagate the labelled atlas to the subject data. Shi et al [184] subsequently

used LARM as an initial step to segment data from different pathological subjects, with the addition of expectation maximisation methods for final refinement.

A comparative study assessing the ability of a number of registration methods for inter-subject cardiac MR registration was presented by Ou et al in [149]. The authors considered 12 popular registration algorithms, utilising a variety of similarity metrics, transformation models and optimisation strategies. Method evaluation was performed on cine-MR data where non-cardiac structures had been removed prior to registration, and algorithm parameters were adjusted such that the aggressiveness of each registration approach was comparable. They concluded from their results, reported as Jacard index overlap scores, that non-intensity-difference-based approaches were generally the better performing algorithms working more consistently; their results showed that their own approach, DRAMMS, which utilises an attribute vector in its similarity measure performed best.

Validation of registration

In order for newly developed algorithms to be accepted for clinical utilisation, validation is an essential component of method development. However, validation of medical image registration is difficult due to the lack of ground truth solutions. As a result, a number of different approaches are reported in the literature.

Without ground truths, a common approach for validation is to utilise the expert knowledge and experience of clinicians to define key regions and landmarks in the image data whose correspondence after registration can be used as a quantitative measure of algorithm performance. Depending on the nature of the annotations given, quantification may take the form of mask overlap scores, contour/surface distances or landmark registration errors.

In [154], where masks of the LV, RV and myocardium were available, qualitative evaluation of image alignment and temporal volume curve correspondences, were supported by quantitative results presented in the form of volume overlap measures as well as the distances measured between these surface representations.

In [183], to quantify how well cardiac motion was captured, manually delineated endo- and epi-cardial surfaces were propagated through the cardiac cycle and the resultant positions compared to surfaces manually defined in each phase. This quantification was coupled with a tracking of 16 landmarks identified based on the intersection of tag lines. The definition of landmarks and tissue surfaces is a subjective process meaning the final results may be influenced by the expert delineations provided; therefore, Shi et al went on to assess their results taking into account inter-observer variability - here their results showed that due to poor definition of the endocardial border a greater variability was observed in the quality of the endo-cardial tracking results as the user-defined surface was less consistently defined.

Many of the cardiac cine registration methods in the state-of-the-art are applied to the definition of cardiac motion; here researchers can utilise existing knowledge regarding the physiological behaviour of cardiac tissue during contraction and relaxation. This is a particularly attractive approach when validating methods on healthy volunteer data where the patterns of motion and strain are generally reproducible and well described. This is one approach taken by DeCrane et al in [46] where strain measurements derived from healthy volunteer data across basal and mid AHA segments were largely in agreement with those previously recorded in the literature.

In [200], where cardiac motion recovery results were presented using algorithms from a number of groups, the algorithms were compared and validated using a combination of qualitative image and mesh assessment, quantitative landmark tracking, and finally strain curve estimation. Here, the utilisation of a number of validation approaches was effective as the different measures used highlighted specific strengths and weaknesses in the algorithms.

4.2 Reconstruction

As previously discussed in section 3.2.1, one consequence of the limitations associated with MR data acquisition is that the obtained images are highly anisotropic,

and suffer from notable partial volume effect. While multi-oriented acquisitions compensate somewhat for the resultant loss of information, it often remains favourable during clinical post-acquisition analysis to deal with isotropic volumetric data representation. While acquisition advancements are consistently made that improve the resolution, much work has also been focused on the post-acquisition software side. In this section, we will initially discuss general methodologies for data up-sampling, before we give specific consideration to the methods proposed in the literature to generate isotropic high-resolution data from the low-resolution MR image acquisitions.

Intensity-based interpolation

Interpolation is a primary tool used for up-sampling low-resolution data, where new data points are estimated based on the small set of known values. Interpolation requires the estimation of a continuous function to model the relationship between input data points, from which an improvement in data resolution can be achieved by the sampling of this function at a finer spatial scale.

The observed result depends on the selection of interpolation model. Nearest neighbour is the most basic approach, where the intensity values of new data points are assigned according to the closest of the known input values. While preserving the intensity values provided by the original data, the resultant reconstruction has the most "blocky" appearance. An improvement can be seen when using a linear interpolation, where in the multi-dimensional case, the value of new data locations can be calculated as a linearly-weighted mean of neighbourhood points.

An improvement in the smoothness of the result, in regards to the continuity between neighbouring intervals in particular, is possible by increasing the degree of the interpolating polynomial. It is important to note that these are piecewise polynomials, defined within limited intervals of the total signal - the final smooth interpolating functions are referred to as splines. Cubic splines for example, ensure the continuity of both first and second derivatives - the acquired results will be smoother than those achieved using the piecewise linear approach above.

Object-based interpolation

The most notable drawback of intensity-based interpolation is the potential blurring of structural boundaries. It is proposed that such issues may be avoided by using object-based interpolation strategies. These methods utilise information regarding the objects present in the images to help guide the interpolation performed. Shape-based interpolation is an example of an object-based approach, and has been extensively used for the interpolation of binary segmentation masks [170]. The approach has since been generalised to allow for the shape-based interpolation of gray-level data [73].

Multiple-image super-resolution

Multiple-image super-resolution reconstruction is another category of techniques found in the literature. The basic premise of the approach is the combination of information available from several distinct low-resolution images of the same object. The underlying hypothesis is that each low resolution image captures a somewhat different field-of-view, and thus provides complementary but simultaneously unique information, from which a reconstruction can be generated that contains a greater proportion of the high frequency detail of the original scene. Such methods have proven successful in video-processing applications leading to their investigation in the context of medical data [71] ; a review of super-resolution approaches applied to MRI data is given in [204]

A key component of super-resolution methodologies is the observation model, which describes the relationship between the original high-resolution data and the generated low-resolution output images - for MR data this requires definition of the blurring and sampling operations as well as the addition of noise [53].

Machine learning

Here the objective is to use information contained in a large collection of high-resolution images to improve the resolution of similar images acquired at lower resolution. One such set of approaches are patch-based dictionary schemes. In these

methods, a dictionary is first learned which is representative of patches extracted from high-resolution training data. A sparse representation may then be estimated from this dictionary, that forms a higher quality final reconstruction for a new low-resolution input image [221]. More recently, deep learning approaches have been used; these methods do not require specific dictionary specification, rather multiple levels of convolutions are applied to training data, and the necessary operator weightings that give the best approximation of corresponding high-resolution output are estimated [48].

4.2.1 Volumetric reconstruction methodologies applied to CMR data

Here we will consider some 3D reconstruction methodologies applied to CMR data. In [68], Goshtasby et al propose to improve the accuracy of ventricular volume assessment by combining the information in both the LA and SA cine data acquired. A single volumetric representation is generated by first identifying the voxels of the volume occupied by the image data before filling in the empty voxels using a weighted average of the closest slice intensity values.

In addition to averaging approaches proposed for the combination of LA and SA data, super-resolution methods have been put forward that require the formulation of image acquistion models. Such a method is proposed in [167], where a maximum a posteriori framework is used; a high-resolution reconstruction solution is found that, under the application of the acquisition model, generates the closest approximations to the acquired SA and LA low-resolution data.

A similar super-resolution approach requiring the definition of an acquisition model is presented in [144]. In this case three intersecting image stacks in the SA, HLA and VLA orientations were available.

An alternative approach for reconstruction using only the SA data is proposed in [147], where a registration is performed between adjacent slices. The interpolation process performed is then based on both the calculated deformation field and the intensities in the registered images. This method is based on the assumption that adjacent slices contain consistent anatomical structures. Due to likely differences

in structural content between some adjacent cine images, the authors propose to apply the registration in both directions with the objective of making the result more robust.

In Shi et al [181] the authors propose a multi-atlas patch-based super-resolution approach. By combining patches from a number of high resolution atlas datasets, a new super-resolved dataset is formed that, under the application of an MR imaging model, optimally matches the acquired low resolution data of the new dataset.

In [20], the authors propose to improve the final volume resolution using data contained in a dictionary of cardiac MR patches. Specifically, two correlated dictionaries are learned containing paired high- and low-resolution data; the final reconstruction can be generated from the high-resolution dictionary based on the correspondences of the low-resolution data. More recently a deep learning approach was proposed [145] in the form of both a single-image and multi-image convolutional neural network, where the latter combined the resolution benefits of both the SA and LA acquisitions.

This work has been extended in [146] to include anatomical shape priors, which allows their method to better cope with artefacts such as slice misalignment and improve the results obtained for both segmentation and image enhancement operations.

Super-resolution reconstruction methods have also been applied to LGE images to improve the assessment of myocardial scar tissue; in [50], based on the multi-MR-slice super-resolution method described in [161], the authors combined SA, four-chamber, and two-chamber acquisition stacks to form an isotropically resolved volumetric representation of the infarct damaged LV. As well as showing clear qualitative improvements, tissue characterisation using their super-resolved volume resulted in smaller gray zones compared to those identified using standard SA interpolation methods; this was the result hypothesised and agreed with previous observations made in animal experiments [176] and electrophysiology voltage mapping [215].

In [65], the authors utilise a deep CNN architecture to reconstruct high resolution LGE data from a synthetically generated SA stack with highly anisotropic resolution. The results were compared to the ground truth high resolution 3D data acquired and

showed superior performance compared to standard linear and cubic interpolation, as well as a state-of-the-art patch-match algorithm.

4.3 Local Phase

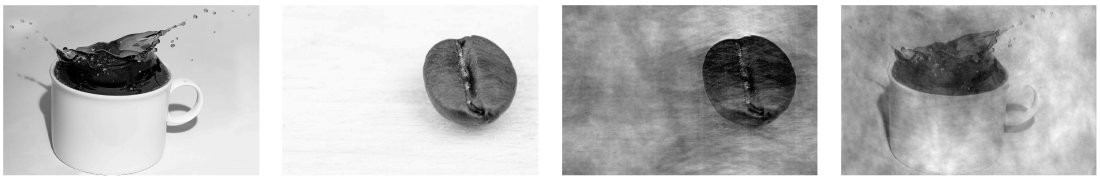
In this section, we introduce the concept of local phase as a descriptor of signal structure, to be used later as a means of assessing image alignment. The local phase property is particularly attractive in this case, as it is invariant to changes in contrast, which means that signal structures (such as lines and edges) have the same characteristic values regardless of the structure amplitudes.

The most frequent use of phase in signal/image processing is by means of Fourier phase, and is obtainable by consideration of the Fourier transform of a signal, $f(x)$. A signal's Fourier transform can be intuitively understood as its decomposition into a weighted sum of an infinite number of sinusoidal components. Expressed in its polar form, indicated in equation 4.24, the Fourier transform of a signal can be compactly described by two components - $F(\omega)$, the amplitude spectrum and $\theta(\omega)$, the phase spectrum. Respectively, they describe the relative contributions made by each frequency component and the zero-offset of each frequency's sinusoidal signal.

$$F[f(x)] = |F(\omega)|\, e^{j\theta(\omega)} \tag{4.24}$$

The significance of Fourier phase in signals has long been established. Oppenheim and Lim [148] observed that many of the important features in a signal - specifically the structural detail - are preserved, even if only the phase of a signal is retained. This is classically exemplified via the inverse Fourier reconstruction of two images. By starting with the Fourier transforms of two images, and substituting their respective phase and magnitude spectra, two new images are generated. The results are given in figure 4.1 and demonstrate that, in both cases, the source image that dominates the output image is that which provided the phase quantity.

The application of the Fourier transform to signals has many associated merits; however, there is a fundamental limitation related to the global nature of its definition. The Fourier transform may tell us which frequency components are



(a) Image A **(b)** Image B **(c)** Magnitude of image A, phase of image B **(d)** Magnitude of image B, phase of image A

Figure 4.1: Images resulting from the substitution of the respective magnitude- and phase-spectra of two images. Original images freely available from https://www.pexels.com/.

present in a signal and their associated amplitudes and phases, but it gives no indication as to where in the signal they are located.

The non-stationary nature of real signals necessitates a definition of variables that can describe the signal locally. If we are to use phase as a feature descriptor for spatial alignment, such a means of localisation is an essential pre-requisite. Alternative phase-based quantities have been developed for this purpose; however, before they can be discussed, it is necessary to introduce the concepts of the Hilbert transform and the analytic signal.

Hilbert transform

The Hilbert transform is defined in the spatial domain by a convolution of the form shown in equation 4.25. Given that a convolution in the spatial domain corresponds to multiplication in the Fourier domain, its frequency domain form can be written as in equation 4.26.

$$\begin{aligned} f_H(x) = H\left[f(x)\right] &= \frac{1}{\pi}\int_{-\infty}^{\infty}\frac{f(\tau)}{\tau - x}d\tau \\ &= f(x) * \frac{1}{\pi x} \end{aligned} \tag{4.25}$$

$$\begin{aligned} F[f_H(x)] &= F\left[f(x) * \frac{1}{\pi x}\right] \\ &= F\left[f(x)\right] F\left[\frac{1}{\pi x}\right] \\ &= -i \cdot sign(u) \cdot F(u) \end{aligned} \tag{4.26}$$

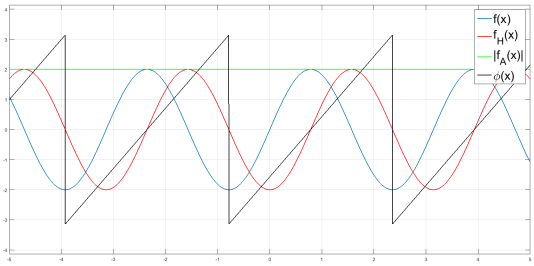


Figure 4.2: Sine wave, $f(x)$, and its: Hilbert transform, $f_H(x)$; instantaneous amplitude, $|f_A(x)|$; instantaneous phase, $\phi(x)$.

where

$$sign(u) = \begin{cases} 1, & \text{if } u > 0. \\ 0, & \text{if } u = 0. \\ -1, & \text{if } u < 0. \end{cases}$$

For the sake of subsequent discussion, the most important observation to be made regarding the Hilbert transform is its effect on signal phase. Positive spectral frequencies see a negative phase shift of $-\frac{\pi}{2}$ ($-i = e^{\frac{-i\pi}{2}}$), while negative frequencies see a positive phase shift of $\frac{\pi}{2}$ ($i = e^{\frac{i\pi}{2}}$). This phase shifting effect can be observed in the spatial domain in figure 4.2. A simple periodic function and its Hilbert transform are shown - we see that they lie out-of-phase by $\frac{\pi}{2}$. This effect is general for all signals; therefore, we can define the Hilbert transform as having a $\frac{\pi}{2}$ phase shifting operation. The original signal and its Hilbert transform are commonly referred to as a quadrature filter pair.

Analytic signal

The analytic signal, first introduced by Gabor in [63], is a complex signal formed by the combination of $f(x)$, the original signal, and its phase-shifted, quadrature form as an imaginary component. As was previously defined, this phase-shifted form is given by the Hilbert transform; therefore, the analytic signal is defined as in equation 4.27. In equation 4.28, the corresponding frequency domain definition is given - we note the operation of the sign function results in a single-sided frequency spectrum - all negative frequency components have been suppressed. However, due to the Hermitian nature of the frequency spectrum of real signals, this does

not result in any loss of information.

$$f_A(x) = f(x) + if_H(x) \tag{4.27}$$

$$\begin{aligned} F[f_A(x)] &= F(u) + i \cdot (-i \cdot sign(u) \cdot F(u)) \\ &= F(u)\,(1 + sign(u)) \end{aligned} \tag{4.28}$$

Significantly, the complex form of the analytic signal facilitates the calculation of a number of meaningful signal quantities. When considered in its polar form, the analytic signal can be expressed as shown in equation 4.29.

$$f_A(x) = |f_A(x)|\, e^{j\phi(x)} \tag{4.29}$$

We saw in equation 4.24 the definition of two quantities, the Fourier amplitude, and Fourier phase; here, in equation 4.29, we have defined their spatial counterparts, referred to as the instantaneous amplitude, $|f_A(x)|$, and the instantaneous phase, $\phi(x)$. Instantaneous amplitude provides a measure of the energy at a given location in the signal, and can be calculated from the complex components of the analytic signal via equation 4.30. The instantaneous phase, $\phi(x)$, is calculated via the argument of the real and imaginary parts of the analytic signal, as defined in equation 4.31, and gives an indication of the type of structure present at each point in the signal - specifically, it reflects the local symmetries of the signal.

$$|f_A(x)| = \sqrt{Re(f_A(x))^2 + Im(f_A(x))^2} \tag{4.30}$$

$$\begin{aligned} \phi(x) &= arg(f_A(x)) \\ &= tan^{-1}\left(\frac{Im(f_A(x))}{Re(f_A(x))}\right) \end{aligned} \tag{4.31}$$

This division of independent signal characteristics is commonly referred to as the split of identity. There are two important properties that relate to this split of identity property, those being invariance and equivariance. The instantaneous phase estimates are said to be invariant to signal energy - meaning that, regardless of the magnitude of the signal change, the phase value will be the same. Similarly, the instantaneous amplitude is invariant to the signal structural change, reflecting only the energy in the signal. The instantaneous phase is also said to be equivariant

with spatial position - this is observed in figure 4.2 (except for locations of phase wrap-around).

We defined the instantaneous phase as a descriptor of signal structure according to the underlying symmetric properties of the signal. To illustrate this, we consider the components of the analytic signal for the two extreme cases of perfectly-symmetric and perfectly-antisymmetric structures. We first note that the Fourier transform of a symmetric structure is purely real-valued (requiring only cosine terms), while that of an antisymmetric structure is purely imaginary (requiring only the sine terms). We saw previously that the operation of the Hilbert transform in the Fourier domain is a multiplication of the Fourier transform of the source signal by the imaginary j unit. Therefore, under the operation of the Hilbert transform, the purely real Fourier representation of a symmetric signal is converted into an exclusively imaginary representation - equivalent to an antisymmetric spatial form. The phase is calculated by the argument of these two terms (see equations 4.31) and, therefore, a result of either 0 or $\pm\pi$, is given, depending on the sign of the symmetric structure. Following the same argument, an antisymmetric structure in the original signal is paired with a symmetric Hilbert component, giving phase values of $\pm\frac{\pi}{2}$. These results are summarised in figure 4.3 and can similarly be observed within the phase relationship of the signal in figure 4.2.

Knutsson H, Westin CF. Monomial phase: A matrix representation of local phase. InVisualization and Processing of Tensors and Higher Order Descriptors for Multi-Valued Data 2014 (pp. 37-73). Springer, Berlin, Heidelberg.

The introduction of the analytic signal has enabled the localisation of amplitude and phase quantification; however, to this point we have considered only the analysis of simple signals with single-frequency spectra. Real signal composition is rarely this basic, rather being formed from a wide range of frequency signals of different amplitudes and phases. In such cases, the meaningful correlation between the instantaneous values and the local signal structure is lost [69]. We can see this in the signal in figure 4.4: the instantaneous amplitude and phase values give little insight into the underlying signal from which they have been calculated.

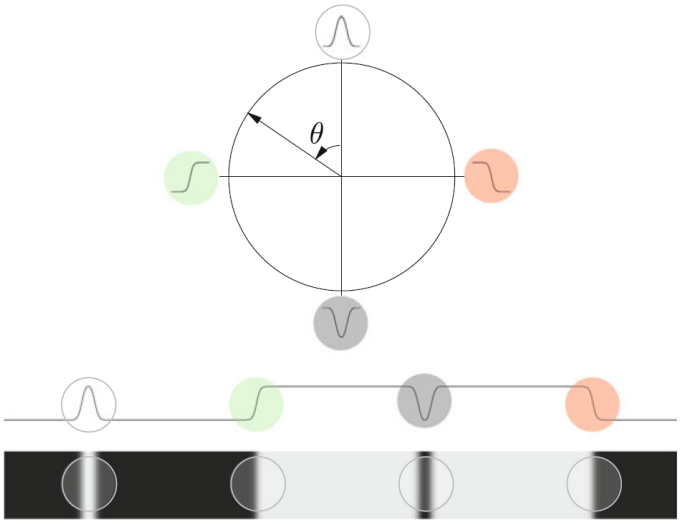


Figure 4.3: Signal structures as they relate to local phase angles [93]. Reprinted by permission from Springer Customer Service Centre GmbH: Springer Nature, Springer eBook, Monomial phase: A matrix representation of local phase, Knutsson H, Westin CF. ©Springer-Verlag Berlin Heidelberg 2014 (2014).

Therefore, the above analysis only holds, if we can decompose our signal into partial, component signals that lie within narrow frequency bands.

Theoretically, to retrieve each partial signal we would require infinitely narrow bandwidth components. However, in the spatial domain, this would require filters with global support, meaning we would again lose any sense of signal localisation. We are familiar with this inverse relationship between spatial and frequency domain signal widths - for the calculation of local phase, we need both frequency and spatial localisation, requiring therefore a trade-off that is satisfactory in both domains. Heisenberg's principle theorises that the uncertainty linked with the

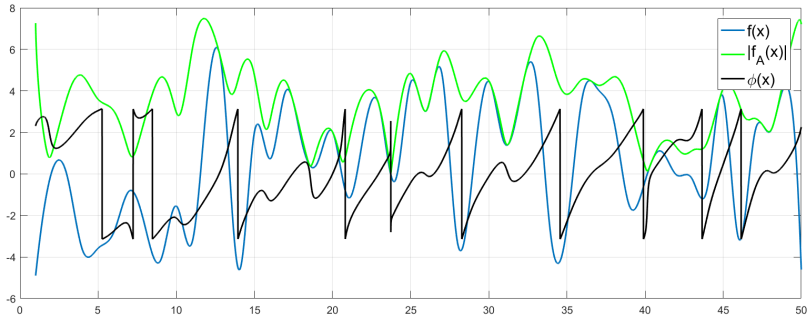


Figure 4.4: A signal composed of a wide frequency spectrum along with its corresponding instantaneous-amplitude, $|f_A(x)|$, and instantaneous-phase, $\phi(x)$.

trade-off between two complementary physical variables such as these is bounded according to the relationship in equation 4.32.

$$\sigma_x \sigma_u \geq \frac{1}{4} \tag{4.32}$$

In addition to the trade-off between spatial and frequency localisation, there are two further requirements of an appropriate filter: firstly, it must be invariant to a global additive shift in the signal, therefore requiring a zero DC component; and secondly, it must have a zero phase-effect on the signal, thus necessitating a symmetric filter. As well as providing an adjustable frequency bandwidth, both of these further requirements are satisfied by appropriate selection of bandpass filters. The above analysis applied to a bandpass-filtered signal now provides a split of identity at a particular spatial scale and the local signal structure is now described according to the two variables referred to as local amplitude and local phase. The analytic signal can be re-written, as in equation 4.33, to incorporate the application of the bandpass filter, *bp*. We see that applying the Hilbert transform to the bandpass filter, and subsequently applying this filter to the original signal, is equivalent to Hilbert transforming the bandpass signal. The bandpass filter and its Hilbert transform are in quadrature; therefore, it is common in the literature to refer to pairs of filters such as these as quadrature filters.

$$\begin{aligned} f_A &= bp(x) * f(x) + iH\left[bp(x) * f(x)\right] \\ &= \left[bp(x) + ibp_H(x)\right] * f(x) \end{aligned} \tag{4.33}$$

While there are a number of possible bandpass quadrature filters here we present only log-Gabor filters - a discussion of respective merits of a range of other filter choices has been presented in [22].

The aim when selecting a bandpass filter is to obtain an acceptable compromise between spatial and frequency uncertainty - we wish to get as close to the equality noted in Heisenberg's principle as possible. The Gabor filter has been shown to be optimal in this respect; however, it does not fulfil all the necessary requirements to be used as a quadrature filter. Specifically, it is neither guaranteed a zero negative frequency spectrum nor a zero DC component, both factors being dependent on

the selected bandwidth. An alternative selection, therefore, is called the log-Gabor filter. Where Gabor filters are characterised by a Gaussian-shaped profile in linear frequency, log-Gabor filters are so formed in the log frequency scale. The key advantage of log-Gabor filters is that, regardless of filter bandwidth, both the DC value and the negative frequency components are zero.

The log-Gabor filter, LG, is defined in the frequency domain by equation 4.34. Here, ω_0 is the centre frequency and k is a parameter, in the range 0 to 1, that is related to the bandwidth and thus controls the filter shape.

$$LG(\omega) = exp\left(-\frac{log^2\left(\frac{\omega}{\omega_0}\right)}{2log^2(k)}\right) \tag{4.34}$$

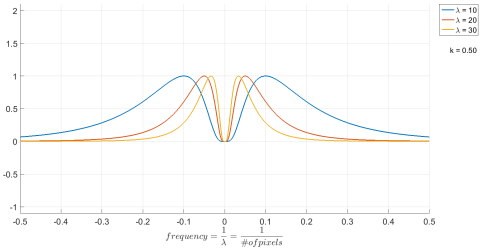
Log-Gabor filters are shown in figures 4.5 and 4.6. The first figure depicts the frequency spectra of filters formed at a range of centre frequencies; in a) we see the bandpass filter spectrum, in b) the Hilbert transform spectrum, and in c) the resultant quadrature filter spectrum - as expected, all negative frequency components have been suppressed. The equivalent plots in the second figure indicate the variation in filter bandwidth with changing k value.

4.3.1 Local phase in nD

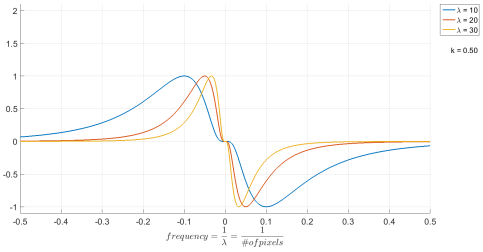
The theoretical description outlined applies to 1D signals. Later in the thesis, we deal with higher dimensional signals which require a generalisation of the above analysis. It has been acknowledged in the literature that this generalisation is not straightforward, as there is no clear nD equivalent to the Hilbert transform or analytic signal.

Previous methods for generating higher-dimensional equivalents of the analytic signal include:

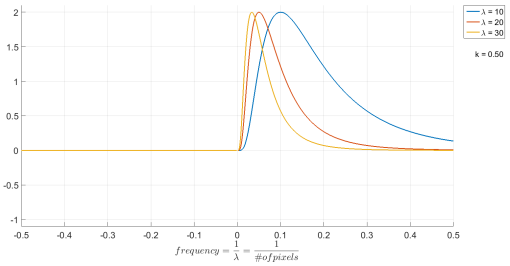
- the partial Hilbert transform [69] - which is entirely dependent on the selection of an axis along which to apply the transform
- the total Hilbert transform [189] - which does not retain the necessary $\frac{\pi}{2}$ phase-shifting property of the Hilbert transform



(a) log-Gabor filter spectra

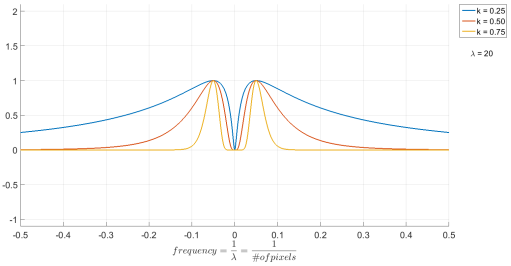


(b) Hilbert transformed log-Gabor filter spectra

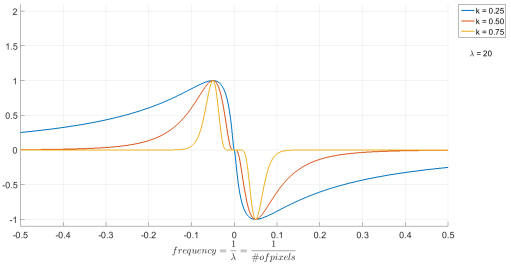


(c) Quadrature filter spectra

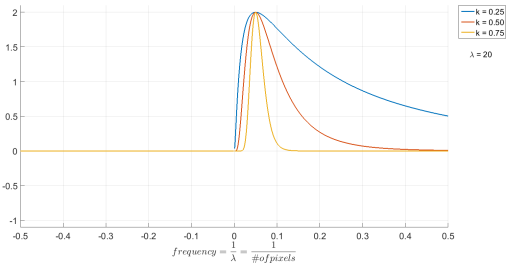
Figure 4.5: Quadrature filter construction in the frequency domain at varying central frequency ($\omega_0 = 1/\lambda_0$) values.



(a) log-Gabor filter spectra

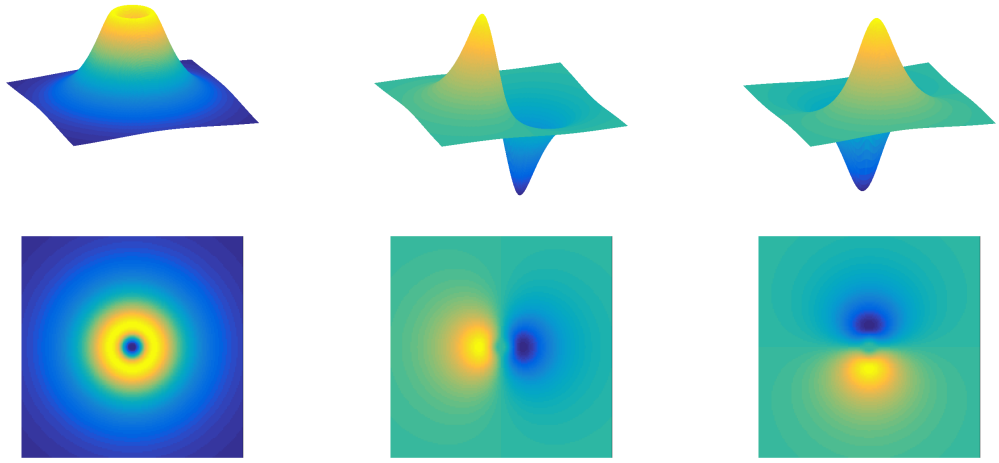


(b) Hilbert transformed log-Gabor filter spectra



(c) Quadrature filter spectra

Figure 4.6: Quadrature filter construction in the frequency domain at varying k values.



(a) Log-Gabor filter spectrum

(b) First Riesz component

(c) Second Riesz component

Figure 4.7: 2D monogenic filter construction in the frequency domain.

- the single orthant analytic approach [76] - which requires multiple analytic signals to adequately represent the signal information
- the hypercomplex signal [28] - which, despite solving many of the above issues, is still not an isotropic transform and as such does not perform satisfactorily under signal rotation.

The monogenic signal

An approach which has gained wide-ranging acceptance, and is thus a commonly-used method in the literature for multi-dimensional local phase calculation, is the monogenic signal [57]. The monogenic signal replaces the Hilbert transform with its vector-valued generalisation, the Riesz transform. The frequency domain definition of the Riesz transform is given in equation 4.35, where in the 2D case $\mathbf{u} = [u, v]$. The monogenic signal is composed of $n + 1$ components, where n represents the dimensionality of the signal in question. Analogously to the analytic signal, the first element of the monogenic signal is the original bandpassed-filtered signal, while, for a 2D image, the second and third elements are given by the application of the Riesz transform to the bandpass-filtered signal. For the analytic signal, we described the possibility of applying the Hilbert transform to the bandpass filter before filtering the image; here, we apply the Riesz transform and generate a pair of directional filters.

These, in addition to the original bandpass filter, are depicted in figure 4.7 - the individual application of all three filters generates the monogenic vector components.

$$I_R(u) = \mathbf{R}[I(\mathbf{u})] = i\frac{\mathbf{u}}{\|\mathbf{u}\|}I_{bp}(\mathbf{u}) \tag{4.35}$$

The spatial domain monogenic signal is defined in vector form in equation 4.36. From this definition the local signal properties - amplitude, phase, and orientation - can be calculated according to equations 4.37, 4.38, and 4.39 respectively.

$$\begin{aligned} I_{monogenic}(x,y) &= \Big[I_{bp}(x,y), \mathbf{I}_R(x,y)\Big] \\ &= \Big[I_{bp}(x,y), I_{R_1}(x,y), I_{R_2}(x,y)\Big] \end{aligned} \tag{4.36}$$

$$A(x,y) = \sqrt{I_{bp}^2(x,y) + I_{R_1}^2(x,y) + I_{R_2}^2(x,y)} \tag{4.37}$$

$$\phi(x,y) = tan^{-1}\left(\frac{\sqrt{I_{R_1}^2(x,y) + I_{R_2}^2(x,y)}}{I_{bp}(x,y)}\right) \tag{4.38}$$

$$\theta(x,y) = tan^{-1}\left(\frac{I_{R_2}(x,y)}{I_{R_1}(x,y)}\right) \tag{4.39}$$

We note that the calculation of local amplitude follows that defined for the 1D analytic case, being equal to the magnitude of the filter outputs. However, the calculation of phase is not so clearly extended; the filter outputs can no longer be simply combined into a complex form from which the phase is given by its argument. Rather, the monogenic representation, in the case of a 2D signal, is a 3D vector, the geometric interpretation of which is illustrated in figure 4.8. Here, the vector must be described by two angles that are defined as the local phase, ϕ, and local orientation, θ. The symmetrical characteristics of image structures are again described by the local phase variable, however the multi-dimensionality of the signal leads to an ambiguity relating to the orientation of structures; the local phase is insufficient to describe the primary orientation of signal change. An estimation of the maximal, intrinsically one-dimensional, signal change direction is given by the local orientation variable.

These signal descriptors can be combined in the form of the phase vector, defined in equation 4.40, which is the local phase value multiplied by the unit

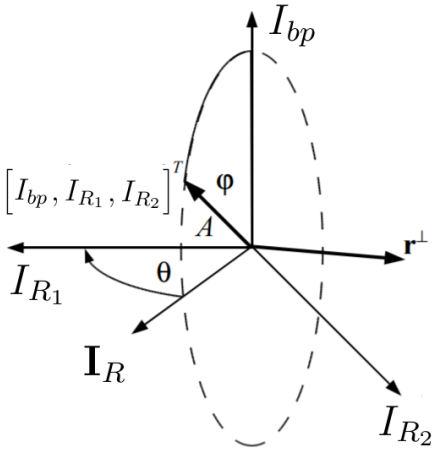


Figure 4.8: The monogenic phase description of a 2D signal, characterised as a vector in 3D space [56]. Adapted by permission from Springer Customer Service Centre GmbH: Springer Nature, Springer eBook, Optical Flow Estimation from Monogenic Phase, Felsberg M. ©Springer Berlin Heidelberg 2007 (2007).

vector associated with the local signal orientation, and which provides a more complete illustration of the local structure.

$$\begin{aligned} r &= \phi(x) \begin{bmatrix} cos(\theta(x)) \\ sin(\theta(x)) \end{bmatrix} \\ &== \phi(x) \begin{bmatrix} \frac{I_{R1}}{\sqrt{I_{R1}^2+I_{R2}^2}} \\ \frac{I_{R2}}{\sqrt{I_{R1}^2+I_{R2}^2}} \end{bmatrix} \\ &== \phi(x) \frac{\mathbf{I}_R}{|\mathbf{I}_R|} \end{aligned} \tag{4.40}$$

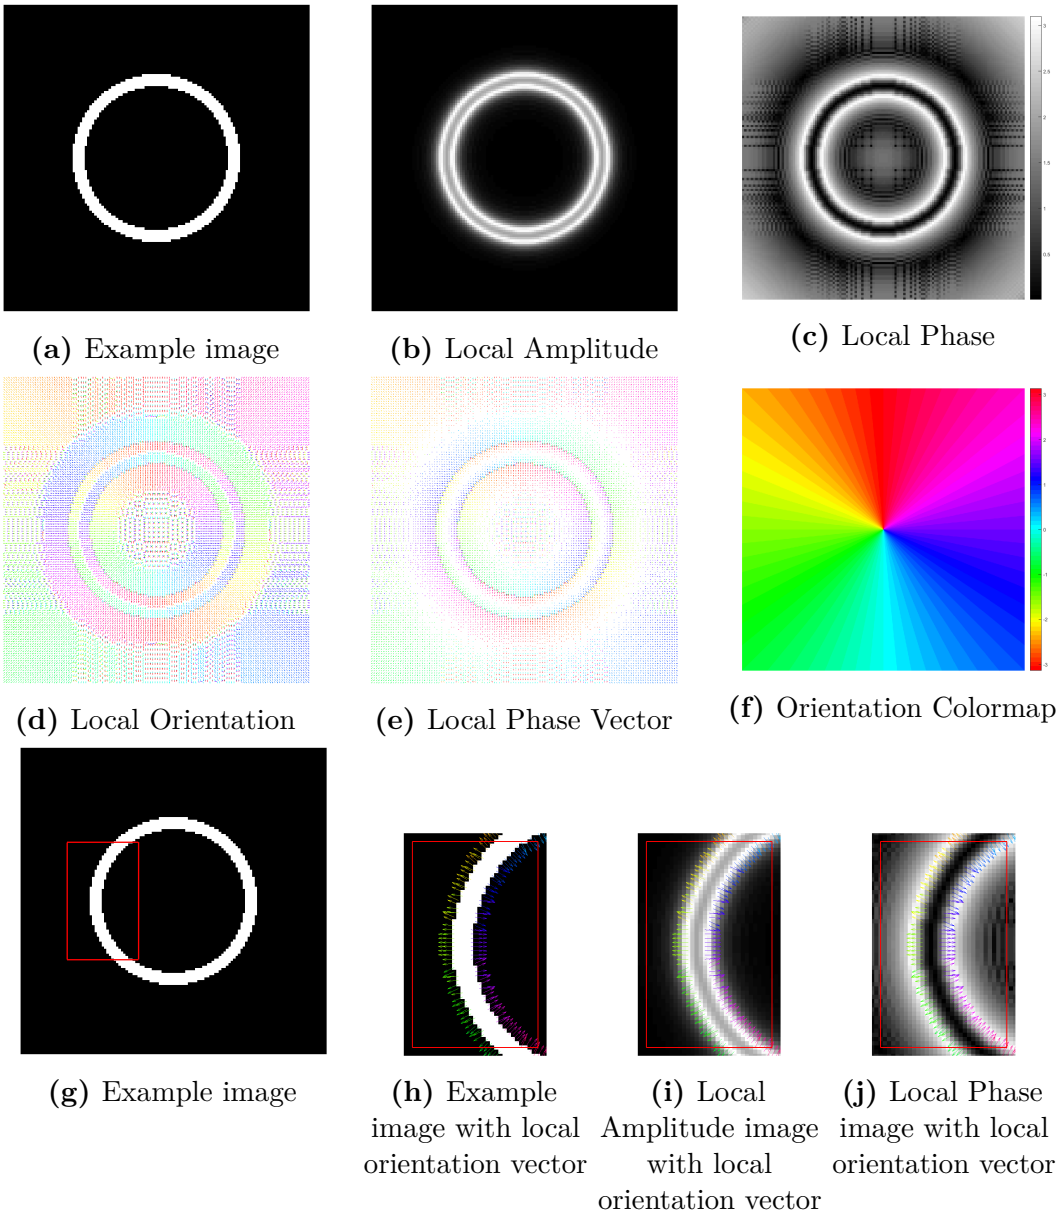
The description of the monogenic signal is given here for 2D signals. Later the concepts will be used for both 2D and 3D data. The extensions to 3D are straightforward, requiring an extra component to the Riesz transform (oriented along the third frequency domain axis), and an extra angle for the definition of the local orientation.

In figures 4.9 and 4.10 the phase-based image features described above are illustrated. In particular, figures 4.9j and 4.10j illustrate the important additional information provided by the calculated local orientation compared to the local phase information alone.

Effect of parameters on local phase estimation in 2D

The local signal properties calculated are primarily dependent on the selection of two parameters; the log-Gabor filter's central frequency and bandwidth (see equation 4.34). As was described for 1D signals, changing these parameters alters the frequency band to which a filter responds.

The central frequency of the filter determines the scale of structures identified in the signal; increasing the frequency (decreasing the wavelength) corresponds to



(a) Example image **(b)** Local Amplitude **(c)** Local Phase

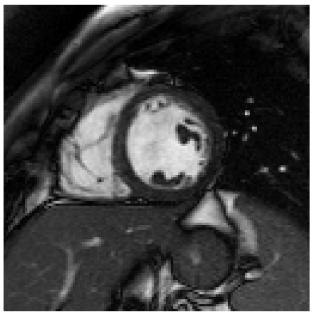
(d) Local Orientation **(e)** Local Phase Vector **(f)** Orientation Colormap

(g) Example image **(h)** Example image with local orientation vector **(i)** Local Amplitude image with local orientation vector **(j)** Local Phase image with local orientation vector

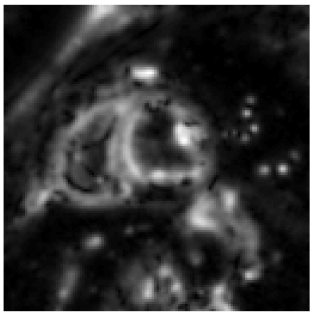
Figure 4.9: Local features of a simple example image.

a decrease in the spatial scale of structures identified. This is clearly illustrated in figure 4.11 where, as the frequency decreases with each row, the structures that can be discerned increase in size.

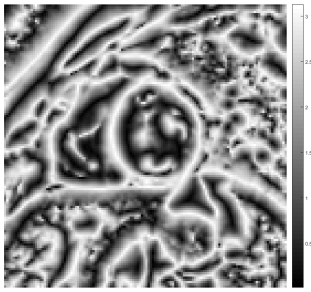
In the same figure, we see the variation with changes in filter bandwidth. The bandwidth determines the spread of frequencies to which the filter responds. Narrower bandwidths, which correspond to larger values of k, result in increasingly periodic phase representation as can be seen towards the left-hand column of figure 4.11. Conversely, with rising bandwidth the phase images increasingly resemble



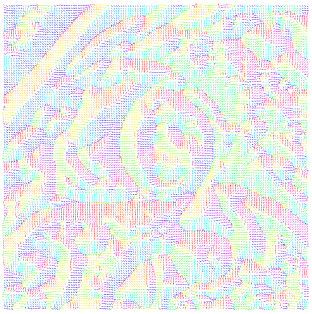
(a) Cine SA image



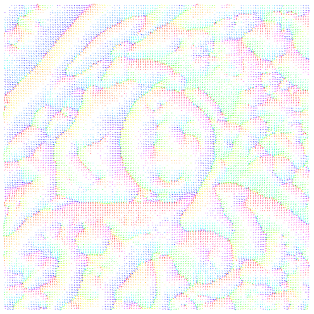
(b) Local Amplitude



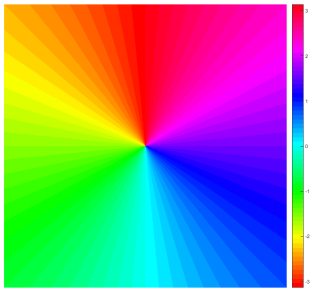
(c) Local Phase



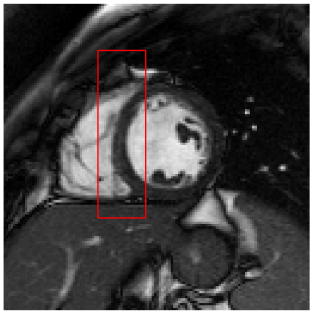
(d) Local Orientation



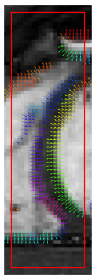
(e) Local Phase Vector



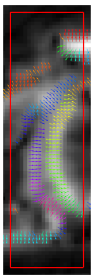
(f) Orientation Colormap



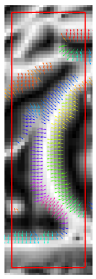
(g) Cine image



(h) Cine image with local orientation vector



(i) Local Amplitude with local orientation vector



(j) Local Phase with local orientation vector

Figure 4.10: Local features of a cardiac cine MR image.

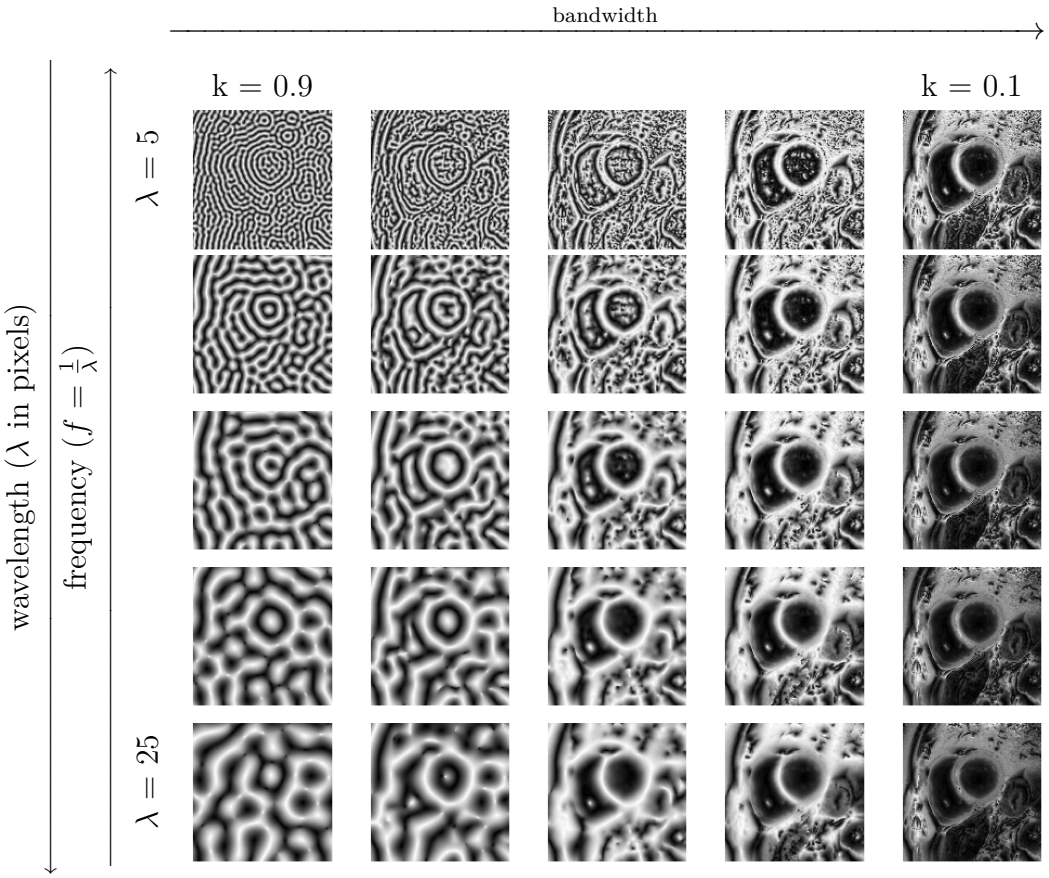


Figure 4.11: Local phase variation with centre frequency and bandwidth.

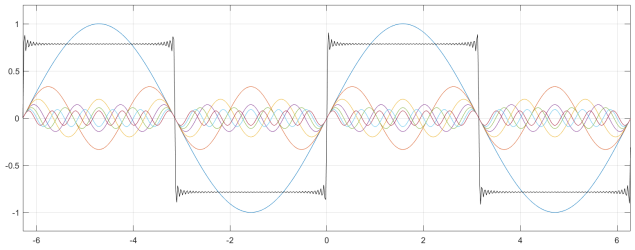
the structures present in the original data as we add frequency components to the representation. It is observable from this figure that, as the bandwidth is increased, the phase image representations become less dependent on the selected scale.

4.3.2 Further phase-based properties

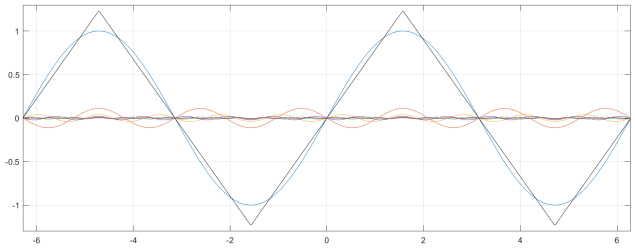
Some further insight can be gained regarding signal structure, if we compare local phase estimates calculated at different frequency bands.

Phase congruency

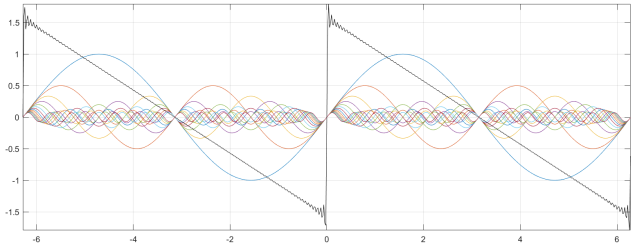
Phase congruency considers the local phase values calculated at different scales. It is a measure that identifies locations at which the local phase values are constant over a wide range of frequencies. It is observable that at salient locations in a signal, such as steps and ridges, the Fourier components that form the signal are all in phase. This is demonstrated if we consider the Fourier components that form the features



(a) Fourier components that form a square wave signal



(b) Fourier components that form a triangular signal



(c) Fourier components that form a sawtooth signal

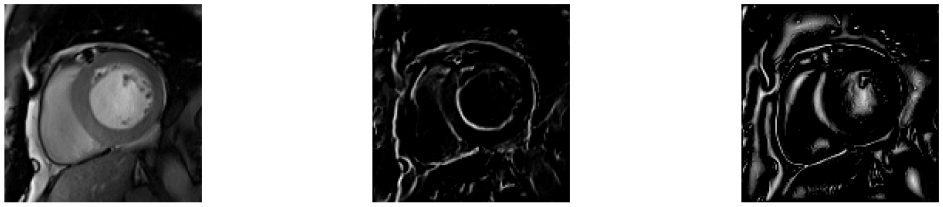
Figure 4.12: The frequency components of the square-, triangular-, and sawtooth-signals which illustrate the phase-congruency, -symmetry and -asymmetry properties.

in signals such as the square wave (figure 4.12a), the triangular wave (figure 4.12b), and the sawtooth signal (figure 4.12c) - in each, we can see that the edges and ridges are located where the phase values are concurrent for each frequency component.

The phase congruency as a feature detector was originally algorithmically proposed in [141]. The measure, which falls within the range 0-1, is defined in equation 4.41, where n frequency components are considered, A defines the amplitude, and $\overline{\phi}$ indicates the mean phase. The measure has subsequently been modified, as detailed in [97], to improve localisation and robustness to noise.

$$PC(x) = \max_{\overline{\phi} \in [0,2\pi]} \frac{\sum_n A_n cos(\phi_n(x) - \overline{\phi})}{\sum_n A_n(x)} \tag{4.41}$$

The phase congruency concept can be extended to higher dimensions and an



(a) **(b)** Phase congruency **(c)** Phase symmetry

Figure 4.13: The phase congruency and phase symmetry calculated for an SA cine MR image.

example result is shown in figure 4.13b.

Local phase symmetry and asymmetry

A second observation that can be made from consideration of the signal's Fourier make-up is related to the symmetric or asymmetric nature of the structure. Symmetric and asymmetric structures appear in a signal at those locations where all the constituent Fourier components are also symmetric or asymmetric respectively. This can again be observed in the plots of figure 4.12.

Symmetry and asymmetry metrics, originally defined in [96], can be calculated given equations 4.42 and 4.43, where the additional parameters T and ϵ are used for noise compensation and to avoid division by zero respectively.

$$Symmetry(x) = \frac{\sum_n \lfloor |cos(\phi_n(x))| - |sin(\phi_n(x))| - T \rfloor}{\sum_n A_n(x) + \epsilon} \tag{4.42}$$

$$Asymmetry(x) = \frac{\sum_n \lfloor |sin(\phi_n(x))| - |cos(\phi_n(x))| - T \rfloor}{\sum_n A_n(x) + \epsilon} \tag{4.43}$$

The above definitions stem from the observation that, as displayed in figure 4.3, symmetric structures have local phase values of 0 and π, while asymmetric structures are characterised by values of $-\frac{\pi}{2}$ and $\frac{\pi}{2}$; therefore the respective metrics above will be maximised accordingly. Again, application of this analysis can be seen in figure 4.13c.

4.3.3 Applications in image registration

Some of the properties of local phase - invariance to contrast change; equivariance to spatial position; confinement to values within the range $-\pi$ and π - have made

it an attractive property for image processing tasks.

Seminal work on the estimation of motion using phase was presented by Fleet and Jepson in [61]. Their work analysed the temporal evolution of contours of constant phase to define oriented component velocities, which are subsequently used to define a full 2D image velocity. Their formulation is equivalent to more classical, differential optical flow techniques, but with intensity values replaced by local phase estimates, as indicated in equation 4.44, where the subindex indicates the phase estimated from a collection of n oriented filters. In [16], this method was compared to other optical flow techniques, with Fleet and Jepson's approach proving to be amongst the most accurate.

$$\phi_{x,n} v_{x,n} + \phi_{y,n} v_{y,n} + \phi_{t,n} = 0 \tag{4.44}$$

Hemmendorff et al [79] similarly applied a bank of quadrature filters to multidimensional signals, and applied a differential solution approach. However, in contrast to Fleet and Jepson, who generate multiple vector fields from each individual filter output, Hemmendorff's 9 filter orientations were embedded into the single problem definition, and a single 2D motion field is calculated. Their formulation provides the flexibility to use a number of different motion models; affine, locally-affine and finite element models were used for multiple medical applications in their presented results. Specifically, they considered the registration of images acquired before and after contrast injection; it is for this reason that they were drawn to the use of local phase, which is invariant to these changes in illumination. Their final presented result shows the qualitatively-improved output of their phase-based approach compared to the classical intensity-based optical flow.

The issue of contrast injection-induced intensity variation was also tackled by the use of a local phase-based representation in [197]. The objective of analysing the perfusion of cardiac tissue is hindered by both cardiac and breathing motion. Quadrature filters were applied, and the deformation between frames calculated for each orientation. An estimate of the deformation field is calculated, based on a weighted sum of the component deformations. A smoothing is subsequently

applied to the deformation field. Their approach was shown to perform better than an NMI-based registration approach.

The application of oriented filter banks has been replaced in many applications by the use of the monogenic signal. Using the monogenic signal, it is possible to accurately define structures both according to type and orientation within the image domain. Felsberg demonstrated the use of what he called the local phase vector within an optical flow framework on synthetic image examples with known ground truths in [56]. This formulation has since been utilised on real data - specifically in multiple cardiac applications. Alessandrini demonstrated its effectiveness in capturing cardiac motion in both echocardiography and tagged MRI for the subsequent calculation of myocardial strain [4, 3]. Both applications suffer from intensity issues: ultrasound image intensities vary as the acoustic angle changes in relation to the tissue, while the tag lines of interest in tagged MRI fade with time - both issues that can be overcome with the use of local phase. Cardiac motion from cine MRI is analysed using a local phase-based optical flow method in [219]. In this paper, the authors revert to using oriented filters for local phase estimates - complex Gabor filters, as detailed in [27], are used and embedded into the optical flow constraint equation. They calculate a deformation based on the Lucas-Kanade local-constraint approach. Their results compare favourably to an implementation of an intensity-based Horn-Schunk approach.

An alternative to directly comparing local phase values was proposed by Mellor and Brady [135, 136]. They suggest the calculation of the MI of phase. They argue that, despite the wide acceptance of MI for multi-modal registration, there are notable drawbacks - specifically, they refer to situations where a non-rigid approach might fail because the perceived intensity relationship between corresponding tissues is in fact poorly defined. They provide a clinical example of this scenario, where reflections from brain ventricle walls in ultrasound images are incorrectly aligned to MR images using an intensity-based approach, while the MI of a phase-based representation recognises the presence of the correct edges in both images, and is not misled by the reflection signal. The key argument is that MI of local phase

searches for relationships between image features rather than image intensities. Their results proved that, as well as having an improved performance in difficult scenarios described above, the phase-MI was able to match the classical MI approach in more standard situations, where intensity-MI would previously have been the metric of choice.

We have already noted the use of local phase in a number of ultrasound applications. This is motivated by the difficult characteristics of ultrasound images - they have a speckled appearance, have low contrast, and their signal to noise ratio is poor. In addition, intensity variations occur as a result of different imaging angles and tissue depths. Therefore, it is commonly accepted that intensity-based processing methods do not perform well. In contrast, it has been shown that local phase methods are effective, enabling the identification of important features in the ultrasound images. As a result, in addition to the above applications, local phase-based methods have been used to register multiple ultrasound imaging planes [70], to align ultrasound and MR data [223], and to derive cardiac motion from ultrasound image frames [198].

5

Image Registration of Cardiac Datasets

Contents

5.1 Introduction 105
5.1.1 Materials 107
5.2 Slice Realignment 107
5.2.1 Methods 107
5.2.2 Experiments and results 114
5.2.3 Discussion 126
5.3 LGE to Cine Registration 129
5.3.1 Methods 131
5.3.2 Experiments and results 133
5.3.3 Discussion 139
5.4 3D CMR Registration - Longitudinal Analysis 142
5.4.1 Methods 142
5.4.2 Experiments and results 150
5.4.3 Discussion 169

5.1 Introduction

In this chapter, a summary is given of the different image registration challenges addressed in order to perform the longitudinal assessment of multi-sequence CMR imaging data.

The main contributions presented are:

- a full 3D rigid slice realignment algorithm for correcting breath-hold related errors. The method proposed evaluates the current slice alignment based on the correspondences of SA-LA intersection profiles as well as the correspondences of adjacent SA slices. These similarity assessments are based on both the images' local phase information, calculated at a number of bandwidths, and the original intensity-based information. For the purposes of result assessment a GUI was designed with interactive slice, contour, and intersection profile viewing capabilities.

- a 3D rigid phase-based slice-to-volume registration between the LGE slices and a cine volume, reconstructed from the SA cine acquisition data. The method proposed compensates for misalignments resulting from breath-hold variation and patient movement by identifying the optimal correlations between local phase-based image representations.

- the longitudinal 3D non-rigid registration of cardiac MR studies using a phase-based metric. A two-step framework is presented, where both the rigid and non-rigid stages utilise phase-based data representations. The non-rigid registration is a multi-resolution, multi-scale approach in which each sequential stage uses smaller local phase scales to refine the deformation calculated. An initial study presenting the application of a phase-based metric to the registration of follow-up cardiac cine data was published at the Computing in Cardiology conference 2014 [88].

- the utilisation of the deformation field result for the co-registration of multiple imaging modalities; we demonstrate application to LGE images for the longitudinal analysis of scar progression in patients with myocardial infarction

Based on results presented in this chapter, a paper is currently under preparation - 'A local phase-based registration of multi-sequence CMR scans for the evaluation of longitudinal tissue changes in myocardial infarction'.

5.1.1 Materials

Before presenting the image registration methods, a brief recap is given of the data used for the results presented (previously discussed in chapter 3). Data was available for 19 patients, each of whom were imaged on two occasions - at the acute stage and a 6 month follow-up assessment. Results are presented for both cine and LGE studies which are composed of an SA stack and a number of intersecting LA slices. As indicated previously in section 3.1, all slices were cropped before processing such that they contain only the heart and close surrounding structures. Endo- and epi-cardial delineations were provided and used to form mesh-based surface representation; in combination, the delineations and meshes are used for multiple validation purposes.

5.2 Slice Realignment

A typical cardiac MR dataset is composed of a number of 2D planar images acquired during multiple patient breath-hold periods. Individually, these images enable the 2D analysis of cardiac tissue state; however, a more comprehensive analysis requires 3D integration of the complementary data. This is made difficult due to differences in lung volume during successive breath-hold acquisitions, leading to variation in tissue position relative to the scanner coordinates.

In the next section, we will propose an automatic method for correcting these slice misalignments driven primarily by the correspondences of image intersection profiles. A full 3D transformation is used, allowing for correction of both in-plane and out-of-plane misalignments.

5.2.1 Methods

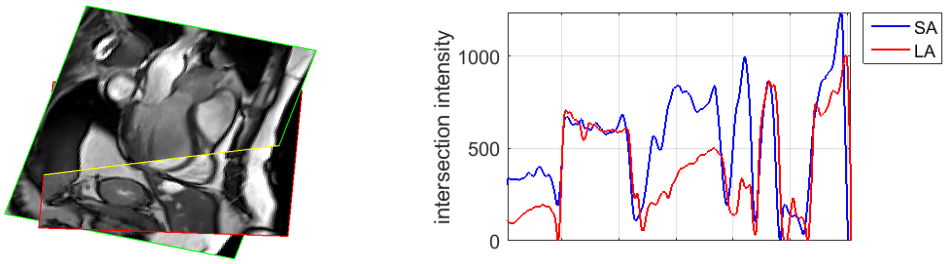
In a similar way to many of the methods discussed in our review of the literature (section 4.1.3), we propose an approach which considers the properties of the intersection lines between images in order to identify the optimal alignment. The hypothesis is that optimal alignment corresponds to the state where the intersection line profiles are most similar. In many of the reviewed methods, the image grayscale values at the intersection lines were considered. In our proposed method, we

additionally consider a local phase representation of the image, therefore providing further intersection profiles.

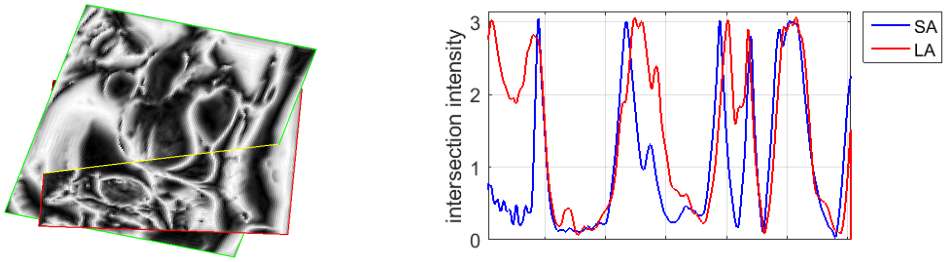
We hypothesise that the local phase representation will better define the structural correspondences of the intersections, as well as providing a contrast-independent measure between the images. Local phase values are calculated using bandpass filters that define the scale and frequency content of the analysed signal. Here we choose to use multiple local phase representations, each calculated at a different filter bandwidth.

The filter bandwidth controls the amount of the signal's frequency spectrum content that is captured. In figure 5.1 we see examples of the local phase outputs acquired with different bandwidths. We observe how wider frequency content, corresponding to more localised spatial filter response, more clearly reflects the image detail and results in sparser phase transitions. As the bandwidth is narrowed the additional peak appearance is concurrent with an increasing linearity of the local phase estimate. The sparseness of the wideband estimate is attractive, as it clearly represents the most important structural components in the signal, while the linearity of the narrowband phase signals can be used in combination with this to identify with high accuracy the final optimal alignment. While we propose the utilisation of local phase profiles, we do not discard the information provided by the intensity images - despite the non-matching intensities expected, the intersecting values should still be correlated.

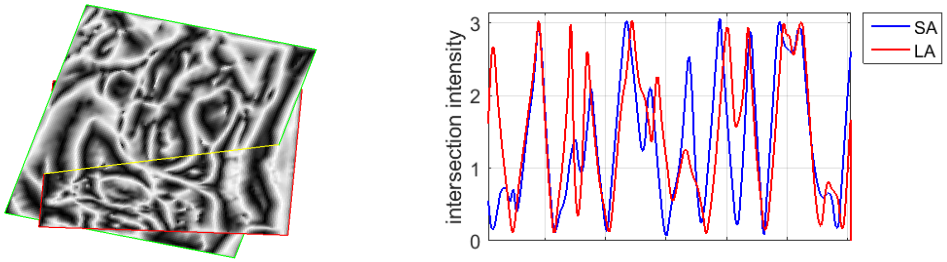
When considering cine data, it is useful to recall that the data has both spatial and temporal components. Cine data are composed of frames captured across the whole cardiac cycle - typically, a sequence contains between 16-25 frames. Significantly, all frames in a cine sequence are acquired during a single breath-hold. As a result, there are no breath-hold related misalignments between frames and therefore the same spatial correction is valid for all image frames. This fact allows the current alignment state to be quantified using information from multiple image frames. It was noted in [113] however, that while the addition of extra frames was initially beneficial to results, no further significant improvement was



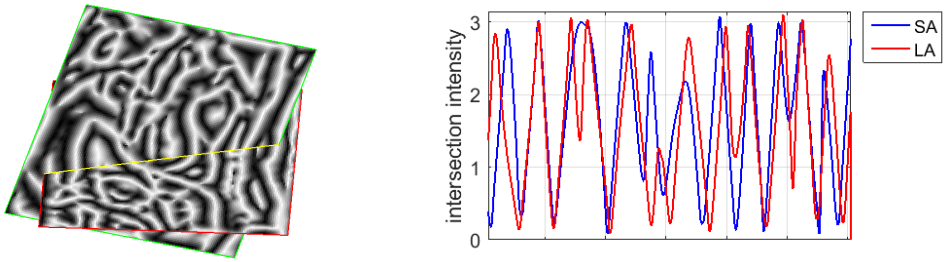
(a) Image Intensity



(b) Wide bandwidth local phase



(c) Medium bandwidth local phase



(d) Narrow bandwidth local phase

Figure 5.1: Illustration of the intensity and local phase image representations used in calculating intersection quality. The data profiles for the intersection line (indicated by the yellow lines in the left hand visualisation) are additionally shown.

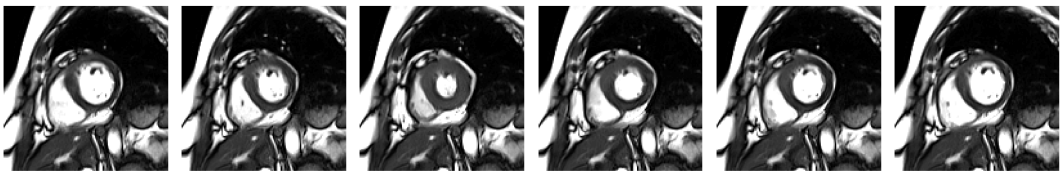


Figure 5.2: The variation in image content observed in a SA slice with cardiac phase.

seen beyond the addition of the first few. Additionally, we note that a significant proportion of the cardiac cycle is occupied by the diastolic phase, during which there is minimal tissue deformation. Given this fact, and with consideration of the increase in computation time required, we choose to use only two image frames - the end-systolic and end-diastolic phases. These two frames are the maximally different phases of the cycle, therefore providing the most distinct information, as illustrated in figure 5.2 where the first and third images capture the approximate diastolic and systolic phases respectively. In our implementation, these frames were identified based on the contours provided by the clinicians, as these are the only two time frames segmented; alternatively, as noted in [154], image content can be used to estimate these frames based on image similarity measures.

As shown in the study by McLeish et al [128], the transformations undergone during respiratory motion are likely to be a combination of both rigid and non-rigid deformation. A notable result from the study is the significant out-of-plane motion undergone during the respiratory cycle; a correction that can capture such motion is essential therefore. The transformation models used in the literature vary; some select simple translational models whilst others have allowed for the full 3D translation and rotation of data. It is notable that there are no proposals to further increase the flexibility of the corrections applied beyond rigid motion, despite the conclusions stated in [128]. However, this is understandable, as the sparsity of the available data would make obtaining accurate results from non-linear approaches very challenging.

The alignment of cardiac MR slices is highly inter-dependent - this may initially encourage us to consider an approach which simultaneously optimises the position of all slices. However, in practice this becomes difficult due to the high dimensional search space and the presence of many local minima. As a result, in this work

we choose an iterative approach, whereby we sequentially optimise the position of each slice, given a consideration of the current slice's correspondences with all other slices. We repeat this process a number of times in order to get closer to the globally optimal positions.

We begin by considering the relative alignment of a single LA slice with all the SA images; each image is considered in turn and allowed to move relative to the others. Three iterations of the realignment are performed before the addition of the next LA slice. Each newly-introduced LA slice is the first to be aligned during that phase of the realignment. To reduce the scale of the initial search spaces the transformation of the SA slices are restricted to be in-plane, while the LA images are given complete freedom to move in 3D. After all LA slices have been introduced, a final iteration gives every slice the freedom to move according to the full 3D rigid transformation model.

Before defining a total cost function to be optimised, we consider how the correspondences of individual intersection profiles are quantified. The observation that the gray value intensities of two intersecting slices may differ has been previously acknowledged, leading to the use of numerous similarity metrics such as mutual information, cross correlation and absolute difference of derivatives. In this work, we use the normalised cross correlation (NCC) metric; this metric is applicable to both intensity and local phase profiles and straightforwardly allows the correspondence measures of both profile forms to be combined to give a single alignment measure.

The initial definition of the full cost function is given in equation 5.1. A single slice, s_T, is selected from the entire set of slices, Ω, to be rigidly transformed according to the transformation, T. The intensity profiles and local phase profiles are denoted by I and ϕ respectively, where the subindex b further indicates the bandwidth of the phase estimate; profiles are calculated for both the end systolic, ES, and end-diastolic, ED, frames. The normalised cross correlation (NCC) of profiles is calculated between the selected slice, s_T and each other image s in the dataset Ω.

$$C_{SA-LA}(T) = \sum_{s\in\Omega|s\neq s_T} \left(\sum_{f=[ED,ES]} \left(NCC\Big(I_{s_T,f}\big(T(x)\big), I_{s,f}(x)\Big) + \sum_{b=1}^{B} NCC\Big(\phi_b\Big(I_{s_T,f}\big(T(x)\big)\Big), \phi_b\Big(I_{s,f}(x)\Big)\Big)\right)\right) \tag{5.1}$$

Additional constraints

The presented cost function looks to improve the image alignments by finding superior correlations between the intersection profiles. However, initial experiments showed that, while improving the intersection correspondences, errors were introduced based on the relative positional changes of adjacent slices. In particular, we observe some erroneous in-plane rotations between adjacent SA slices as well as out-of-plane motions that are clearly not physiologically correct.

The erroneous rotations are likely a result of the small number of intersection lines available to assess the alignment quality, those being the intersections of the SA with the HLA, VLA and LVOT images. Based on only these three intersections, rotations can occur which do not significantly change the structures along these intersections, but do result in inconsistent alignment between adjacent SA slices. This effect is illustrated in figure 5.3a, where three adjacent SA slices are shown with the contours of the middle slice overlain on each. We can see, based on the right ventricle delineation in particular, that the relative rotations of each slice are not consistent. This inconsistency is further illustrated in the mesh visualisation, where there are unexpected distortions in the right ventricle surface.

To reduce the potential of erroneous rotations, a further term was introduced to the cost evaluated, which considers the relative alignment of adjacent SA slices. For this, the cross-correlations of the image intensities are calculated between adjacent slices. This quantity, which is expressed in equation 5.2, is combined with the previously calculated intersection correlation as in equation 5.3. The impact of this additional term is illustrated in figure 5.3b, where the improvement of the relative rotations compared to those shown in 5.3a are qualitatively observable.

Figure 5.3: In (a) results using the intersection-based cost function alone are shown, while in (b) the correspondences between adjacent SA slices and distance between slices were also included. In both cases 3 SA slices are shown, overlain on which are the LV-endo, LV-epi, and RV-endo contours of the middle of the three images. In (a) we see how there is an inconsistency in the correction; the in-plane rotations of the two end slices do not match the rotational position of the central slice - this is most clearly observed from the RV-endo contours. In (b), where a similarity term that takes into account the adjacent slices was used, we see that the relative rotational positions of the slices are more consistent, with the delineations being a closer match to the image data. The meshes displayed give another illustration of the improved alignment of the SA slices where the "bulge" present on the outer wall of the RV mesh in (a) is corrected in (b).

$$C_{SA-SA}(T) = \sum_{s \in SA | s=[s_T-1, s_T+1]} \left(\sum_{f=[ED,ES]} \left(NCC\Big(I_{s_T,f}\big(T(x)\big), I_{s,f}(x) \Big) \right) \right) \quad (5.2)$$

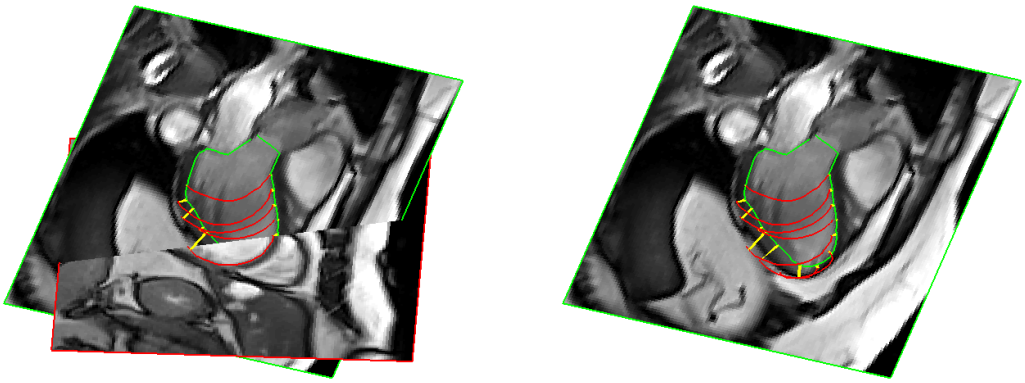
$$C(T) = C_{SA-LA}(T) + C_{SA-SA}(T) \quad (5.3)$$

Our second observation, regarding the out-of-plane motions of SA slices, has been previously noted by Elen et al in [54]. Changes in position in the out-of-plane direction do not always result in clear changes in the cost function; this may be taken as an indication that the local structure of the heart is fairly constant along the direction of the LA. A quantitatively-improved intersection correlation may be found where, qualitatively, the result is clearly incorrect. Such erroneous corrections manifest themselves as significant reductions in distances between SA slices - in the most extreme case SA slice positions were swapped. When qualitatively assessed, it may be clear that new SA slice positions are not anatomically correct, but the current cost function does not convey this. Similar to that proposed in [54], we introduce a penalisation which aims to limit the out-of-plane shift of slices. The ratio of the change in inter-slice distance relative to the original distance is calculated during the optimisation, with any increase above a certain threshold penalised by a comparative reduction in the total correlation value.

The final positions of each slice in the dataset is therefore based on the iterative minimisation of the negative of the defined cost function. Minimisation is performed using the Nelder-Mead simplex algorithm [99], where the optimal parameters are sought for the 3D rigid transformation of each slice.

5.2.2 Experiments and results

The local phase values estimated are dependent on two primary parameters: the central frequency (alternatively specified via the wavelength), which controls the scale of structures identified; and the bandwidth, which indicates the range of frequencies to be included during the estimation of local phase. The final parameter selection was based on qualitative assessment of the filter outputs. A wavelength value of 10 (corresponding to 10 pixels) was chosen to provide a desirable level



(a) **(b)**

Figure 5.4: Images show the closest identified points between SA and LA contours. The yellow lines indicate the distances calculated for the validation of image alignment. In (a) both an SA and LA image are indicated, while (b) shows the positions of all SA endo-cardial contours in relation to the LA data.

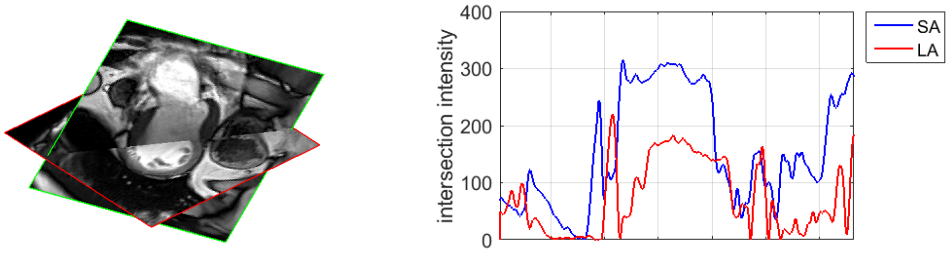
of structural detail. Three radial bandwidth parameters were used, specified via the parameter k which was defined in equation 4.34. Values of 0.3, 0.5, and 0.7 were selected.

As previously indicated, both cine and LGE images were processed. The alignment method used on both forms of data was the same, apart from the availability of only a single frame for LGE realignment.

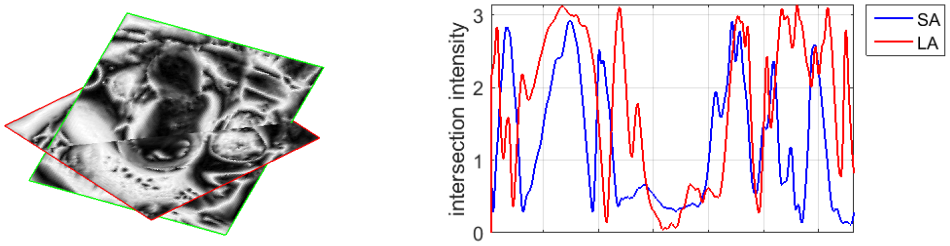
The presented quantitative validation results are based on the expert delineations of the data provided by our clinical collaborators. For each SA-LA slice intersection, the shortest distances between corresponding contours in the SA and LA images are calculated. This analysis provides four distances for each SA-LA intersection - both endo- and epi-cardial distances calculated at either side of the left ventricle. In figure 5.4, an illustrative example is shown of the endocardial distances calculated for a single LA image and all its intersecting SA data before slice realignment.

Cine results

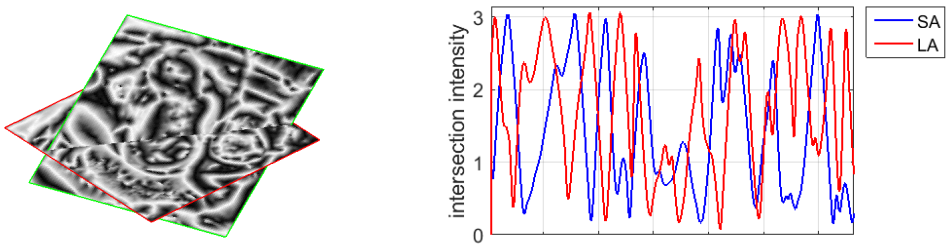
Figure 5.5 shows an example of the initial misalignment between an SA slice and an LA slice. In addition to the original image data, the figure displays the phase-



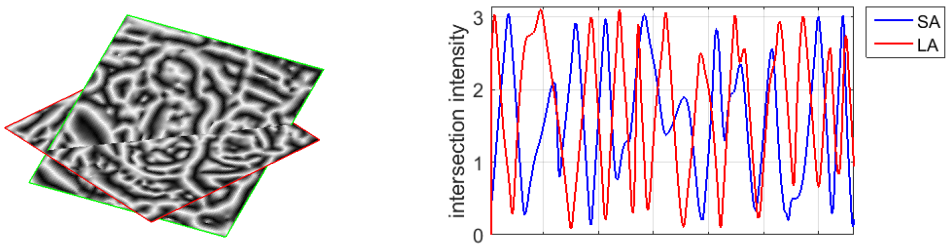
(a) Intersecting SA and LA intensity data with corresponding intersection profiles.



(b) Intersecting SA and LA wideband local phase with corresponding intersection profiles.

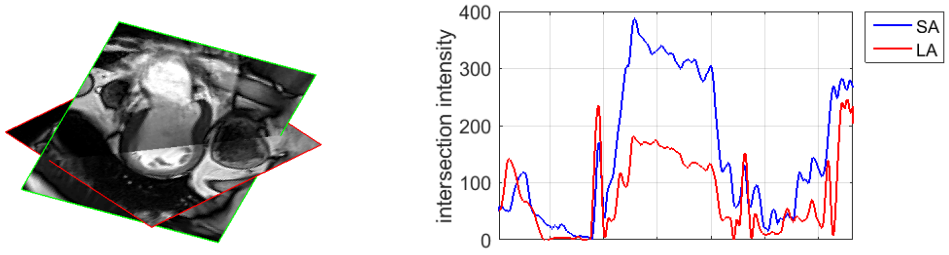


(c) Intersecting SA and LA intermediate bandwidth local phase with corresponding intersection profiles.

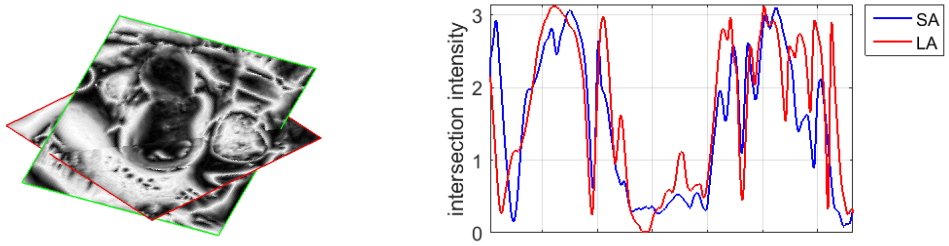


(d) Intersecting SA and LA narrowband local phase with corresponding intersection profiles.

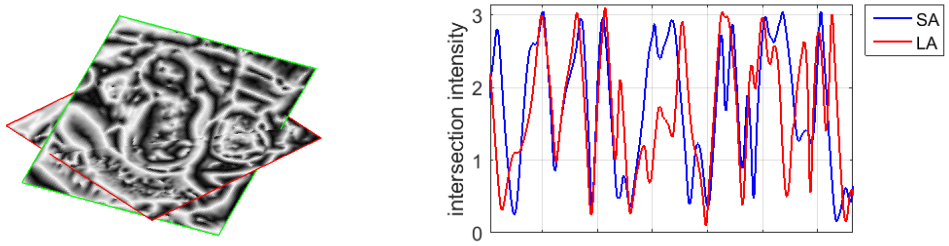
Figure 5.5: Sample SA and LA slices before alignment.



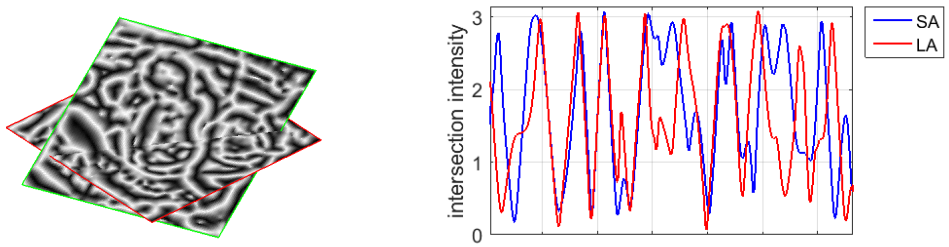
(a) Intersecting SA and LA intensity data with corresponding intersection profiles.



(b) Intersecting SA and LA wideband local phase with corresponding intersection profiles.

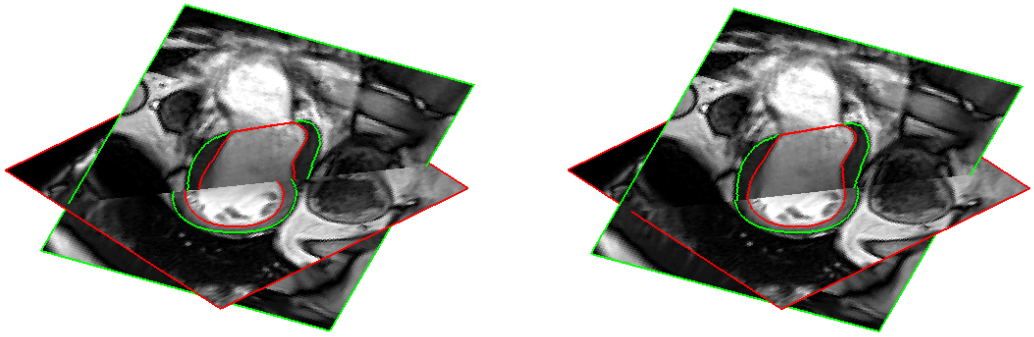


(c) Intersecting SA and LA intermediate bandwidth local phase with corresponding intersection profiles.



(d) Intersecting SA and LA narrowband local phase with corresponding intersection profiles.

Figure 5.6: Sample SA and LA slices after alignment.



(a) **(b)**

Figure 5.7: The position of an SA slice in relation to a single LA image is shown. Their respective endo- and epi-cardial delineations are also displayed. (a) and (b) indicate the correspondence of the data before and after realignment respectively.

based representations used; the calculated profiles used for quantifying the current alignment are also plotted. Figure 5.6 shows the same SA-LA slice pair after alignment, where the improved correspondence of the data is clear. In figure 5.7 the same slice pair is shown with their endo- and epi-cardial delineations overlain. It is clear to see how the correspondence of these contours improves after alignment.

In figure 5.8 the LA image is shown with the delineated endocardial contour overlain in green. The location of the SA endocardial contours that cross this surface are indicated by the red markers. The distances are calculated both before (5.8a) and after (5.8b) alignment, and are plotted in figure 5.8c in green and red respectively. By considering the two distances calculated for each slice (indicated by two green/red points per slice number), we observe that there is a notable reduction in mean contour distance in each case. Based on the nature of the positional corrections observed, we can interpret the misalignment in this case specifically as being caused primarily by a shift of the LA.

Further results are illustrated in 5.9. The acquired cine LA slices are compared to cross-sections taken at the same coordinates from volumetric nearest neighbour interpolations, generated from the cine SA data. The first column contains the LA slices overlain with their endo- and epi-cardial delineations. The second and third columns show the cross-sections generated from nearest neighbour volume interpolations before and after alignment respectively. The same LA delineations are

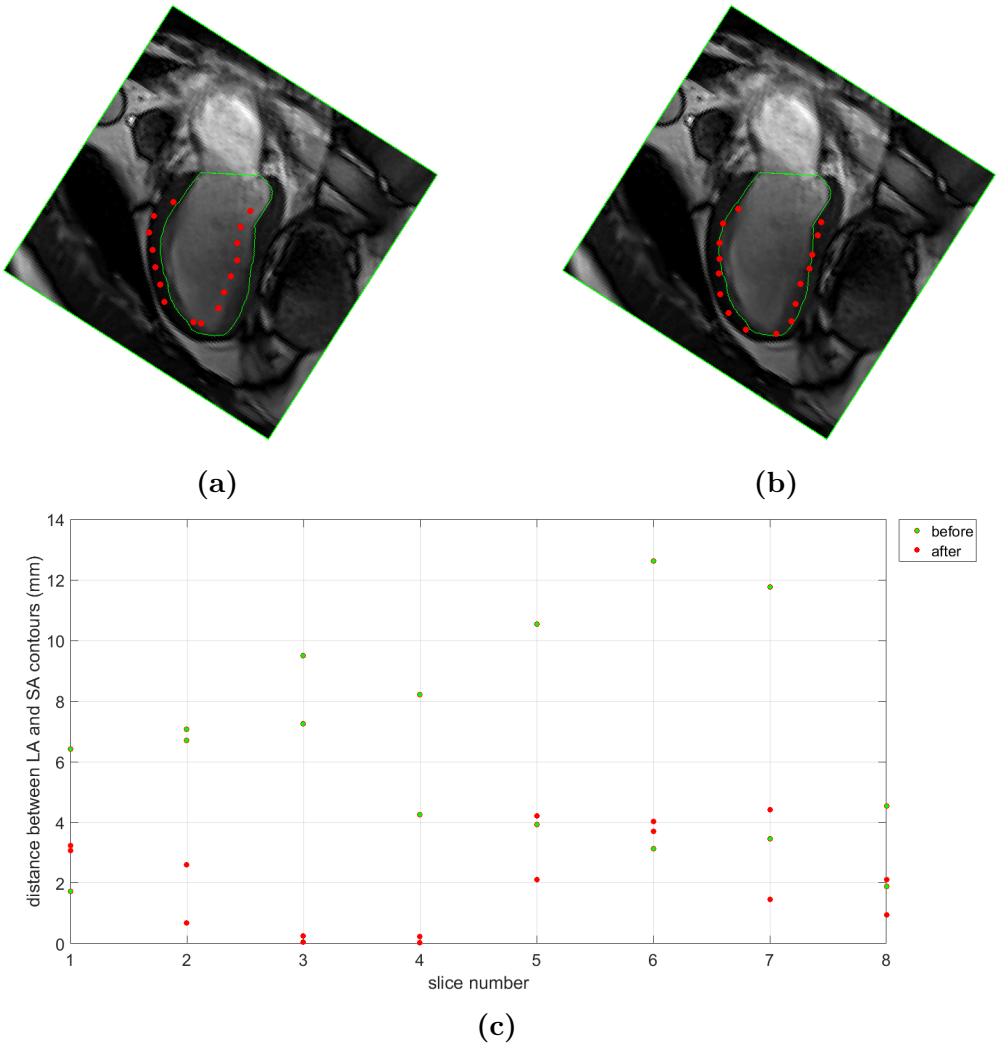
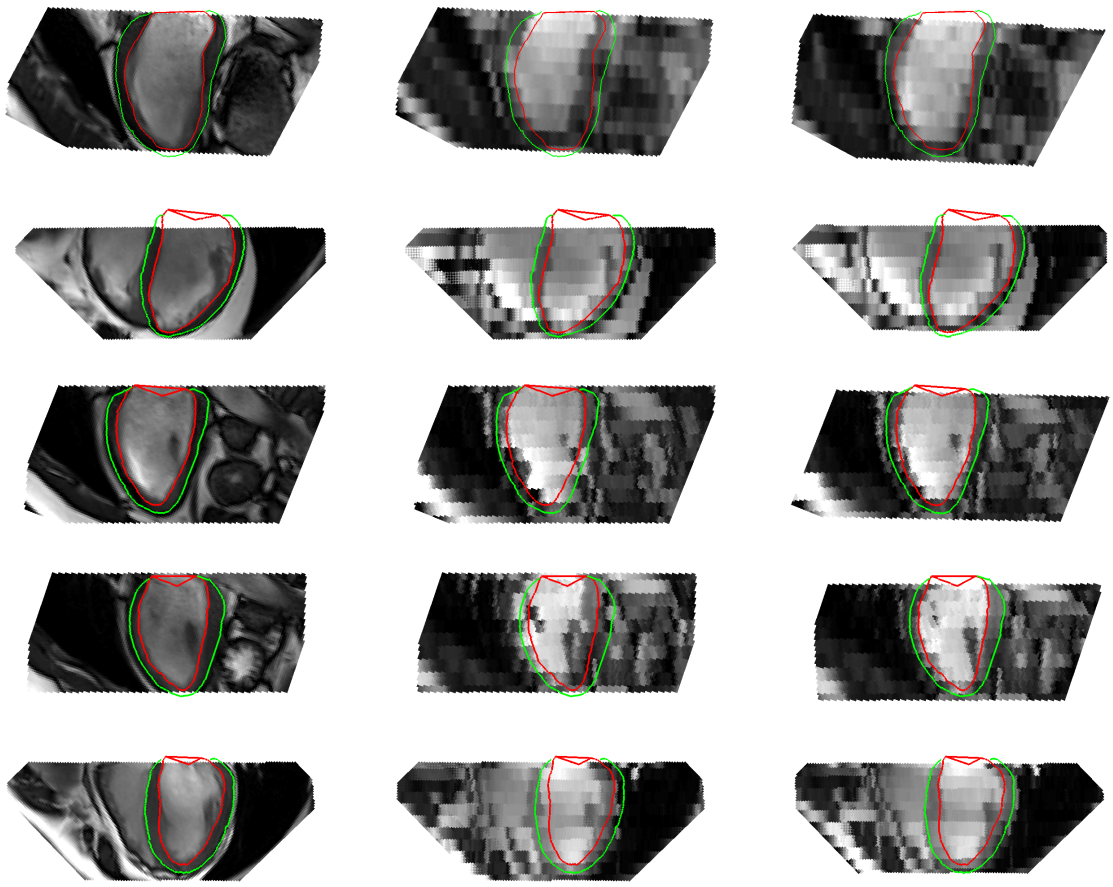


Figure 5.8: An example result showing the position of SA intersections with the LA data. (a) and (b) show these intersections before and after alignment respectively. In (c) the closest distances between SA and LA endocardial contours are plotted. The distances before and after alignment are plotted in green and red respectively.

overlain to give a clearer illustration of the improved slice positions. The recovery of the correct anatomical shape of the left ventricle after slice realignment is apparent. This improvement in shape is also clearly evident in figure 5.10, where an example of endo- and epi-cardial meshes generated using the original slice positions and the corrected positions respectively are shown.

Finally, in figure 5.11, we provide a summary of the contour distances calculated across all cine datasets. The results are shown in the form of histograms indicating the number of SA-LA closest point pairs that fall into each distance bin of width 0.5mm. In all of the plots the distances before correction are indicated in blue,



Cine LA slice | Nearest neighbour interpolation from SA data before realignment | Nearest neighbour interpolation from SA data after realignment

Figure 5.9: A series of results is shown, where the first column contains cine LA slices, which can be compared to the nearest neighbour interpolations generated from SA data before and after realignment in the second and third columns respectively.



(a) **(b)**

Figure 5.10: An example mesh generated from the endo- and epi-cardial contours delineated on the SA slices, where (a) and (b) illustrate the cases before and after realignment respectively.

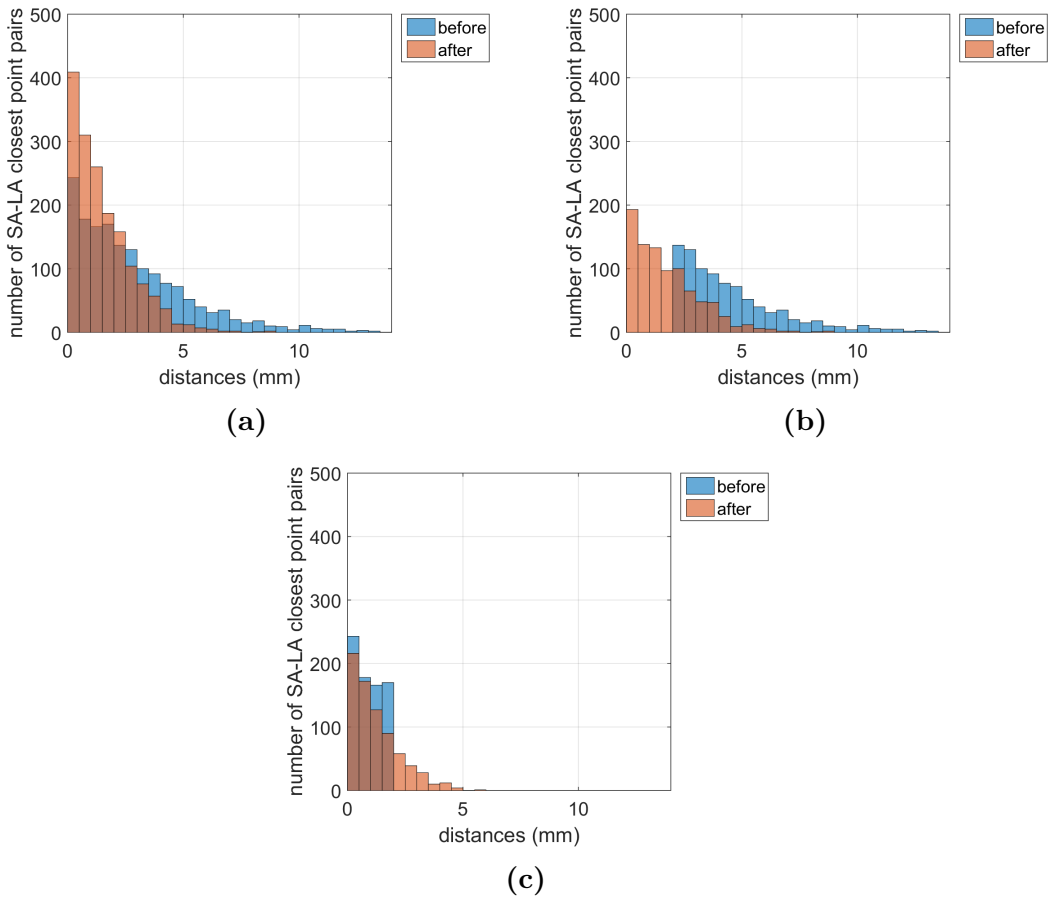


Figure 5.11: These histograms summarise the distances calculated between SA and LA endocardial contours. The number of closest point pairs for each 0.5mm distance bin is plotted in blue before alignment and red after alignment. In (a) all pairs are shown, while (b) and (c) show the results of only those pairs that had an initial inter-contour distance of more than 2mm and less that 2mm respectively.

while the distances after realignment are in red. The distances displayed here are based on endocardial contours; the epicardial results show the same trends.

The achievable minimum contour distances are limited by inconsistencies in the image segmentations (this is discussed by Villard et al [205] where they showed that, using a contour-based realignment technique, they were limited to results with a mean contour distance of 0.43mm). Due to this, it is more illustrative if we present the results as two separate sets. Here we have based this division of results on an assumption that smaller initial contour distances are caused by delineation error while, contrastingly, the larger distances are expected to be actual alignment errors. Here we have chosen a 2mm threshold for this division, as this is approximately the size of a pixel diagonal (2.18mm).

Therefore, three histogram plots are given: 5.11a summarises all distances; 5.11b summarises distances that were initially greater than the 2mm threshold; and 5.11c summarises the distances that were initially less than 2mm.

The initial misalignment of the endocardial contours is calculated as 2.91mm ± 2.60mm. After realignment, this was reduced to 1.51mm ± 1.29mm. The corresponding result for the epicardium was 2.93mm ± 2.65mm before and 1.58mm ± 1.34mm after.

The number of SA-LA closest point pairs that had an endocardial misalignment of greater than 2mm before correction was 885, of which 561 were reduced below the 2mm threshold, while the mean and standard deviation were reduced from 4.60mm ± 2.45mm to 1.76mm ± 1.44mm. Correspondingly for the epicardium the mean and standard deviation dropped from 4.72mm ± 2.50mm to 1.70mm ± 1.40mm.

As is evident from figure 5.11c, the distance between contours does not consistently decrease; those that do, drop from 3.17mm ± 2.18mm to 1.22mm ± 1.20mm, while the remaining point pairs change from 1.04mm ± 1.10mm to 2.03mm ± 1.45mm.

Components of the algorithm

The results presented above, based on the algorithm described, are used in the remainder of this thesis for the reconstruction of volumetric representations and mesh formation. However, as a number of metrics and terms are included as part of the algorithm, in this short section we collate a number of further quantitative results obtained using only parts of the algorithm, thus giving an indication as to the influence of the components to the final alignment result.

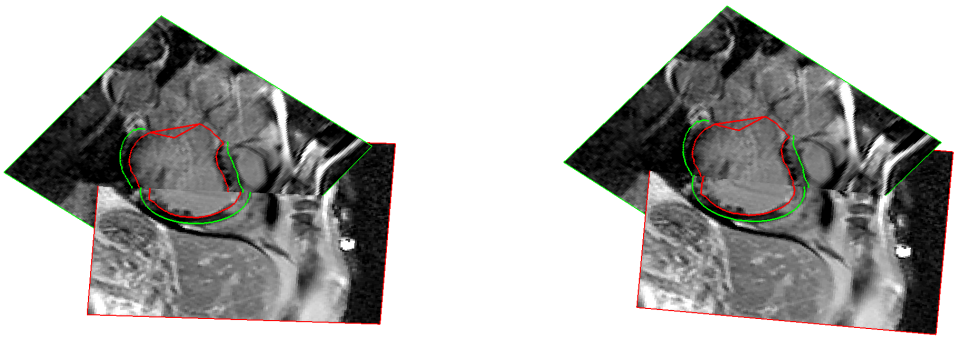
In table 5.1 contour distances computed using the different components of the algorithm are shown. Using the intensity information only we observe that, while the algorithm is successful in reducing the errors in slice alignment, the reported final contour distances are significantly larger than those achieved by each other variant of the algorithm where phase-based information is included. Utilising phase-based information we benefit from a significant improvement in algorithm

	Endo	Epi
Original	2.91mm ± 2.60mm	2.93mm ± 2.65mm
Intensity	1.83mm ± 1.52mm	1.86mm ± 1.58mm
Local phase	1.55mm ± 1.39mm	1.61mm ± 1.37mm
Intensity & Local Phase	1.58mm ± 1.38mm	1.62mm ± 1.37mm
Intensity, Local Phase & Adjacent Slice Term	1.52mm ± 1.30mm	1.58mm ± 1.35mm
Intensity, Local Phase & Slice Distance Term	1.57mm ± 1.37mm	1.61mm ± 1.34mm
Final Presented Results	1.51mm ± 1.29mm	1.58mm ± 1.34mm

Table 5.1: Results obtained using different components of the cost function

performance, reducing contour distances to 1.55mm and 1.61mm for endo- and epi-cardial borders respectively. We note that no gain in performance is obtained by the inclusion of intensity in addition to local phase.

The addition of the final two components, slice-to-slice correlation and inter-slice distances, improves the quantitative results slightly, and when both are combined in the final presented approach we obtain the best results in the reported table; however, we see that these extra improvements are an order of magnitude smaller than that gained by the introduction of phase-based information. While the quantitative improvements are not significant we note that qualitative observations showed that these extra metrics satisfied our objectives as illustrated previously in figure 5.3.



(a) **(b)**

Figure 5.12: Sample result for LGE data. (a) and (b) show the slice positions before and after realignment respectively.

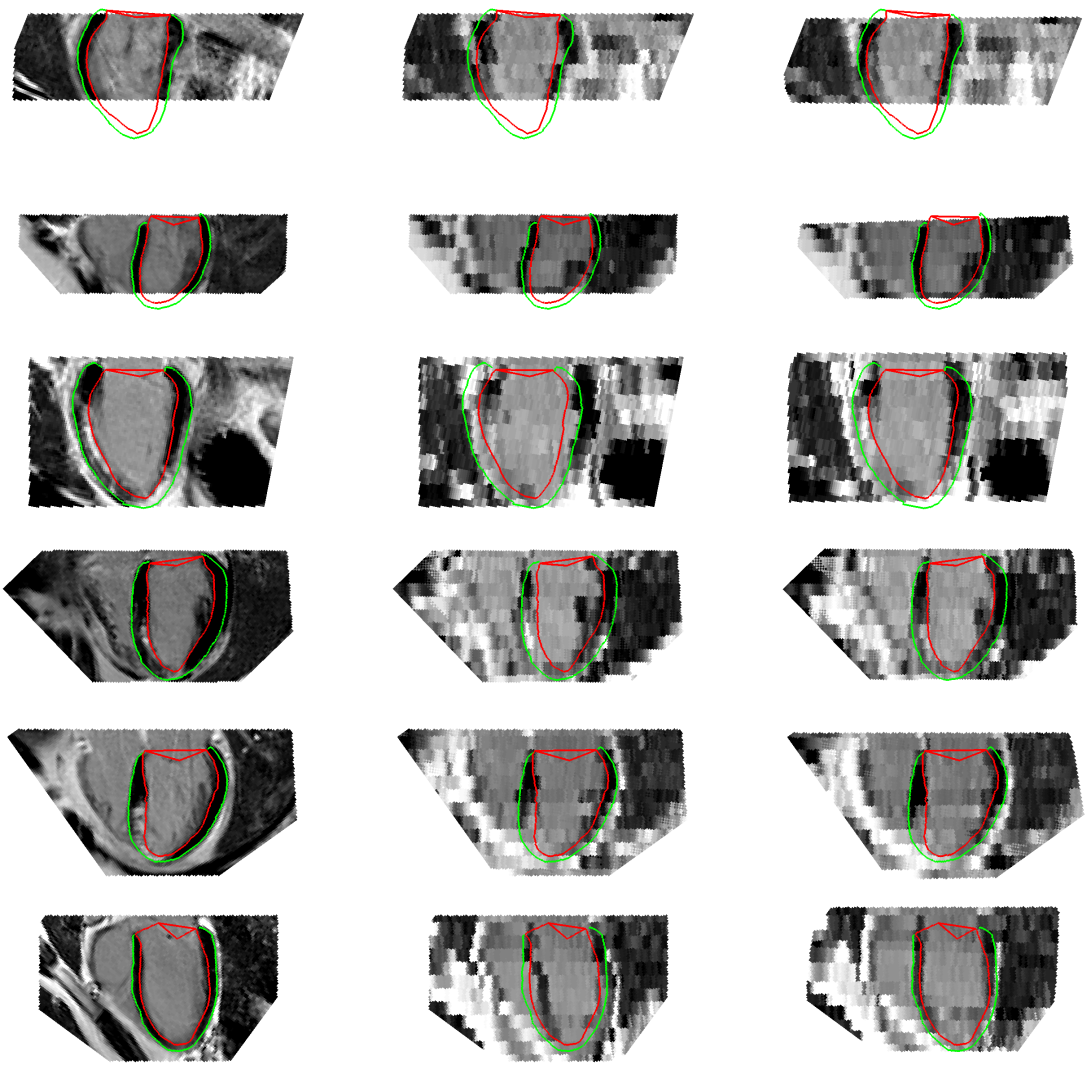


Figure 5.13: A series of results is shown, where the first column contains LGE LA slices, which can be compared to the nearest neighbour interpolations generated from SA data before and after realignment in the second and third columns respectively.

LGE results

The alignment algorithm was applied in the same way to the LGE datasets. Figure 5.12 provides an example of the alignment between the SA and LA LGE data before and after registration. As for the cine results, a series of nearest neighbour cross-sections are shown in figure 5.13 to illustrate the improvement in left ventricular shape observed, while the endocardial contour distances calculated are summarised in figure 5.14.

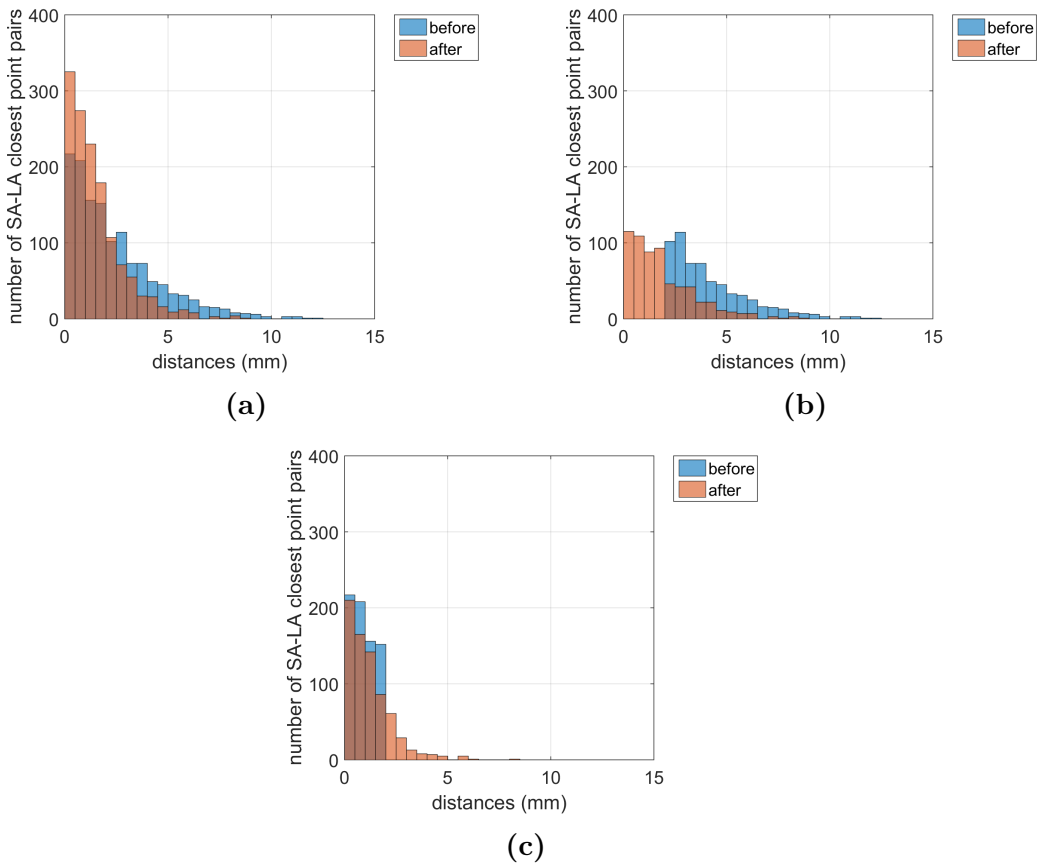


Figure 5.14: These histograms summarise the distances calculated between SA and LA endocardial contours for the LGE data. The number of closest point pairs for each 0.5mm distance bin is plotted in blue before alignment and red after alignment. In (a) all pairs are shown, while (b) and (c) show the results of only those pairs that had an initial inter-contour distance of more than 2mm and less that 2mm respectively.

The initial misalignment of the endocardial contours is calculated as 2.41mm ± 2.13mm. After realignment, this was reduced to 1.51mm ± 1.35mm. The corresponding result for the epicardium was 2.59mm ± 2.32mm before and 1.67mm ± 1.57mm after.

The number of SA-LA closest point pairs that had an endocardial misalignment of greater than 2mm before correction was 621, of which 405 were reduced below the 2mm threshold, while their mean and standard deviation were reduced from 4.18mm ± 1.92mm to 1.86mm ± 1.56mm. Correspondingly for the epicardium the mean and standard deviation dropped from 4.40mm ± 2.15mm to 2.21mm ± 1.92mm.

5.2.3 Discussion

Breath-hold acquisition is the primary routine used for imaging cardiac infarct patients. Unfortunately, this can lead to notable misalignments between images obtained during an acquisition session. Such misalignments cause clear errors during integration of data, and can make further post-processing difficult.

In the previous section, results were presented that demonstrated the ability of the proposed realignment algorithm to compensate for a significant proportion of errors introduced during cardiac MRI acquisition. Results were both qualitatively and quantitatively demonstrated, indicating an improvement in the consistency of slice alignment after processing.

The quality of the alignments was quantified via the distances between endo- and epi-cardial contours. The general trend of the calculated values supported visual results, indicating reductions in the distances between contours after alignment. Despite clear improvements, notable distances do in some cases remain. However, these outstanding differences between contours are not solely indicative of persisting slice misalignment. Due to the lack of a ground truth, we have based our own validation on the segmentations performed by clinical collaborators. The segmentation process is highly dependent on image quality and is understandably a very subjective process; as a result, it suffers significantly from both inter- and intra-expert variation. For example, delineation of the SA data becomes particularly difficult in slices nearer the apex due to the lack of definition of the myocardial-blood-pool surface. Conversely, the apex is likely well-defined in the LA, but recognition of other cardiac structures such as the papillary muscles can become less clear. Some examples are shown in figure 5.15 of cases where we see good image alignment, but poor agreement of the delineations.

This observation regarding difficulties nearer the apex is reinforced, if we consider the specific distribution of the alignment improvements achieved. In table 5.2 the mean and standard deviation of the calculated differences are expressed before and after registration, according to the relative position of the slice from the apex of the heart.

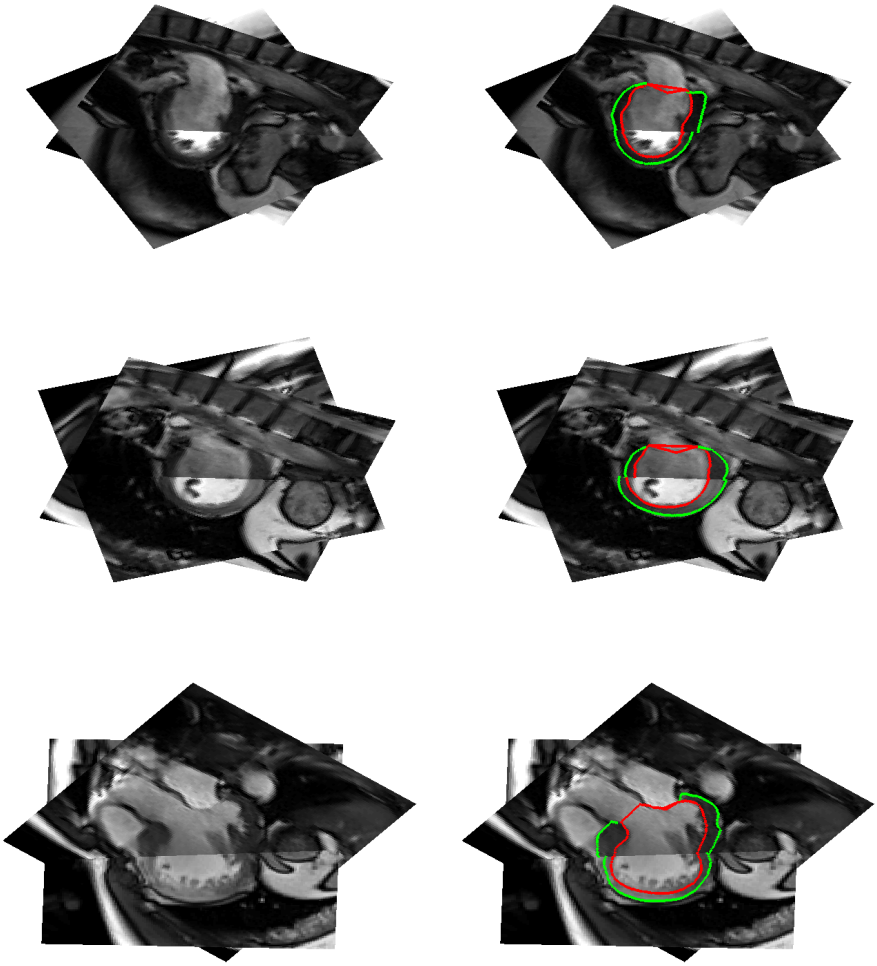


Figure 5.15: Example results that illustrate good image alignment on the left-hand side, but poor contour correspondence due to differing image delineation on the right.

The primary observation here is that the errors nearer the apex appear, in general, to be greater than those observed in the mid-ventricle and basal regions; this is true both before and after registration. A comparison of the relative improvements in the mean distances observed also indicates that these are somewhat less in the apical region. These results suggest therefore that, near to the apex, both the contouring and the alignment process itself is more difficult.

Significantly however, the numbers recorded in the table indicate that, in the mid-ventricular region in particular, the mean error is reduced to below the approximate in-plane pixel dimensions 1.5mm x 1.5mm.

A number of methods were presented in the literature which proposed using data

delineations to estimate the slice realignment correction. However, the observations that we have just made regarding the accuracy of the delineation process is a key reason that we choose not to utilise such an approach. It is perhaps worth noting that, in many of those papers, the final desired outcome is a mesh- or model-based representation of the cardiac surfaces, and therefore the accurate relative alignment of the contours was of primary importance.

The quantitative results can be considered in conjunction with a qualitative assessment. The qualitative results, which were presented in the form of slice intersections, LA cross-sectional views, and meshes generated before and after alignment, give a good indication of the improvement in the underlying anatomical shape captured by the realigned images.

In addition to the intersection profile information, our cost function gave consideration to the correspondence of adjacent slices and their proximity to each other. A qualitative evaluation of the results indicated that the occurrences of notable inaccurate rotations were reduced and excessive out-of-plane adjustment was suppressed. While erroneous rotations were the motivation for introducing the additional adjacent slice correspondence term it has likely also improved the overall ability of the algorithm to correctly align the SA stack.

Distance from apex (mm)	Inter-contour distances	
	Before	After
< 10	3.13 ± 2.36	1.96 ± 1.52
20	3.07 ± 2.72	1.94 ± 1.54
30	2.90 ± 2.63	1.57 ± 1.28
40	2.92 ± 2.74	1.37 ± 1.21
50	2.81 ± 2.51	1.23 ± 1.08
60	2.78 ± 2.64	1.32 ± 1.06
70	2.99 ± 2.84	1.35 ± 1.12
80	2.52 ± 2.11	1.38 ± 1.18
90	2.26 ± 2.01	1.08 ± 1.13

Table 5.2: Distances between SA and LA endocardial contours grouped according to the distance of the slice from the apex of the heart.

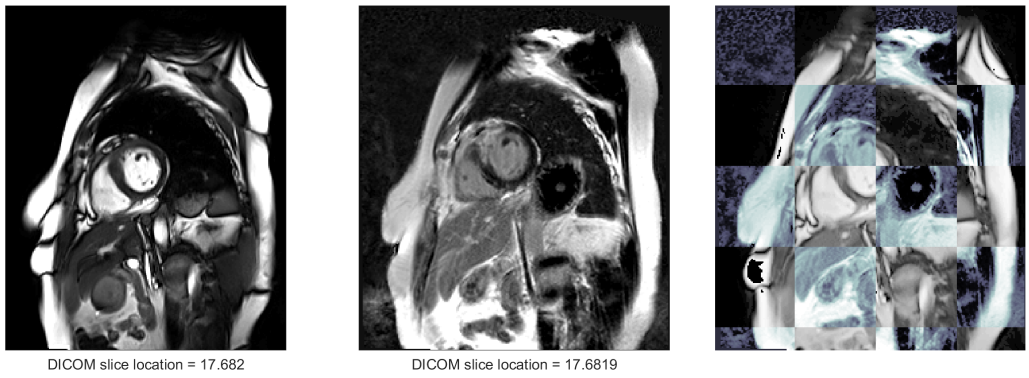
In contrast to many of the papers on this topic in the literature, we proposed to use local phase representations of the data in addition to the raw intensity data. We believe that using local phase-based data representations would improve the accuracy of alignment, in particular due to the strong characterisation of significant image features such as the myocardial borders, and papillary muscles. The local phase-based metric is particularly attractive here due to the intensity differences that we regularly observe between the intersecting LA and SA data. Specifically in this work we used local phase representations that were calculated at different bandwidths - the differing bandwidths provide complementary representations with the wider-bandwidths providing qualitatively higher degrees of image detail, while the narrower bandwidths generated local phase estimates that were linear over wider spatial widths. An alternative approach would be to use a number of different central frequency representations, as this would provide information at a number of scales and therefore characterisation of different structures for alignment.

5.3 LGE to Cine Registration

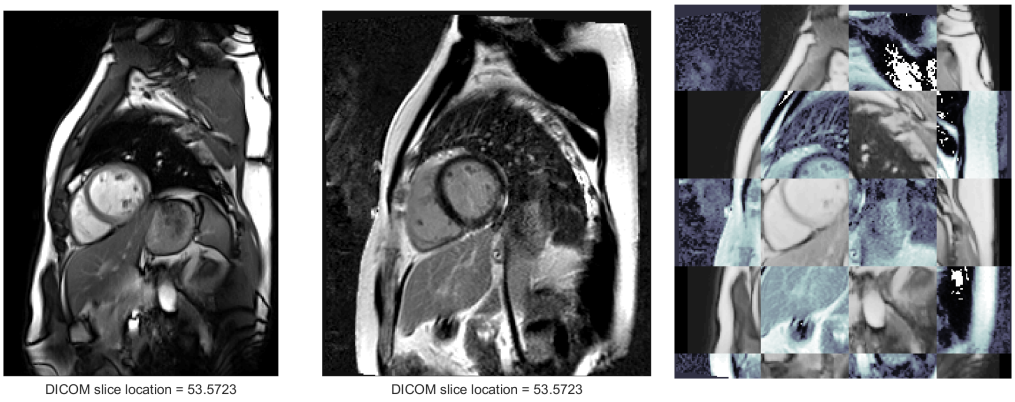
LGE MR imaging has become the gold-standard modality for identification of scar tissue. It is acquired in the same imaging session as the cine data, with the same image locations selected; however, these acquisitions are similarly susceptible to patient breath-hold- and movement-related difficulties. This is of marked importance here due to the acquisition of LGE images significantly later than cine images, increasing the likelihood of patient movements.

Misalignments between intersecting LGE images were compensated for using the slice realignment method discussed in the previous section. Having performed this realignment, it it is necessary in this section to discuss further processing steps that align this data to the corresponding cine data.

In figures 5.16a and 5.16b, we illustrate two cases of slice misalignments that appear in the data. The two examples show cine and LGE images that, according to their DICOM headers, were acquired in the same physical space; however, in both cases, the imaged tissue certainly do not match. The nature of the misalignments



(a)



(b)

Figure 5.16: Two examples of misalignment between the acquired cine and LGE data are shown in (a) and (b) The left-hand and centre column show the cine and LGE data respectively, while the right-hand column displays a checkerboard visualisation of the misalignments.

observed differ; in the first case, the misalignment appears to be primarily an in-plane motion, while, in the second, a significant out-of-plane displacement is evident. The scale of the shift in the latter suggests a significant adjustment in patient position.

In the clinical setting, the comparison between cine and LGE data is based on the selection of the most similar slices acquired from both modalities, in order that a segmental analysis of the myocardium may be performed on each. To enable a fully 3D assessment of the clinical data, in this section we will propose a slice-to-volume registration approach for aligning our LGE data to volumetric cine data; such an approach allows for the compensation of both the in-plane motions illustrated in figure 5.16a, and the out-of-plane shifts, such as in figure 5.16b.

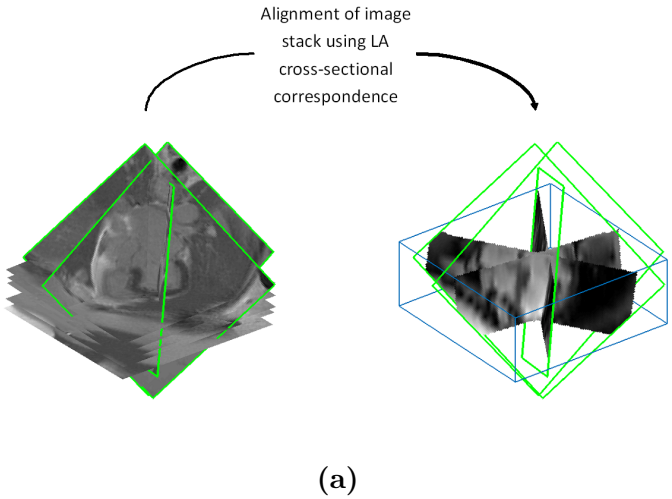
5.3.1 Methods

To find the optimal position of the LGE image data, a 3D volumetric representation of the cine data is used. An isotropic cine volume is generated from the realigned slices using MATLAB's linear interpolation of scattered data. With this data presentation, a full 3D rigid slice-to-volume registration is performed between the LGE slices and the cine volume.

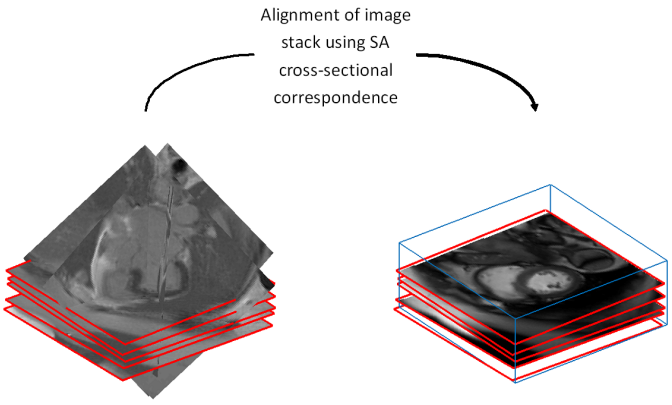
A three stage process is applied, illustrated in figure 5.17:

- the whole LGE image stack, treated as a single moving dataset, is registered, based on the correspondence evaluated between LA LGE images and the volume. See figure 5.17a where the LGE LA slices outlined in green on the left hand side are compared to their cross-sections taken from the cine volumes shown outlined in green on the right hand side;
- the whole LGE image stack is registered, based on the correspondence evaluated between SA LGE images and the volume. See figure 5.17b where the LGE SA slices outlined in red on the left hand side are compared to their cross-sections taken from the cine volumes shown outline in red on the right hand side;
- LGE slices are individually adjusted to align with the cine volumetric data. See figure 5.17c where all LGE SA and LA slices are outlined respectively in red and green on the left hand side are compared to their cross-sections taken from the cine volumes shown on the right hand side.

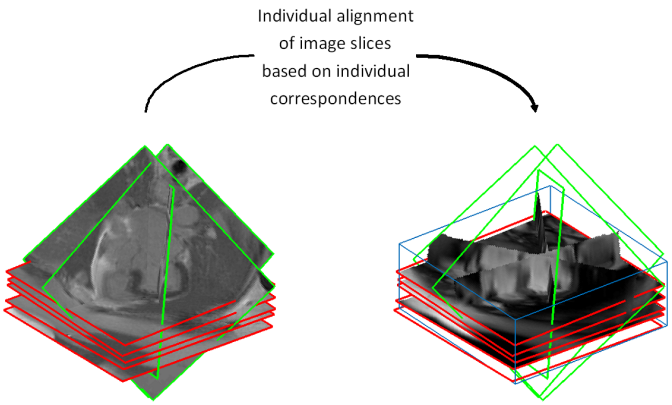
The first stage is proposed in order to align the structures of the LA images without the extra local minima introduced by the SA slice data; this step is expected to allow for any significant out-of-plane motions between the data. In the second step, we again treat the entire image stack as a single entity and calculate a rigid correction based on a comparison of LGE SA slices and the volume's plane intersections. Finally, an individual realignment of each slice to the volume is performed. This final step is taken such that any remaining misalignments, that



(a)



(b)



(c)

Figure 5.17: The three steps of the LGE slice realignment are illustrated. In the first two steps the entire LGE image stack is moved together. Firstly the LA LGE data are used to assess the alignment with the volume - the right-hand images indicate the plane cross-sections calculated from the cine volume. In the second step only the SA data are considered. Finally, each slice is moved independently to allow for any small remaining misalignments.

may exist due to inconsistent or erroneous realignment that occurred between SA and LA slices in the previous stage, may be accounted for - in the first two steps the global registration of the entire image stack will not have accounted for this. In each of the three stages a 3D rigid transformation is calculated. The first two steps require the calculation of a single 3D transformation, while in the final step a 3D rigid transformation is calculated for each slice independently.

Due to the different gray scale appearances of cine and LGE images we chose to utilise local phase-based representation of the data for quantification of data similarity. This choice was based on the understanding that, while the intensity distributions vary significantly, the underlying structural content of both image modalities is equivalent and, therefore, a similarity criterion that looks for correspondences in these structures is favourable. The normalised cross-correlation between the image local phase and the volume local phase cross-section is calculated. As was described in the previous slice realignment section, the local phase was estimated at a number of bandwidths, which therefore combines structural information at a range of frequency spreads.

5.3.2 Experiments and results

The slice-to-volume registration was performed on LGE data acquired from each of the patients for both hospital visits. Similar to the previous slice realignment step, bandwidths of 0.3, 0.5 and 0.7 were used at a scale of 10 pixels for local phase estimation.

The validation of the results is based on the delineations of the endo- and epi-cardial borders provided by the clinicians. However, since the cine and LGE delineations are unlikely to lie in the same planes, we consider the distances between the LGE delineated contours and the meshes generated from the cine delineations.

Figure 5.18 shows an example of the input data used. An LA LGE slice is shown in 5.18a and the corresponding cine volume cross-section in 5.18c. However, it is the local phase data indicated in 5.18b and 5.18d that is used for assessing the current data alignment.

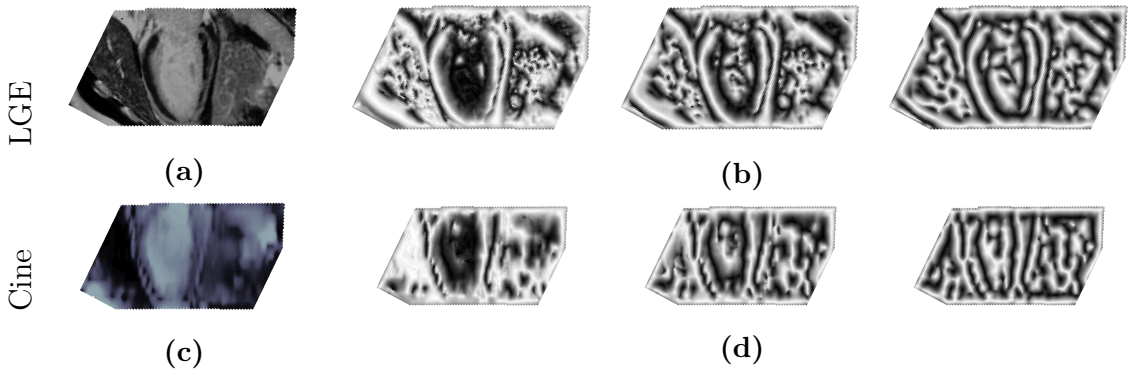


Figure 5.18: Example data used for LGE slice to volume registration. A sample LGE slice and cine volume cross-section before alignment are shown in (a) and (c). Local phase representations at a range of bandwidths are shown for the LGE slice in (b) and for the cine volume cross-section in (d).

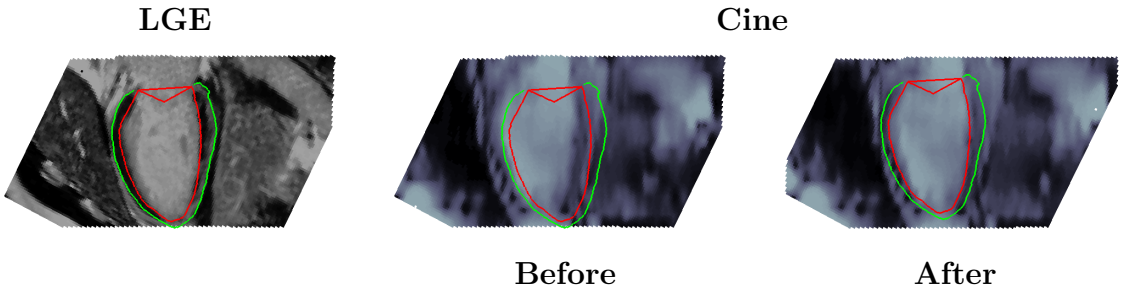


Figure 5.19: A sample result showing the alignment of an LA LGE image with the cine volume. In the left-hand image an LGE slice is shown, while the middle and right-hand images show cross-sections of this slice with a cine volume before and after registration respectively. Overlain on each are the endo- and epi-cardial contours drawn on the LGE data.

Figure 5.19 indicates the relative alignment of that same LGE slice to the volumetric data in the form of checkerboard illustrations before and after registration. Additionally, the LGE myocardial delineations are overlain, aiding in the illustration of the improved alignment of the LGE slice after registration.

Figures 5.20 and 5.21 illustrate the equivalent input data and registration results for a SA slice.

In figure 5.22 a number of further LA alignment results are shown, while in 5.23 example SA results are given. In all cases the improved alignment of the results is clear, based on both the qualitative appearance of the image data and the position of the delineations.

The improvement in data alignment is quantified via the calculation of distances between the LGE endo- and epi-cardial delineations and the corresponding meshes

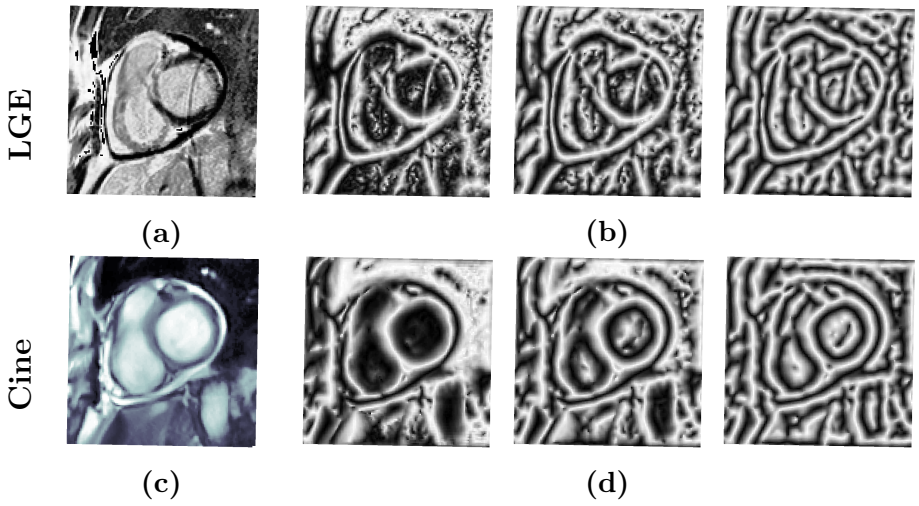


Figure 5.20: Example data used for LGE slice to volume registration. A sample LGE slice and cine volume cross-section before alignment are shown in (a) and (c). Local phase representations at a range of bandwidths are shown for the LGE slice in (b) and for the cine volume cross-section in (d).

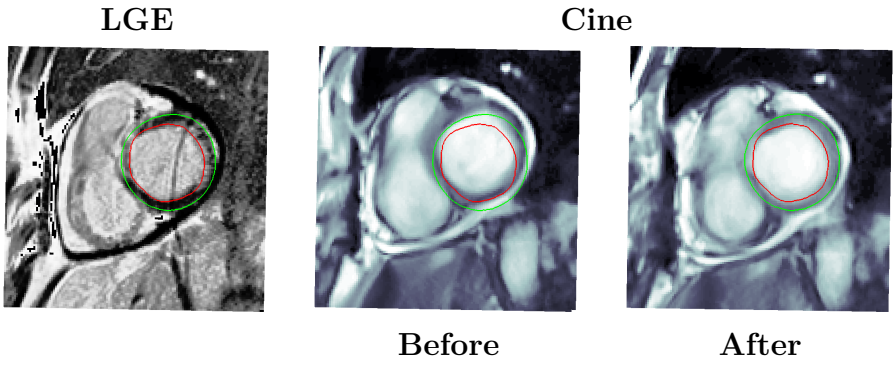


Figure 5.21: A sample result showing the alignment of a SA LGE image with the cine volume. In the left-hand image an LGE slice is shown, while the middle and right-hand images show cross-sections of this slice with a cine volume before and after registration respectively. Overlain on each are the endo- and epi-cardial contours drawn on the LGE data.

generated for the cine data. In figure 5.24 a number of illustrative examples are given; in both the endo- and epi-cardial cases the calculated distances are represented by the black lines.

Finally, the calculated distances are summarised in figure 5.25. There is an observable reduction in the calculated distances for both the endo- and epi-cardial borders. Specifically, the endocardial distances reduced from a mean and standard deviation of 3.30mm $\pm$ 3.15mm to 1.80mm $\pm$ 1.52mm. 56% of contour points that were greater than 2mm before alignment have become less than 2mm after alignment,

	Endo Distances	**Epi Distances**
Original	3.31mm ± 3.16mm	3.26mm ± 3.07mm
Step 1	2.34mm ± 2.08mm	2.33mm ± 1.98mm
Step 2	2.13mm ± 1.77mm	2.04mm ± 1.68mm
Final	1.80mm ± 1.52mm	1.77mm ± 1.41mm

Table 5.3: Distances calculated between the LGE SA contours and the cine meshes after each stage of the algorithm illustrated in figure 5.17

and their distribution has reduced from 5.10mm ± 3.13mm to 2.13mm ± 1.69mm.

The algorithm developed here is a three stage process; intermediate results obtained for each stage are presented in table 5.3. These numerical results, which summarise the contour-to-mesh distances calculated after each stage, clearly demonstrate the benefit of each additional component. The greatest improvement is observed after the first stage, thus demonstrating that the registration based on the LA data performs well and is capable of recovering much of the initial misalignment. Using the SA data, for which there are more slices as compared to the LA, allows for a further improvement in the alignment results, while finally, the adjustments allowed when each of the slices is given the freedom to move independently provides the best numerical results.

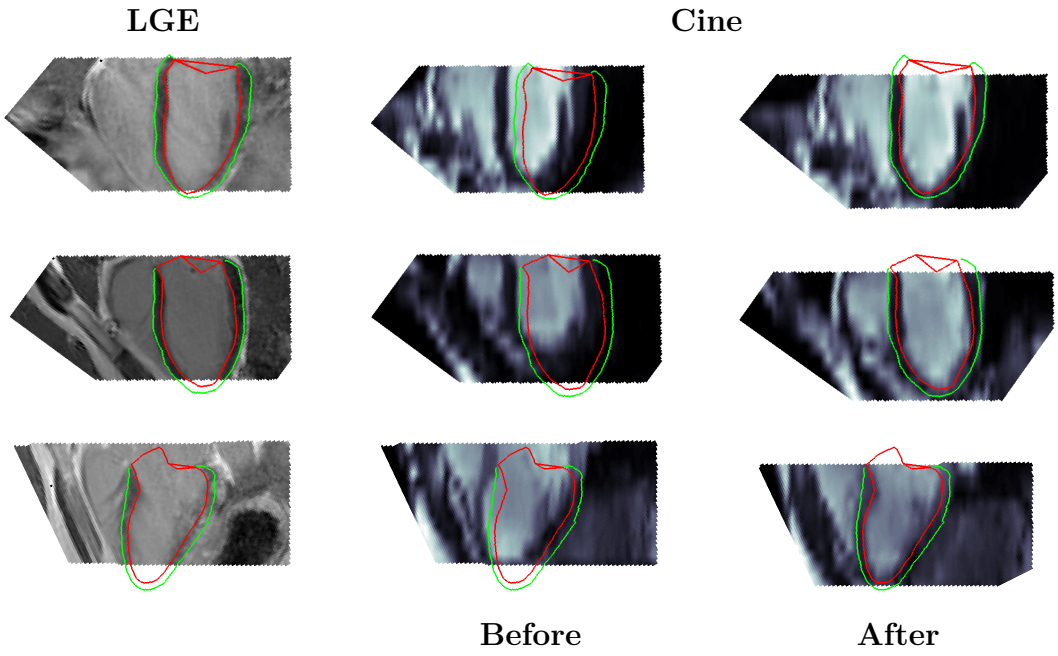


Figure 5.22: A series of example results showing the alignment of LA LGE slices with the cine volumes. In the left-hand column LGE slices are shown, while the middle and right-hand columns show cross-sections of these slices with cine data before and after registration respectively. Overlain on each are the endo- and epi-cardial contours drawn on the LGE data.

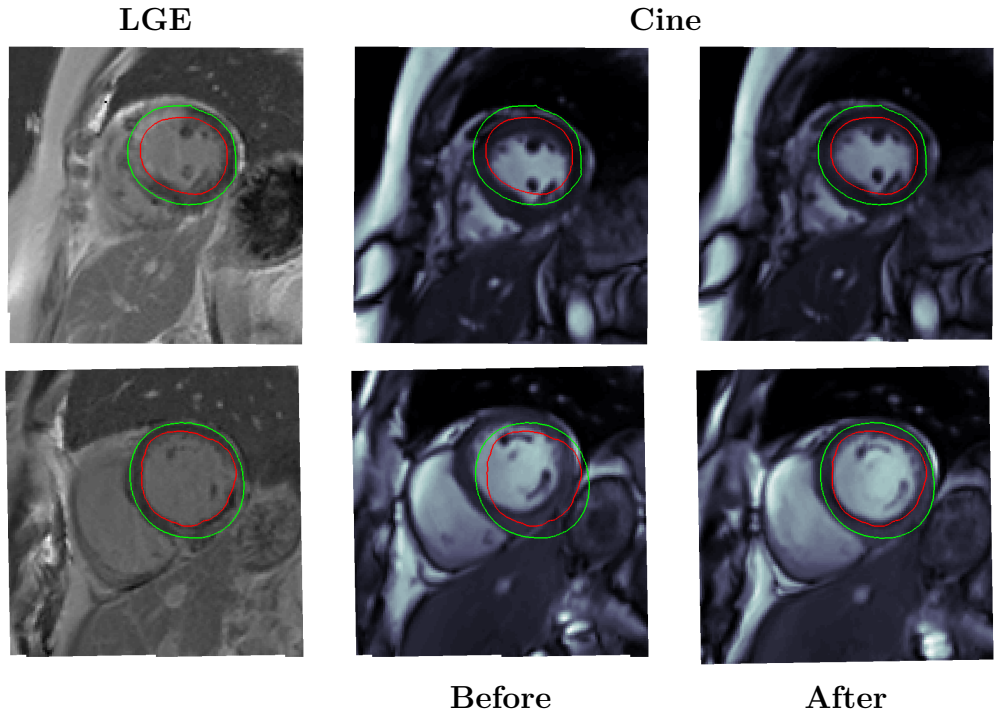


Figure 5.23: A series of example results showing the alignment of SA LGE slices with the cine volumes. In the left-hand column LGE slices are shown, while the middle and right-hand columns show cross-sections of these slices with cine data before and after registration respectively. Overlain on each are the endo- and epi-cardial contours drawn on the LGE data.

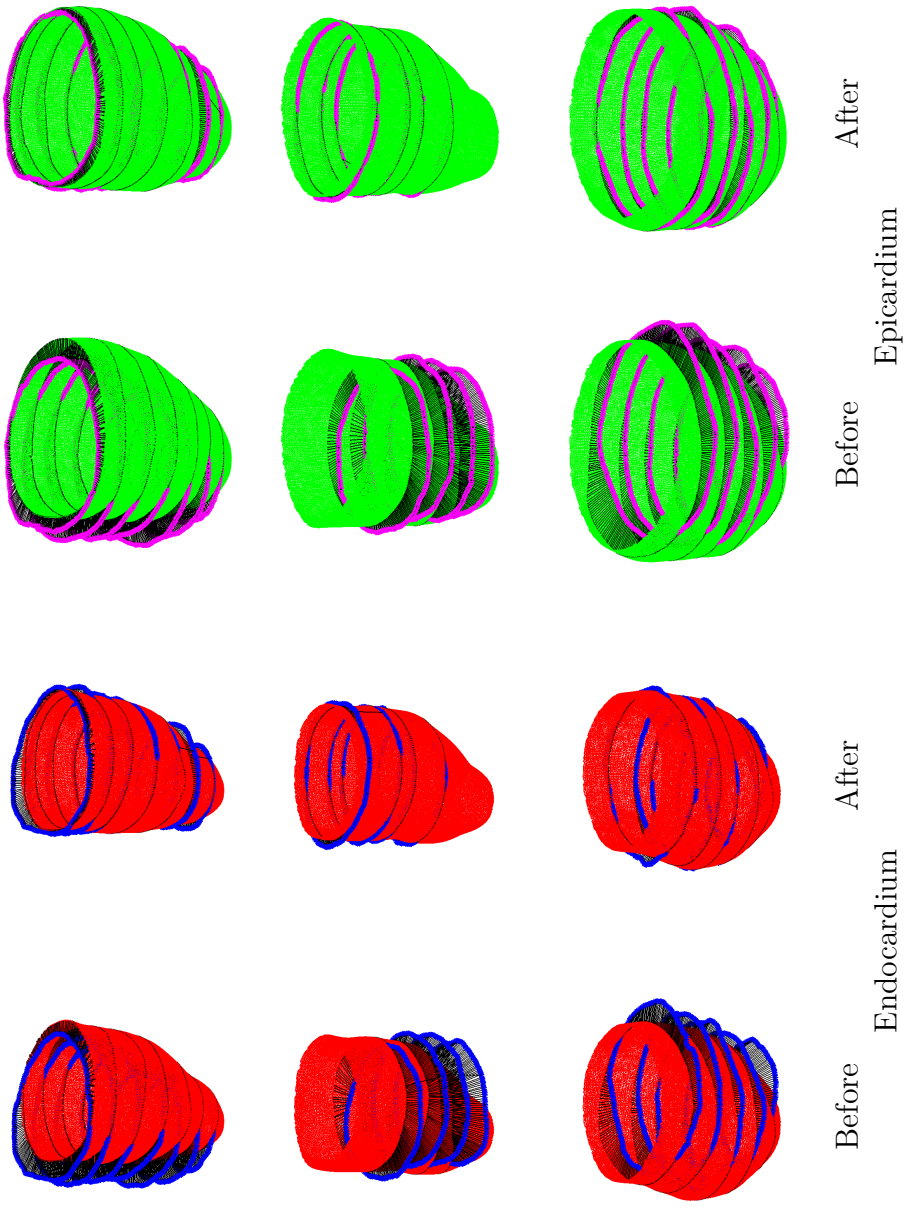


Figure 5.24: The relative positions of LGE data delineations and cine meshes are shown both before and after registration; the results for both the endocardium (in red) and epicardium (in green) are given. The black lines drawn between the contours and meshes indicate the distances used for validation of the results.

5.3.3 Discussion

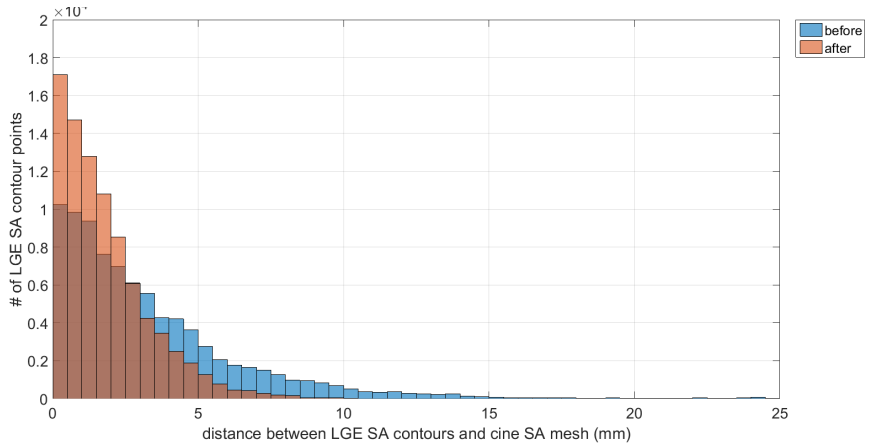
The alignment of LGE data to the corresponding cine acquisitions is undermined by the time required between their respective acquisitions. The period required increases the likelihood that the patient may move within the scanner, while it may also further reduce the likely consistency of breath-hold volumes compared to those held previously during cine acquisition. While breath-hold variations observed between LGE SA and LA images are corrected using the method described in section 5.2, this still leaves the misalignment between LGE and cine data. In the previous pages we have discussed the methodology for re-establishing data alignment and have presented the acquired results.

A local phase-based approach was presented that registered the 2D planar slice data to an isotropic cine volume interpolated from SA slices. Since the LGE SA and LA images had previously been co-aligned, the majority of the registration procedure calculated the best 3D rigid transformation for the entire image dataset as a single entity.

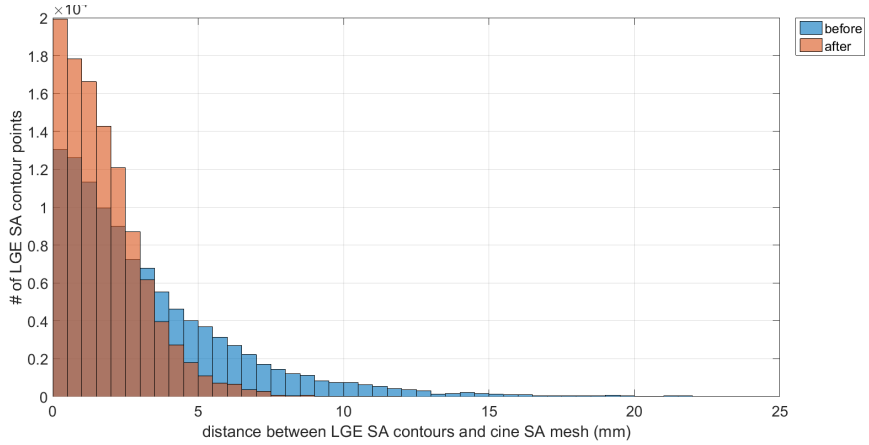
A number of qualitative results were presented; these included examples of both LA and SA alignment between the LGE and cine data, each indicating a clear improvement. The LA results in particular illustrate the requirement for a 3D registration procedure, due to the potentially significant shifts in the LA plane between cine and LGE acquisitions.

Quantitative validation was performed by considering the distances between LGE SA endo- and epi-cardial delineations and the corresponding surface meshes generated from cine SA contours. The improved alignment after registration can be clearly seen in the example results in figure 5.24. These examples illustrate corrections in response to different forms of misalignment - the first case for example indicates an in-plane correction of contour positions, while the second case shows clear misalignment in both the in-plane and out-of-plane directions.

The distances calculated were summarised in the subsequent histogram plots in figure 5.25. Evidently, there is a significant reduction in the errors observed. However, here again we see that misalignments remain.



(a)



(b)

Figure 5.25: These histograms summarise the distances calculated between the LGE SA contours and the cine meshes. The number of closest point pairs for each 0.5mm distance bin is plotted in blue before alignment and red after alignment. In (a) the results for the endocardium are shown, while (b) indicate the corresponding epicardial results.

These discrepancies can be partly explained by the differences in the respective delineations of the cine and LGE data. We previously alluded to the subjective nature of the segmentation process; here, the difficulties are exacerbated by the differing qualitative characteristics of cine and LGE data. Furthermore, the cine delineations used are mesh-based, meaning that between the original cine slice positions the fidelity of the endo- and epi-cardial surfaces to the interpolated image data may be somewhat reduced, when compared to manual delineations.

In addition to the issues of segmentation, the registration method used here does not account for non-rigid misalignment between the cine and LGE images. Both sets of images are acquired during a single imaging session using a triggering

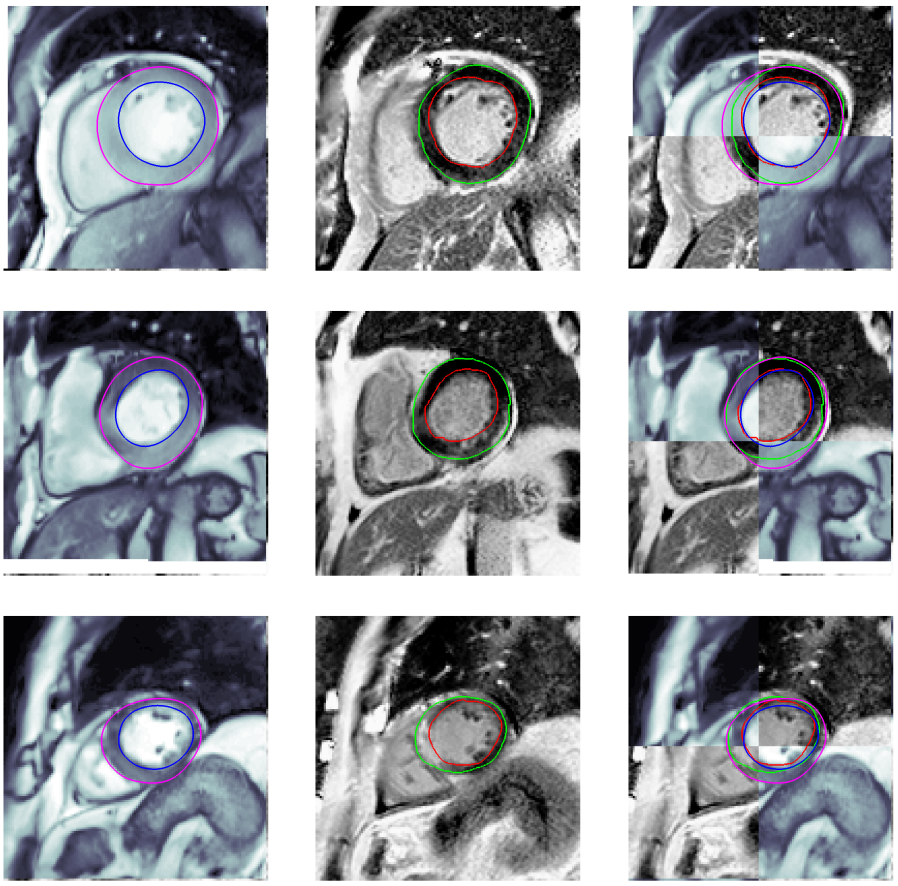


Figure 5.26: Discrepancies in the segmentation of cine and LGE image data. In the first and second column the cine and LGE images are shown with their respective endo- and epi-cardial delineations. In the final column the images are shown using a checkerboard display. With both sets of delineations overlain, the differences are clearly visible.

system such that the same cardiac phase is imaged; despite this some degree of non-linear deformation may exist. However, as we expect the scale of these non-linear deformations to be relatively small, and that compensation of these differences would require a significant increase in the number of degrees of freedom in the transformation model, we chose not to introduce this into the discussed approach as it would complicate the process of finding a reasonable approximation to the global minimum.

A series of examples is shown in figure 5.26 that illustrate some cases where, despite image correspondence, the respective delineations in the cine and LGE cases can vary notably.

5.4 3D CMR Registration - Longitudinal Analysis

Follow-up cardiac assessments are essential for the clinician, providing insight into the progression of disease, from which a more accurate prognosis may be attained. The objective of this section is to discuss a framework for the identification of 3D correspondences between cardiac MR datasets acquired from a single patient on two separate visits to the hospital. We will then consider the utilisation of the resultant alignment (in combination with the results of the cine-LGE registration in the previous section) to assess changes that may be observed in the tissue between image acquisitions.

5.4.1 Methods

Isotropic volumetric data representations are generated from the original anisotropic slice data via simple linear interpolation. The interpolation of the scattered data is performed using MATLAB's in-built *scatteredInterpolant* function. This approach is based on a Delaunay triangulation of input data points, followed by a weighted averaging of the intensities at the tetrahedron vertices, local to each new data point. Due to the inherent intensity differences present between SA and LA image data, the results presented here are acquired based on volumes generated from the SA slice data only. A comprehensive approach to 3D image reconstruction is presented in chapter 6, but its development was completed after the work presented in this chapter, and thus was not employed here.

A two stage registration is performed, with an initial rigid realignment, followed by a refinement using a variational optical flow approach.

Rigid registration

The original volume positions in the MRI scanner are calculated according to the DICOM header patient position information. Due to the acquisition of images during two entirely distinct sessions, there is likely a significant initial data offset, such as that illustrated in figures 5.27a and 5.27b. One means of compensating for

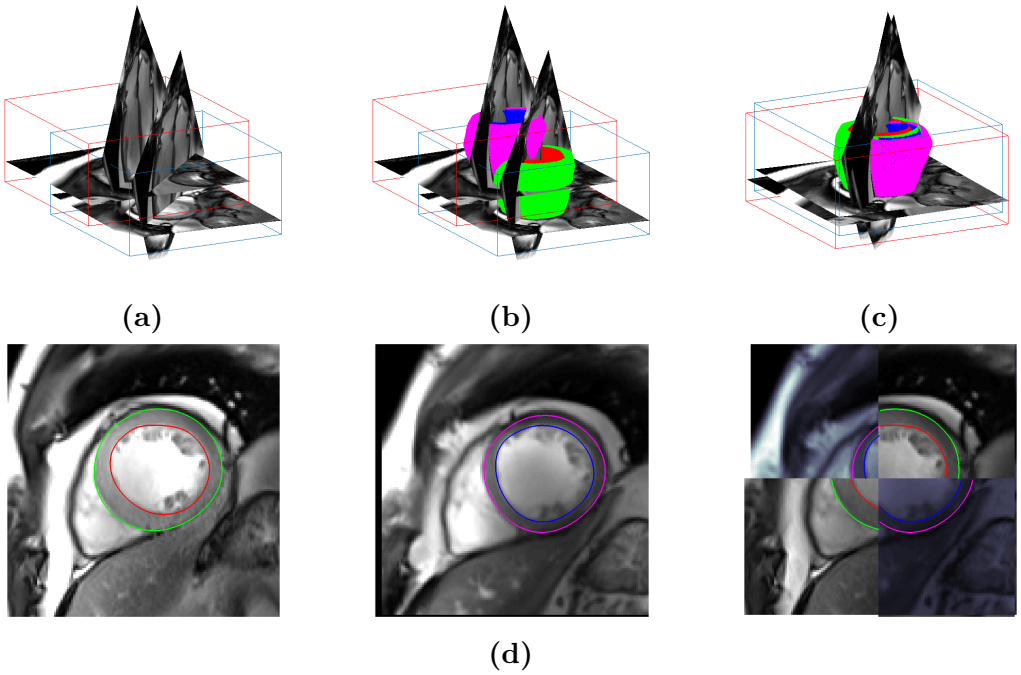


Figure 5.27: (a) and (b) show the positions of the cardiac cine MR data acquired during each imaging session as derived from the DICOM header information - a single SA and LA slice from each session is shown. (c) indicates the resulting positions if the coordinates of the most similar SA slice from each set are matched (the SA slices shown are not the slices that are matched). (d) shows the correspondence between volume cross-sections based on this slice matching.

the misalignment may be to select the most similar SA slices in each of the two datasets and match their coordinates, however this approach still leaves notable tissue misalignments as illustrated in figure 5.27d, where an SA cross-sectional image from both volumes is shown after slice matching. Rather, we require a registration approach that aligns the data based on image content. A preliminary rigid registration step will provide good initialisation for a subsequent non-rigid alignment, which will capture the more localised tissue differences.

Due to the previously-described intensity variation between MRI datasets, we choose to again utilise a local phase-based representation of the data, illustrated in figure 5.28, rather than the raw intensity form. Utilising this data representation, optimal alignment is achieved when a maximum correspondence of structures is found - structures which are characterised by the local phase values. Using the volumetric local phase-based representations, the quality of the correspondence is calculated via the mean squared difference metric. The optimal 3D rigid

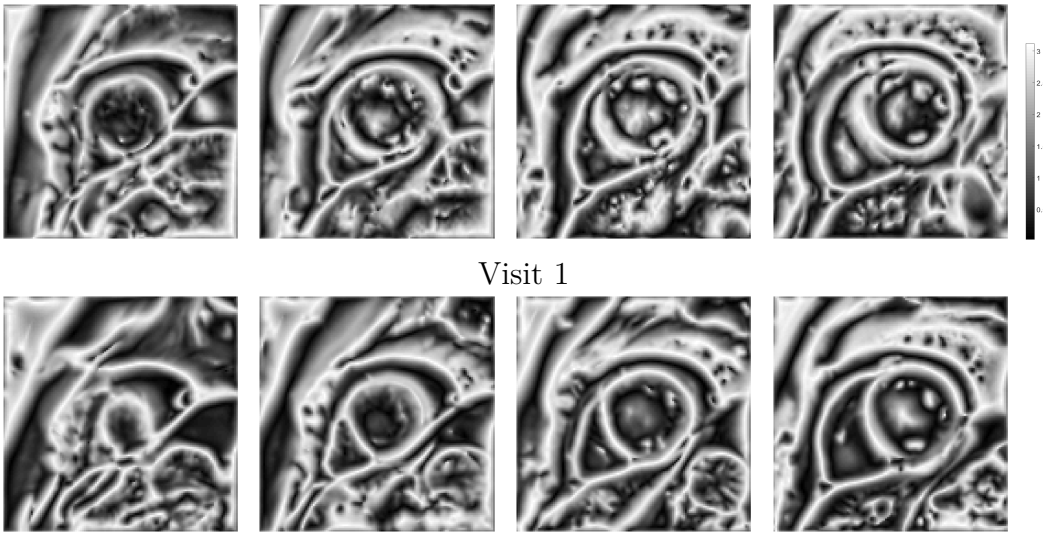


Figure 5.28: An illustration of the local phase used for calculating the optimal rigid alignment between the cine volumes.

transformation is calculated in a three-tiered multi-resolution framework, using the adapted gradient descent approach described in [91]. The above initialisation registration was implemented using the elastix software package [92, 180].

Non-rigid registration - Phase-based optical flow

While the rigid registration provides a good initial alignment between the cardiac MR volumes, it is insufficient to capture all the observed differences. In order to compensate for more localised differences that occur due to both the highly deformable nature of the tissue and the changes that occur in tissue over time, a non-rigid registration method is required. In this work we utilised the combined local-global variational optical flow framework discussed previously whilst using a local phase-based data representation - this modification is described in the next section

Local-global optical flow with phase-vector data term

The optical flow methods outlined in the introduction to image registration were all based on an intensity constancy assumption. However, in many applications this assumption is violated: therefore, there is considerable interest in adapting optical flow formulations to better cope with these situations. As has been discussed in previous sections local phase-based metrics are a particularly attractive alternative

to intensity-based metrics. However, while in the previous sections local phase values were used directly to evaluate correspondence, here we use the local phase vector (discussed in section 4.3.1). The local phase vector is superior to the local phase alone as it additionally distinguishes image structure based on the local orientation; this is particularly important for the deformable deformations applied here, where we wish to avoid incorrect correspondences of opposed edges. In this section we describe the formulation of the combined local-global optical flow method using the phase vector as the data term.

The final energy function may be expressed precisely as in section 4.1.2, apart from the substitution of the intensity signal by the local phase vector, r.

$$E(\mathbf{u}) = \int \mathbf{u}_{\mathbf{x}t}^T \mathbf{J}_\rho(\nabla_{\mathbf{x}t}\mathbf{r})\mathbf{u}_{\mathbf{x}t} + \alpha \left(\|\nabla \mathbf{u}\|^2 \right) d\mathbf{x} \tag{5.4}$$

where, as previously defined, there are multiple components to the phase vector, $\mathbf{r} = [r_1, r_2, r_3]$, which are combined in the structure tensor according to:

$$\mathbf{J}_\rho(\nabla_{\mathbf{x}t}\mathbf{r}) = \begin{bmatrix} K_\rho * (r_{1x}^2 + r_{2x}^2 + r_{3x}^2) & K_\rho * (r_{1x}r_{1y} + r_{2x}r_{2y} + r_{3x}r_{3y}) \\ K_\rho * (r_{1x}r_{1y} + r_{2x}r_{2y} + r_{3x}r_{3y}) & K_\rho * (r_{1y}^2 + r_{2y}^2 + r_{3y}^2) \\ K_\rho * (r_{1x}r_{1z} + r_{2x}r_{2z} + r_{3x}r_{3z}) & K_\rho * (r_{1y}r_{1z} + r_{2y}r_{2z} + r_{3y}r_{3z}) \\ K_\rho * (r_{1x}r_{1t} + r_{2x}r_{2t} + r_{3x}r_{3t}) & K_\rho * (r_{1y}r_{1t} + r_{2y}r_{2t} + r_{3y}r_{3t}) \end{bmatrix}$$

$$\begin{bmatrix} K_\rho * (r_{1x}r_{1z} + r_{2x}r_{2z} + r_{3x}r_{3z}) & K_\rho * (r_{1x}r_{1t} + r_{2x}r_{2t} + r_{3x}r_{3t}) \\ K_\rho * (r_{1y}r_{1z} + r_{2y}r_{2z} + r_{3y}r_{3z}) & K_\rho * (r_{1y}r_{1t} + r_{2y}r_{2t} + r_{3y}r_{3t}) \\ K_\rho * (r_{1z}^2 + r_{2z}^2 + r_{3z}^2) & K_\rho * (r_{1z}r_{1t} + r_{2z}r_{2t} + r_{3z}r_{3t}) \\ K_\rho * (r_{1z}r_{1t} + r_{2z}r_{2t} + r_{3z}r_{3t}) & K_\rho * (r_{1t}^2 + r_{2t}^2 + r_{3t}^2) \end{bmatrix} \tag{5.5}$$

Directly calculating the derivatives of phase-based quantities is not appropriate due to its inherent wrap-around property. Therefore, the phase vector derivatives are calculated directly from the filter outputs according to the equations defined in [56] and expressed for reference here.

$$\theta_1 = tan^{-1}\left(\frac{I_{R_2}}{I_{R_1}}\right) \tag{5.6}$$

$$\theta_2 = tan^{-1}\left(\frac{\sqrt{I_{R_1}^2 + I_{R_2}^2}}{I_{R_3}}\right) \tag{5.7}$$

$$\begin{aligned}\nabla^T \mathbf{r}(\mathbf{x}) &= \begin{bmatrix} r_{1x} & r_{1y} & r_{1z} \\ r_{2x} & r_{2y} & r_{2z} \\ r_{3x} & r_{3y} & r_{3z} \end{bmatrix} \\ &= \begin{bmatrix} cos(\theta_1)sin(\theta_2) \\ sin(\theta_1)sin(\theta_2) \\ cos(\theta_2) \end{bmatrix} \begin{bmatrix} cos(\theta_1)sin(\theta_2) \\ sin(\theta_1)sin(\theta_2) \\ cos(\theta_2) \end{bmatrix}^T \frac{I_{bp}\nabla^T \mathbf{I}_R - \mathbf{I}_R^T \nabla I_{bp}}{I_{bp}^2 + |\mathbf{I}_R|^2}\end{aligned} \tag{5.8}$$

$$r_t = \frac{(I_{bp})_{old}(\mathbf{I}_R)_{new} - (\mathbf{I}_R)_{old}(I_{bp})_{new}}{\left|(I_{bp})_{old}(\mathbf{I}_R)_{new} - (\mathbf{I}_R)_{old}(I_{bp})_{new}\right|} tan^{-1}\left(\frac{(I_{bp})_{old}(\mathbf{I}_R)_{new} - (\mathbf{I}_R)_{old}(I_{bp})_{new}}{(I_{bp})_{old}(I_{bp})_{new} + (\mathbf{I}_R)_{old}^T(\mathbf{I}_R)_{new}}\right) \tag{5.9}$$

Optical flow approaches assume small motions between images, as the initial linearisation via the Taylor expansion only holds if the displacement vectors are small. Therefore, in order to capture greater differences, it is common for optical flow methods to be implemented within multi-resolution strategies, beginning at a coarse resolution and updating the results with finer resolution refinements. Such an approach is adopted here; each iteration generates a new small displacement field that is combined with that calculated at the previous scale. With each estimate, the original moving image is deformed towards the target image. In addition to using multiple resolutions, we adopt a multi-scale approach where sequential registration steps use local phase values estimated at decreasing spatial scales. The registration is performed initially using a large spatial scale that will identify broader image structures, before utilising a smaller spatial scale which will capture further signal detail and refine the final registration result.

Parameter selection

We briefly discuss the effect of three parameters on the registration performance as these are considered to have the most influence on the final result. The first two parameters, ρ and α, relate directly to the registration framework, modulating the local and global smoothing respectively. The third parameter affects the nature of the input data via the scale of the local phase estimation.

Local and global smoothing - α and ρ

As described in the introduction to optical flow, both the α and ρ parameters, indicated in equation 5.4, are introduced to enable the calculation of the flow field solution, $\mathbf{u}$. In the context of the combined local-global method they perform complementary tasks; ρ controlling the local averaging of flow, and α influencing the global smoothness of the entire motion field. The effect of both parameters is illustrated in figures 5.29 and 5.30, where the registration of two cardiac frames, and two separate cardiac studies are demonstrated respectively.

It is clear that smaller parameter values give the algorithm increased freedom to find a minimiser to the data term of the cost function. However, such a result is not always desirable as the result is an over-fitting and the deformation field is too irregular - this is of particular importance where the registration of medical data is considered, requiring a deformation that is physiological. Insight into the nature of the computed transformation is given by considering the Jacobian matrix, defined in equation 5.10 [171], where $\mathbf{u}$ is the calculated displacement vector and I is the identity matrix.

$$J(I+\mathbf{u}) = \begin{bmatrix} 1+\frac{\delta u}{\delta x} & \frac{\delta u}{\delta y} & \frac{\delta u}{\delta z} \\ \frac{\delta v}{\delta x} & 1+\frac{\delta v}{\delta y} & \frac{\delta v}{\delta z} \\ \frac{\delta w}{\delta x} & \frac{\delta w}{\delta y} & 1+\frac{\delta w}{\delta z} \end{bmatrix} \tag{5.10}$$

The determinant of the Jacobian (which is also commonly referred to directly as the Jacobian) is an important quantity, as it describes the volume-changing characteristics of the deformation; values greater than 1 indicate volumetric expansion, while those less than 1 are indicative of contraction. These characteristics of the Jacobian are clearly shown in figure 5.29c, which illustrates the deformation calculated between two cardiac cine phases. Of similar importance is the recognition of negative determinant values, which indicate voxel disappearance associated with folding of the image domain. The example in figure 5.30 is more representative of the deformation expected in the current registration task. Negative Jacobian values are seen if small α and ρ parameters are used, as indicated in figure 5.30c. We choose

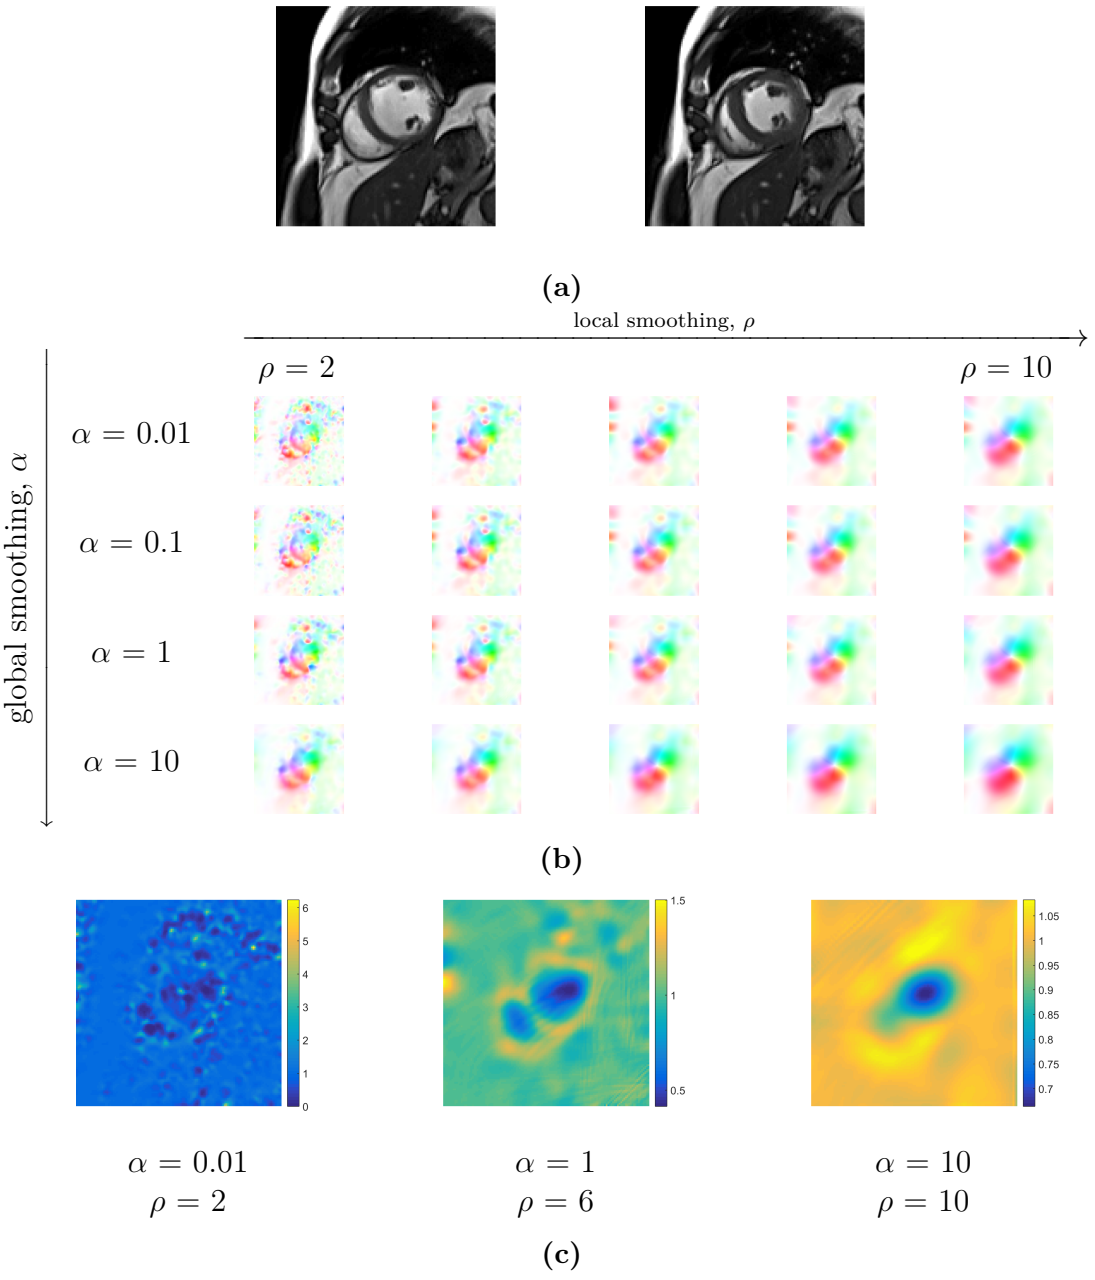


Figure 5.29: The effect of the local and global smoothing parameters, ρ and α is illustrated. In (a) the two cardiac frames that are registered are shown. (b) illustrates the variation in the motion field computed, where the colour indicates the vector direction, and hue the magnitude. In (c) examples of the computed Jacobian determinants are shown.

our parameters based on the observed smoothness of the field and found during tests that an α value of approximately 5, and a ρ of between 6 and 10 worked well.

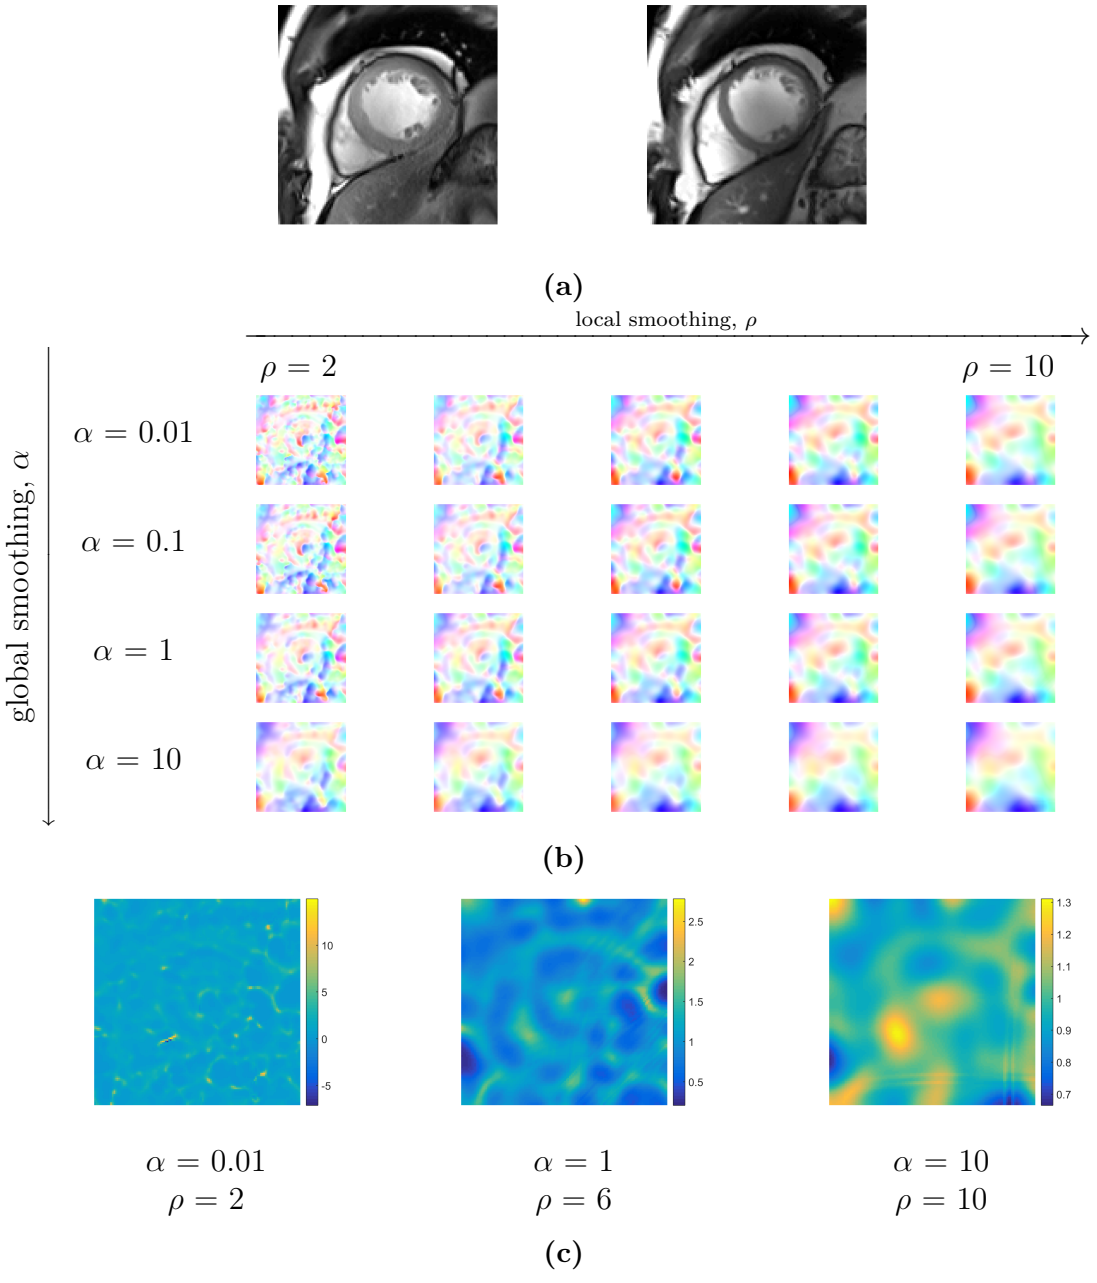


Figure 5.30: The effect of the local and global smoothing parameters, ρ and α is illustrated. In (a) the two SA cardiac images that are registered are shown. (b) illustrates the variation in the motion field computed, where the colour indicates the vector direction, and hue the magnitude. In (c) examples of the computed Jacobian determinants are shown.

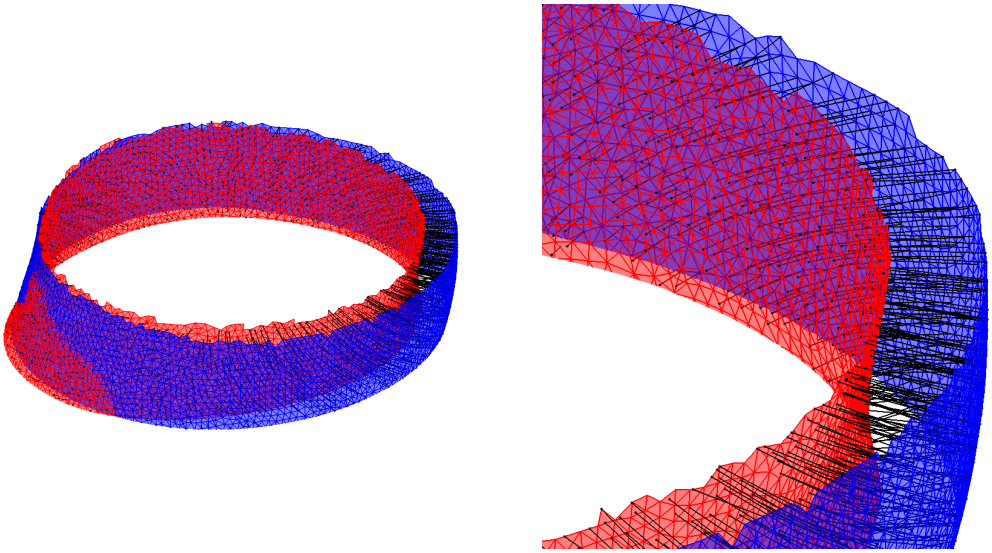


Figure 5.31: Illustration of the distances calculated between meshes for registration validation. The black lines connect the vertices of one mesh to the closest point on the second mesh.

Local phase scale - λ

As previously indicated, the λ parameter determines the scale of the structures identified by the local phase estimation. Due to the increasing detail available as the scale is lowered, the parameter is utilised as a means of iteratively refining the result. A first registration is performed using a comparatively large λ value ($\lambda = 10$), so that larger tissue structures are broadly aligned. With a further two registration steps, smaller values ($\lambda = 6, \lambda = 3$) are utilised that contain finer image detail and thus lead to improved registration accuracy.

5.4.2 Experiments and results

Here, we present the results acquired using the discussed registration framework on the 19 patient datasets available. Firstly, the results of the rigid initialisation will be qualitatively presented, before we go on to show non-linear registration results. The latter non-linear results are quantified by means of the distances calculated between the endo- and epi-cardial surface meshes generated for each visit's data, as illustrated in figure 5.31.

Rigid

The rigid results are presented qualitatively in figure 5.32. A series of results is shown that illustrate the relative alignment of the data from the two hospital visits before and after the rigid registration. In each case we see that the data has been brought into broadly the same coordinate-space with the majority of the rigid translational and rotational misalignments accounted for. On the right-hand side of the figure a number of SA slices are shown after registration; these images indicate that after the rigid registration there is still a need for further non-linear refinement of the results.

Non-linear

The results for the non-rigid registration were quantified by means of distances calculated between meshes, generated from the SA endo- and epi-cardial delineations. In addition to the quantitative results, figures are shown to indicate the improved alignment of cardiac structures after non-rigid registration.

Firstly, some examples of LA cross-sectional data are shown in figure 5.33. The images show LA cross-sections taken from the first visit volumes. Overlain on the image data are the intersections of the generated endo- and epi-cardial meshes from both visits with the given slice. The red and green contours indicate the segmentations for the first visit data (that corresponding to the image data shown), while the blue and magenta contours are taken from the fourth visit data. Comparing the left and right columns of the figure, which show the mesh positions before and after registration respectively, we can see that in each case the correspondence of the mesh positions is improved after registration. This also results in an improved fidelity of the fourth visit meshes (blue and magenta) to the image data shown.

The alignment of the data can be illustrated by considering the SA cross-sections from the two visits. These results are illustrated in figures 5.34 and 5.35. Both figures are divided into four columns, displaying checkerboard illustrations of results, where each line represents a single example. The first two columns show the correspondences of a single slice before registration (with inverted checkerboard display), while the third and fourth show the same data after registration. Both

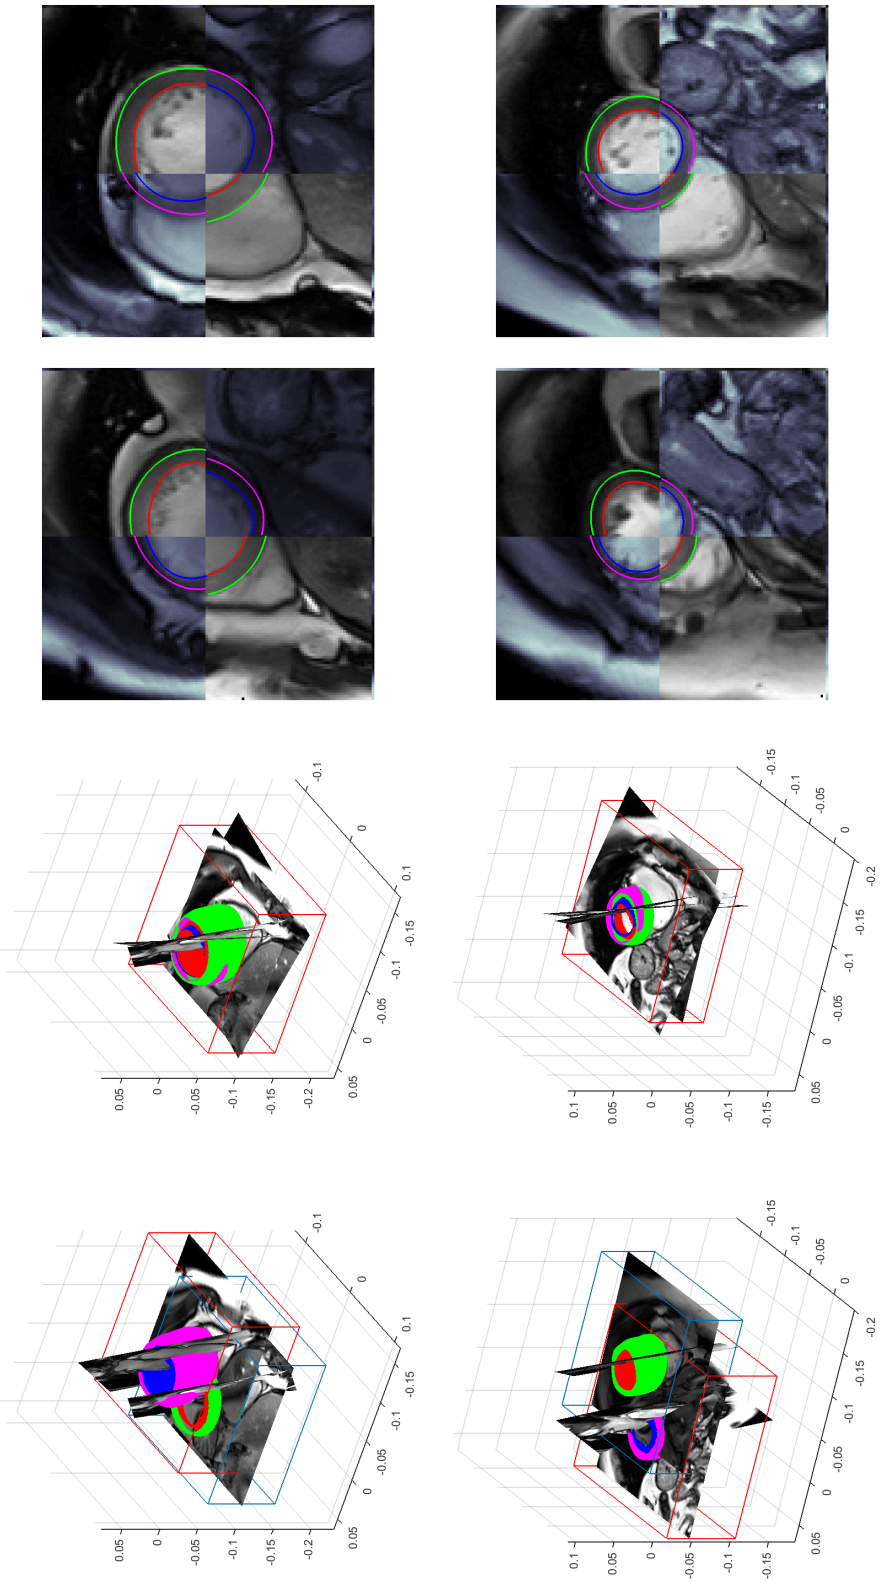


Figure 5.32: Results obtained after the rigid registration step. The first column indicates the original data positions and the second column the resulting, rigidly-registered positions. Two example SA slices are displayed with checkerboard appearance indicating the alignment of visit 1 and visit 4 data.

figures provide similar illustrations, but in the second the mesh positions are additionally overlain. These delineations further illustrate the improved alignment of the data after registration.

Quantification of the results is accomplished via the closest distances between meshes. In figure 5.36, an example of the relative position of the endocardial surfaces both before and after registration is given. The distances calculated are displayed on the mesh surfaces, illustrating the local positional discrepancies between the corresponding surfaces. Finally, these calculated distances are displayed in the histogram plot; in blue, the distances calculated using the initial mesh positions are shown, while in red the improved alignment after registration is indicated.

In figure 5.37 a number of further examples are given showing the post rigid- and post non-rigid-registration distances calculated for either the endo- or epicardial surfaces. In each case, the improvements after application of the non-rigid registration are clear to see.

A summary of all the results is given in figure 5.38. The first plot indicates the distances before and after non-rigid registration for the endocardial surface, while the second plot shows the corresponding results for the epicardium. The mean and standard deviation of the endocardial distances are reduced from 1.98mm $\pm$ 1.59mm to 1.40mm $\pm$ 1.32mm, while the epicardial distances dropped from 1.66mm $\pm$ 1.33mm to 1.34mm $\pm$ 1.14mm.

Comparison with other methods

In this section, to place our results in context, we present registration results using two alternative algorithms.

The first is an implementation of the local-global optical flow algorithm using intensity in place of local phase for similarity quantification. This follows therefore the original formulation of the method, and allows us to investigate the importance of our choosing a local phase-based metric for image alignment quantification. Similar to the local phase-based implementation, a three-level multi-resolution

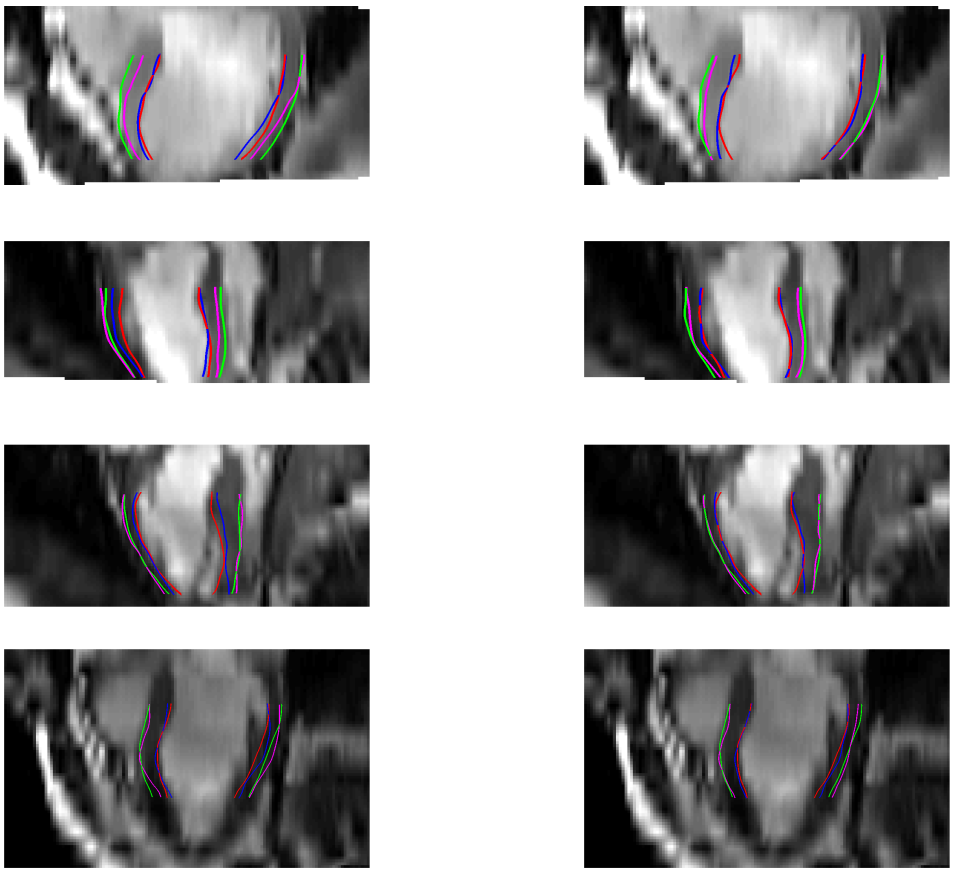


Figure 5.33: Examples of the non-rigid registration results. All images show cross-sections taken from the first visit cine volumes. Overlain are the intersections of the endo- and epi-cardial meshes from both visits with this plane. The red and green delineations are from the first visit, while the blue and magenta delineations are from the fourth visit. The left-hand column indicates the mesh positions after rigid registration. In the right-hand column the fourth visit meshes (blue and magenta) have been deformed according to the calculated non-rigid deformation.

scheme is used; with local smoothing $\rho = 6$, we present quantitative results using a range of global smoothness parameter values, α.

The second approach is an FFD-based approach with NMI similarity metric. Such methods have been commonly utilised in the literature for cardiac registration applications [154, 10, 164] and therefore the results give an indication as to the performance of our method in relation to the state-of-the-art. This B-spline-based approach has been implemented using the elastix toolbox; a three-level multi-resolution scheme is used, optimised via a gradient descent approach. Here we investigate the algorithm's performance across a range of grid spacings, which are

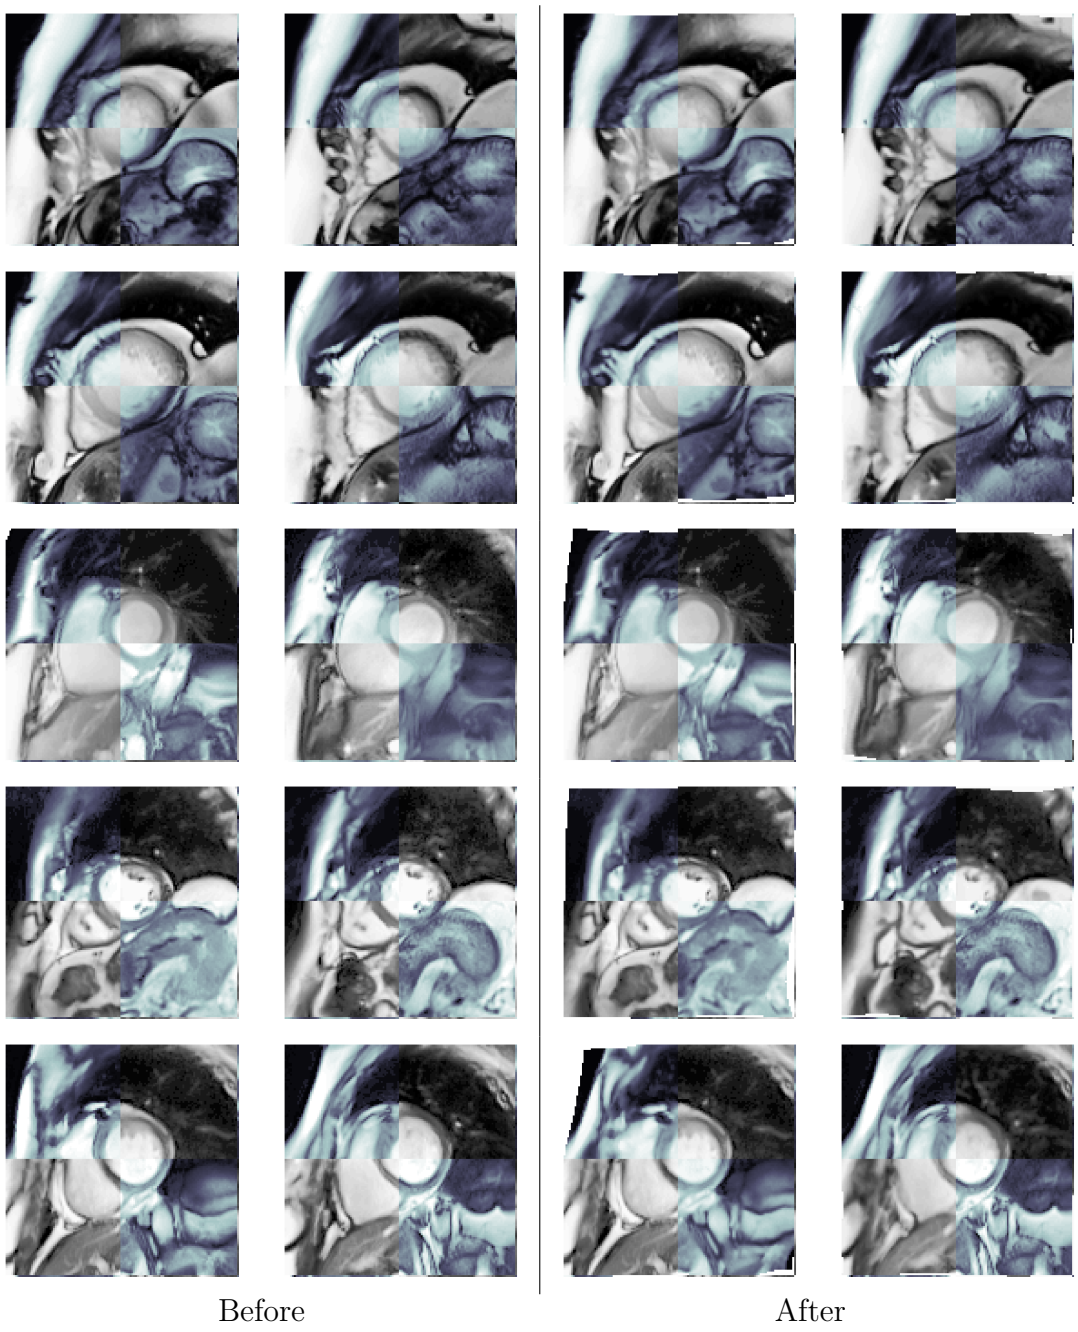


Figure 5.34: Example results of SA alignment before and after non-rigid registration. The checkerboard displays show data from both the first and fourth visits. The first two columns indicate the alignment before registration, while the third and fourth columns show the same data after registration.

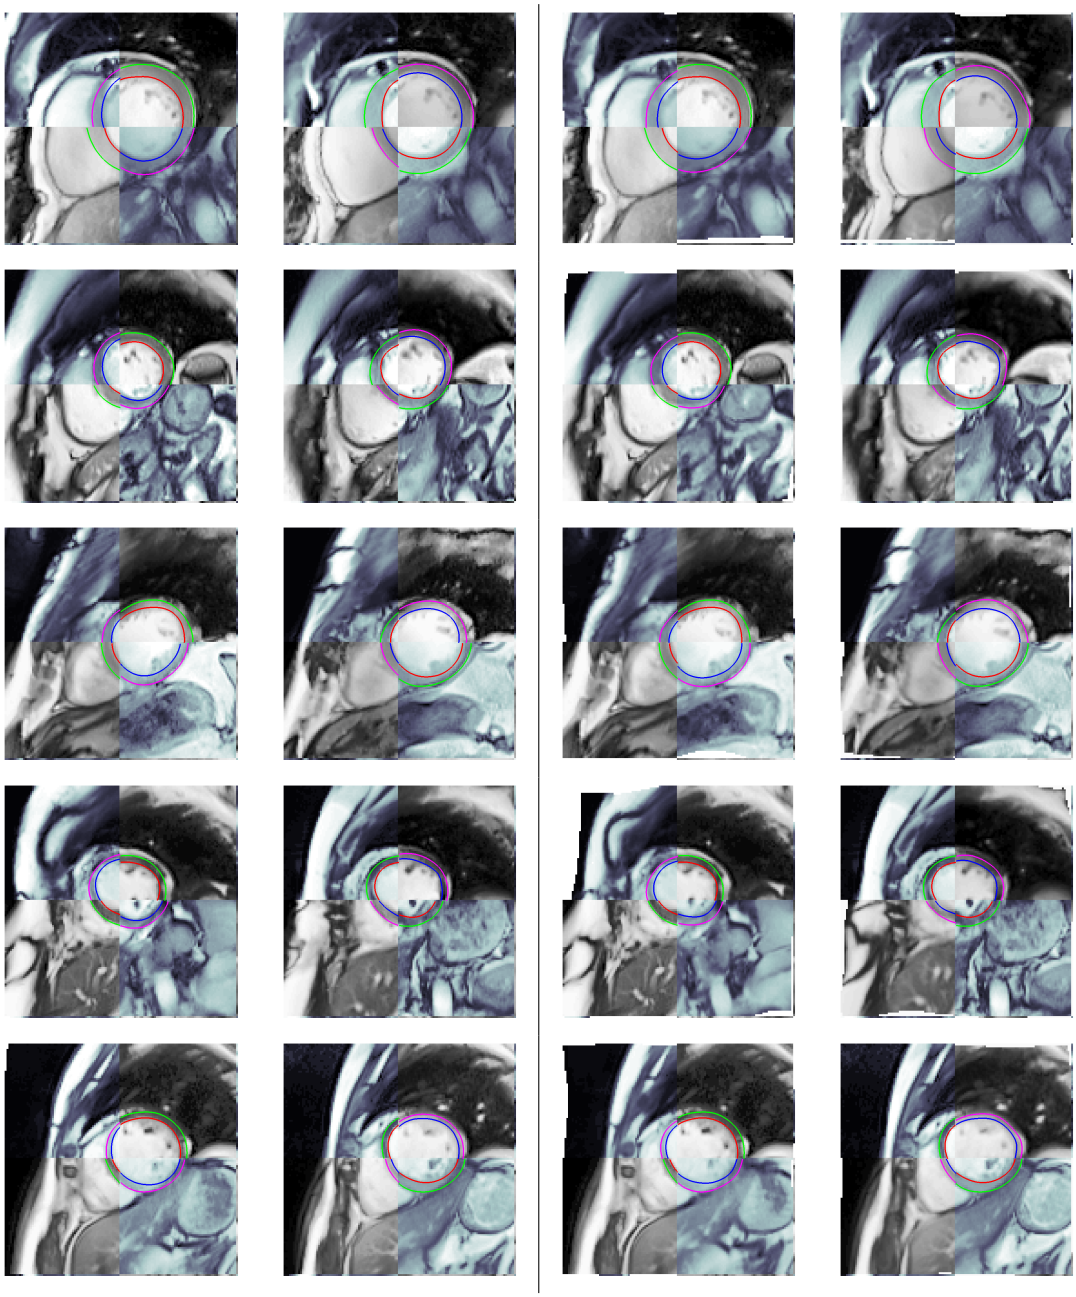
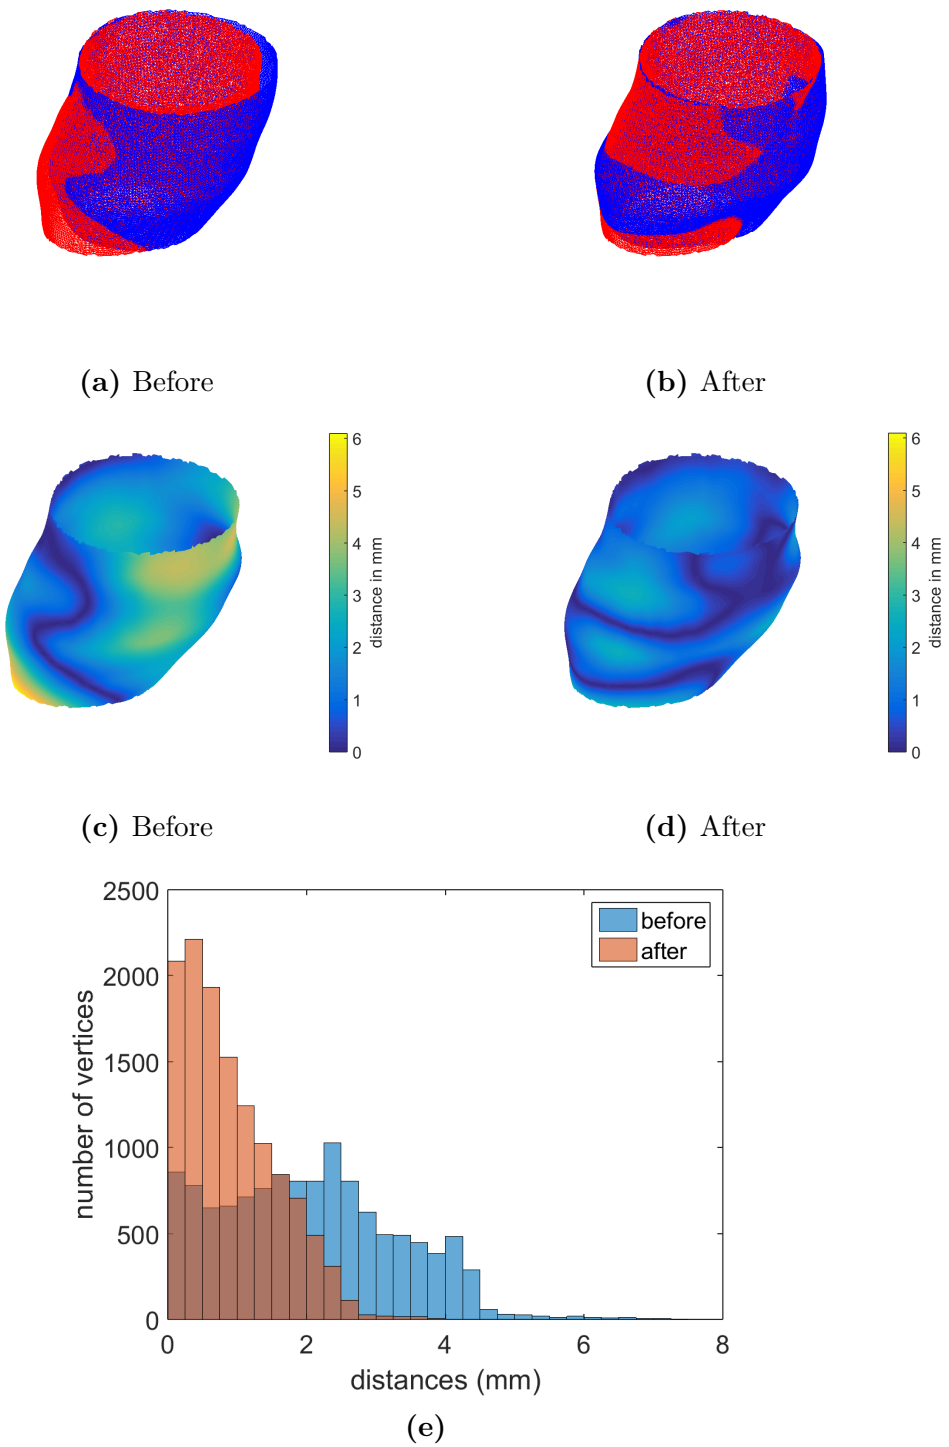


Figure 5.35: Example results of SA alignment before and after non-rigid registration. The checkerboard displays show data from both the first and fourth visits. The intersections of the respective endo- and epi-cardial meshes are indicated in red and green for the first visit, and in blue and magenta for the fourth visit. The first two columns indicate the alignment before registration, while the third and fourth columns show the same data after registration.

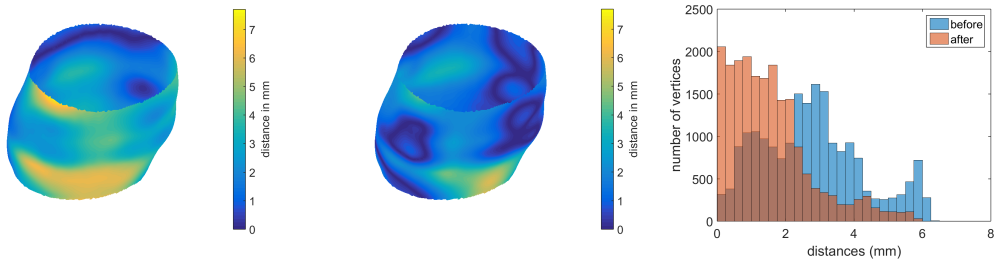


(a) Before **(b)** After

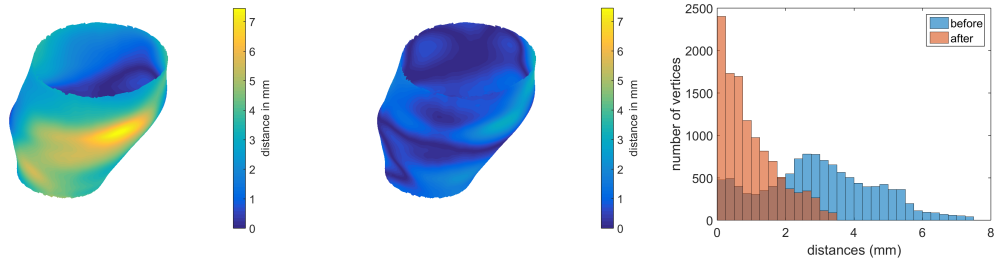
(c) Before **(d)** After

(e)

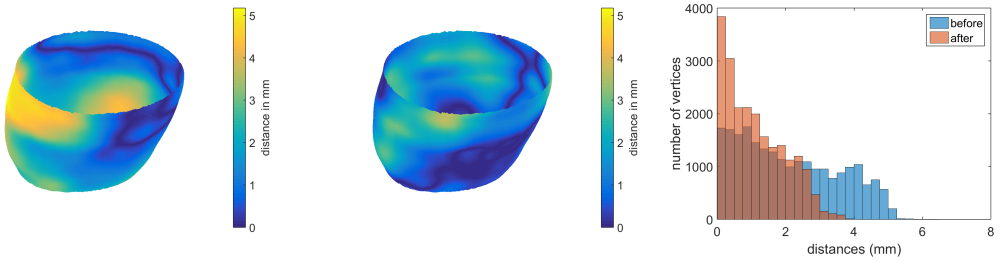
Figure 5.36: An example result showing the endocardial surfaces from both visits. (a) illustrates the relative positions of the first visit in red and the fourth visit in blue before registration. In (b) the blue mesh is deformed based on the registration result. The colour surface in (c) and (d) indicates the local distance to the other mesh before and after registration respectively. The histogram in (e) displays the calculated distances before registration in blue and after registration in red.



(a)



(b)



(c)

Figure 5.37: In (a), (b), and (c), a series of example results is given. The first and second columns show colour surfaces that indicate the distances calculated to the corresponding mesh from the fourth visit before and after registration respectively. In the third column the calculated distances before and after registration are shown by the blue and red histograms.

reported according to their specification at the final, highest resolution level.

Figures 5.39 and 5.40 summarise the distances calculated between corresponding endo- or epi-cardial surfaces according to alpha value or grid spacing, where the blue line indicates the initial distances and the red line indicates the distances after registration. In both cases we are varying the tradeoff between matching image content according to the respective metrics, intensity or NMI, and preserving the smoothness of the final transformation calculated; from left to right in both figures the smoothness of the transformations is increased. These results are complemented by qualitative examples of the registration results in figures 5.41 and 5.42.

It is clear from figure 5.39 that the intensity-based optical flow approach behaves

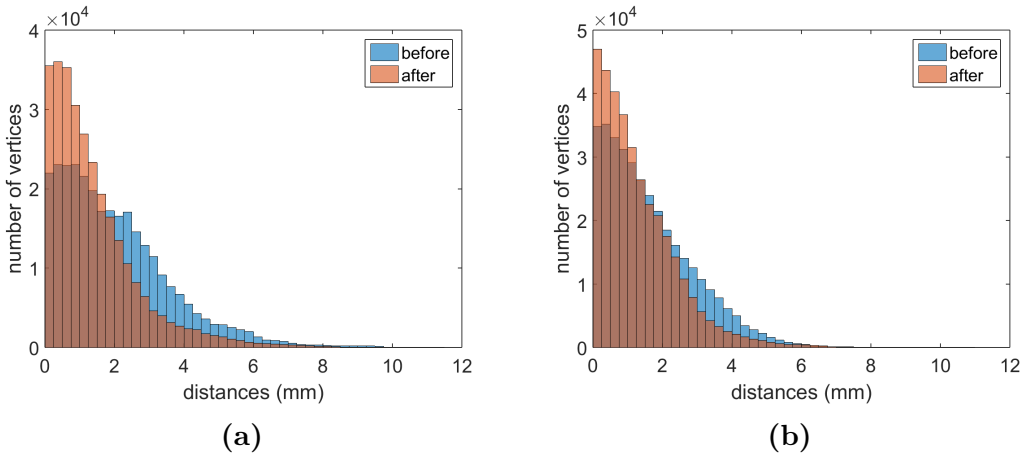
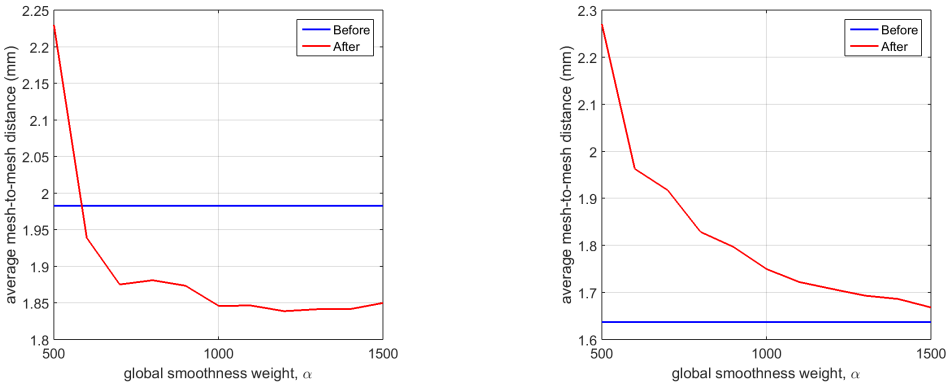


Figure 5.38: These histograms summarise the distances calculated between the cine meshes from each visit. The number of mesh vertices per distance bin are plotted in blue before registration and in red after registration. The distances calculated for the endocardial and epicardial surfaces are plotted in (a) and (b) respectively.

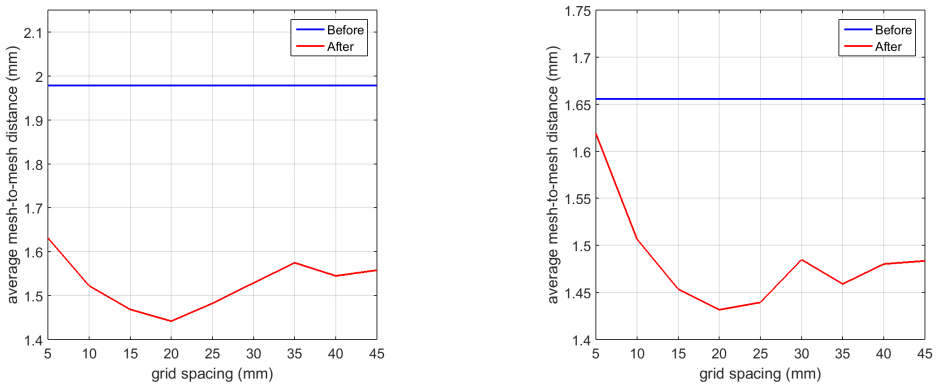
poorly. With less global smoothness, as a result of small α values, the final mesh distances are worse which suggests that the metric is not able to guide the algorithm to the correct data correspondence. At these values of α we further observed qualitatively that the resulting deformation fields are too loose and become physiologically implausible. While the distances improve with increasing global smoothness, we do not observe a significant reduction compared to the non-registered mesh distances. The example illustrated in figure 5.41c shows how the registration fails to align the data, this is particularly clear when considering the left ventricle at the lowest global smoothness which is shown in the left-most image of figure 5.41c. This does not improve at higher values of α, when there is more global smoothness, as the deformation becomes more restricted.

The second algorithm based on B-splines and NMI has been commonly used in the literature, and thus we expected good results in both endo- and epi-mesh alignments. In the literature, results have been reported using a grid spacing of 5 x 5 x 5mm [10], however, here results at this grid spacing were substantially worse than ours and further to this we observed that, using this grid resolution, the transformations calculated again became non-physiological, see for example the left-most image in figure 5.42c. We subsequently did a systematic study of the effect of grid size as well as other aspects of the algorithm such as the optimiser and sampling



(a) Endocardium **(b)** Epicardium

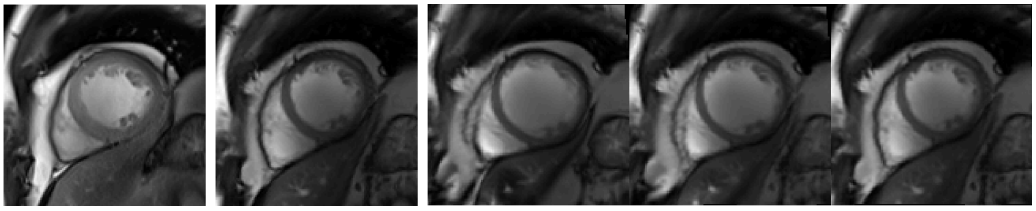
Figure 5.39: Distances calculated between cine meshes from each visit using the intensity-based optical flow approach for a range of global smoothness parameters, α, where the blue line indicates the initial errors and the red line indicates the distances after registration.



(a) Endocardium **(b)** Epicardium

Figure 5.40: Distances calculated between cine meshes from each visit using the NMI-based B-spline registration approach for a range of grid spacings, where the blue line indicates the initial errors and the red line indicates the distances after registration.

protocol. The best results were obtained at a grid spacing of 20 x 20 x 20mm with the endo- and epi-cardial distances reduced to 1.44mm ± 1.30mm and 1.43mm ± 1.21mm respectively. We observe therefore that our results compare favourably with the state-of-the-art literature with slight improvements in the final results for both the endo- and epi-cardial surfaces at 1.40mm and 1.34mm respectively.



(a) **(b)** **(c)**

Figure 5.41: (a) and (b) show the first and fourth visit data respectively. In (c) registration results using the intensity-based optical flow approach are shown at α values of 500, 1000, and 1500 respectively. We see how the algorithm fails to find the correct alignment.



(a) **(b)** **(c)**

Figure 5.42: (a) and (b) show the first and fourth visit data respectively. In (c) registration results using the B-spline-based NMI approach are shown at grid spacings of 5 x 5 x 5mm, 25 x 25 x 25mm, and 45 x 45 x 45mm. We note in particular the non-physiological deformation in the left-most image of (c).

Application of deformation fields to LGE data

Previously in this chapter, the rigid alignment of the LGE image slices to the cine volumetric datasets from each visit was performed. Based on this pre-alignment, here we use the transformations calculated via the registration of the cardiac cine volumetric data to align corresponding LGE interpolated, volumetric data. These processes, indicated in figure 5.43, bring all the CMR data into a single coordinate frame enabling their simultaneous clinical assessment.

A qualitative indication of the results attainable is given in figure 5.44, where a number of checkerboard display SA slices are shown. As in the previous assessment of 3D registration, a quantitative analysis considers the alignment of the respective first and fourth visit meshes. Figure 5.45 illustrates how the reduction in distances of cine meshes can be reflected in a reduction in LGE mesh distances upon application of the cine registration result.

In figure 5.46 the combined results from all datasets are shown. While some degree of improvement is seen in the distances calculated for the endocardium,

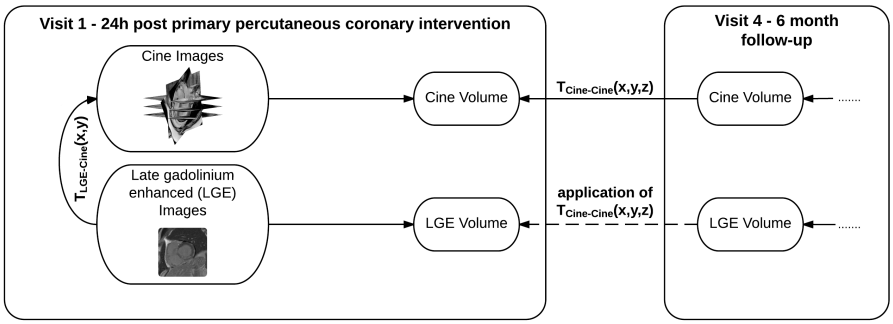


Figure 5.43: The described framework for multi-CMR-image registration. The LGE images have been registered to the cine data in each visit dataset. The 3D registration, $T_{cine-cine}(x, y, z)$ is calculated between the cine volumes from each visit. Finally, the registration result, $T_{cine-cine}$, is applied to the reconstructed LGE volumes.

the results are less clear than those acquired for the cine data. The inaccuracies that the quantitative results suggest are largely a consequence of the issue of segmentation differences, in addition to some non-rigid mis-registration that still exists between the LGE and cine data.

Analysis of scar

The alignment of LGE CMR data from longitudinal acquisitions allows for the comparison of scar tissue identified during both visits. Using the results of the 3D registration described in the previous section, here we present some initial results which illustrate the 3D analysis possible.

As a result of segmentation variation, the myocardial masks for the data from each visit do not match. For the purposes of the scar analysis carried out, a new myocardial mask is calculated which is common to both visit datasets; the edges of the new mask lie at equal distances from the visit delineations. Some examples are indicated in figure 5.47: the respective visit delineations are overlain in the first two columns, while the third and fourth columns show the edges of the newly-calculated, common segmentation on the same image data.

In figures 5.48, 5.49, and 5.50 example results are given of hyper-enhancement identified from LGE data after 3D registration for a number of patients. In each case, the respective hyper-enhanced tissue in the initial 24h acquisition and the 6 month follow-up images are shown. Firstly, these figures give an indication as to the

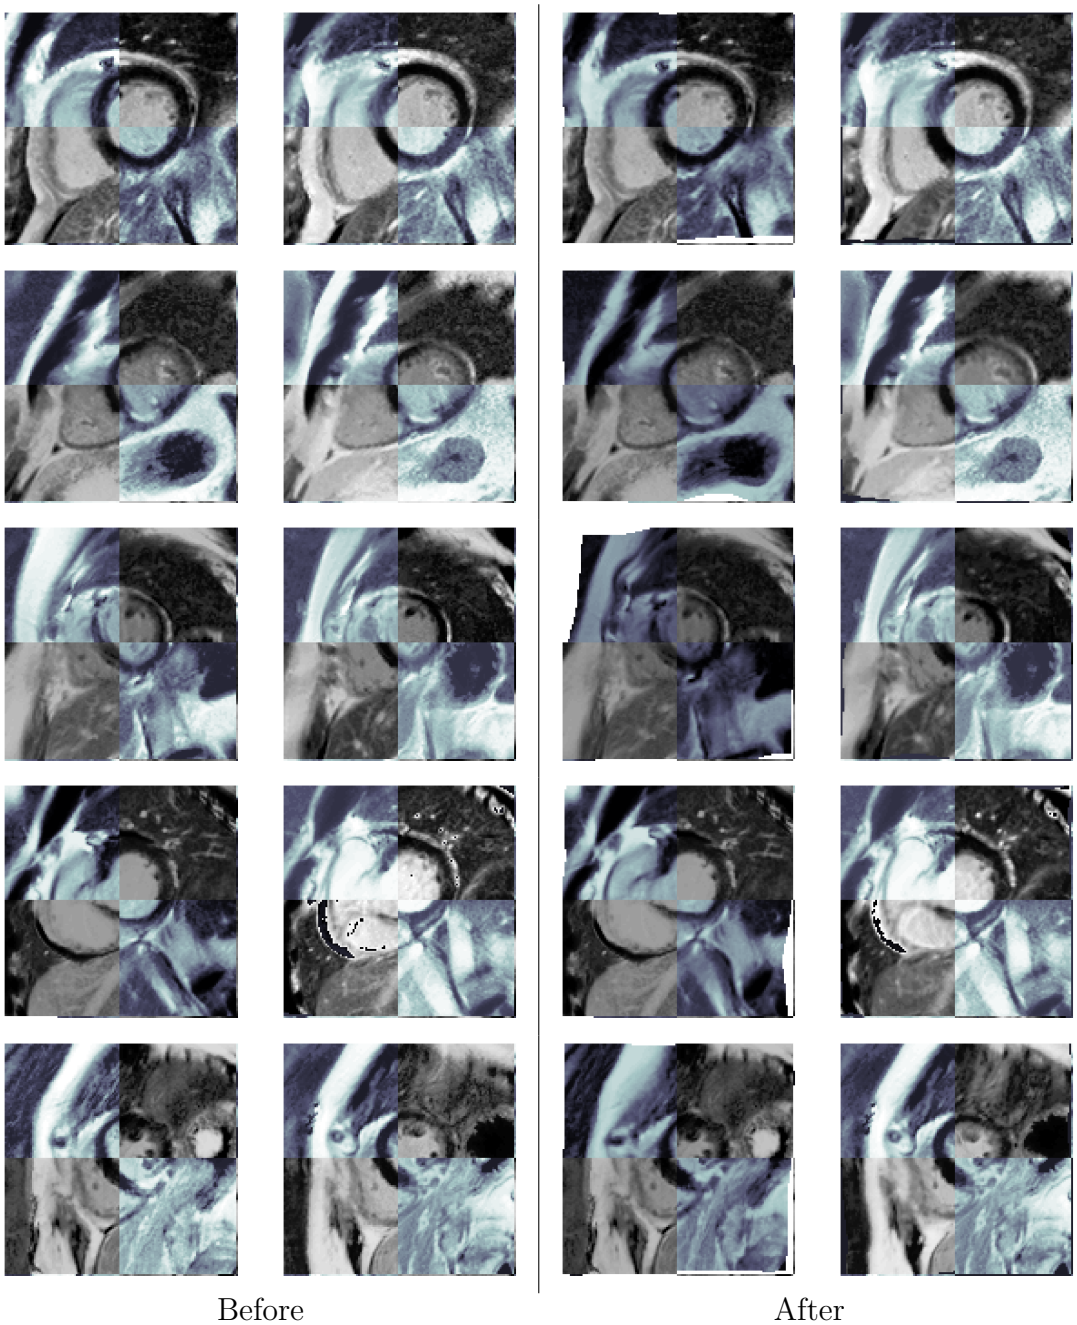
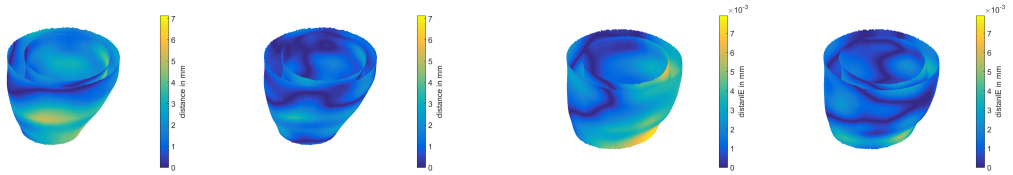
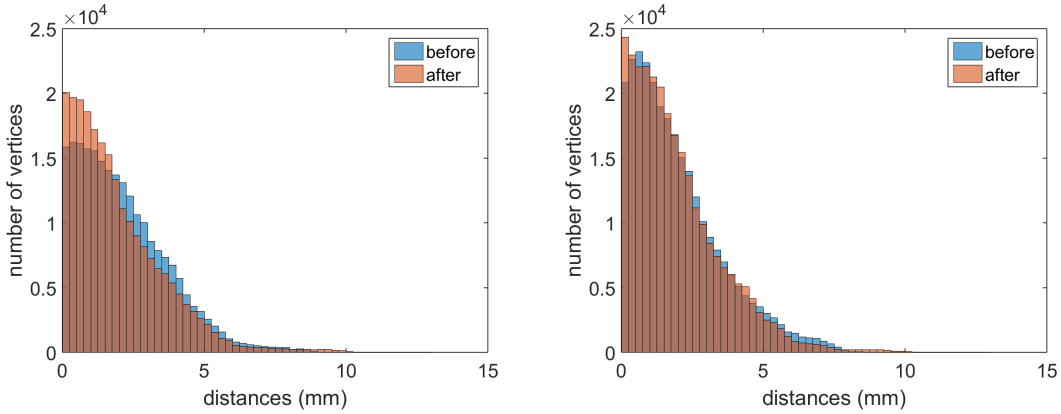


Figure 5.44: Example results of SA alignment before and after non-rigid registration. The checkerboard displays show data from both the first and fourth visits. The first two columns indicate the alignment before registration, while the third and fourth columns show the same data after registration.



(a) Cine endocardial distances **(b)** LGE endocardial distances

Figure 5.45: The meshes of the cine data compared to the meshes of the LGE data; the same registration field is applied to both sets of meshes.



(a) endocardium **(b)** epicardium

Figure 5.46: These histograms summarise the distances calculated between the LGE meshes from each visit. The number of mesh vertices per distance bin are plotted in blue before registration and in red after registration. The distances calculated for the endocardial and epicardial surfaces are plotted in (a) and (b) respectively.

respective size and location of scar tissues identified in the longitudinal acquisitions. In addition, we look to provide some quantitative illustration of the changes in tissue characterisation between the first and fourth visit. Here we show results in terms of myocardial hyper-enhancement transmurality. Transmurality is estimated by considering the myocardial mask in segments of 5 degrees - for each segment the number of hyper-enhanced pixels is calculated as a ratio of the total number of pixels in the segment and reported as a percentage measure. Due to the narrow angular width of each segment this calculation is taken as an estimation of the transmurality.

3D renderings of the myocardial regions under consideration, formed from the myocardial masks, and the hyper-enhanced regions are shown in the figures; in sub-figures (a) the myocardial regions are illustrated in relation to LGE LA data; sub-figures (b) and (c) illustrate the regions of hyper-enhancement identified within the myocardial region for each case. The volumes of hyper-enhancement

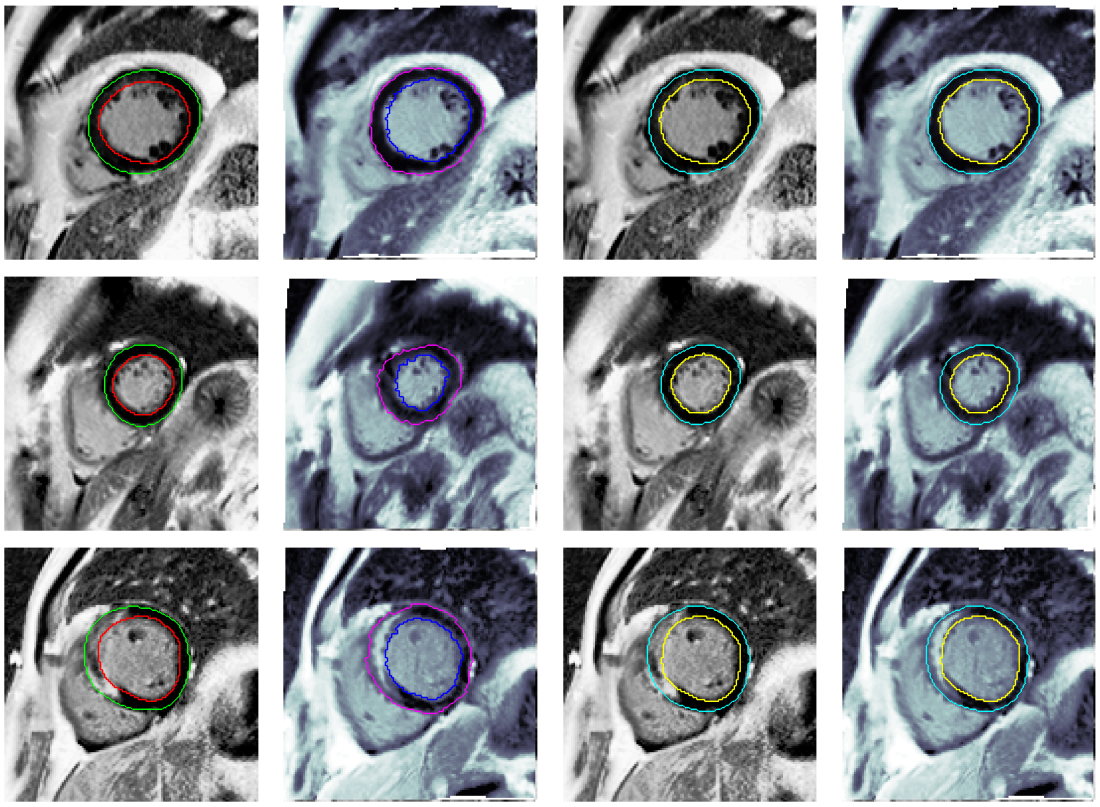


Figure 5.47: Illustration of the new myocardial delineations that lie at equal distances from the contours of each visit. The new delineations are displayed in the third and fourth columns, overlain on images from both visits in yellow and cyan.

are colour-coded according to the local degree of transmurality, from 0% in blue to 100% in yellow. Finally, the change in transmurality is reported in histogram form in (d); the proportions of myocardial mask undergoing positive or negative transmurality changes are shown.

The changes can be clearly seen in the 3D visualisations, and the observations can be related to the imaged tissue by considering the depicted LGE slices where yellow coloured pixels are used to indicate the regions of hyper-enhancement. The histogram plots provide a quantitative assessment of the changes: the x-axis reports the percentage change in hyper-enhanced myocardial transmurality, while the y-axis indicates the percentage of the myocardial-mask that undergoes these changes. We see a trend in these examples, namely that there is reduced hyper-enhancement in the follow up assessments.

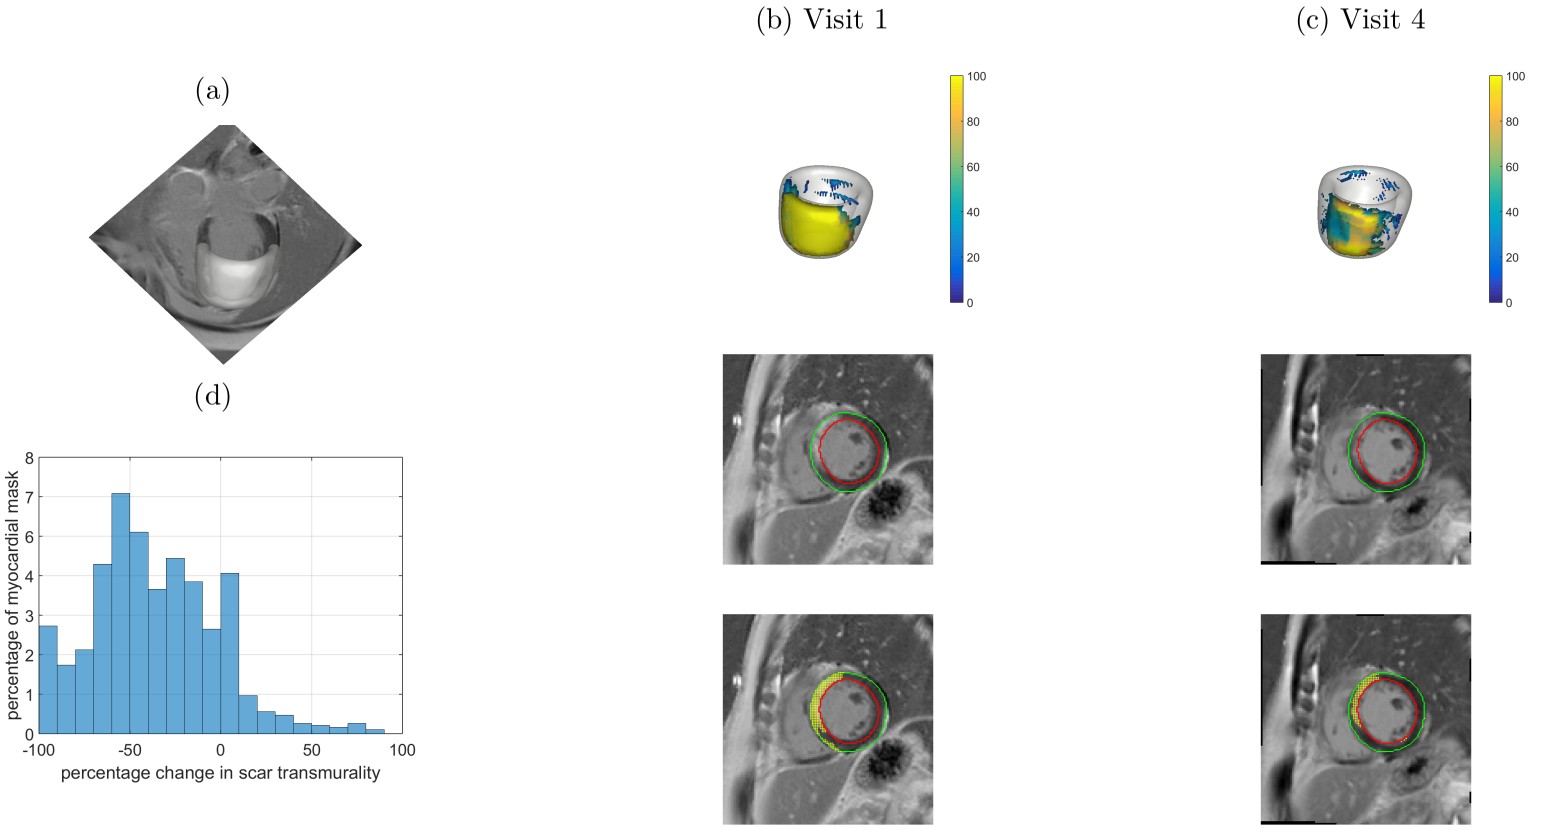


Figure 5.48: Example result of longitudinal scar assessment showing hyper-enhancement identified in follow-up acquisitions from the same patient after myocardial infarction. In (a) an LGE slice is shown with a 3D rendering of the myocardium assessed. (b) and (c) show the scar identified in both the first and fourth visits; the 3D illustration shows the myocardial volume assessed with the blue-yellow coloured rendering indicating the regions of hyper-enhancement identified; the colour-axis indicates ranges of hyper-enhancement transmurality from 0% in blue to 100% in yellow. In the bottom most image of both (b) and (c) the pixels identified as enhanced are indicated in yellow. The histogram plot in (d) shows the change in transmurality identified in the follow-up data (shown in (c)) compared to the first visit (shown in (b)).

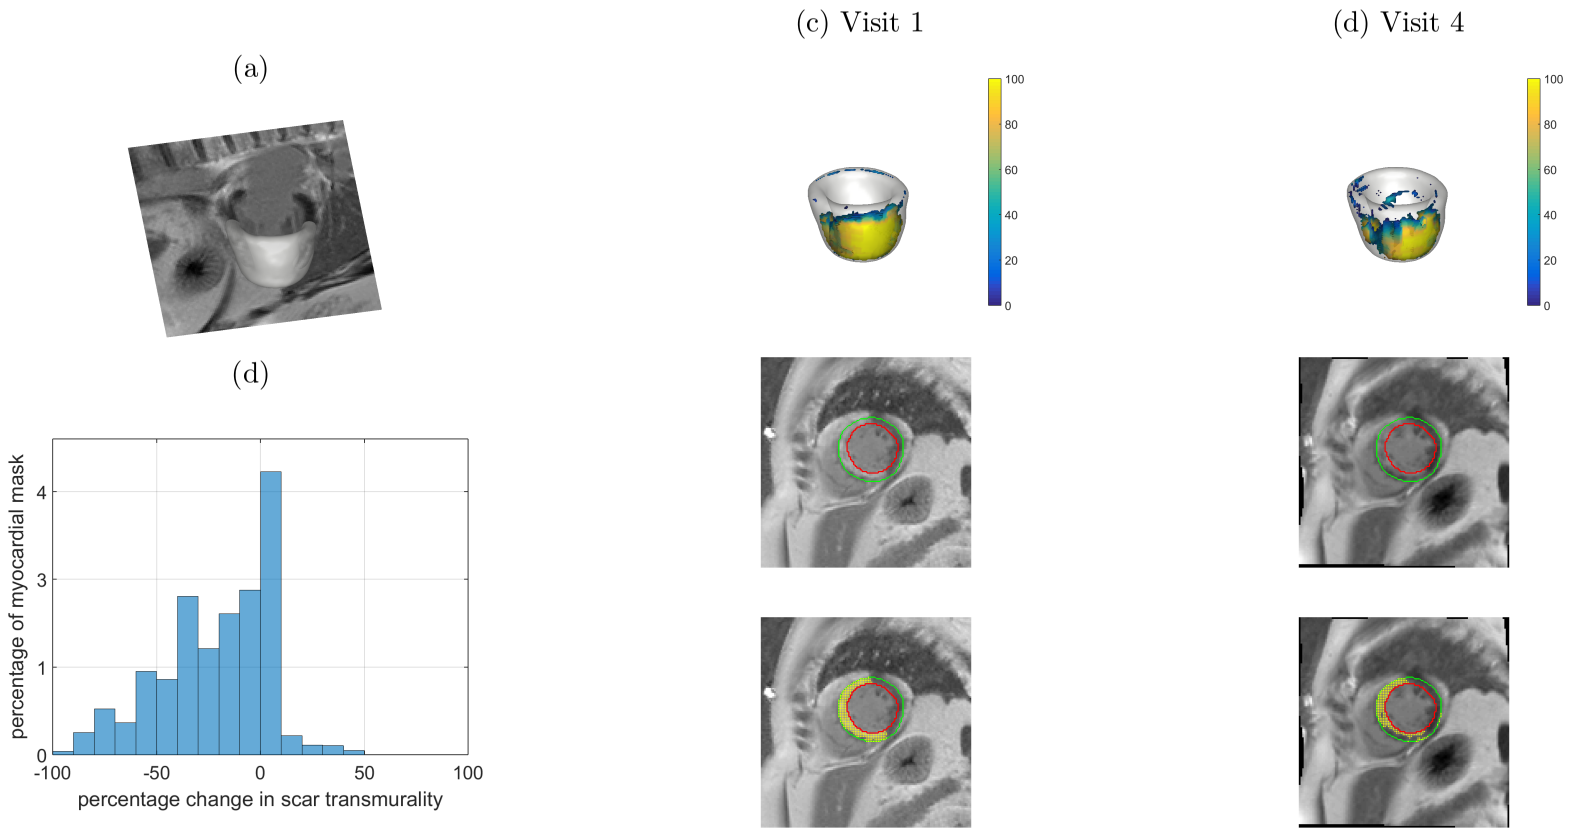


Figure 5.49: Example result of longitudinal scar assessment showing hyper-enhancement identified in follow-up acquisitions from the same patient after myocardial infarction. In (a) an LGE slice is shown with a 3D rendering of the myocardium assessed. (b) and (c) show the scar identified in both the first and fourth visits; the 3D illustration shows the myocardial volume assessed with the blue-yellow coloured rendering indicating the regions of hyper-enhancement identified; the colour-axis indicates ranges of hyper-enhancement transmurality from 0% in blue to 100% in yellow. In the bottom most image of both (b) and (c) the pixels identified as enhanced are indicated in yellow. The histogram plot in (d) shows the change in transmurality identified in the follow-up data (shown in (c)) compared to the first visit (shown in (b)).

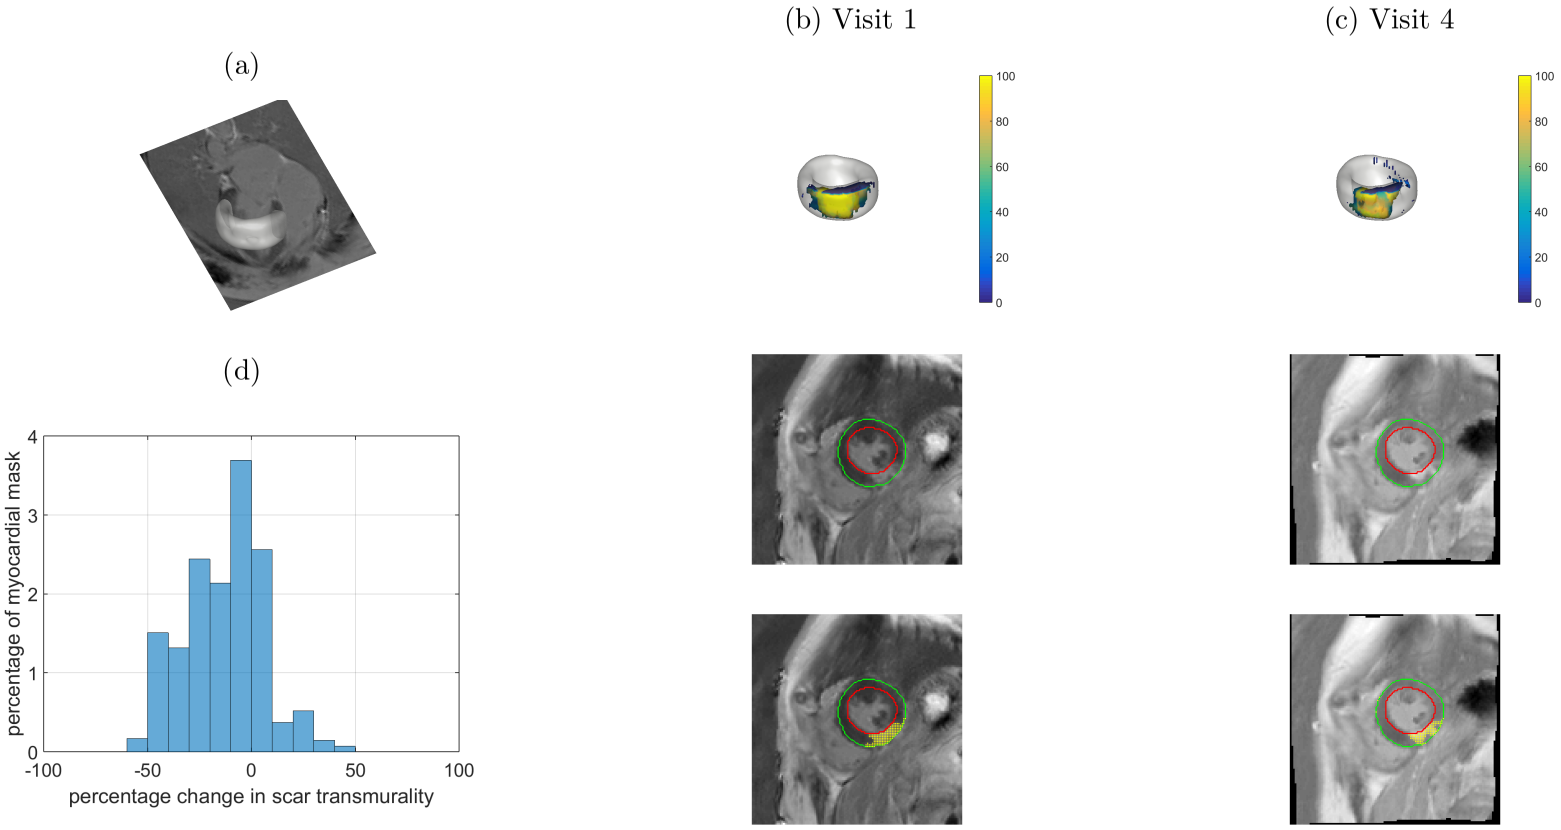


Figure 5.50: Example result of longitudinal scar assessment showing hyper-enhancement identified in follow-up acquisitions from the same patient after myocardial infarction. In (a) an LGE slice is shown with a 3D rendering of the myocardium assessed. (b) and (c) show the scar identified in both the first and fourth visits; the 3D illustration shows the myocardial volume assessed with the blue-yellow coloured rendering indicating the regions of hyper-enhancement identified; the colour-axis indicates ranges of hyper-enhancement transmurality from 0% in blue to 100% in yellow. In the bottom most image of both (b) and (c) the pixels identified as enhanced are indicated in yellow. The histogram plot in (d) shows the change in transmurality identified in the follow-up data (shown in (c)) compared to the first visit (shown in (b)).

5.4.3 Discussion

To analyse changes in tissue structure and composition in response to cardiac injury, the correspondences between the acquired images is required. As was noted in our initial study, presented in section 3.2.4, there are likely differences in the tissue planes imaged on each occasion. If tissue changes are assessed according to "corresponding" 2D slices, the analysis is likely to be affected by the fact that these slices are actually at different locations, capturing different planes of cardiac tissue. Therefore an analysis of tissue changes requires a registration that accounts for 3D spatial differences between the data. In this section, we have presented a two-stage, local phase-based framework for the 3D non-linear alignment of cardiac cine MR data.

Our assessment of the rigid registration was performed qualitatively; a representative series of results was shown in figure 5.32 which clearly illustrated the proficiency of the rigid phase-based registration to provide initial rough alignment between the longitudinally acquired datasets. The selection of the scale parameter ($\omega = \frac{1}{\lambda}$) had an effect on the acquired result; using too fine a scale ($\lambda \approx 5$), which captures a high proportion of small structural detail in the data, the algorithm was unable to reliably identify an acceptable initial alignment. Rather, we found that a scale parameter of $\lambda = 10$ worked consistently in the datasets processed.

The remaining non-rigid misalignments were compensated for using the optical flow approach. The phase-based approach used was very successful in enabling the non-rigid alignment of the data, as qualitatively and quantitatively illustrated in the results section. The qualitative results, presented in the form of cross-sections taken from both datasets illustrate this most clearly by looking at the correspondences of the left ventricle myocardium. In figure 5.35 attention is drawn to this, by consideration on the alignment of the endo- and epi-cardial delineations.

The quantitative validation was based on the distances calculated between myocardial surface meshes. The initial distances calculated before non-rigid registration do in fact give us some indication of the performance of the rigid step - specifically, the mean distance was calculated at approximately 2mm, a little

greater than the in-plane pixel dimension. The non-rigid registration was able to further reduce this mean to below the pixel resolution of 1.5mm.

The non-rigid method outlined is an iterative procedure, composed of 3 successive registrations performed with decreasing local phase scales. As intended, the increase in structural detail observed with each successive scale improves the accuracy of the registration result: the initial mean error of 1.98mm, as was reported in the results section, was reduced through 1.59mm, to 1.49mm and finally 1.40mm.

As with many of the previous discussions there are outliers from the mean distance errors after registration. Similarly to those previous analyses we find that a number of these notable distances occur in the apical regions of the meshes, as indicated in figure 5.51, for which we previously indicated that both segmentation and the registration itself is made difficult due to poor image definition.

Due to the variation in intensities seen between similar tissue in different datasets, we have proposed using a local phase-based registration approach in place of a more standard intensity-based optical flow scheme. Preliminary results obtained using the intensity-based approach showed the final mean error after registration was only reduced to 1.86mm from the original 1.98mm, a notably worse result than that achieved with the local phase-based metric. This comparative result could be strengthened if we were to perform the registrations on more cases, using a larger range of algorithm parameters.

It is notable in the literature that few of the proposed approaches utilise metrics based on the differences in intensity values, rather often selecting mutual information-based similarity terms. This results from an appreciation for the non-quantitative nature of the MRI signal, meaning the same tissues can demonstrate significant intensity differences between scans. One example where intensity values were used is the method proposed by De Craene et al [46, 200]; however, they note that intensity variation is still an observed phenomena across the cycle, and to improve robustness they require an additional frame-by-frame intensity similarity term which is more consistent as they expect the signal to vary smoothly in accordance with the changing magnetization state with time. Our decision to utilise a non-intensity,

local phase-based metric is therefore consistent with the literature. However, while mutual information accounts for the lack of an intensity matching between the data, as a measure it lacks spatial information which could therefore reduce its effectiveness in non-rigid image registration applications. Another potential issue with the mutual information metric occurs if the intensity values of adjacent but distinct structures are similar - depending on the number of histogram bins selected, this could also have a detrimental effect on the registration result. Using local phase as a similarity measure alleviates these issues - as discussed, the local phase-based metric is robust to intensity and contrast differences. The recognition of boundaries between structures is no longer dependent on the extent of the intensity difference between the structures, but rather on the scale of the filter selected, and will identify the same image structures in both datasets to be aligned.

The other primary difference between our method and those discussed in the literature is the transformation model employed. B-spline based approaches have been extensively utilised in the literature: this approach has the advantage of providing a smooth motion field defined by an adaptable number of parameters, in contrast, the transformation defined by our optical flow-based method is a result of a parametrisation for each voxel in the image constrained by local and global components for overall smoothness. We feel that this parametrisation for each point is advantageous in allowing a greater flexibility for local registration.

A number of the methods in the literature utilised the physiological properties of cardiac tissue to further constrain their registration - in particular incompressibility terms were added to the registration cost calculations. We note however that this was due to the application of these methods to cardiac motion estimation - in such cases the incompressible behaviour of the tissue is valid, and similarly may be valid if registration were to be performed on multiple acquisitions in a short space of time (i.e. the same day). However, in our work we were looking to register data acquired over a period of six months, a period in which some degree of tissue remodelling was expected in response to the cardiac injury; as a result, such an assumption becomes invalid.

For comparative purposes we have presented additional results utilising two alternative registration algorithms. In our own method we utilised local phase information as we hypothesised that differences between intensities observed in the cine MR volumes would make intensity-based methods ineffective. The results obtained here using the intensity-based approach support this hypothesis. While in specific cases where the intensity characteristics of the two volumes were favourable we observed improvements in the alignment of data, we found that across all datasets the registration approach was not successful, leading in many cases to worsening registration results.

The second algorithm utilised was an FFD-based approach with the NMI similarity metric. This approach was selected as it has been regularly reported in the literature for registration of cardiac MR data. We collected results at a range of different grid spacings and found that the best results were obtained using a grid spacing of $20mm^3$. In contrast to the intensity-based approach, the results obtained here do show improvement in alignment with a reduction in both the endo- and epi-cardial mesh distances to 1.44mm ± 1.30mm and 1.43mm ± 1.21mm for endo- and epi-cardial surfaces respectively. The positive results obtained here further illustrate that non-intensity-based approaches are preferable for the registration of cine MR data. We have noted that when using a B-spline approach we require smaller grid spacings to achieve local deformation accuracy, however, our results showed that reducing the grid spacing to $5mm^3$, which has been previously reported in the literature, we observed worse results and in some cases in-plausible deformation fields.

Our results show that the phase-based optical flow approach was the most accurate, however the measured improvement compared to the B-spline approach is small.

Application of deformation fields to LGE data

The registration results were applied to the pre-aligned LGE data such that the complementary CMR studies were brought into a common, aligned space. While positive qualitative results were obtained, the quantitative error measurements,

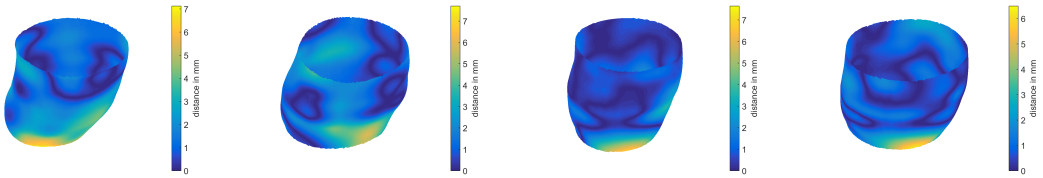


Figure 5.51: Meshes on which the distances to the corresponding surface in the other visit are displayed. These examples illustrate how in a number of cases the largest error values are observed near the apex.

shown in figure 5.46, were less convincing - particularly when compared to the corresponding cine results. There are two primary reasons for the inconsistencies in the acquired results: firstly, there are segmentation differences between the single-visit cine and LGE images as well as the final registered multi-visit LGE data; and secondly, any remaining misalignments that exist between cine and LGE data, lead to an exacerbation of the errors observed in the final LGE-LGE registration result. In particular this result demonstrates that an additional non-rigid registration of the LGE data to the cine data may be a beneficial addition. However, as we indicated earlier, this is not a straightforward addition as the complexity of the registration will be notably increased, with a non-exact similarity metric combined with a non-linear deformation model requiring minimisation. Therefore, addressing this challenge is a potential for future work that can further improve the acquired results.

The changes in scar tissue

Finally, we presented examples of the analysis of scar tissue changes observed between the longitudinal acquired CMR data. In each of examples shown, we see reductions in the transmurality of the scar identified. These results are consistent with those presented in [43], where LGE hyper-enhancement was assessed via volume estimates and 2D segmental analysis in data acquired from similar ST-elevation infarction patients. Therefore, it is possible to similarly conclude that the hyper-enhancement identified during the first visit acquisitions is an overestimation of actual scar tissue, and that in reality the degree of salvageable tissue is greater than that which would be concluded based on the original assessment of LGE and oedema data. The 3D registration that we have presented here enables the fully

quantitative geometric analysis of hyper-enhanced regions, overcoming limitations associated with assessment of non-corresponding slices. The insight gained from this 3D evaluation of changes in scale and transmurality of damaged tissue may be used by clinicians, in combination with other clinical biomarkers, to evaluate intervention efficacy and patient prognosis.

6

Volume Reconstruction of Cardiac Datasets from Multiple Slices

Contents

6.1 Introduction 175
6.2 Diffusion-based Processing 178
6.3 Methods 188
6.3.1 Reaction-diffusion interpolation 188
6.3.2 Multi-frame reconstruction - Combining information from all cardiac phases 194
6.3.3 Registration-based reconstruction 196
6.3.4 Integration of LA and SA data 199
6.4 Experiments and Results 204
6.5 Discussion 213

6.1 Introduction

As alluded to in previous chapters, due to the limited imaging time available and inherent acquisition restrictions, the clinical data gathered is both highly anisotropic and non-contiguous - issues that are depicted in figure 6.2. However, isotropic descriptions of the tissue are desirable, as they can lead to improvements in estimates of anatomical and functional indices, which are important for the characterisation of a patient's condition.

Intensity-based interpolation methods such as the nearest neighbour, linear, and cubic spline-based schemes are all standard approaches for upsampling from low to high resolution data. In the previous chapter we used linear interpolation to generate the volumes for registration; however, while an intensity-based approach such as this has the advantage of being computationally efficient, the generated images suffer from issues such as the blurring of object edges

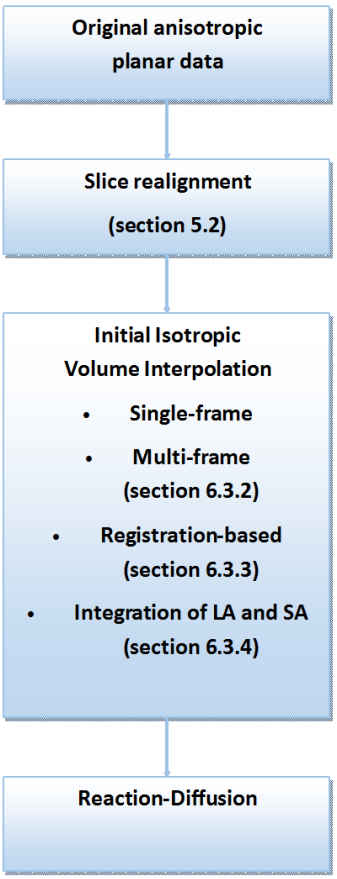


Figure 6.1: The steps involved in the isotropic reconstruction of cardiac MR volumes using anisotropic reaction-diffusion.

In this chapter we discuss an interpolation approach that is based on an anisotropic diffusion process. The diffusion system employed is coupled with a reaction component that simultaneously looks to retain the important in-plane detail available in the input 2D acquisitions in the final isotropic reconstruction. In addition to presenting the reaction-diffusion-based interpolation, a series of potential data initialisation schemes are explored that consider alternative means of utilising the input planar MR data.

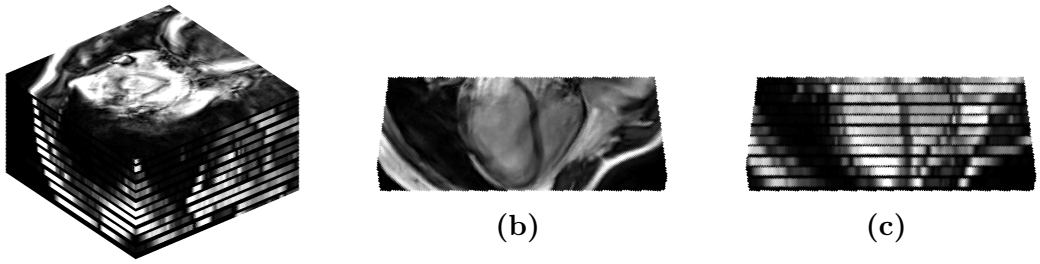
The stages of the reconstruction process are illustrated in figure 6.1. After slice realignment, an initial reconstruction estimate is generated using one of a number of methods that will be presented during the course of this chapter. The reaction-diffusion procedure is then applied to this initial reconstruction to generate a final volume representation.

Before discussing the various components of the proposed methods, in the next section we discuss the basic theory of diffusion-based image processing that motivates the use of a reaction-diffusion-based interpolation scheme.

The work presented in this chapter is a continuation of work undertaken by previous group members during their masters projects. In their work they explored the utilisation of registration for improved interpolation of cine MR data, including the potential for using data deformed according to cardiac cycle registration as input to a reaction-diffusion-based interpolation scheme; their work explored the utilisation of these approaches via experiments carried out on a small number of volunteer datasets.

Here we further develop and adapt these methods. We have introduced the phase-based image registration approach discussed in the previous chapter in place of an intensity-based approach for aligning cardiac frames and for the registration between slices that enables registration-based interpolation. A key development of the results presented in this chapter is the introduction of LA data - this has two positive effects on the results obtained; firstly, taking into account LA data allows for SA slice positions to be corrected, a pre-processing step that was previously discounted and relied on data from consistent breath-holders; and secondly, using the intensity correction method presented in this chapter, we are able to include LA data directly in the interpolation approach, resulting in improvements in the definition of the apical region. The algorithms discussed have been adapted therefore to deal with slices of varying orientation. Finally, we show first results of the discussed interpolation methods on patient data; these results illustrate some of the limitations of the approaches presented, therefore giving insight into which of the method variations may be best utilised according to the particular data available. In addition to these extensions, a more comprehensive validation of the presented methods has been undertaken including a phase-based metric that gives an indication to endo- and epi-cardial surface definition.

A paper, 'Anisotropic reaction-diffusion interpolation for cardiac cine MR volume reconstruction', is in preparation based on the reconstruction methods and results presented in this chapter.



(a) (b) (c)

Figure 6.2: Illustration of CMR slice thickness and inter-slice gaps. In (c) the volume cross-section corresponding to the LA slice in (b) is shown.

6.2 Diffusion-based Processing

Diffusion is the physical process by which a material is observed to move down its concentration gradient. This concentration equilibration process has a well-known mathematical formulation expressed in equation 6.1. Known as Fick's law, it describes how a flux (or flow), j, is driven by the existence of a concentration gradient, ∇I. The relationship between the two is described by the diffusion tensor, D, which may take one of a number of forms depending on the nature of the diffusion occurring.

$$j = -D \cdot \nabla I \tag{6.1}$$

In a purely diffusion-based system, there is a conservation of material mass; this allows the change in concentration with time to be expressed via the continuity equation in 6.2. Here, the rate of change at a given location is equal to the net flux that enters or leaves the space. Combining Fick's law and the continuity equation, an expression for the diffusion equation is defined as in 6.3.

The theoretical realisation of the diffusion process has found application in many image processing tasks [214]. In this context, it may be used to describe the redistribution of pixel intensities, I, according to the image intensity gradients. The continuity equation therefore describes this movement, while simultaneously maintaining the image's average intensity.

$$\frac{\delta I}{\delta t} = -divj \tag{6.2}$$

$$\frac{\delta I}{\delta t} = div(D \cdot \nabla I) \tag{6.3}$$

Homogeneous isotropic diffusion

Diffusion processes are characterised by the selection of the diffusion tensor, D. Diffusion that is equal in all directions and does not vary according to location is referred to as homogeneous isotropic diffusion. Under these conditions, we need not consider diffusion in the form of a tensor, but could refer to the term instead as the diffusion coefficient - that being a scalar value. However, for comparison with the methodological variations discussed subsequently, we note that homogeneous isotropic diffusion may be interpreted as a diffusion system characterised by a scalar multiplied identity matrix that is equal at all image locations. The major drawback of this homogeneous diffusion is that, while reducing the noise present in the image, it also induces the less favourable effect of blurring important structures such as edges. The issue is illustrated in figures 6.3b.

$$D(\mathbf{x}) = D = c * \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \tag{6.4}$$

Inhomogeneous isotropic diffusion

Such limitations may be addressed by the introduction of inhomogeneous isotropic methods. These approaches apply diffusion in a spatially-dependent way; while the diffusion remains equal in all directions, the rate at which it occurs now varies according to the image's local characteristics. In particular, the objective is to reduce diffusion at salient locations in the domain, thus avoiding the excessive blurring observed previously. Such an approach was proposed by Perona and Malik [152], who referred to their presented approach as anisotropic diffusion. However, this is not consistent with the definition for anisotropic diffusion given in the next section; rather, the spatial dependency of the diffusion coefficient makes their method inhomogeneous isotropic only. We will thus refer to this method as "inhomogeneous isotropic" in this thesis, and reserve the term anisotropic for the techniques described in the next section. The diffusion coefficient was made

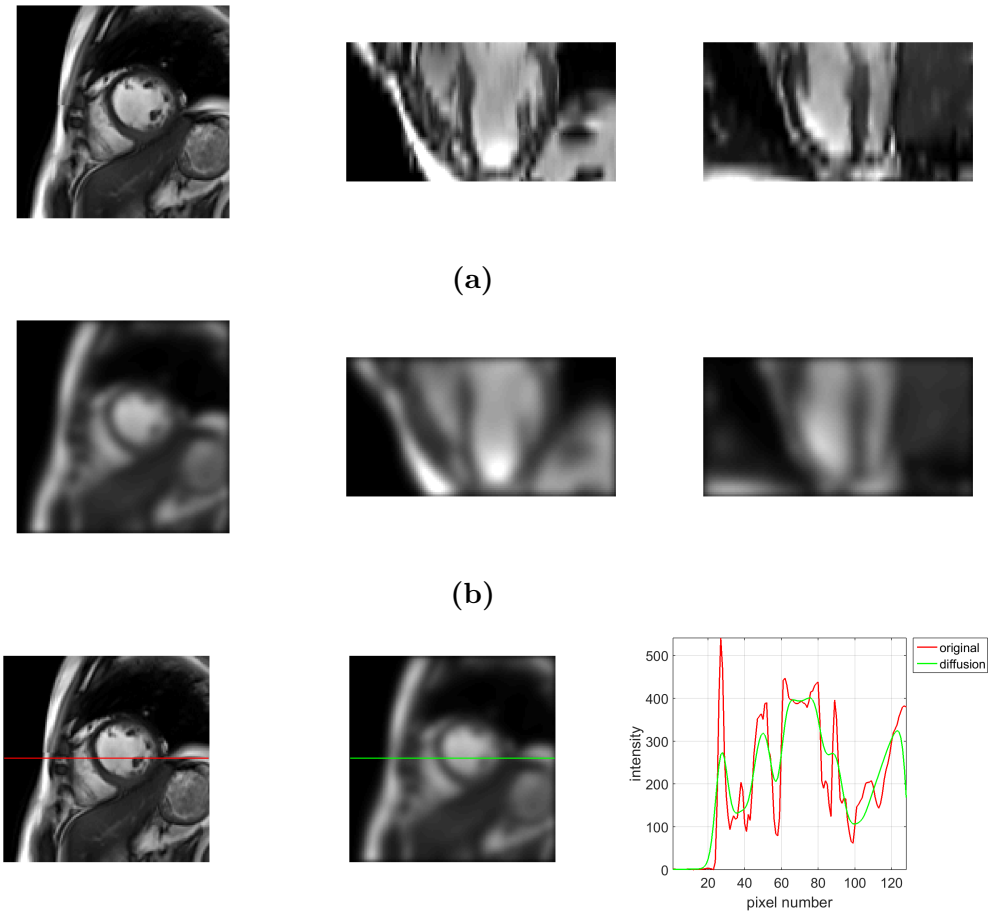
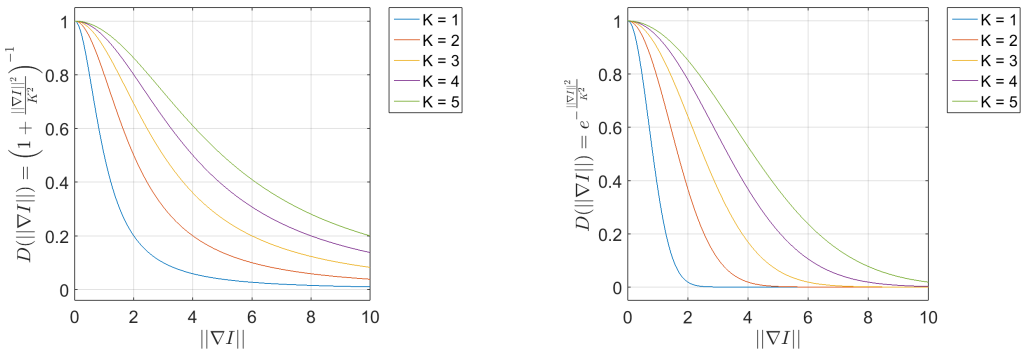


Figure 6.3: Illustration of the blurring of edges seen when using a pure isotropic diffusion model. The original data and the diffused data are shown in (a) and (b) respectively. In (c) an intensity profile taken from each is plotted showing the reduced definition of structure edges.

a function of the image gradient, $D(x) = f(\|\nabla I(x)\|)$; by forming a function for which $D(x) \to 0$ as $\|\nabla I(x)\| \to \infty$, the diffusion across edges can be reduced.

$$\frac{\delta I}{\delta t} = div(D(\|\nabla I\|) \cdot \nabla I) \tag{6.5}$$

In [152], the two diffusion functions expressed in equations 6.6 and 6.7 were presented. Both are formulated such that diffusion is reduced at locations that have a higher likelihood of being edges; positions that are identified by considering the magnitude of the image gradients. This reduction in diffusivity with increasing gradient magnitudes is illustrated for both formulations in figure 6.4. We can observe here also the effect of the K parameter in controlling how fast the diffusivity decays. Often referred to as the contrast parameter, it can additionally be interpreted



(a) **(b)**

Figure 6.4: Diffusion coefficients plotted as functions of image gradient magnitude. The formulations of equations 6.6 and 6.7 are shown in (a) and (b) respectively.

as a scale selecting parameter, influencing the features that will be preserved during the diffusion.

$$D(\mathbf{x}) = \frac{1}{1 + \frac{\left\|\nabla I(\mathbf{x})\right\|^2}{K^2}} \tag{6.6}$$

$$D(\mathbf{x}) = e^{-\left(\frac{\left\|\nabla I(\mathbf{x})\right\|}{K}\right)^2} \tag{6.7}$$

As the Perona-Malik formulation looks to avoid diffusion of strong image edges, characterised by large gradients, this would also result in a retention of image noise. Therefore it is common to pre-filter the image data by a Gaussian smoothing kernel, $I_\sigma = K_\sigma * I$. The resulting diffusion coefficients are then estimated from the regularised set of gradient magnitudes.

$$D(\mathbf{x}) = f(\|\nabla I_\sigma(\mathbf{x})\|) \tag{6.8}$$

The Perona-Malik formulation is further capable of enhancing structure definition; this effect is a result of a reverse diffusion process that occurs for gradient magnitudes larger than the contrast parameter. The observed effect - a result of moving material against the concentration gradient - is an enhancement of the feature boundary. However, this reverse diffusion also makes the system ill-posed and can lead to singularities in the solution [31].

Anisotropic diffusion

Despite their benefits, these inhomogeneous approaches are still limited to isotropic diffusion; the degree of diffusion has simply been regulated according to the local image characteristics. A preferable approach is one that, in addition to being spatially dependent, is variably diffusive according to direction.

Such a method was discussed by Weickert in [213], and was referred to as an anisotropic diffusion system. While the diffusion relationships for all previously discussed methods used scalar coefficients, here tensors are required that encode the directionality of the diffusion process. By manipulating the eigenvalues and eigenvectors of this diffusion tensor, the rates of diffusion can be adjusted according to orientation. Anisotropic diffusion has the significant advantage therefore of enabling the reduction of noise along notable structures, but not across them.

In order to define appropriate diffusion tensors, a description of the image structure is required, thus necessitating a consideration of the local neighbourhood to each pixel. While a pixel's gradient information indicates the direction of signal change, it contains no further information regarding the image's local construction. However, if the gradient information for an entire neighbourhood is taken into account, it is possible to distinguish between edges and corners, in addition to reducing the effect of misleading signal noise. Direct averaging of image gradients over significant regions can be problematic, due to the risk of cancelling gradients that differ by sign alone (for example the opposite gradients on either side of thin lines). This issue is illustrated in figure 6.5c, where the structure orientation is shown, estimated via the argument of the gradient components (figure 6.5b). With increasing local smoothing, we see that the definition of the inner circles is lost. For this reason it is commonplace, when describing image structure, to use the structure tensor.

The structure tensor is formed by the gradient vector outer product. By this process, antipodal gradients (i.e. gradients that have the same orientation, but opposite directions) are mapped to identical matrix forms. This has the beneficial outcome that, under a neighbourhood averaging operation, there is no longer a

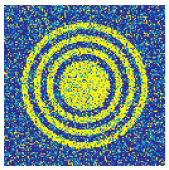
cancellation effect, which previously led to a loss in orientation information for gradient vectors alone. The structure tensor is expressed in equation 4.15; we note here that the original gradient vectors are still locally smoothed, such that the effect of spurious noise is negated; however, the degree of smoothing may be curtailed to minimise the loss of structural information.

$$\begin{aligned} J_\rho(\nabla I_\sigma) &= K_\rho * \left(\nabla I_\sigma \nabla I_\sigma^T\right) \\ &= K_\rho * \left(\begin{bmatrix}(I_\sigma)_x \\ (I_\sigma)_y\end{bmatrix}\begin{bmatrix}(I_\sigma)_x & (I_\sigma)_y\end{bmatrix}\right) \\ &= K_\rho * \begin{bmatrix}(I_\sigma)_x^2 & (I_\sigma)_x(I_\sigma)_y \\ (I_\sigma)_x(I_\sigma)_y & (I_\sigma)_y^2\end{bmatrix} \end{aligned} \tag{6.9}$$

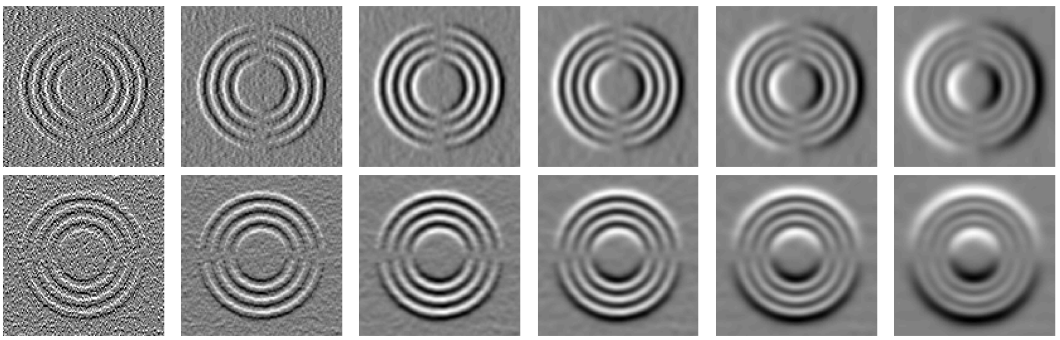
This formulation can be referred to as the multi-scale structure tensor, due to the use of two separate scale parameters: σ, which controls the initial image smoothing, is referred to as the local scale parameter, while ρ is referred to as the integration scale parameter, as it determines the spatial extent of the structure tensor averaging operation [109]. Figure 6.5d illustrates the benefit of using the structure tensor over the gradient-based orientation estimation, whereby no structural information is lost with increasing local smoothing. In figure 6.5e we see the further improvement that results from using the multi-scale structure tensor, where the gradient information is pre-filtered to reduce the effect of noise.

The significance of the structure tensor is realised by its eigendemcomposition. The eigenvalues and eigenvectors describe the distribution of gradients within the local neighbourhood, thus providing insight into the underlying structural construction. In figure 6.6 a series of example 3D structures is shown, below which, the eigenvalues and eigenvectors that characterise these structures are indicated. The first eigenvector is aligned in the direction of the mean local gradient, while the third is aligned to the orientation of maximal signal coherence. The eigenvalues of the system reflect the corresponding signal changes along each vector orientation; therefore, the eigenvalues are ordered such that $\lambda_1 \geq \lambda_2 \geq \lambda_3$.

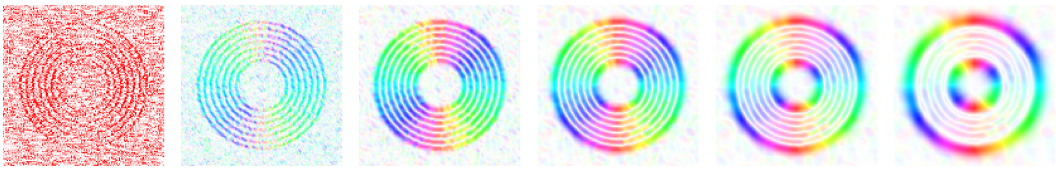
Specifically, in the first example of figure 6.6, where there is no significant image gradient, we see how all three eigenvalues are equal resulting in an approximately spherical ellipsoidal description. In contrast, the second example shows one clear



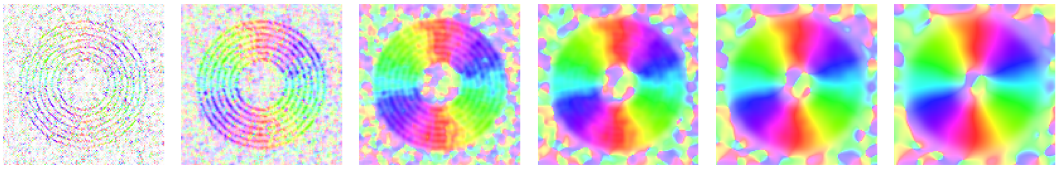
(a)



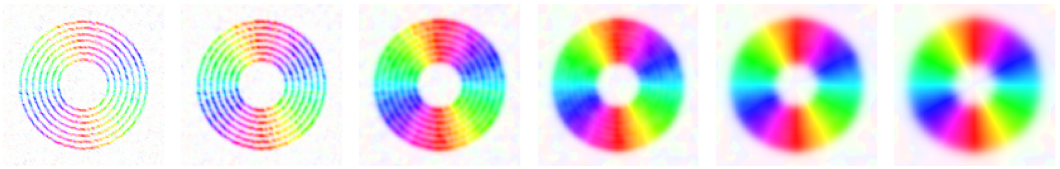
(b)



(c) Orientation estimated from neighbourhood averaged gradient, calculated for different values of σ (increasing from left to right).



(d) Orientation of the primary eigenvector of the structure tensor $(K_\rho * \left(\nabla I \nabla I^T\right))$, scaled by its corresponding eigenvalue, calculated for different values of ρ (increasing from left to right).



(e) Orientation of the primary eigenvector of the multi-scale structure tensor $(K_\rho * \left(\nabla I_\sigma \nabla I_\sigma^T\right))$, scaled by its corresponding eigenvalue, calculated for different values of ρ (increasing from left to right) with $\sigma = 1$.

Figure 6.5: Illustration of the local orientation estimation using the gradient, structure tensor, and multi-scale structure tensor. The example image is shown in (a), while its gradient, computed at different local averaging scales, is shown in (b).

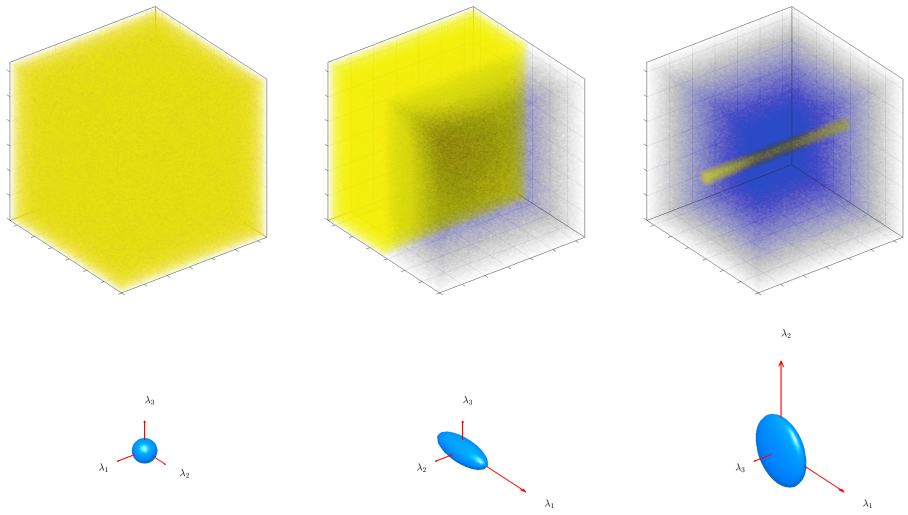


Figure 6.6: Example 3D image neighbourhoods with corresponding eigenvectors below; the eigenvectors are scaled by their eigenvalues. In the first example, where there is no primary orientation of intensity change, all three eigenvalues are equal, in contrast to the second and third examples where the primary orientations of intensity change are characterised by large eigenvalues.

3D boundary. The ellipsoid here is very anisotropic having one eigenvalue, whose eigenvector is oriented across this intensity boundary, that is much larger than the other two. In the final example we see a structure characterised by large intensity changes in two orientations, while in the third orientation the image intensity is comparatively constant. Therefore, this structure is characterised by two large eigenvalues, while the third, corresponding to the direction along the image structure, is small. The eigenvalues can therefore be used to gain insight into the local description of the signal.

By utilising information gained from the structure tensor regarding the local construction of a signal, systems can be designed such that the diffusion process is guided towards an enhancement of image structures, coupled with a de-noising process. The diffusion tensors may be constructed from a set of orthonormal vectors and corresponding eigenvalues. Utilising the eigenvectors calculated for the structure tensor, the eigenvalues are chosen according to the desired diffusion action.

Reaction-diffusion systems

In some cases it may be necessary to introduce extra terms to the above diffusion formulations, and the resulting systems are referred to as reaction-diffusion algorithms. The new system equation is defined in 6.10, where μ indicates the reaction component and α is a parameter used to weigh the relative influence of diffusion and reaction. In such formulations the reaction term acts as a material source or sink, introducing notable changes to the total composition of the system - the systems are no longer governed solely by the spatial redistribution of material.

$$\frac{\delta I}{\delta t} = (1-\alpha)div(D \cdot \nabla I) + \alpha\mu \tag{6.10}$$

There are many possible motivations for introducing such terms. An inherent limitation of the previous diffusion-only systems is that, with increasing time, important structural detail will be lost, despite the introduction of the anisotropic diffusion process. A reaction term which considers the difference between the image at the current iteration compared to the original has been proposed as one way to minimise this deviation from the original data. A second motivation may be the existence of gaps in data. Diffusion systems are capable of filling such gaps; however, during this process there will be a decrease in the average value of the signal. If a reaction term is used, it can act as a source for extra signal quantity, that can be used to fill in the gaps without reducing the average signal value.

Solving the diffusion system

Applying the above theory to the discrete image volumes in our own work requires a discretisation of the continuous domain formulation.

The time derivative is approximated straightforwardly according to equation 6.11, leading to the image update equation of the form in 6.12.

$$\frac{\delta I}{\delta t} \approx \frac{I^t - I^{t-1}}{\Delta t} \tag{6.11}$$

$$I^t = I^{t-1} + \Delta t((1-\alpha)div(D \cdot \nabla I) + \alpha\mu) \tag{6.12}$$

The expansion of the divergence term, outlined in equation 6.13, shows that there will be a need to calculate both single- and mixed-axis second-order partial derivative terms. These may be approximated by repeated application of the central finite difference scheme. Examples of each are given in equations 6.14 and 6.15. The remaining terms in the diffusion matrix can be calculated in the same way.

$$\begin{aligned} div(D\nabla I) &= div\left(\begin{bmatrix} D_{11} & D_{12} & D_{13} \\ D_{21} & D_{22} & D_{23} \\ D_{31} & D_{32} & D_{33} \end{bmatrix}\begin{bmatrix} I_x \\ I_y \\ I_z \end{bmatrix}\right) \\ &= div\left(\begin{bmatrix} D_{11}I_x + D_{12}I_y + D_{13}I_z \\ D_{21}I_x + D_{22}I_y + D_{23}I_z \\ D_{31}I_x + D_{32}I_y + D_{33}I_z \end{bmatrix}\right) \\ &= \frac{\delta}{\delta x}(D_{11}I_x + D_{12}I_y + D_{13}I_z) + \frac{\delta}{\delta y}(D_{21}I_x + D_{22}I_y + D_{23}I_z) \\ &\quad + \frac{\delta}{\delta z}(D_{31}I_x + D_{32}I_y + D_{33}I_z) \end{aligned} \tag{6.13}$$

$$\begin{aligned} \frac{\delta}{\delta x}(D_{11}(x,y,z)I_x(x,y,z)) &= \frac{\delta}{\delta x}D_{11}(x,y,z)\left(\frac{I(x+h_x,y,z) - I(x-h_x,y,z)}{2h_x}\right) \\ &= \frac{1}{2h_x}\left(D_{11}(x+h_x,y,z)\left(\frac{I(x+2h_x,y,z) - I(x,y,z)}{2h_x}\right) - \right. \\ &\quad \left. D_{11}(x-h_x,y,z)\left(\frac{I(x,y,z) - I(x-2h_x,y,z)}{2h_x}\right)\right) \end{aligned} \tag{6.14}$$

$$\begin{aligned} \frac{\delta}{\delta x}(D_{12}(x,y,z)I_y(x,y,z)) &= \frac{\delta}{\delta x}D_{12}(x,y,z)\left(\frac{I(x,y+h_y,z) - I(x,y-h_y,z)}{2h_y}\right) \\ &= \frac{1}{2h_x}\left(D_{12}(x+h_x,y,z)\left(\frac{I(x+h_x,y+h_y,z) - I(x+h_x,y-h_y,z)}{2h_y}\right) - \right. \\ &\quad \left. D_{12}(x-h_x,y,z)\left(\frac{I(x-h_x,y+h_y,z) - I(x-h_x,y-h_y,z)}{2h_y}\right)\right) \end{aligned} \tag{6.15}$$

Using a standard single pixel step size, $h = h_x = h_y = h_z = 1$, requires a wider spatial region of pixels for approximating the divergence. Instead, we can use the half-step central difference, $h = h_x = h_y = \frac{1}{2}$, such that the spatial extent is limited to a $3 \times 3 \times 3$ window. Using the half-step difference generates terms

that do not lie on actual image grid points; these values can be approximated using a one-sided average value: $f(x+\frac{1}{2}) = \frac{1}{2}(f(x+1)+f(x))$.

6.3 Methods

6.3.1 Reaction-diffusion interpolation

Having discussed the general theory of diffusion-based image processing in the previous section, here we will look more specifically at the proposed formulation for isotropic interpolation of thick slice cardiac MR data.

The general equation for the reaction-diffusion system that will be described is defined in equation 6.16, where the α term introduces a weighting between the relative contributions of the diffusion and reaction components.

$$\frac{\delta I}{\delta t} = (1-\alpha) div(\mathbf{D} \cdot \nabla I) + \alpha\mu \tag{6.16}$$

Firstly, we consider the diffusion component. The diffusion tensor is calculated according to analysis of the structure tensor. The eigenvectors of our diffusion tensor, $D(x)$, are those calculated from the structure tensor. However, the eigenvalues are adapted according to the desired diffusion properties. Similar to previous methods, we employ eigenvalues that will minimise the loss of important structural detail during the diffusion process. As previously noted, the eigenvalues are ordered such that $\lambda_1 \geq \lambda_2 \geq \lambda_3$, i.e. the first (and largest) eigenvalue is associated with the direction of greatest signal change, whereas the last is affiliated with the least. We require an inversion of this relationship, such that the orientation of greatest signal change corresponds to the orientation of minimal diffusion. Therefore, the selected eigenvalues of our diffusion tensor, λ_D, are calculated by a straightforward inversion and normalisation of the structure tensor eigenvalues λ_{ST}; these new eigenvalues are defined in equation 6.17.

$$\lambda_{D_i} = \frac{1}{\lambda_{ST_i}} \left(\sum_{j=1}^{3} (\frac{1}{\lambda_{ST_j}}) \right)^{-1} \tag{6.17}$$

For a structure tensor eigenvalue (λ_{ST}) equal to 0 the corresponding diffusion tensor eigenvalue (λ_D) is set to 1 .

Diffusion parameter selection

The structure tensor integration scale, ρ, and the number of update iterations, t, are two parameters that have a significant effect on the final result of the diffusion process. In this section, results are shown that illustrate their effect using the anisotropic system that has been described to this point. Diffusion is applied to a cine MR volume generated by linearly interpolating the cine SA slice data.

The integration scale is responsible for defining the neighbourhood from which local orientation is estimated; this is significant, as the identified local structure is subsequently used to define the tensor components of the diffusion process. Averaging over larger areas reduces the influence of smaller structures on the orientation estimation, and the resulting diffusion process will be governed according to the characterisation of the larger scale local structure. The second parameter is the number of iterations; since a newly-estimated diffusion term is added with each iteration, the cumulative diffusion is dependent on the total number of iterations applied.

The effect of both parameters on the diffusion result are illustrated in figures 6.7 and 6.8, where we see the qualitative change in the volumes generated according to the magnitude of the integration scale selected and the number of iterations of the diffusion process run. The results show how, as both parameters are increased, more diffusion is observed. Therefore, the parameter selection is a compromise between the ability to smooth and connect coherent structures without simultaneously losing significant structural detail.

If we consider the integration scale in particular, it is responsible for defining the neighbourhood from which local orientation is estimated; this is significant for the operation of this algorithm, as the identified local structure is subsequently used to define the tensor components of the diffusion process. Averaging over larger areas reduces the influence of noise and smaller structures on the orientation estimation. This stabilises the orientation estimation, and the resulting diffusion process will be governed according to the characterisation of larger scale local structures.

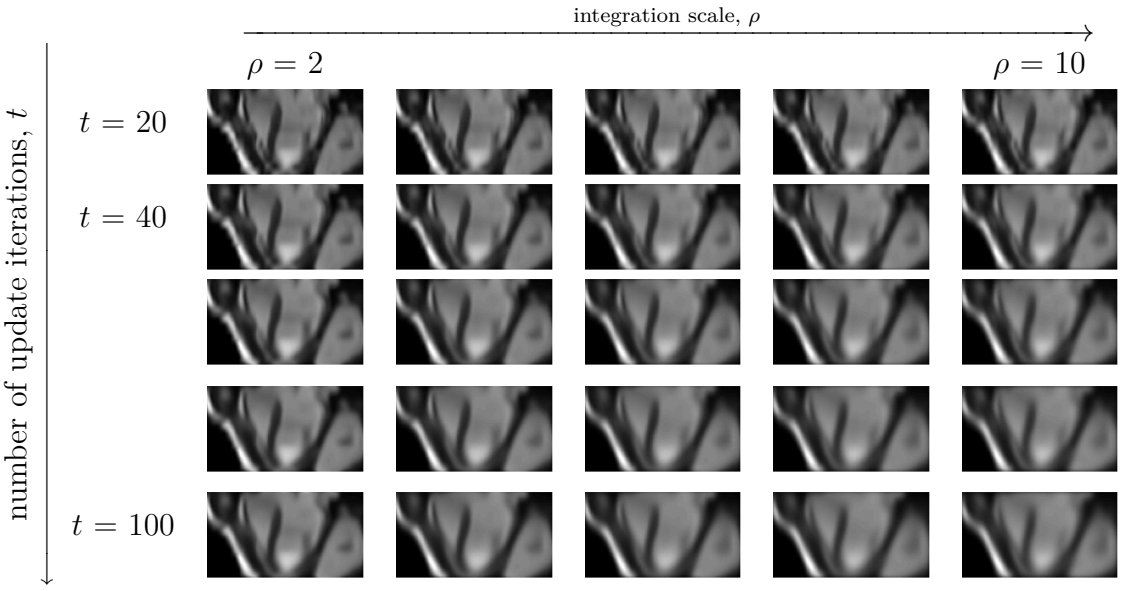


Figure 6.7: LA cross-sections taken from volumes resulting from anisotropic diffusion over a range of integration scales, ρ, and number of iterations, t.

The desirable outcome of this effect is to increase the coherence of image details, enabling in certain cases, where the integration scale is large enough, the closure of interrupted structures. However, if the integration scale is too large, the diffusion process will result in a loss of important image detail as the diffusion process lacks specificity to the structure.

These conflicting outcomes are demonstrated more clearly if we look at the most extreme result in figure 6.9, where an integration scale $\rho = 10$ was used with 100 iterations of the algorithm; there is a clear reduction in noise and an improvement in connecting coherent structures observed in the LA representation; however, we see clearly in the SA view how an excessive amount of diffusion has resulted in a loss of structural definition.

Reaction term

In addition to the appropriate selection of ρ and t parameters, structural detail may also be retained in the final result if we use a reaction term. Here the reaction term is used to limit deviation away from the original 2D planar slices, while simultaneously taking into account the anisotropic nature of the CMR acqusitions.

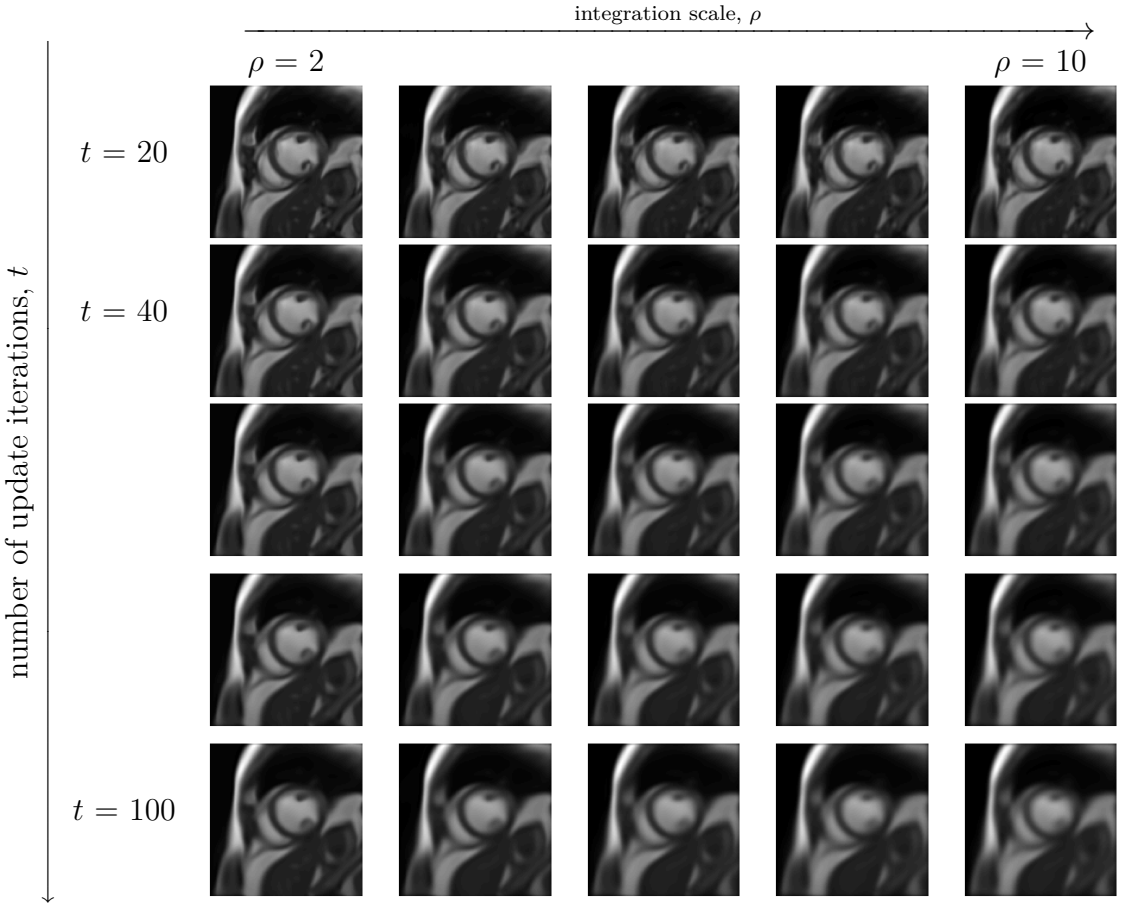


Figure 6.8: SA cross-sections taken from volumes resulting from anisotropic diffusion over a range of integration scales, ρ, and number of iterations, t.

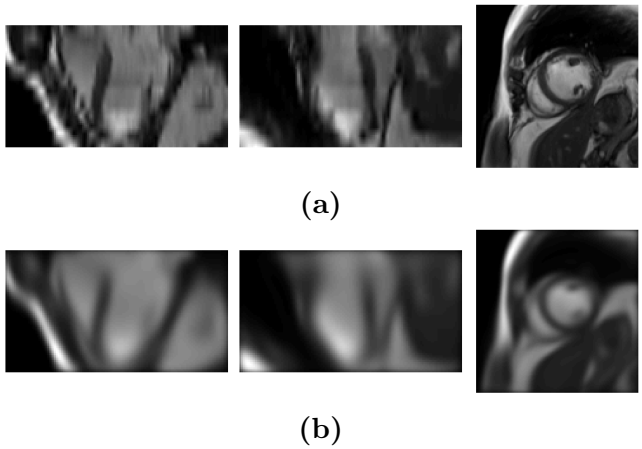


Figure 6.9: The original data are shown in (a), while the result acquired using anisotropic diffusion with parameters $\rho = 10$ and $t = 100$ is given in (b). Using these parameters an excessive amount of structural detail is lost.

We consider that a voxel's intensity value represents the averaged signal received from a thick block of tissue. While in the theoretical, ideal case, the RF pulse applied excites a rectangular out-of-plane signal profile, the reality is a signal profile which gradually decays nearer the borders. In this work we will approximate the slice profile as that of a Gaussian whose standard deviation is equal to half the slice thickness. Using this model, the difference between the current interpolated volume and the acquired slices can be calculated according to equation 6.18, where G indicates the Gaussian function defined by equation 6.19. The σ_x, σ_y, and σ_z parameters are calculated such that both the slice orientation and slice resolution are taken into account.

$$I_{Diff}(\mathbf{x}) \begin{cases} = I_0(\mathbf{x}) - G(\mathbf{x}) * I_t(\mathbf{x}), & \text{if } \mathbf{x} = \text{centre of slice voxels.} \\ = 0, & \text{if } \mathbf{x} \neq \text{centre of slice voxels.} \end{cases} \tag{6.18}$$

$$\begin{aligned} G(\mathbf{x}) &= \frac{1}{(2\pi)^{\frac{3}{2}} |\mathbf{\Sigma}|^{\frac{1}{2}}} e^{-\frac{1}{2}\mathbf{x}^T \mathbf{\Sigma}^{-1} \mathbf{x}} \\ \mathbf{\Sigma} &= \begin{bmatrix} \sigma_x^2 & 0 & 0 \\ 0 & \sigma_y^2 & 0 \\ 0 & 0 & \sigma_z^2 \end{bmatrix} \end{aligned} \tag{6.19}$$

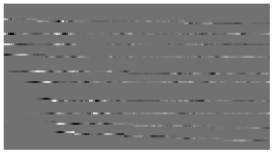
The differences calculated provide a sparse distribution of reaction points located at the centres of the thick voxels as illustrated in figure 6.10a. A reaction term can be generated for all voxels by using the same Gaussian acquisition model; the final reaction term is calculated according to equation 6.20, and is illustrated in figure 6.10b.

$$\mu = G(\mathbf{x}) * I_{diff}(\mathbf{x}) \tag{6.20}$$

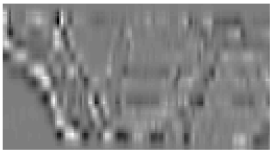
In the reaction-diffusion equation (equation 6.16) it was indicated that the system components were weighted according to the α parameter; its effect on the final result is shown in figure 6.11.

Improving the initialisation image

The interpolation method described requires an initial image that is iteratively updated according to the reaction-diffusion system. In the previously presented results, the initial image provided was a linear reconstruction of the SA slice data.



(a) I_{diff} - differences calculated at slice centres



(b) μ - reaction term

Figure 6.10: In (a) the result of the calculation indicated in equation 6.18 is displayed, where the differences between the volume intensities and the originally acquired planar slices are calculated. In (b) the reaction term that is generated by applying the Gaussian acquisition model to the result in (a) is shown.

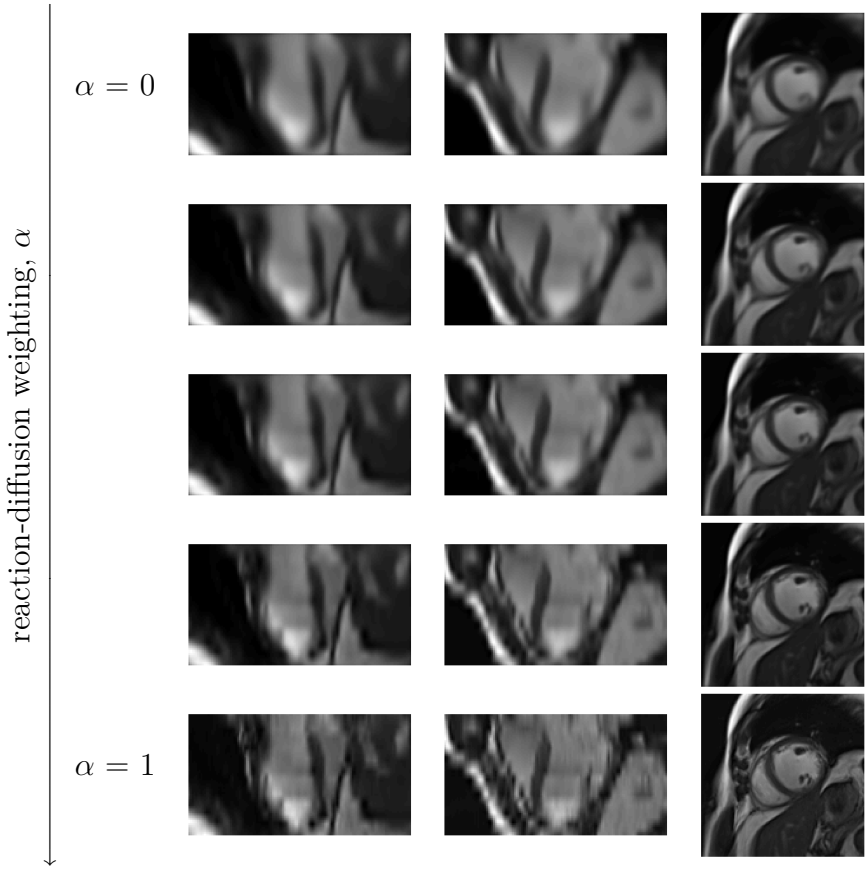


Figure 6.11: The reaction and diffusion components are weighted using the α parameter as indicated in equation 6.16. Here, results are shown that indicate the variation in outcome according to the applied weighting, α. The indicated results were acquired using the parameters $\rho = 5$ and $t = 50$.

In these next sections, we explore possible alternative initialisation images that may further improve the final interpolated result.

6.3.2 Multi-frame reconstruction - Combining information from all cardiac phases

Cardiac cine datasets are multi-frame acquisitions that span the entire cardiac cycle. While the acquired slices are at fixed positions in relation to the MRI scanner, the heart itself is in motion. As a result, successively acquired cardiac frames capture slightly different regions of the heart tissue. In particular, the motion causes perceptible out-of-plane tissue movement. This is illustrated in figure 6.12 where the partial appearance, and subsequent disappearance of the left ventricle outflow tract is recognisable in this basal cine slice. Based on such an observation, we propose to utilise the complementary information available from all cardiac frames

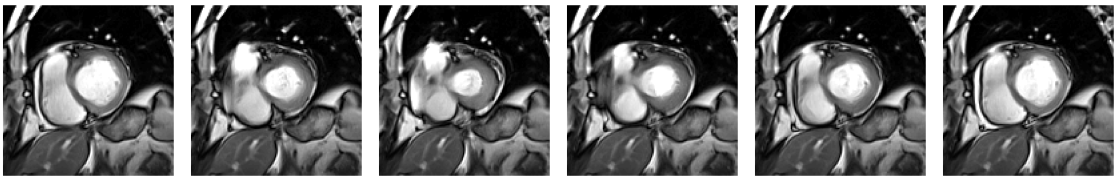


Figure 6.12: Illustration of the changing cardiac structures captured across the cardiac cycle.

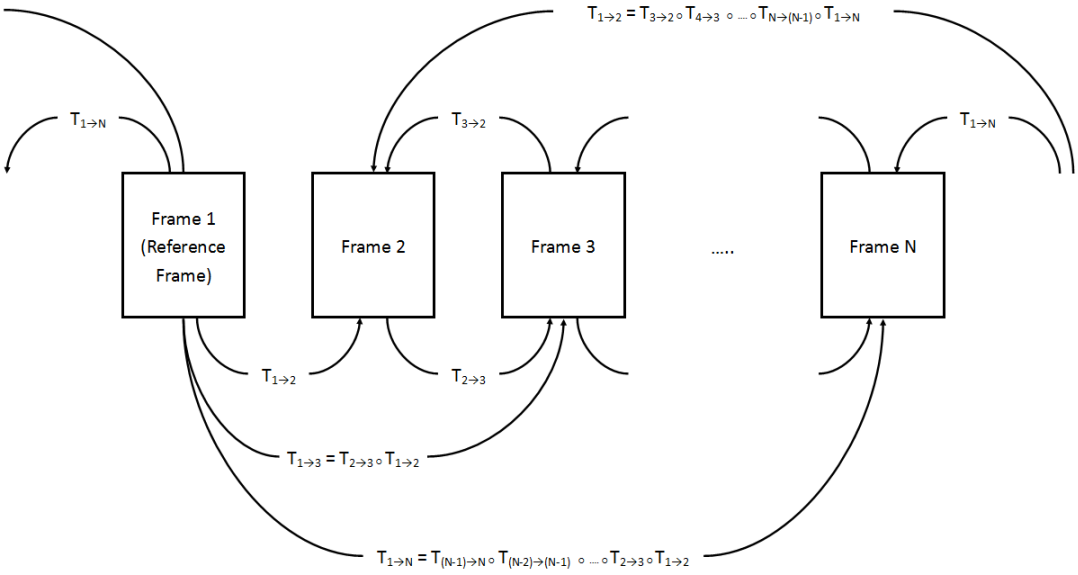


Figure 6.13: The framework for registering frames of the cardiac cycle. The individual correspondences of adjacent frames are composed to enable to alignment of distal cardiac frames. The registration process is performed in both temporal directions such that a final deformation between frames can be estimated using a weighted average based on the distance between frames.

to improve the quality of a single-cardiac-phase 3D reconstruction.

To map the data from all time frames to the single selected reference frame, their correspondences must first be identified; for this we use the combined local-global optical flow scheme with phase-vector data term discussed at the end of the previous chapter. Due to significant differences between distal points of the cardiac cycle, the registrations can prove difficult and can lead to unreliable results. To reduce the likely impact of poor registration between distal points, we first consider the individual alignments between adjacent cardiac frames. These more reliable transformations, $T_{i\rightarrow(i+1)}$, may then be composed to provide initial alignment between data in distal cardiac phases, as expressed in equation 6.21, where $\mathbf{u}$ indicates the deformations calculated by the optical flow registration between image frames I.

$$\begin{aligned}
T_{1\rightarrow2}(\mathbf{x}) &= \mathbf{x} + \mathbf{u}_{1\rightarrow2}(\mathbf{x}) \quad \text{such that, } I_1(\mathbf{x}) \approx I_2(T_{1\rightarrow2}(\mathbf{x})) \\
&\qquad\qquad\qquad\qquad\qquad\qquad \approx I_2(\mathbf{x} + \mathbf{u}_{1\rightarrow2}(\mathbf{x})) \\
T_{2\rightarrow3}(\mathbf{x}) &= \mathbf{x} + \mathbf{u}_{2\rightarrow3}(\mathbf{x}) \\
T_{1\rightarrow3}(\mathbf{x}) &= \left(T_{2\rightarrow3} \circ T_{1\rightarrow2}\right)(\mathbf{x}) \\
&= T_{2\rightarrow3}(T_{1\rightarrow2}(\mathbf{x})) \\
&= \mathbf{x} + \mathbf{u}_{1\rightarrow2}(\mathbf{x}) + \mathbf{u}_{2\rightarrow3}(\mathbf{x} + \mathbf{u}_{1\rightarrow2}(\mathbf{x})) \\
T_{1\rightarrow N}(\mathbf{x}) &= (T_{(N-1)\rightarrow N} \circ T_{(N-2)\rightarrow(N-1)} \circ .... \circ T_{2\rightarrow3} \circ T_{1\rightarrow2})(\mathbf{x})
\end{aligned} \tag{6.21}$$

To further improve the reliability of these composed transformations, we exploit the cyclical nature of the heart's deformation by computing the registrations in both temporal directions. The final estimated deformation for each frame is then given by a weighted average of the deformations along each cycle direction; the weightings are calculated based on the distance of each respective frame from the reference. This bi-directional registration process is illustrated in figure 6.13.

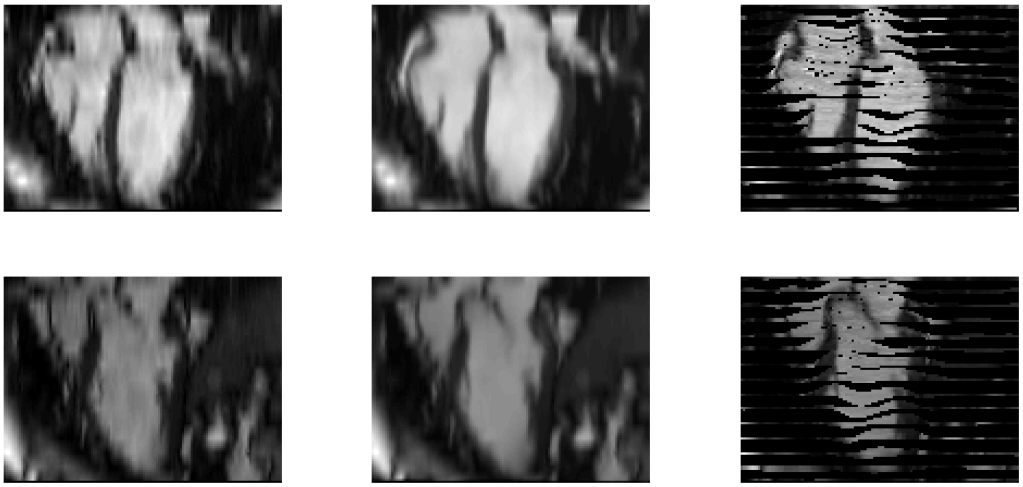
Using the registration results, the data contained in each frame of the cardiac cycle is deformed to match the reference frame. In figure 6.14 we illustrate how these registration results can be used for reconstruction in a number of ways. Figure 6.14a shows the original linear reconstruction that we used in the last chapter and which generates the volumes used to derive the deformation field between frames

here. In figure 6.14b we show the result obtained when we take the mean of all frames of the cardiac cycle registered to this specific frame - each of the deformed 3D frames contributes to the final volume generated, using weightings that are set according to the distance of each frame from the reference. In figure 6.14c we show the points contributed from all slices and all cardiac frames after registration to this one frame - it shows out-of-plane deformation that fills in the gaps between original slice centres. In figure 6.14d the new linear interpolation result using all the points in figure 6.14c is shown. This volume is generated to be used as a new initialisation to the reaction-diffusion system. Using this volume, in combination with a reaction term generated from the deformed slice data (figure 6.14c), the final result is illustrated in figure 6.14e.

6.3.3 Registration-based reconstruction

An alternative method that has been proposed in the literature is registration-based interpolation [147]. This approach is based on the identification of correspondences between adjacent slices via image registration methods. Subsequently, the registration results may be used to generate new interpolated image planes. As acknowledged in [151], such an approach to interpolation is dependent on the assumption that adjacent image slices contain similar structures, and that the prescribed registration method is capable of identifying the necessary transformation for their alignment. In this section, the necessary alignment between adjacent SA slices is computed using a 2D implementation of the combined local-global optical flow approach with local phase vector that we have previously used for 3D registration of cine data. Using the computed 2D deformations we generate intermediate planes of data points between the two originally-defined slices, as indicated in figure 6.15, where the original planes are shown to be non-parallel as a result of the slice realignment step previously performed.

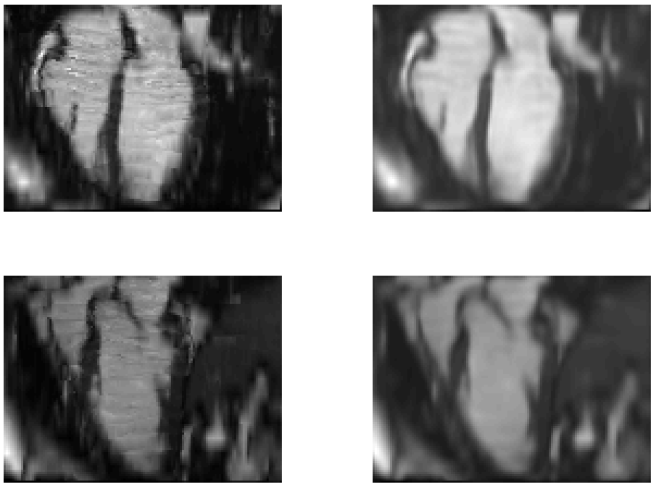
Figure 6.15 indicates a single new intermediate slice. In practice, the motion field is partitioned into a number of small segments and each partition is used to generate a new set of data points. These deformed data points no longer lie



(a) Single-frame linear reconstruction

(b) Mean data from all registered frames

(c) Deformed slice data



(d) Multi-frame linear reconstruction

(e) Reaction-diffusion reconstruction

Figure 6.14: Illustration of the different results generated using frame-to-frame registration.

on a perfect regular planar grids, but rather form a scattered collection of data points in a 3D space. Specifically, the positions of new data points are defined according to equation 6.22, where the two original adjacent slices are $I_A(x_A, y_A, z_A)$ and $I_B(x_B, y_B, z_B)$, and the calculated deformation field is defined by $\mathbf{u}_{A\rightarrow B}$.

$$\begin{aligned} x_{New}(\mathbf{x}) &= x_A(\mathbf{x}_A) + f \cdot (x_B(\mathbf{x}_A + \mathbf{u}_{A\rightarrow B}) - x_A(\mathbf{x}_A)) \\ y_{New}(\mathbf{x}) &= y_A(\mathbf{x}_A) + f \cdot (y_B(\mathbf{x}_A + \mathbf{u}_{A\rightarrow B}) - y_A(\mathbf{x}_A)) \\ z_{New}(\mathbf{x}) &= z_A(\mathbf{x}_A) + f \cdot (z_B(\mathbf{x}_A + \mathbf{u}_{A\rightarrow B}) - z_A(\mathbf{x}_A)) \end{aligned} \tag{6.22}$$

In addition to the newly-defined voxel locations, the intensity values of these data points are also linearly interpolated, according to the difference in intensities between the corresponding positions in the two registered images. This intensity interpolation is similarly expressed in equation 6.23

$$I_{New}(\mathbf{x}) = I_A(\mathbf{x}_A) + f \cdot (I_B(\mathbf{x}_A + \mathbf{u}_{A\rightarrow B}) - I_A(\mathbf{x_A})) \tag{6.23}$$

From the dense collection of scattered intensity values we generate an isotropic volume representation of these intensities, by using the same linear scattered interpolation which we have used previously; the interpolant is subsequently sampled at an isotropic set of coordinates where the resolution is equal to the original slice in-plane resolution.

It has previously been proposed that the robustness of the results can be increased if the registration is performed in both directions between adjacent slices [147]. We therefore generate a second set of results, in which the scattered points

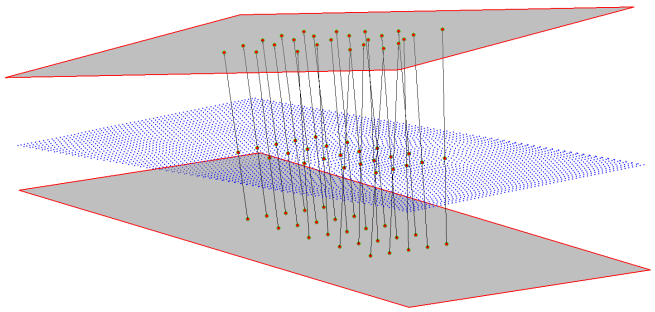


Figure 6.15: Illustration of a deformed set of points generated using a fraction of the 2D registration field calculated between two planes.

are acquired by registering image B to image A. A further scattered interpolation from these points can be used to form a second volume.

The final result is formed by taking a weighted average of the two generated volumes. Each point in the volumes is weighted, based on the distances to the adjacent slices; for the first volume, the voxel weightings are higher nearer the deformed image, B, while in the second volume the weightings are reversed such that they are higher near to image A. The final volume calculation is expressed in equation 6.24.

$$I_F(\mathbf{x}) = w(\mathbf{x}) \cdot V_{A \to B}(\mathbf{x}) + (1 - w(\mathbf{x})) \cdot V_{B \to A}(\mathbf{x}) \tag{6.24}$$

Figure 6.16 illustrates the benefit of using the registration-based approach compared to the standard, linear interpolation of the acquired 2D slices.

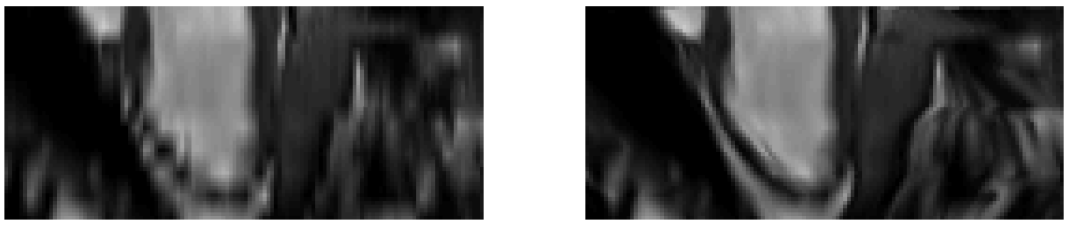
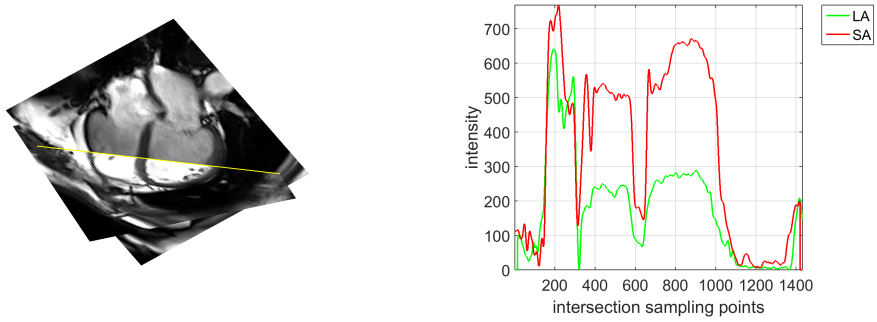


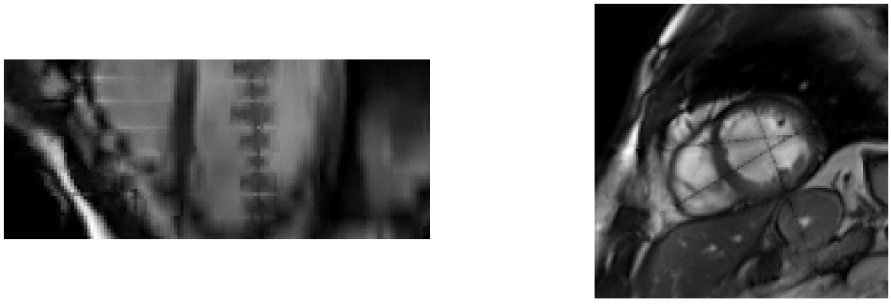
Figure 6.16: On the left, an example cross-section is shown, taken through a linearly interpolated volume generated from the original SA slices. On the right, the same cross-section is taken from a volume interpolated from all the points acquired using registration between slices.

6.3.4 Integration of LA and SA data

All of the above interpolations use only the SA data. However, LA data plays an important role in characterising the shape of the heart, in particular near the apex where the definition provided by the SA is poor. Integrating the complementary data available in the SA and LA is desirable; however, this is not straightforward due in large part to the intensity differences that exist between the slices as a result of the lack of a specific MR intensity scale. The potential intensity differences are shown in figure 6.17a, while the effect on the volumes that can be interpolated is illustrated by the volumetric cross-sections in figure 6.17b. The dark lines are visible at the



(a)



(b)

Figure 6.17: In (a) an illustration of the intensity differences observed between SA and LA data is given. In (b) a volume generated using the combination of SA and LA data illustrates the issue caused by the inconsistent intensities.

locations of the LA data due to their lower relative intensity values. Integration of both SA and LA data therefore requires an initial intensity correction process.

The method of intensity correction presented here is based on that proposed by Zhen et al in [224]. Their algorithm was used for correcting intensities between CT and cone-beam CT images, on the assumption that these differences can be corrected based on small patch-based calculations, rather than a global intensity correction mapping. For our own application, we use the same linear correction model expressed in equation 6.25, but now the intensity differences are defined only at the points of intersection between images rather than over entire image domains.

$$I_{new}(\mathbf{x}) = a(\mathbf{x})I_{orig}(\mathbf{x}) + b(\mathbf{x}) \tag{6.25}$$

The intensity differences are defined at the intersection locations between SA and LA data. While either the SA or LA data could be corrected, here the necessary coefficients are calculated for the LA slices due to the greater amount of intersection information available; each LA image intersects with all SA slices in the stack.

The intersection intensities are used to calculate the necessary coefficients for the mapping in equation 6.25, specifically they are calculated by considering the intensities of local sections of the intersection. Based on short, corresponding sets of the intersection line intensities, $\mathbf{I}_{SA}(\mathbf{x})$ and $\mathbf{I}_{LA}(\mathbf{x})$ centred on $\mathbf{x}$, the $a(\mathbf{x})$ and $b(\mathbf{x})$ coefficients are approximated such that the intensity distributions of the two sets are made to match. The intensity distributions are characterised according to the first and second order intensity moments. The general definition for an intensity moment is given in equation 6.26, where M^p denotes an image moment of order p, and N indicates the number of pixels within each local set of the intersection line. The distributions are matched therefore by matching these two moments, as expressed in equation 6.27.

$$M^p(\mathbf{I}(\mathbf{x})) = \frac{1}{N}\sum_{i=1}^{N} I(\mathbf{x}_i)^p \tag{6.26}$$

$$\begin{aligned} M(\mathbf{I}_{SA}(\mathbf{x})) &= M(a(\mathbf{x})\mathbf{I}_{LA}(\mathbf{x}) + b(\mathbf{x})) \\ M^2(\mathbf{I}_{SA}(\mathbf{x})) &= M^2(a(\mathbf{x})\mathbf{I}_{LA}(\mathbf{x}) + b(\mathbf{x})) \end{aligned} \tag{6.27}$$

If the first and second order moments are matched, then the standard deviation of the distributions will also be matched, giving the relationship shown in equation 6.28. From this relation, an expression for the $a(\mathbf{x})$ coefficient can be derived, given in equation 6.29. Using the estimated $a(\mathbf{x})$ coefficient, and the first moment relationship defined in 6.27, the $b(\mathbf{x})$ coefficient can also be calculated and is expressed in 6.30.

$$\begin{aligned} std\left(\mathbf{I}_{SA}(\mathbf{x})\right) &= std\left(a(\mathbf{x})\mathbf{I}_{LA}(\mathbf{x}) + b(\mathbf{x})\right) \\ &= a(\mathbf{x}) std\left(\mathbf{I}_{LA}(\mathbf{x})\right) \end{aligned} \tag{6.28}$$

$$\begin{aligned} a(\mathbf{x}) &= \frac{std\left(\mathbf{I}_{SA}(\mathbf{x})\right)}{std\left(\mathbf{I}_{LA}(\mathbf{x})\right)} \\ &= \frac{\sqrt{M^2(\mathbf{I}_{SA}(\mathbf{x})) - (M(\mathbf{I}_{SA}(\mathbf{x}))^2)}}{\sqrt{M^2(\mathbf{I}_{LA}(\mathbf{x})) - (M(\mathbf{I}_{LA}(\mathbf{x}))^2)}} \end{aligned} \tag{6.29}$$

$$b(\mathbf{x}) = M(\mathbf{I}_{SA}(\mathbf{x})) - a(\mathbf{x})M(\mathbf{I}_{LA}(\mathbf{x})) \tag{6.30}$$

These calculations provide coefficients for each intersection location in a LA image. It is necessary, however, to correct the entire LA image; to do this we

extrapolate the $a(\mathbf{x})$ and $b(\mathbf{x})$ correction coefficients to all other pixels in the image using the expressions given in equations 6.31, 6.32, and 6.33, which were proposed in [224]. The new values of $a(\mathbf{x})$ and $b(\mathbf{x})$ are approximated by taking weighted averages of the $a(\mathbf{x})$ and $b(\mathbf{x})$ values of surrounding pixels, where T represents the neighbourhood of a pixel $\mathbf{x}$. Weightings are allocated to each neighbourhood pixel, based on the similarity of its intensity to that of the pixel $\mathbf{x}$. The hypothesis here is that pixels that have similar original intensities to $\mathbf{x}$ are likely from the same tissue type, requiring similar correction coefficients, $a(\mathbf{x})$ and $b(\mathbf{x})$; these pixels are therefore given a higher weighting. Dissimilar intensities are indicative of different tissues for which the $a(\mathbf{x})$ and $b(\mathbf{x})$ values will differ. The final parameter, h, is used to adjust the degree to which we want to enforce the intensity similarity constraint.

$$a(\mathbf{x}) = \frac{\sum_{i \in T} w(\mathbf{x}_i) a(\mathbf{x}_i)}{\sum_{i \in T} w(\mathbf{x}_i)} \tag{6.31}$$

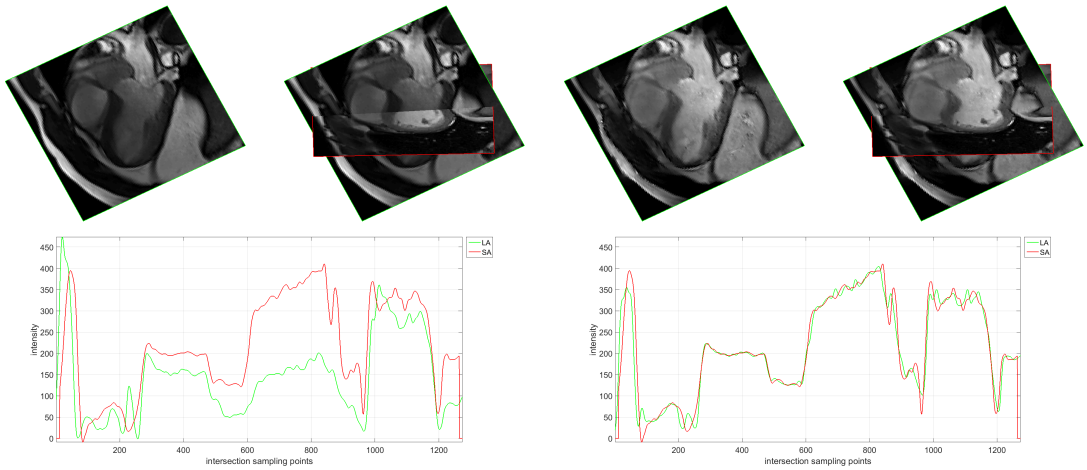
$$b(\mathbf{x}) = \frac{\sum_{i \in T} w(\mathbf{x}_i) b(\mathbf{x}_i)}{\sum_{i \in T} w(\mathbf{x}_i)} \tag{6.32}$$

$$w(\mathbf{x}_i) = e^{-\left(\frac{(I_{LA}(\mathbf{x}_i) - I_{LA}(\mathbf{x}))^2}{h^2}\right)} \tag{6.33}$$

The $a(\mathbf{x})$ and $b(\mathbf{x})$ correction maps are initially very sparse, consisting of only a few lines of known correction coefficients; therefore, an iterative process must be used to fill the unknown values. In order to reach a satisfactory extrapolation over the entire LA image domain, 1000 iterations of the calculations were performed with the h parameter set to 30.

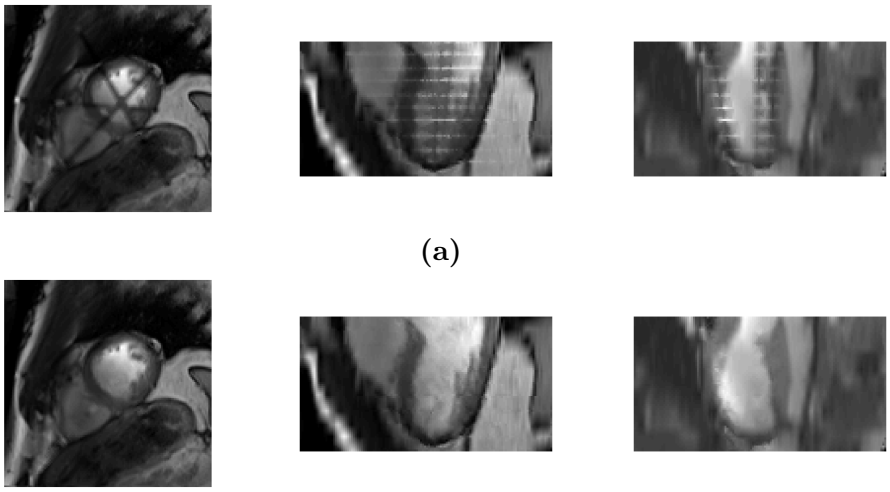
For each dataset, intensity correction was performed on the three LA slices for all cardiac cine time frames. In figure 6.18a, an LA image along with a single intersecting SA slice is shown. We see from the intensity profile indicated that the intensities differ but, more significantly, that the mode of variation is not consistent across tissue type. In figure 6.18b, the same LA-SA image pair and its intersection profile is shown after intensity correction.

In figure 6.19, cross-sections of volumes generated with the original uncorrected image data and, subsequently, the corrected data are shown.



(a) **(b)**

Figure 6.18: On the left-hand side, in (a), the respective intensities of the LA and SA images are shown before correction. On the right-hand side the same slices are indicated after correction.



(a)

(b)

Figure 6.19: The interpolated volumes using both SA and LA data, with and without intensity correction are shown in (a) and (b) respectively.

Introduction of LA data to previous reconstructions

Having corrected for the intensity discrepancies between SA and LA data we can now introduce this additional information into the previously discussed reconstruction approaches; in each case, this provides an increase in the number of available scattered data points for interpolation.

6.4 Experiments and Results

For the frame-to-frame reconstruction, the 3D phase-based optical flow registration approach previously described in chapter 5 was used. Due to the similarity of adjacent frames, and in order to save overall computation time, the registration was performed at a single local phase scale using the parameter $\lambda = 3$. The global smoothing parameter was set to $\alpha = 1$, and local smoothing parameter $\rho = 6$.

For the 2D registration between slices, we used the same set of parameters as previously applied for the registration of cine volumes from different hospital visits. Specifically, this set of parameters has an initial local phase scale of $\lambda = 10$; this larger scale enables the alignment of structures that differ more significantly between the adjacent SA images.

Reaction-diffusion results were acquired for volumes initialised using each of the methods discussed, both with and without LA data. The results are presented qualitatively in the form of cross-sections from the final volumes, alongside the original linear reconstructions. The reaction-diffusion results shown were acquired using a local smoothing parameter $\sigma = 1$, an integration scale $\rho = 5$, $t = 30$ iterations, and reaction-diffusion component weightings $\alpha = 0.5$.

In addition to the visual results, a series of quantitative measures is calculated in order to assess the quality of the reconstructions; these are the signal-to noise ratio (SNR), contrast-to-noise ratio (CNR), and a phase-based feature point estimator. As indicated in their respective definitions below, these quantities are estimated based on the left ventricular region only - *myo* indicates the myocardial region defined by the generated endo- and epi-cardial meshes, and *bp* indicates the blood pool region enclosed by the endocardial mesh. We consider the reconstructions specifically in this region, as the definition of these structures for the purposes of clinical delineation is an important appraisal of reconstruction quality. The calculated metrics are reported alongside the first set of results to give qualitative context to the numbers. The final quantitative results are summarised in tables 6.1 and 6.2. The first table reports the changes in metric values after reaction-diffusion

in comparison to their respective initialisations, while the second table reports the changes in each case compared to the basic linear interpolation approach.

Percentage change in signal-to-noise ratio (SNR): The SNR is calculated in both the myocardial and blood pool regions for both the initialisation and reaction-diffusion reconstructions. The SNR is calculated based on the mean and standard deviation of the intensities in these regions respectively, and the final $\Delta(SNR)$ measure is taken as the average of the two.

$$\begin{aligned}\Delta(SNR) &= \frac{1}{2}\left(\Delta(SNR^{myo}) + \Delta(SNR^{bp})\right) * 100 \\ &= \frac{1}{2}\left(\left(\frac{SNR^{myo}_{react-diff}}{SNR^{myo}_{linear}} - 1\right) + \left(\frac{SNR^{bp}_{react-diff}}{SNR^{bp}_{linear}} - 1\right)\right) * 100 \\ \text{where, } SNR &= \frac{\mu}{\sigma}\end{aligned} \tag{6.34}$$

Percentage change in contrast-to-noise ratio (CNR): We consider the contrast between the myocardium and blood pool in each of the interpolations and quantify its increase relative to the noise level of the image.

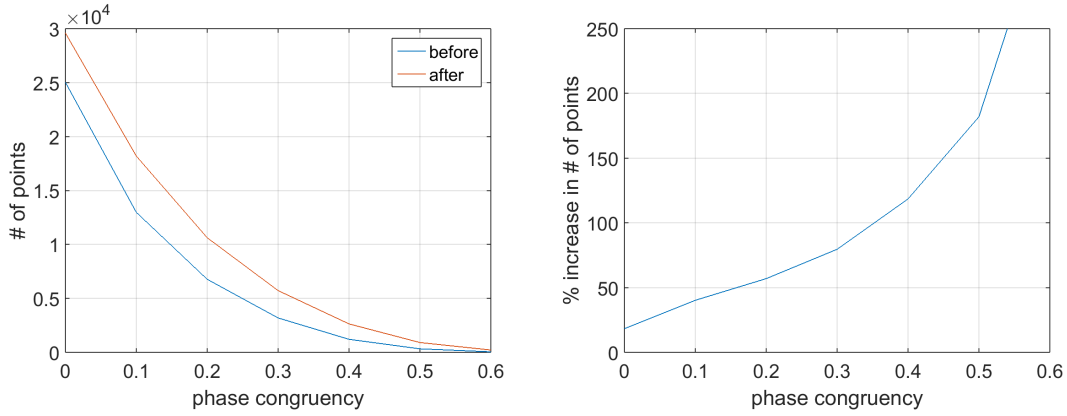
$$\begin{aligned}\Delta(CNR) &= \left(\frac{CNR_{react-diff}}{CNR_{linear}} - 1\right) * 100 \\ \text{where, } CNR &= \frac{\mu^{myo} - \mu^{bp}}{\sqrt{(\sigma^{myo})^2 + (\sigma^{bp})^2}}\end{aligned} \tag{6.35}$$

Percentage change in phase congruency-based feature points: A desirable characteristic of an improved volumetric interpolation is the superior definition of important structural surfaces - here in particular we look for better definitions of both the endo- and epi-cardial layers. Phase congruency is a measure used for feature identification, with values that fall in the range 0 to 1. An example of the phase congruency calculated for a volume before and after reaction-diffusion is shown in figure 6.20. Here we hypothesise that, after the application of the reaction-diffusion interpolation, a greater number voxels in the endo- and epi-cardial surface regions will be characterised by higher phase congruency values. Therefore, the reconstruction quality is quantified via the percentage increase in the number of phase congruency-based feature



(a) **(b)**

Figure 6.20: Phase congruency values calculated for a volume before, (a), and after, (b), application of reaction-diffusion.



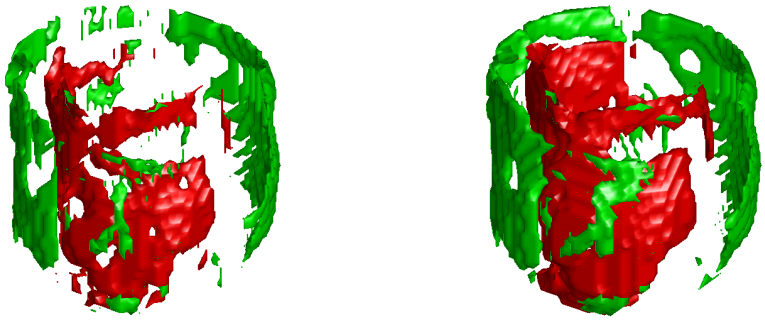
(a) **(b)**

Figure 6.21: The variation in feature points detected as a function of the phase-congruency threshold value used. In (a) the number of feature points is plotted before and after reaction-diffusion in blue and red respectively, while in (b) the percentage change is plotted.

points identified above a given threshold, within constrained regions near the endo- and epi-cardial mesh surfaces.

$$\Delta(F) = \left(\frac{\#F_{react-diff}}{\#F_{initial}} - 1 \right) * 100 \tag{6.36}$$

In figure 6.21 we see the variation based on this selected threshold. What is clear is that in each case the number of feature points is increased in the reaction-diffusion case. Figure 6.22 provides a visual illustration of the increase in the number of identified feature points at the endo- and epi-cardial surfaces, based on a phase-congruency threshold of 0.2.



(a) **(b)**

Figure 6.22: 3D illustration of the phase-congruency based feature points detected in the vicinity of the endo- (red) and epi-cardial (green) surfaces before and after application of the reaction-diffusion method.

Single-frame linear initialisation

Figure 6.23 shows a series of example results where reaction-diffusion has been applied to the original linear reconstructions generated from the original SA slice data. In each case a LA cross-section is shown taken through the left and right ventricles.

Multi-frame linear initialisation

In figure 6.24, the results acquired using the multi-frame initialisation are shown. A number of LA cross-sections are shown for each example result, which illustrate: the single-frame linear interpolation; the position of deformed SA slices; the linear interpolation from all deformed points; and the final reaction-diffusion result initialised with this interpolation.

Registration-based linear initialisation

In figure 6.25 a number of examples are shown of cross-sections taken from volumes generated using a registration-based interpolation technique. Specifically, these are indicated in the middle column of the figure. The first column shows the equivalent linear reconstruction, while the last column indicates the result obtained from the reaction-diffusion system when the registration-based initialisation image is used.

Integrated LA and SA initialisation

In figure 6.26 results acquired using the combination of both SA and LA data are shown. Specifically, these are the results when both LA and SA data are used as part of the multi-frame reaction-diffusion approach. These results are shown as they most clearly illustrate the additional benefits of using LA data as part of the reconstruction process.

		$\Delta(SNR)$	$\Delta(CNR)$	$\Delta(\#F)$
Single-frame	SA	$6.83\% \pm 2.51\%$	$3.90\% \pm 1.69\%$	$58.57\% \pm 18.86\%$
	LA	$7.23\% \pm 3.19\%$	$4.13\% \pm 1.97\%$	$71.16\% \pm 21.98\%$
Multi-frame	SA	$6.55\% \pm 3.23\%$	$3.44\% \pm 2.15\%$	$77.61\% \pm 25.50\%$
	LA	$6.48\% \pm 2.80\%$	$2.90\% \pm 1.67\%$	$99.60\% \pm 33.28\%$
Registration-based	SA	$5.32\% \pm 2.37\%$	$3.81\% \pm 2.20\%$	$47.48\% \pm 14.31\%$
	LA	$5.67\% \pm 2.53\%$	$4.04\% \pm 2.18\%$	$51.58\% \pm 14.71\%$

Table 6.1: Changes in signal-to-noise ratio (SNR), contrast-to-noise ratio (CNR), and phase congruency-based features ($\#F$) calculated for the reaction-diffusion results compared to the respective initialisation volumes.

		$\Delta(SNR)$	$\Delta(CNR)$	$\Delta(\#F)$
Single-frame	SA	$6.83\% \pm 2.51\%$	$3.90\% \pm 1.69\%$	$58.57\% \pm 18.86\%$
	LA	$6.30\% \pm 3.45\%$	$3.67\% \pm 3.10\%$	$62.28\% \pm 18.29\%$
Multi-frame	SA	$6.10\% \pm 5.55\%$	$-3.40\% \pm 6.33\%$	$44.45\% \pm 20.68\%$
	LA	$6.21\% \pm 5.35\%$	$-2.63\% \pm 7.33\%$	$49.62\% \pm 20.92\%$
Registration-based	SA	$4.74\% \pm 3.50\%$	$4.66\% \pm 2.76\%$	$62.31\% \pm 16.64\%$
	LA	$4.87\% \pm 3.54\%$	$4.74\% \pm 2.67\%$	$63.91\% \pm 17.30\%$

Table 6.2: Changes in signal-to-noise ratio (SNR), contrast-to-noise ratio (CNR), and phase congruency-based features ($\#F$) calculated for the reaction-diffusion results compared to the basic linear interpolation volumes.

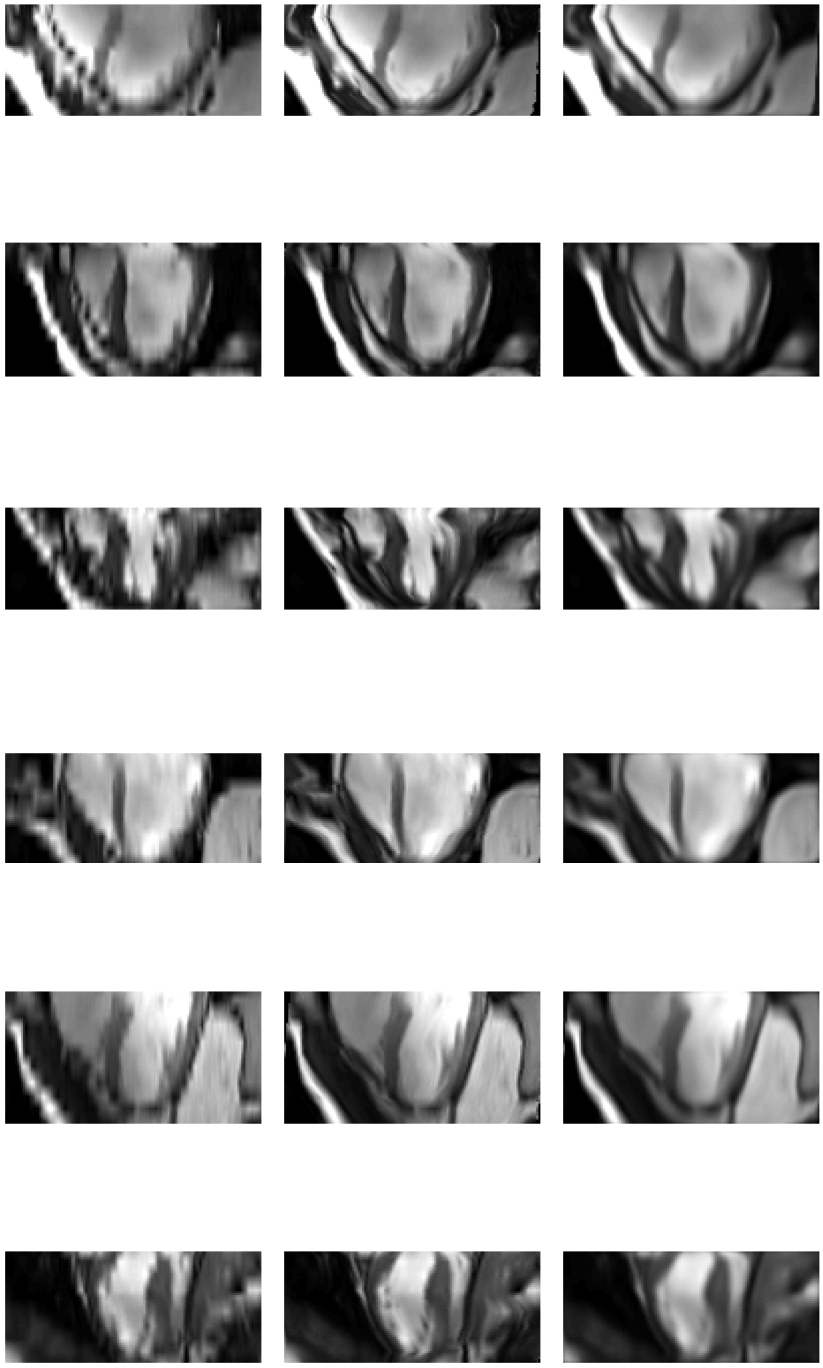


Figure 6.23: Example reconstruction results using the reaction-diffusion approach on single-frame linear initialisation volumes.



(a)	**(b)**	**(c)**	**(d)**
Single-frame linear reconstruction	Deformed slice data	Multi-frame linear reconstruction	Reaction-diffusion reconstruction

Figure 6.24: Example reconstruction results using the reaction-diffusion approach on multi-frame linear initialisation volumes.

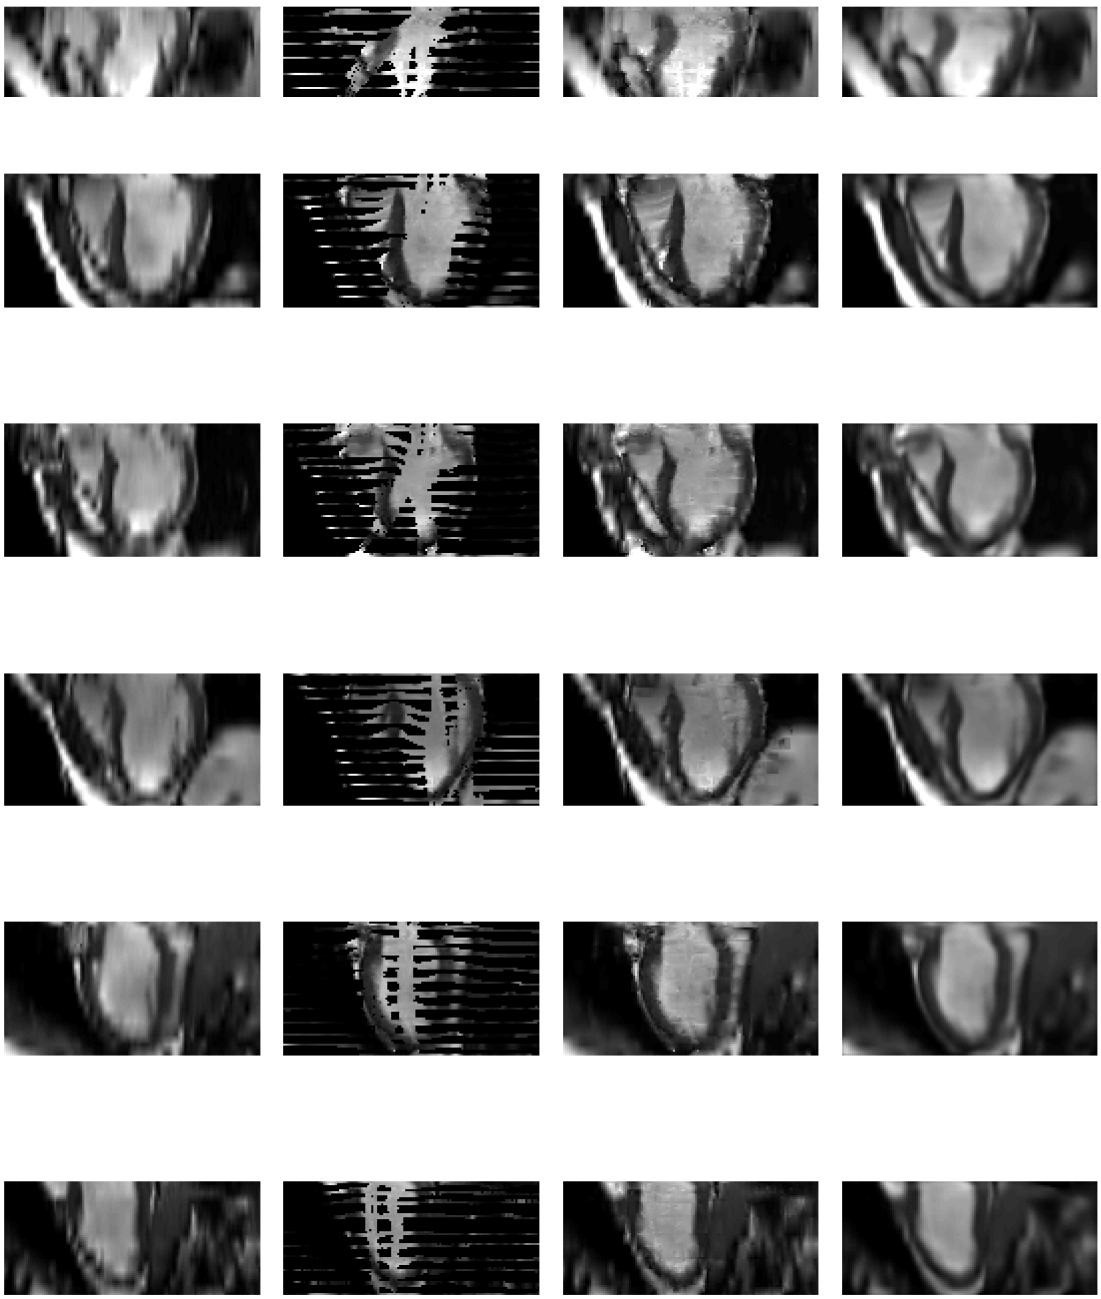


(a) Single-frame linear reconstruction

(b) Registration-based linear reconstruction

(c) Reaction-diffusion reconstruction

Figure 6.25: Example reconstruction results using the reaction-diffusion approach on registration-based linear initialisation volumes.



(a)	**(b)**	**(c)**	**(d)**
Single-frame linear reconstruction	Deformed slice positions	Multi-frame linear reconstruction	Reaction-diffusion reconstruction

Figure 6.26: Example reconstruction results using the reaction-diffusion approach on multi-frame linear initialisation volumes using both LA and SA information. (See figure 6.14 for an equivalent example result using SA data only. See the end of section 6.3.2 for the corresponding description of the different images shown.)

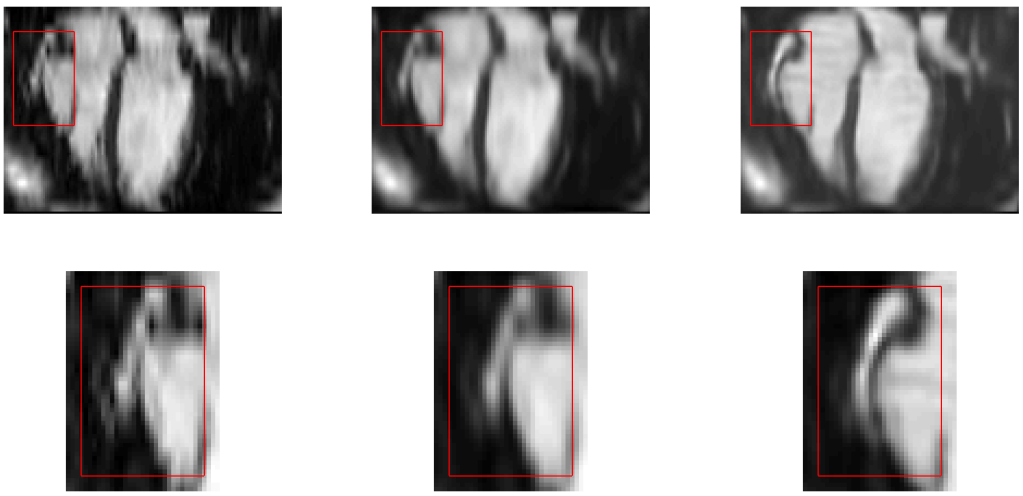
6.5 Discussion

We have presented a series of reconstruction methods, specifically designed for their use in cardiac MR sequences, that are based on the reaction-diffusion process. In the first part of the chapter we discussed the theoretical concepts of the reaction-diffusion process, and introduced its application to the interpolation of thick slice cardiac MR data. The reaction-diffusion process requires an initial volumetric reconstruction. A number of variants were explored which looked to best use the initially available data; in this discussion we will consider the merits and limitations of each.

The anisotropic reaction-diffusion process alone is capable of enhancing the qualitative appearance of the reconstructions. In figures 6.23, 6.24, 6.25, and 6.26, we saw that compared to their respective initialisations the final reconstructions showed visible improvements in the continuity of structures, the definition of boundaries such as those that exist between the myocardium and blood pool, and the homogeneity of intensities within regions. The improvement in the definition of boundaries is reflected in particular in the phase-based features identified, where as reported in table 6.1 there are notable increases in all cases after reaction-diffusion. We similarly see improvements in both the SNR and CNR.

The final reaction-diffusion result is dependent on the initialisation image provided. Due to the slice thickness and the gaps between acquired images, there is a significant distance between slice centres; as a result, direct linear interpolation of the slice data can be inaccurate. A number of alternative initialisations were therefore explored.

The first of these we referred to as the multi-frame initialisation. We proposed the combination of complementary information contained in all of the acquired cardiac cine frames, by performing a 3D registration of volumes generated for each frame. The estimated deformations were used to calculate the resulting positions of the SA slice voxels from each frame. The deformation of the slices from different frames results in a filling in of gaps between the acquired slice centres.



(a) Linear single-frame reconstruction

(b) Reaction-diffusion single-frame reconstruction

(c) Reaction-diffusion multi-frame reconstruction

Figure 6.27: Illustration of the reconstruction improvements possible - firstly using the single-frame reaction-diffusion method, and secondly the multi-frame reaction-diffusion method - compared to the basic linear reconstruction.

The potential for this method was previously exemplified in figure 6.14; here in figure 6.27, an enlarged view of the right ventricle's basal region from this result is shown. The initial single-frame linear interpolation, the single-frame reaction-diffusion, and the multi-frame reaction-diffusion reconstructions are displayed. The single-frame reaction-diffusion result is qualitatively superior to the basic linear interpolation, but the multi-frame method improves further upon this, capturing a greater proportion of the high frequency detail near the right ventricle base.

While this example illustrated the potential of the method, it is worth noting that the dataset on which the method was tested was acquired from a healthy volunteer. The results when applied to real infarct patient data showed in general a lower proportion of filling in between the SA slices. This is illustrated in figure 6.28 where the position of all deformed points is shown for a number of different cases. The first line contains examples of volunteer patient datasets and in each case we see a significant proportion of the gaps is filled in. In contrast, the second line indicates some examples from real patients where the proportion of filled in points is visibly reduced. This result reflects the compromised contractile action



(a) Volunteer patient data

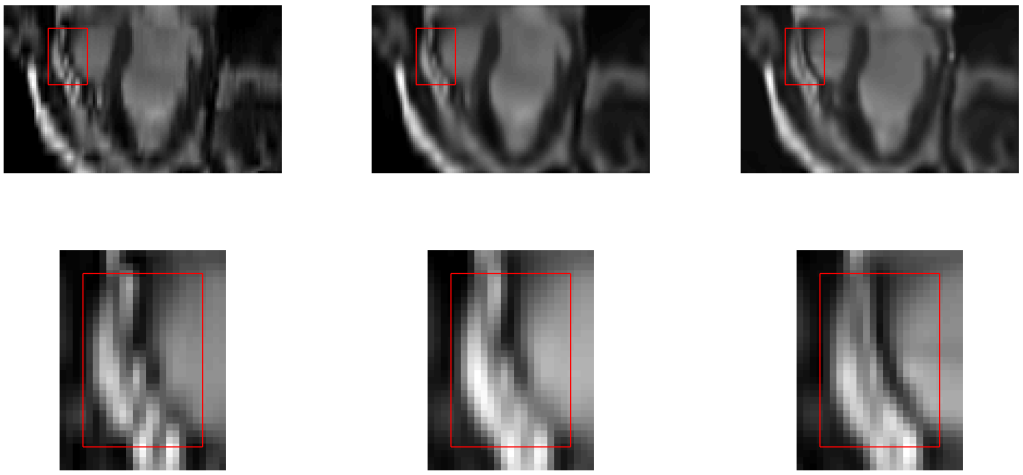


(b) Infarction patient data

Figure 6.28: Illustration of results obtained using healthy volunteer data (a) compared to patient data (b). Resultant positions of SA data points from each cardiac frame after frame-to-frame registration are shown. Due to the reduction in contractile capability of the heart in infarcted patients, these results tend to display less filling in of the gaps between acquired slices compared to that observed in healthy volunteers.

in many of the patients. While the filling in may be reduced in a number of cases, this does not mean that no benefit was observed. Figure 6.24 in the results section illustrated a number of example cases for real patient data where the definition of the cardiac ventricle anatomy is improved using the multi-frame reaction-diffusion approach, while figure 6.29 shows an enlarged region taken from an infarct patient dataset, where, in an equivalent way to the volunteer healthy patient data, we see in particular the improvement of the right ventricle endocardial surface.

In table 6.2 we reported that, as with all of the approaches, the SNR is improved for the multi-frame reaction-diffusion results compared to the basic linear reconstruction; however, the CNR is slightly reduced. This result is likely due to an increase in the intensity variation and noise in the initial multi frame reconstruction - some variation in intensity values is introduced by the other cine frames, in particular in the blood pool, while simultaneously the spatial distribution of the scattered input data results in a poorer quality interpolation. Similarly, while there is an increase in the number of phase-based feature points identified, the degree of improvement is less than that reported for the single-frame reaction-diffusion and registration-based



(a) Linear single-frame reconstruction

(b) Reaction-diffusion single-frame reconstruction

(c) Reaction-diffusion multi-frame reconstruction

Figure 6.29: Illustration of the reconstruction improvements possible - firstly using the single-frame reaction-diffusion method, and secondly the multi-frame reaction-diffusion method - compared to the basic linear reconstruction.

reaction-diffusion approaches. This again we put down to the quality of the multi-frame initialisation image; a greater number of iterations of the reaction-diffusion approach may be needed in order to achieve comparable feature-based improvements in the multi-frame case. Despite the quantitative results, it is important to note the improved qualitative appearance of ventricle shape in the multi-frame results.

The second proposed initialisation was a registration-based approach, the results of which were indicated in figure 6.25. The improvements in this case are qualitatively very clear; the connectivity of image structures contained in adjacent slices is improved and the sharpness of the structure boundaries is increased. The former observation is particularly clear if we consider right ventricular walls. The improvement in interpolated volume quality is backed up by increases in each of the SNR, CNR and phase-based feature points identified.

The primary limitation of the registration-based approach is the risk introduced by mis-registration. If the registration method employed fails to capture the differences between two slices, the final reconstruction result will be erroneous. This was observed in two cases in our own results and is illustrated in figure 6.30,



Figure 6.30: Illustration of slice-to-slice mis-registration, leading to errors in the subsequent registration-based linear reconstructions.

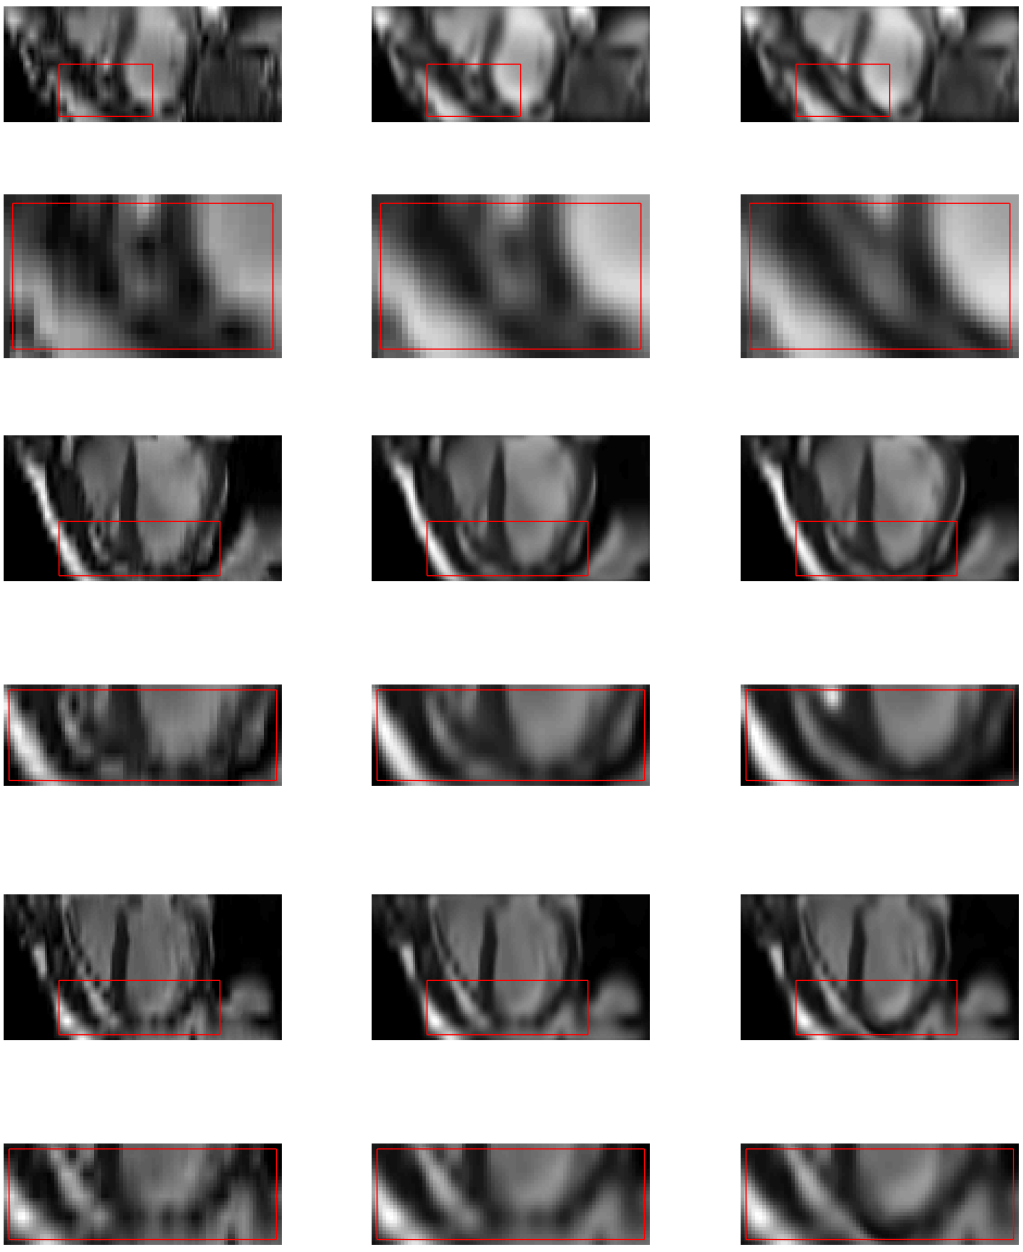
where we see that a region of the right ventricle, and a region of the left ventricle are incorrectly registered respectively. While it is possible that other registration methods might avoid these particular mis-registrations, we cannot ensure that no case-specific registration issues arise.

The final addition to the reconstructions was the introduction of LA data; this provides high resolution information in an orientation that is approximately perpendicular to that of the SA. To enable this integration of information, the intensities of the LA data first had to be adjusted to match the SA data.

With the adjusted intensities, the introduction of the LA data is most clearly illustrated by considering the multi-frame interpolation approach. Figure 6.26 gave some examples of these reconstructions, and we saw the increase in the total number of points from which the interpolated volume could be generated. A primary desired outcome from the introduction of LA data is the improved definition of the apical region of the left ventricle; in figure 6.31 we see examples of this. We see that utilisation of the reaction-diffusion approach on the basic SA-only reconstruction shows some improvement, but the addition of the LA data provides clear advantages for the definition of the apical region.

In table 6.2, where the final reconstructions are all compared to the basic linear interpolation, we see that there is little quantitative difference between the improvements observed when using both the LA and SA data, compared to the SA alone - the benefits of combined utilisation of LA and SA are best appreciated qualitatively.

Therefore, in figure 6.32 we show further examples of the visual improvements provided by the inclusion of LA data. Here we compare the reconstructions generated



(a) Linear single-frame reconstruction

(b) Reaction-diffusion single-frame reconstruction

(c) Reaction-diffusion single-frame reconstruction with LA

Figure 6.31: Illustration of the reconstruction improvements using the reaction-diffusion method both with and without LA data.

with each of the initialisation approaches. In these cases the best definition of the apical region is provided by the multi-frame approach, as the deformation of the LA planes during the contractile action of the heart provides a greater coverage of this area. In contrast we see that the registration-based approach does not see as much benefit from using the LA data - this is possibly due to the higher density of scattered points in the apical region, generated via the registration procedure, meaning that the LA slices introduced have a lesser effect on the final reconstructed volume.

We have introduced a series of volume reconstruction techniques in this chapter. While there are benefits to each technique, they also have limitations, meaning that the selection of which approach to use must be subject to a few considerations.

The basic reaction-diffusion technique works consistently well with all datasets. For the initialisation of the reaction-diffusion procedure we would conclude that both the multi-frame and registration-based approaches are superior to using a basic linear interpolation.

While the multi-frame initialisation improves the definition of certain regions by filling in the gaps between slices, the process of registering all frames is quite time-consuming. In addition, we see that in patients whose contractile function is compromised, the filling-in process - which is the primary benefit of this technique - is limited; in such cases there is minimal gain from this approach. However, it is worth noting that, when both SA and LA data are combined, it is this multi-frame approach that sees the most benefit and provides the best definition of the apical regions.

If a faster reconstruction is required, the registration-based initialisation is a better choice as it only requires 2D registrations between SA slices. However, it is important with this approach to be aware of any potential mis-registrations between SA planes that lead to errors in the reconstructions.

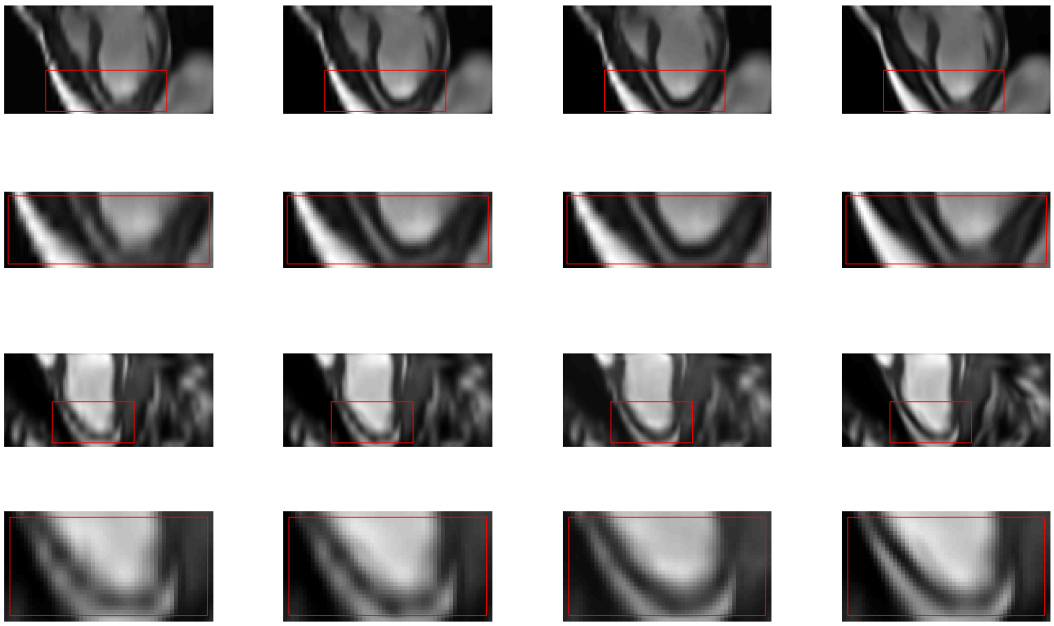
Multi-image super-resolution techniques, which utilise complementary information of the same scene from different viewpoints, are one of the primary categories of algorithms in the literature. In the context of cardiac MR data, these viewpoints can be the different SA and LA slices acquired during a typical MR examination. While we have similarly combined these complementary views in our approach,

we have gone further and in addition utilised the extra information contained in the other frames of the cardiac cine sequence. This approach takes advantage of the fact that, while data are acquired at stationary planar positions, the cardiac structures are in motion and therefore the specific tissues imaged across the cycle change. By estimating the motion via image registration techniques, the imaged tissues can be propagated to their positions at one temporal time-point.

An important aspect of many of the papers in the literature is the MR imaging model, which accounts in particular for the anisotropy of the acquired data; this model looks to identify the super-resolution reconstructions that, under its operation, provide optimal correspondence to the true low-resolution MR data acquired. The thickness, orientation and profile of the acquired slice data were incorporated into our own approach via the reaction component of the reaction-diffusion interpolation method, which ensured that the isotropic reconstruction was reflective of the true slices acquired.

There are a number of papers in the literature that propose super-resolution reconstruction of cardiac MR data by utilising machine learning methods; these take advantage of the similar structural data available in higher resolution MR images and have obtained some qualitatively impressive results. Some of these approaches are based on combined low- and high-resolution dictionaries; the most notable drawback of such techniques is the requirement for large amounts of corresponding data at both resolution levels from which the dictionaries are formed. More recently published are deep-learning processes that are able to learn directly the models that allow for up-sampling from the low- to high-resolution data; here again there is a requirement initially for a significant amount of training data from which the relationships can be learned before being applied to generate new reconstructions.

Whilst these methods have proven very successful in the literature, we feel that an advantage of the approaches described in this chapter is that they are based solely on the data acquired for an individual patient, not being reliant on the availability of a prior database of images. This may be particularly important in cases of severe or abnormal remodelling where previous representative high-resolution data



(a)	**(b)**	**(c)**	**(d)**
Reaction-diffusion single-frame reconstruction with SA only	Reaction-diffusion single-frame reconstruction with LA	Reaction-diffusion multi-frame reconstruction with LA	Reaction-diffusion registration-based reconstruction with LA

Figure 6.32: Illustration of the reconstruction improvements using the reaction-diffusion method with LA data for each of the initialisation approaches proposed.

may not be available. A further advantage is that key components of the methods such as multi-frame alignment or slice-to-slice correspondence can potentially be accomplished using easily available registration methods.

7 Conclusions and Future Work

Contents

7.1 Conclusions 222
7.2 Future Work 224

7.1 Conclusions

Cardiac remodelling following ischemic injury is a key indicator of patient prognosis. The analysis of local tissue changes over time has the potential to offer new insight into the ventricular remodelling that occurs in response to clinical intervention. In this thesis we have presented a framework for the longitudinal registration of CMR data; local phase-based methods were proposed that accomplish the steps necessary to take 2D planar cine and LGE acquisitions through to the identification of 3D correspondences on interpolated data. The final registration results were used to exemplify the potential for clinical analysis through the identification and comparison of scar tissue.

Firstly, it is worth commenting on the reasons for undertaking the 3D analysis presented. Conventional analyses in the clinical setting consider the developments in diseased state via the acquisition of corresponding images during follow-up hospital sessions. However, acquisition of the same images is very difficult, and

therefore the actual planes recovered do not match. We gave some indication to these errors in the study presented in chapter 3; in that work we estimated plane discrepancies via the registration of an LA slice and showed that these slice misalignments are not insignificant, concluding therefore that these discrepancies should not be neglected during analysis.

Having identified the need for a registration that accounts for deformations both in- and out-of-plane, in chapter 5 we presented a framework capable of bringing into alignment cine and LGE image data acquired during multiple hospital visits in a 3D domain. As a result of patient motion and breath-hold acquisition, a number of registration steps are required to account for spatial discrepancies in the data.

Due to differences in intensities between acquisitions, we opted for registration methods that were based on local phase data. These algorithms proved effective in improving the correspondence of data in all applications, with the benefit of local phase information being of particular note in the results reported for SA-LA slice realignment and 3D registration when compared to intensity-based methods.

Finally, we demonstrated how the framework outlined allowed for the analysis of how regions of scar, or specifically regions of hyper-enhancement, in the LGE images change between the two acquisitions.

In the second part of the thesis we addressed the fundamental anisotropy limitation of MR imaging. The 2D planar images that are acquired as standard during a cardiac imaging examination result in a representation of the data that suffers from partial volume effects and, when integrated to 3D, results in poor visual reconstructions. We presented a series of interpolation results using a reaction-diffusion-based scheme. The diffusion process, which is anisotropic, reduces the influence of noise and improves the connectivity of structures due to its use of local structure description, while the reaction term ensures the fidelity to the original input data in the final result. Tests were carried out using a number of different volume initialisations, each of which combined the available planar input data in different ways.

While the post-processing-based approaches discussed were able to improve somewhat on the reconstructed volumes generated, they are fundamentally limited according to the image sets acquired. Clearly, any reduction in slice thickness and inter-slice gaps will improve the final reconstruction; the partial-volume effect will be decreased and, in particular, the multi-frame reconstruction approach will be enhanced due to an increased proportion of gap that can be filled. However, implementing such changes may not be possible due, for example, to sequence restrictions. Alternatively, notable benefits were observed by the inclusion of LA data in the reconstructions; therefore, procurement of a greater number of LA slices may be a more straightforward modification for improving the results obtained. In addition to the improved definition of ventricle boundaries, especially at the apex, the additional information provided by extra LA slices in the form of further slice intersections could improve the robustness of the earlier SA-LA realignment operation.

We note that the primary objective of our reconstruction methods has been to enhance image quality so that the accuracy of subsequent processing steps, such as segmentation and registration can be improved - in particular in the apical region. Using the reaction-diffusion-based interpolation schemes presented we have improved the continuity of many of these salient structural surfaces such as the ventricle borders, reducing the appearance of anisotropic-resolution-induced blurring. This outcome, which in our quantitative results is best exemplified by the increase in phase-based feature points, could lead to improvements in both segmentation and registration post-processing tasks. In order to further improve the results using the methods presented here, we believe that acquisition of some additional image data will be of greatest benefit.

7.2 Future Work

The requirement for the data registration steps outlined in this thesis originate from a lack of correspondence in the slices acquired during the distinct scanning sessions; this effect might be reduced if a registration step were introduced into

the acquisition process followed by clinicians. The LA images are acquired in advance of the SA stack; if, at this stage, a simple and fast registration of LA data were performed, it may be possible to improve the correspondence of the SA data acquired, and thus reduce errors introduced in subsequent analysis.

The primary objective of the presented framework is the possible benefit to clinical analysis from bringing complementary CMR protocols into the same coordinate frame. Further work is required therefore that correlates changes in functional information, interpreted via the analysis of the cine sequences from each visit, with the corresponding changes in scar tissue observed. Work currently being undertaken by fellow group members explores this functional information, with emphasis on the stresses and strains experienced in healthy and pathological myocardial regions. In future work, this analysis, which is currently undertaken on the single-visit data, could be extended to the co-registered multi-visit data in order to assess the corresponding changes in strains experienced.

As was outlined in chapter 2, the efficient operation of the heart is a result of complex electro-mechanical processes. Therefore, it is of additional interest to see how the mechanical deficiencies may be related to underlying electrical conductivities. By utilisation of computational models - which have become important tools in cardiac research and are increasingly employed in combination with image data [100] - it would be interesting to see how the specific scar tissues identified affect the propagation of the electrical signal; how this is indicated in conventional clinical signals such as the ECG; and, finally, how these electrical characterisations change during follow-up examinations. Based on observations in regard to the localised scar development, it may be interesting to see if there are indicators in the mode of electrical signal propagation, characterised in the ECG signal, that can be recognised as signs of the expected cardiac remodelling process.

As a longer-term objective, the combination between images and models can further the advance of personalised medicine; computational models, personalised by means of the images, can help us to understand the condition of a particular patient, predict the evolution of the disease and design personalised treatments.

To improve upon and advance the work presented relating to volumetric reconstruction of cine-MR data, there are two clear future steps, both of which relate to further validation of the work. A primary clinical goal of the super-resolution cardiac cine-MR interpolation is to improve the segmentations of the ventricles, leading to improved estimation of clinically important cardiac indices, for example the EF. Therefore, a quantification of the reconstruction results based on segmentation approaches would perhaps give clinicians a more illustrative measure of the algorithm's performance. Such a study could consider both the inter- and intra-user reproducibility of post-processing cine analysis, in addition to the performance of state-of-the-art automated segmentation techniques on different reconstructions produced. Coupled with this validation work, a comparison of our results with other state-of-the-art methods, in particular machine learning approaches that are becoming common-place in the literature, would be a valuable addition. However, this work is currently hindered by a lack of easily available algorithms and the need for extensive databases of high resolution images on which to train the required dictionaries or networks.

References

[1] Global health estimates 2015: Deaths by cause, age, sex, by country and by region, 2000-2015. Geneva, World Health Organization, 2016.

[2] D. Adair, M. Walton, A. Kroker, K. Rivard, S. Sudhindra, and S. Pandey. Physics of MRI. MediaWiki. [Online; accessed 12-May-2017].

[3] M. Alessandrini, A. Basarab, H. Liebgott, and O. Bernard. Myocardial motion estimation from medical images using the monogenic signal. *IEEE Transactions on Image Processing*, 22(3):1084–1095, 2013.

[4] 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, M. Pop, K. Rhode, M. Sermesant, and A. Young, editors, *Statistical Atlases and Computational Models of the Heart. Imaging and Modelling Challenges, STACOM 2012*, volume 7746 of *Lecture Notes in Computer Science*, pages 159–168. Springer, Berlin, Heidelberg, 2013.

[5] S. C. Amartu and H. J. Vesselle. A new approach to study cardiac motion: The optical flow of cine MR images. *Magnetic Resonance in Medicine*, 29(1):59–67, 1993.

[6] F. Andreu, C. Ballester, V. Caselles, J. Mazón, et al. Minimizing total variation flow. *Differential and Integral Equations*, 14(3):321–360, 2001.

[7] L. Axel and L. Dougherty. MR imaging of motion with spatial modulation of magnetization. *Radiology*, 171(3):841–845, 1989.

[8] L. Axel, A. Montillo, and D. Kim. Tagged magnetic resonance imaging of the heart: a survey. *Medical Image Analysis*, 9(4):376 – 393, 2005.

[9] W. Bai, D. Peressutti, S. Parisot, O. Oktay, M. Rajchl, D. O'Regan, S. Cook, A. King, and D. Rueckert. Beyond the AHA 17-segment model: motion-driven parcellation of the left ventricle. In *International Workshop on Statistical Atlases and Computational Models of the Heart*, pages 13–20. Springer, 2015.

[10] W. Bai, W. Shi, A. de Marvao, T. J. 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.

[11] W. Bai, W. Shi, C. Ledig, and D. Rueckert. Multi-atlas segmentation with augmented features for cardiac MR images. *Medical image analysis*, 19(1):98–109, 2015.

[12] W. Bai, W. Shi, D. P. O'Regan, T. Tong, H. Wang, S. Jamil-Copley, N. S. Peters, and D. Rueckert. A probabilistic patch-based label fusion model for multi-atlas segmentation with registration refinement: application to cardiac MR images. *IEEE Transactions on Medical Imaging*, 32(7):1302–1315, 2013.

[13] J. Barajas, K. L. Caballero, J. G. Barnés, F. Carreras, S. Pujadas, and P. Radeva. Correction of misalignment artifacts among 2-D cardiac MR images in 3-D space. In *Proceedings of the 2006 1st International Workshop on Computer Vision for Intravascular and Intracardiac Imaging*, pages 114–121. IEEE, 2006.

[14] D. Barbosa, D. Friboulet, J. D'hooge, and O. Bernard. Fast tracking of the left ventricle using global anatomical affine optical flow and local recursive block matching. *Proceedings of the MICCAI Challenge on Endocardial Three-dimensional Ultrasound Segmentation-CETUS*, pages 17–24, 2014.

[15] J. Baron. Experience with 3D optical flow on gated MRI cardiac datasets. In *Proceedings of the 2004 1st Canadian Conference on Computer and Robot Vision*, pages 370–377. IEEE, 2004.

[16] J. L. Barron, D. J. Fleet, S. S. Beauchemin, and T. Burkitt. Performance of optical flow techniques. In *Proceedings of the 1992 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 236–242. IEEE, 1992.

[17] S. S. Beauchemin and J. L. Barron. The computation of optical flow. *ACM Computing Surveys (CSUR)*, 27(3):433–466, 1995.

[18] A. M. Beek, H. P. Kühl, O. Bondarenko, J. W. Twisk, M. B. Hofman, W. G. van Dockum, C. A. Visser, and A. C. van Rossum. Delayed contrast-enhanced magnetic resonance imaging for the prediction of regional functional improvement after acute myocardial infarction. *Journal of the American College of Cardiology*, 42(5):895 – 901, 2003.

[19] J. Betts, O. College, P. Desaix, J. Johnson, E. Johnson, O. Korol, D. Kruse, B. Poe, J. Wise, M. Womble, et al. *Anatomy and Physiology*. Open Textbooks. OpenStax College, 2013.

[20] K. K. Bhatia, A. N. Price, W. Shi, J. V. Hajnal, and D. Rueckert. Super-resolution reconstruction of cardiac MRI using coupled dictionary learning. In *Proceedings of the 2014 IEEE 11th International Symposium on Biomedical Imaging (ISBI)*, pages 947–950. IEEE, 2014.

[21] R. Bonow, R. Gibbons, D. Berman, L. Johnson, J. Rumberger, M. Schwaiger, K. Taubert, F. Wackers, J. Ritchie, B. Brundage, L. Dreifus, C. Higgins, S. Nissen, H. Schelbert, J. Seward, B. Zaret, E. Garcia, S. Bacharach, and T. Boyce. Standardization of cardiac tomographic imaging. *Circulation*, 86(1):338–339, 1992.

[22] D. Boukerroui, J. A. Noble, and M. Brady. On the choice of band-pass quadrature filters. *Journal of Mathematical Imaging and Vision*, 21(1):53–80, 2004.

[23] C. Broit. *Optimal registration of deformed images*. PhD thesis, Computer and Information Science, University of Pennsylvania, 1981.

[24] J. J. Brown and C. B. Higgins. Myocardial paramagnetic contrast agents for MR imaging. *American Journal of Roentgenology*, 151(5):865 – 871, 1988.

[25] L. G. Brown. A survey of image registration techniques. *ACM Computing Surveys (CSUR)*, 24(4):325–376, 1992.

[26] A. Bruhn, J. Weickert, and C. Schnörr. Lucas/Kanade meets Horn/Schunck: Combining local and global optic flow methods. *International Journal of Computer Vision*, 61(3):211–231, 2005.

[27] E. Bruno and D. Pellerin. Robust motion estimation using spatial Gabor-like filters. *Signal Processing*, 82(2):297–309, 2002.

[28] T. Bulow and G. Sommer. Hypercomplex signals-a novel extension of the analytic signal to the multidimensional case. *IEEE Transactions on Signal Processing*, 49(11):2844–2852, 2001.

[29] M. Cabezas, A. Oliver, X. Lladó, J. Freixenet, and M. B. Cuadra. A review of atlas-based segmentation for magnetic resonance brain images. *Computer Methods and Programs in Biomedicine*, 104(3):e158–e177, 2011.

[30] M. C. Carminati, F. Maffessanti, and E. G. Caiani. Nearly automated motion artifacts correction between multi breath-hold short-axis and long-axis cine CMR images. *Computers in Biology and Medicine*, 46(1):42–50, 2014.

[31] F. Catté, P.-L. Lions, J.-M. Morel, and T. Coll. Image selective smoothing and edge detection by nonlinear diffusion. *SIAM Journal on Numerical Analysis*, 29(1):182–193, 1992.

[32] A. G. Chandler, R. J. Pinder, T. Netsch, J. A. Schnabel, D. J. Hawkes, D. L. Hill, and R. Razavi. Correction of misaligned slices in multi-slice cardiovascular magnetic resonance using slice-to-volume registration. *Journal of Cardiovascular Magnetic Resonance*, 10:13, 2008.

[33] R. Chandrashekara, R. Mohiaddin, and D. Rueckert. Cardiac motion tracking in tagged MR images using a 4D B-spline motion model and nonrigid image registration. In *Proceedings of the 2004 2nd IEEE International Symposium on Biomedical Imaging: Nano to Macro*, pages 468–471. IEEE, 2004.

[34] R. Chandrashekara, R. H. Mohiaddin, and D. Rueckert. Analysis of myocardial motion in tagged MR images using nonrigid image registration. In *Proceedings of SPIE Medical Imaging 2002: Image Processing*, volume 4684, pages 1168–1179. International Society for Optics and Photonics, 2002.

[35] R. Chandrashekara, R. H. Mohiaddin, and D. Rueckert. Analysis of 3-D myocardial motion in tagged MR images using nonrigid image registration. *IEEE Transactions on Medical Imaging*, 23(10):1245–1250, 2004.

[36] R. Chandrashekara, R. H. Mohiaddin, and D. Rueckert. Comparison of cardiac motion fields from tagged and untagged MR images using nonrigid registration. In A. F. Frangi, P. I. Radeva, A. Santos, and M. Hernandez, editors, *Functional Imaging and Modeling of the Heart, FIMH 2005*, volume 3504 of *Lecture Notes in Computer Science*, pages 425–433. Springer, Berlin, Heidelberg, 2005.

[37] Y. Chenoune, C. Constantinides, R. El Berbari, E. Roullot, F. Frouin, A. Herment, and E. Mousseaux. Rigid registration of delayed-enhancement and cine cardiac MR images using 3D normalized mutual information. In *Proceedings of 2010 Computing in Cardiology (CinC)*, pages 161–164. IEEE, 2010.

[38] A. Chiribiri, R. M. Botnar, and E. Nagel. Magnetic resonance coronary angiography: Where are we today? *Current Cardiology Reports*, 15(328):1–9, 2013.

[39] G. E. Christensen. *Deformable shape models for anatomy*. PhD thesis, Sever Institute of Technology, Washington University, 1994.

[40] C. Ciofolo, M. Fradkin, B. Mory, G. Hautvast, and M. Breeuwer. Automatic myocardium segmentation in late-enhancement MRI. In *Proceedings of the 2008 5th IEEE International Symposium on Biomedical Imaging: Nano to Macro*, pages 225–228. IEEE, 2008.

[41] A. Collignon, F. Maes, D. Delaere, D. Vandermeulen, P. Suetens, and G. Marchal. Automated multi-modality image registration based on information theory. *Information Processing in Medical Imaging*, 3(6):263–274, 1995.

[42] W. R. Crum, T. Hartkens, and D. Hill. Non-rigid image registration: Theory and practice. *The British Journal of Radiology*, 77:140–153, 2004.

[43] E. Dall'Armellina, N. Karia, A. C. Lindsay, T. D. Karamitsos, V. Ferreira, M. D. Robson, P. Kellman, J. M. Francis, C. Forfar, B. D. Prendergast, et al. Dynamic changes of edema and late gadolinium enhancement after acute myocardial infarction and their relationship to functional recovery and salvage index. *Circulation: Cardiovascular Imaging*, 4(3):228–236, 2011.

[44] E. Dall'Armellina, S. K. Piechnik, V. M. Ferreira, Q. L. Si, M. D. Robson, J. M. Francis, F. Cuculi, R. K. Kharbanda, A. P. Banning, R. P. Choudhury, T. D. Karamitsos, and S. Neubauer. Cardiovascular magnetic resonance by non contrast T1-mapping allows assessment of severity of injury in acute myocardial infarction. *Journal of Cardiovascular Magnetic Resonance*, 14:15, 2012.

[45] M. De Craene, G. Piella, O. Camara, N. Duchateau, E. Silva, A. Doltra, J. D'hooge, J. Brugada, M. Sitges, and A. F. Frangi. Temporal diffeomorphic free-form deformation: Application to motion and strain estimation from 3D echocardiography. *Medical image analysis*, 16(2):427–450, 2012.

[46] M. De Craene, C. Tobon-Gomez, C. Butakoff, N. Duchateau, G. Piella, K. S. Rhode, and A. F. Frangi. Temporal diffeomorphic free form deformation (TDFFD) applied to motion and deformation quantification of tagged MRI sequences. In *International Workshop on Statistical Atlases and Computational Models of the Heart*, pages 68–77. Springer, 2011.

[47] E. Dikici, T. O'Donnell, R. Setser, and R. White. Quantification of delayed enhancement MR images. In C. Barillot, D. R. Haynor, and P. Hellier, editors, *Medical Image Computing and Computer-Assisted Intervention, MICCAI 2004*, volume 3216 of *Lecture Notes in Computer Science*, pages 250–257. Springer, Berlin, Heidelberg, 2004.

[48] C. Dong, C. C. Loy, K. He, and X. Tang. Learning a deep convolutional network for image super-resolution. In *European Conference on Computer Vision*, pages 184–199. Springer, 2014.

[49] L. Dougherty, J. C. Asmuth, A. S. Blom, L. Axel, and R. Kumar. Validation of an optical flow method for tag displacement estimation. *IEEE Transactions on Medical Imaging*, 18(4):359–363, 1999.

[50] O. Dzyubachyk, Q. Tao, D. H. Poot, H. J. Lamb, K. Zeppenfeld, B. P. Lelieveldt, and R. J. van der Geest. Super-resolution reconstruction of late gadolinium-enhanced MRI for improved myocardial scar assessment. *Journal of Magnetic Resonance Imaging*, 42(1):160–167, 2015.

[51] R. L. Ehman and J. P. Felmlee. Adaptive technique for high-definition MR imaging of moving structures. *Radiology*, 173(1):255–263, 1989.

[52] R. L. Ehman, M. McNamara, M. Pallack, H. Hricak, and C. Higgins. Magnetic resonance imaging with respiratory gating: techniques and advantages. *American Journal of Roentgenology*, 143(6):1175–1182, 1984.

[53] M. Elad and A. Feuer. Restoration of a single superresolution image from several blurred, noisy, and undersampled measured images. *IEEE Transactions on Image Processing*, 6(12):1646–1658, 1997.

[54] 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 on Medical Imaging*, 29(3):868–878, 2010.

[55] J. Felblinger and C. Boesch. Amplitude demodulation of the electrocardiogram signal (ECG) for respiration monitoring and compensation during MR examinations. *Magnetic Resonance in Medicine*, 38(1):129–136, 1997.

[56] M. Felsberg. Optical flow estimation from monogenic phase. In B. Jähne, R. Mester, E. Barth, and H. Scharr, editors, *Complex Motion*, volume 3417 of *Lecture Notes in Computer Science*, pages 1–13. Springer, Berlin, Heidelberg, 2007.

[57] M. Felsberg and G. Sommer. The monogenic signal. *IEEE Transactions on Signal Processing*, 49(12):3136–3144, 2001.

[58] J. P. Finn, K. Nael, V. Deshpande, O. Ratib, and G. Laub. Cardiac MR imaging: State of the technology. *Radiology*, 241(2):338 – 354, 2006.

[59] J. M. Fitzpatrick, J. B. West, and C. R. Maurer. Predicting error in rigid-body point-based registration. *IEEE Transactions on Medical Imaging*, 17(5):694–702, 1998.

[60] D. Fleet and Y. Weiss. Optical flow estimation. In N. Paragios, Y. Chen, and O. Faugeras, editors, *Handbook of Mathematical Models in Computer Vision*, pages 237–257. Springer, Boston, MA, 2006.

[61] D. J. Fleet and A. D. Jepson. Computation of component image velocity from local phase information. *International Journal of Computer Vision*, 5(1):77–104, 1990.

[62] D. Fortun, P. Bouthemy, and C. Kervrann. Optical flow modeling and computation: a survey. *Computer Vision and Image Understanding*, 134:1–21, 2015.

[63] D. Gabor. Theory of communication. Part 1: The analysis of information. *Journal of the Institution of Electrical Engineers - Part III: Radio and Communication Engineering*, 93(26):429–441, 1946.

[64] L. L. Geyer, U. J. Schoepf, F. G. Meinel, J. W. Nance Jr, G. Bastarrika, J. A. Leipsic, N. S. Paul, M. Rengo, A. Laghi, and C. N. De Cecco. State of the art: iterative CT reconstruction techniques. *Radiology*, 276(2):339–357, 2015.

[65] A. Giannakidis, O. Oktay, J. Keegan, V. Spadotto, I. Voges, G. Smith, I. Pierce, W. Bai, D. Rueckert, S. Ernst, et al. Super-resolution Reconstruction of Late Gadolinium Enhancement Cardiovascular Magnetic Resonance Images using a Residual Convolutional Neural Network. In *Proceedings of the 25th Annual Meeting of ISMRM, Honolulu, HI, USA*, 2017.

[66] D. T. Ginat, M. W. Fong, D. J. Tuttle, S. K. Hobbs, and R. C. Vyas. Cardiac imaging: Part 1, MR pulse sequences, imaging planes, and basic anatomy. *American Journal of Roentgenology*, 197(4):808–815, 2011.

[67] G. Glover and N. Pelc. A rapid-gated cine MRI technique. *Magnetic Resonance Annual*, 21(2):299–334, 1988.

[68] A. A. Goshtasby and D. A. Turner. Fusion of short-axis and long-axis cardiac MR images. *Computerized Medical Imaging and Graphics*, 20(2):77–87, 1996.

[69] G. H. Granlund and H. Knutsson. *Signal processing for computer vision*. Springer US, Boston, MA, 1995.

[70] V. Grau, H. Becher, and J. A. Noble. Registration of multiview real-time 3-D echocardiographic sequences. *IEEE Transactions on Medical Imaging*, 26(9):1154–1165, 2007.

[71] H. Greenspan. Super-resolution in medical imaging. *The Computer Journal*, 52(1):43–63, 2009.

[72] J. Greupner, E. Zimmermann, A. Grohmann, H.-P. Dübel, T. Althoff, A. C. Borges, W. Rutsch, P. Schlattmann, B. Hamm, and M. Dewey. Head-to-head comparison of left ventricular function assessment with 64-row computed tomography, biplane left cineventriculography, and both 2-and 3-dimensional transthoracic echocardiography: comparison with magnetic resonance imaging as the reference standard. *Journal of the American College of Cardiology*, 59(21):1897–1907, 2012.

[73] G. J. Grevera and J. K. Udupa. Shape-based interpolation of multidimensional grey-level images. *IEEE Transactions on Medical Imaging*, 15(6):881–892, 1996.

[74] F. Grothues, G. C. Smith, J. C. Moon, N. G. Bellenger, P. Collins, H. U. Klein, and D. J. Pennell. Comparison of interstudy reproducibility of cardiovascular magnetic resonance with two-dimensional echocardiography in normal subjects and in patients with heart failure or left ventricular hypertrophy. *The American Journal of Cardiology*, 90(1):29–34, 2002.

[75] S. N. Gupta and J. L. Prince. On variable brightness optical flow for tagged MRI. In Y. Bizais, C. Barillot, and R. Di Paola, editors, *Information Processing in Medical Imaging*, volume 3 of *Computational Imaging and Vision*, pages 323–334. Springer Netherlands, 1995.

[76] S. L. Hahn. Multidimensional complex signals with single-orthant spectra. *Proceedings of the IEEE*, 80(8):1287–1300, 1992.

[77] J. V. Hajnal and D. L. G. Hill. *Medical image registration*. CRC Press, 2001.

[78] T. Hartkens, D. L. Hill, A. D. Castellano-Smith, D. J. Hawkes, C. Maurer, A. J. Martin, W. A. Hall, H. Liu, and C. L. Truwit. Measurement and analysis of brain deformation during neurosurgery. *IEEE Transactions on Medical Imaging*, 22(1):82–92, 2003.

[79] M. Hemmendorff, M. T. Andersson, T. Kronander, and H. Knutsson. Phase-based multidimensional volume registration. *IEEE Transactions on Medical Imaging*, 21(12):1536–1543, 2002.

[80] A. Hennemuth, A. Seeger, O. Friman, S. Miller, B. Klumpp, S. Oeltze, and H.-O. Peitgen. A comprehensive approach to the analysis of contrast enhanced cardiac MR images. *IEEE Transactions on Medical Imaging*, 27(11):1592–1610, 2008.

[81] D. L. Hill, P. G. Batchelor, M. Holden, and D. J. Hawkes. Medical image registration. *Physics in Medicine and Biology*, 46(3):R1, 2001.

[82] B. K. Horn and B. G. Schunck. Determining optical flow. *Artificial Intelligence*, 17(1-3):185–203, 1981.

[83] E.-S. H. Ibrahim. Myocardial tagging by cardiovascular magnetic resonance: Evolution of techniques - pulse sequences, analysis algorithms, and applications. *Journal of Cardiovascular Magnetic Resonance*, 13:36, 2011.

[84] T. D. Karamitsos, E. Dall'Armellina, R. P. Choudhury, and S. Neubauer. Ischemic heart disease: Comprehensive evaluation by cardiovascular magnetic resonance. *American Heart Journal*, 162(1):16–30, 2011.

[85] M. Kazhdan, M. Bolitho, and H. Hoppe. Poisson surface reconstruction. In *Proceedings of the 4th Eurographics Symposium on Geometry Processing*, pages 61–70, 2006.

[86] P. Kellman and A. E. Arai. Cardiac imaging techniques for physicians: late enhancement. *Journal of magnetic resonance imaging*, 36(3):529–542, 2012.

[87] C. Kelly, V. Grau, S. Neubauer, R. Choudhury, and E. Dall'Armellina. Reproducibility of short axis slice locations in longitudinal cardiovascular magnetic resonance (CMR) studies. *Journal of Cardiovascular Magnetic Resonance*, 16(1):228, 2014.

[88] C. Kelly, S. Neubauer, R. Choudhury, E. Dall'Armellina, and V. Grau. A local phase-based algorithm for registration of CMR scans from multiple visits. In *Proceedings of 2014 Computing in Cardiology (CinC)*, pages 937–940. IEEE, 2014.

[89] R. J. Kim, D. S. Fieno, T. B. Parrish, K. Harris, E.-L. Chen, O. Simonetti, J. Bundy, J. P. Finn, F. J. Klocke, and R. M. Judd. Relationship of MRI delayed contrast enhancement to irreversible injury, infarct age, and contractile function. *Circulation*, 100(19):1992–2002, 1999.

[90] S. Kleijn and O. Kamp. Clinical application of three-dimensional echocardiography: past, present and future. *Netherlands Heart Journal*, 17(1):18–24, 2009.

[91] S. Klein, J. P. Pluim, M. Staring, and M. A. Viergever. Adaptive stochastic gradient descent optimisation for image registration. *International Journal of Computer Vision*, 81(3):227, 2009.

[92] S. Klein, M. Staring, K. Murphy, M. A. Viergever, and J. P. Pluim. elastix: A toolbox for intensity-based medical image registration. *IEEE Transactions on Medical Imaging*, 29(1):196–205, 2010.

[93] H. Knutsson and C.-F. Westin. Monomial phase: A matrix representation of local phase. In C.-F. Westin, A. Vilanova, and B. Burgeth, editors, *Visualization and Processing of Tensors and Higher Order Descriptors for Multi-Valued Data*, Mathematics and Visualization, pages 37–73. Springer, Berlin, Heidelberg, 2014.

[94] J. Knuuti, F. Bengel, J. J. Bax, P. A. Kaufmann, D. Le Guludec, P. Perrone Filardi, C. Marcassa, N. Ajmone Marsan, S. Achenbach, A. Kitsiou, et al. Risks and benefits of cardiac imaging: an analysis of risks related to imaging for coronary artery disease. *European heart journal*, 35(10):633–638, 2013.

[95] M. A. Konstam, D. G. Kramer, A. R. Patel, M. S. Maron, and J. E. Udelson. Left ventricular remodeling in heart failure. *JACC: Cardiovascular Imaging*, 4(1):98–108, 2011.

[96] P. Kovesi. Symmetry and asymmetry from local phase. In *Proceedings of the 10th Australian Joint Conference on Artificial Intelligence*, pages 185–190, 1997.

[97] P. Kovesi. Image features from phase congruency. *Videre: Journal of Computer Vision Research*, 1(3), 1999.

[98] D.-J. Kroon and C. H. Slump. MRI modalitiy transformation in demon registration. In *Biomedical Imaging: From Nano to Macro, 2009. ISBI'09. IEEE International Symposium on*, pages 963–966. IEEE, 2009.

[99] J. C. Lagarias, J. A. Reeds, M. H. Wright, and P. E. Wright. Convergence properties of the Nelder–Mead simplex method in low dimensions. *SIAM Journal on Optimization*, 9(1):112–147, 1998.

[100] P. Lamata, R. Casero, V. Carapella, S. A. Niederer, M. J. Bishop, J. E. Schneider, P. Kohl, and V. Grau. Images as drivers of progress in cardiac computational modelling. *Progress in biophysics and molecular biology*, 115(2):198–212, 2014.

[101] P. Lanzer, E. Botvinick, N. Schiller, L. Crooks, M. Arakawa, L. Kaufman, P. Davis, R. Herfkens, M. Lipton, and C. B. Higgins. Cardiac imaging using gated magnetic resonance. *Radiology*, 150(1):121–127, 1984.

[102] L. Le Folgoc, H. Delingette, A. Criminisi, and N. Ayache. Sparse Bayesian registration of medical images for self-tuning of parameters and spatially adaptive parametrization of displacements. *Medical image analysis*, 36:79–97, 2017.

[103] N. Lee and T. Hyeon. Designed synthesis of uniformly sized iron oxide nanoparticles for efficient magnetic resonance imaging contrast agents. *Chemical Society Reviews*, 41(7):2575–2589, 2012.

[104] G. W. Lenz, E. M. Haacke, and R. D. White. Retrospective cardiac gating: A review of technical aspects and future directions. *Magnetic Resonance Imaging*, 7(5):445–455, 1989.

[105] H. Lester and S. R. Arridge. A survey of hierarchical non-linear medical image registration. *Pattern Recognition*, 32(1):129–149, 1999.

[106] J. R. Levick. *An introduction to cardiovascular physiology*. Butterworth-Heinemann, 2013.

[107] L. S. Lilly. *Pathophysiology of heart disease: A collaborative project of medical students and faculty*. Lippincott Williams & Wilkins, 2012.

[108] J. A. Lima and M. Y. Desai. Cardiovascular magnetic resonance imaging: Current and emerging applications. *Journal of the American College of Cardiology*, 44(6):1164 – 1171, 2004.

[109] T. Lindeberg and J. Gårding. Shape-adapted smoothing in estimation of 3-d depth cues from affine distortions of local 2-d brightness structure. *Image and Vision computing*, 15(6):415–434, 1997.

[110] S. Loncaric and T. Macan. Point-constrained optical flow for LV motion detection. In *Proceedings of SPIE Medical Imaging 2000: Physiology and Function from Multidimensional Images*, volume 3978, pages 521–529. International Society for Optics and Photonics, 2000.

[111] C. Loop. Smooth subdivision surfaces based on triangles. Master's thesis, University of Utah, Department of Mathematics, 1987.

[112] M. Lorenzo-Valdés, G. I. Sanchez-Ortiz, R. Mohiaddin, and D. Rueckert. Atlas-based segmentation and tracking of 3D cardiac MR images using non-rigid registration. In T. Dohi and R. Kikinis, editors, *Medical Image Computing and Computer-Assisted Intervention, MICCAI 2002*, volume 2488 of *Lecture Notes in Computer Science*, pages 642–650. Springer, Berlin, Heidelberg, 2002.

[113] J. Lötjönen, M. Pollari, S. Kivistö, and K. Lauerma. Correction of motion artifacts from cardiac cine magnetic resonance images. *Academic Radiology*, 12(10):1273–1284, 2005.

[114] J. M. Lötjönen, V. M. Järvinen, B. Cheong, E. Wu, S. Kivistö, J. R. Koikkalainen, J. J. Mattila, H. M. Kervinen, R. Muthupillai, F. H. Sheehan, and K. Lauerma. Evaluation of cardiac biventricular segmentation from multiaxis MRI data: A multicenter study. *Journal of Magnetic Resonance Imaging*, 28(3):626–636, 2008.

[115] Y. Lu, G. Wright, and P. E. Radau. Automatic myocardium segmentation of LGE MRI by deformable models with prior shape data. *Journal of Cardiovascular Magnetic Resonance*, 15(1):14, 2013.

[116] Y. Lu, Y. Yang, K. A. Connelly, G. A. Wright, and P. E. Radau. Automated quantification of myocardial infarction using graph cuts on contrast delayed enhanced magnetic resonance images. *Quantitative Imaging in Medicine and Surgery*, 2(2):81–86, 2012.

[117] B. D. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *Proceedings of the 7th International Joint Conference on Artificial intelligence (IJCAI)*, Vancouver, BC, Canada, 1981.

[118] G. K. Lund, A. Stork, K. Muellerleile, A. A. Barmeyer, M. P. Bansmann, M. Knefel, U. Schlichting, M. Müller, P. E. Verde, G. Adam, T. Meinertz, and M. Saeed. Prediction of left ventricular remodeling and analysis of infarct resorption in patients with reperfused myocardial infarcts by using contrast-enhanced MR imaging. *Radiology*, 245(1):95–102, 2007.

[119] T. Macan and S. Loncaric. 3D cardiac motion estimation by point-constrained optical flow algorithm. In *Proceedings of the 2001 2nd International Symposium on Image and Signal Processing and Analysis (ISPA)*, pages 255–259. IEEE, 2001.

[120] J. A. Maintz and M. A. Viergever. A survey of medical image registration. *Medical Image Analysis*, 2(1):1–36, 1998.

[121] T. Makela, P. Clarysse, O. Sipila, N. Pauna, Q. C. Pham, T. Katila, and I. E. Magnin. A review of cardiac image registration methods. *IEEE Transactions on Medical Imaging*, 21(9):1011–1021, 2002.

[122] S. B. Malik, N. Chen, R. A. Parker III, and J. Y. Hsu. Transthoracic echocardiography: pitfalls and limitations as delineated at cardiac CT and MR imaging. *Radiographics*, 37(2):383–406, 2017.

[123] A. Mang, W. R. Crum, O. Camara, J. A. Schnabel, G. B. Caseiras, T. M. Buzug, H. R. Jäger, T. A. Yousry, and D. J. Hawkes. Modelling tumour growth patterns with non-rigid image registration. In T. M. Buzug, D. Holz, J. Bongartz, M. Kohl-Bareis, U. Hartmann, and S. Weber, editors, *Advances in Medical Engineering*, volume 114 of *Springer Proceedings in Physics*, pages 139–144. Springer, Berlin, Heidelberg, 2007.

[124] T. Mansi, J.-M. Peyrat, M. Sermesant, H. Delingette, J. Blanc, Y. Boudjemline, and N. Ayache. Physically-constrained diffeomorphic demons for the estimation of 3D myocardium strain from cine-MRI. In N. Ayache, H. Delingette, and M. Sermesant, editors, *Functional Imaging and Modeling of the Heart, FIMH 2009*, volume 5528 of *Lecture Notes in Computer Science*, pages 201–210. Springer, Berlin, Heidelberg, 2009.

[125] S. Marchesseau, N. Duchateau, and H. Delingette. Segmentation and registration coupling from short-axis Cine MRI: application to infarct diagnosis. In *International Workshop on Statistical Atlases and Computational Models of the Heart*, pages 48–56. Springer, 2016.

[126] N. B. Mascarenhas, R. Muthupillai, B. Cheong, M. Pereyra, and S. D. Flamm. Fast 3D cine steady-state free precession imaging with sensitivity encoding for assessment of left ventricular function in a single breath-hold. *American Journal of Roentgenology*, 187(5):1235 – 1239, 2006.

[127] J. R. McClelland, D. J. Hawkes, T. Schaeffter, and A. P. King. Respiratory motion models: A review. *Medical Image Analysis*, 17(1):19–42, 2013.

[128] K. McLeish, D. L. 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.

[129] K. McLeod, A. Prakosa, T. Mansi, M. Sermesant, and X. Pennec. An incompressible log-domain demons algorithm for tracking heart tissue. In O. Camara, E. Konukoglu, M. Pop, K. Rhode, M. Sermesant, and A. Young, editors, *Statistical Atlases and Computational Models of the Heart. Imaging and Modelling Challenges, STACOM 2011*, volume 7085 of *Lecture Notes in Computer Science*, pages 55–67. Springer, Berlin, Heidelberg, 2012.

[130] K. McLeod, C. Seiler, M. Sermesant, and X. Pennec. A near-incompressible poly-affine motion model for cardiac function analysis. In *International Workshop on Statistical Atlases and Computational Models of the Heart*, pages 288–297. Springer, 2012.

[131] K. Mcleod, C. Seiler, M. Sermesant, and X. Pennec. Spatio-temporal dimension reduction of cardiac motion for group-wise analysis and statistical testing. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 501–508. Springer, 2013.

[132] K. Mcleod, C. Seiler, N. Toussaint, M. Sermesant, and X. Pennec. Regional analysis of left ventricle function using a cardiac-specific polyaffine motion model. In *International Conference on Functional Imaging and Modeling of the Heart*, pages 483–490. Springer, 2013.

[133] K. McLeod, M. Sermesant, P. Beerbaum, and X. Pennec. Spatio-temporal tensor decomposition of a polyaffine motion model for a better analysis of pathological left ventricular dynamics. *IEEE Transactions on Medical Imaging*, 34(7):1562–1575, 2015.

[134] P. Medrano-Gracia, B. R. Cowan, B. Ambale-Venkatesh, D. A. Bluemke, J. Eng, J. P. Finn, C. G. Fonseca, J. A. 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*, 16:56, 2014.

[135] M. Mellor and M. Brady. Non-rigid multimodal image registration using local phase. In C. Barillot, D. R. Haynor, and P. Hellier, editors, *Medical Image Computing and Computer-Assisted Intervention, MICCAI 2004*, volume 3216 of *Lecture Notes in Computer Science*, pages 789–796. Springer, Berlin, Heidelberg, 2004.

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

[137] C. Metz, S. Klein, M. Schaap, T. van Walsum, and W. J. Niessen. Nonrigid registration of dynamic medical imaging data using nD+ t B-splines and a groupwise optimization approach. *Medical image analysis*, 15(2):238–249, 2011.

[138] D. D. Miller, G. Holmvang, J. B. Gill, D. Dragotakes, H. L. Kantor, R. D. Okada, and T. J. Brady. MRI detection of myocardial perfusion changes by gadolinium-DTPA infusion during dipyridamole hyperemia. *Magnetic Resonance in Medicine*, 10(2):246 – 255, 1989.

[139] W. Moody, N. Edwards, C. Chue, R. Taylor, C. Ferro, J. Townend, and R. Steeds. Variability in cardiac MR measurement of left ventricular ejection fraction, volumes and mass in healthy adults: Defining a significant change at 1 year. *The British Journal of Radiology*, 88(1049), 2015.

[140] J. C. Moon, C. H. Lorenz, J. M. Francis, G. C. Smith, and D. J. Pennell. Breath-hold FLASH and FISP cardiovascular MR imaging: Left ventricular volume differences and reproducibility. *Radiology*, 223(3):789–797, 2002.

[141] M. C. Morrone and R. A. Owens. Feature detection from local energy. *Pattern Recognition Letters*, 6(5):303–313, 1987.

[142] S. G. Myerson, J. Francis, and S. Neubauer. *Cardiovascular magnetic resonance*. Oxford University Press, 2010.

[143] N. M. I. Noble, D. L. G. Hill, M. Breeuwer, J. A. Schnabel, D. J. Hawkes, F. A. Gerritsen, and R. Razavi. Myocardial delineation via registration in a polar coordinate system. *Academic Radiology*, 10(12):1349–1358, 2003.

[144] F. Odille, A. Bustin, B. Chen, P.-A. Vuissoz, and J. Felblinger. 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 of *Lecture Notes in Computer Science*, pages 435–442. Springer, Cham, 2015.

[145] O. Oktay, W. Bai, M. Lee, R. Guerrero, K. Kamnitsas, J. Caballero, A. de Marvao, S. Cook, D. O'Regan, and D. Rueckert. Multi-input cardiac image super-resolution using convolutional neural networks. In S. Ourselin, L. Joskowicz, M. R. Sabuncu, G. Unal, and W. Wells, editors, *Medical Image Computing and Computer-Assisted Intervention, MICCAI 2016*, volume 9902 of *Lecture Notes in Computer Science*, pages 246–254. Springer, Cham, 2016.

[146] O. Oktay, E. Ferrante, K. Kamnitsas, M. Heinrich, W. Bai, J. Caballero, S. A. Cook, A. de Marvao, T. Dawes, D. P. O'Regan, et al. Anatomically constrained neural networks (ACNNs): application to cardiac image enhancement and segmentation. *IEEE Transactions on Medical Imaging*, 37(2):384–395, 2018.

[147] H. Ólafsdóttir, H. Pedersen, M. S. Hansen, M. Lyksborg, M. F. Hansen, S. Darkner, and R. Larsen. Registration-based interpolation applied to cardiac MRI. In *Proceedings of SPIE Medical Imaging 2010: Image Processing*, volume 7623. International Society for Optics and Photonics, 2010.

[148] A. V. Oppenheim and J. S. Lim. The importance of phase in signals. *Proceedings of the IEEE*, 69(5):529–541, 1981.

[149] Y. Ou, D. H. Ye, K. M. Pohl, and C. Davatzikos. Validation of DRAMMS among 12 popular methods in cross-subject cardiac MRI registration. In *International Workshop on Biomedical Image Registration*, pages 209–219. Springer, 2012.

[150] M. R. Paling and J. R. Brookeman. Respiration artifacts in MR imaging: reduction by breath holding. *Journal of Computer Assisted Tomography*, 10(6):1080–1082, 1986.

[151] G. P. Penney, J. A. Schnabel, D. Rueckert, M. A. Viergever, and W. J. Niessen. Registration-based interpolation. *IEEE Transactions on Medical Imaging*, 23(7):922–926, 2004.

[152] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–639, 1990.

[153] D. Perperidis, M. Lorenzo-Valdes, R. Chandrashekara, A. Rao, R. Mohiaddin, G. Sanchez-Ortiz, and D. Rueckert. Building a 4D atlas of the cardiac anatomy and motion using MR imaging. In *Proceedings of the 2004 2nd IEEE International Symposium on Biomedical Imaging: Nano to Macro*, pages 412–415. IEEE, 2004.

[154] D. Perperidis, R. H. Mohiaddin, and D. Rueckert. Spatio-temporal free-form registration of cardiac MR image sequences. *Medical Image Analysis*, 9(5):441–456, 2005.

[155] D. Perperidis, A. Rao, M. Lorenzo-Valdés, R. Mohiaddin, and D. Rueckert. Spatio-temporal alignment of 4D cardiac MR images. In I. E. Magnin, J. Montagnat, P. Clarysse, J. Nenonen, and T. Katila, editors, *Functional Imaging and Modeling of the Heart, FIMH 2003*, volume 2674 of *Lecture Notes in Computer Science*, pages 205–214. Springer, Berlin, Heidelberg, 2003.

[156] D. Perperidis, A. Rao, R. Mohiaddin, and D. Rueckert. Non-rigid spatio-temporal alignment of 4D cardiac MR images. In J. C. Gee, J. B. A. Maintz, and M. W. Vannier, editors, *Biomedical Image Registration, WBIR 2003*, volume 2717 of *Lecture Notes in Computer Science*, pages 191–200. Springer, Berlin, Heidelberg, 2003.

[157] C. Petitjean and J.-N. Dacher. A review of segmentation methods in short axis cardiac MR images. *Medical Image Analysis*, 15(2):169–184, 2011.

[158] R. I. Pettigrew, J. N. Oshinski, G. P. Chatzimavroudis, and W. T. Dixon. MRI techniques for cardiovascular imaging. *Journal of Magnetic Resonance Imaging*, 10(5):590–601, 1999.

[159] J. P. Pluim, J. A. Maintz, and M. A. Viergever. Mutual-information-based registration of medical images: a survey. *IEEE Transactions on Medical Imaging*, 22(8):986–1004, 2003.

[160] M. Poon, V. Fuster, and Z. Fayad. Cardiac magnetic resonance imaging: a one-stop-shop evaluation of myocardial dysfunction. *Current Opinion in Cardiology*, 17(6):663–670, 2002.

[161] D. H. Poot, V. Van Meir, and J. Sijbers. General and efficient super-resolution method for multi-slice MRI. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 615–622. Springer, 2010.

[162] I. A. Popescu, B. Irving, A. Borlotti, E. Dall'Armellina, and V. Grau. Myocardial Scar Quantification Using SLIC Supervoxels-Parcellation Based on Tissue Characteristic Strains. In *International Workshop on Statistical Atlases and Computational Models of the Heart*, pages 182–190. Springer, 2016.

[163] J. L. Prince and E. R. McVeigh. Motion estimation from tagged MR image sequences. *IEEE Transactions on Medical Imaging*, 11(2):238–249, 1992.

[164] E. Puyol-Antón, M. Sinclair, B. Gerber, M. S. Amzulescu, H. Langet, M. De Craene, P. Aljabar, P. Piro, and A. P. King. A multimodal spatiotemporal cardiac motion atlas from MR and ultrasound data. *Medical image analysis*, 40:96–110, 2017.

[165] S. Queirós, D. Barbosa, B. Heyde, P. Morais, J. L. Vilaça, D. Friboulet, O. Bernard, and J. D'hooge. Fast automatic myocardial segmentation in 4D cine CMR datasets. *Medical image analysis*, 18(7):1115–1131, 2014.

[166] S. Queirós, J. L. Vilaça, P. Morais, J. C. Fonseca, J. D'hooge, and D. Barbosa. Fast left ventricle tracking using localized anatomical affine optical flow. *International Journal for Numerical Methods in Biomedical Engineering*, 33(11), 2017.

[167] S. U. Rahman and S. Wesarg. Combining short-axis and long-axis cardiac MR images by applying a super-resolution reconstruction algorithm. In *Proceedings of SPIE Medical Imaging 2010: Image Processing*, volume 7623. International Society for Optics and Photonics, 2010.

[168] A. Rao, G. I. Sanchez-Ortiz, R. Chandrashekara, M. Lorenzo-Valdés, R. Mohiaddin, and D. Rueckert. Comparison of cardiac motion across subjects using non-rigid registration. In T. Dohi and R. Kikinis, editors, *Medical Image Computing and Computer-Assisted Intervention, MICCAI 2002*, volume 2488 of *Lecture Notes in Computer Science*, pages 722–729. Springer, Berlin, Heidelberg, 2002.

[169] A. Rao, G. I. Sanchez-Ortiz, R. Chandrashekara, M. Lorenzo-Valdés, R. Mohiaddin, and D. Rueckert. Construction of a cardiac motion atlas from MR using non-rigid registration. In I. E. Magnin, J. Montagnat, P. Clarysse, J. Nenonen, and T. Katila, editors, *Functional Imaging and Modeling of the Heart, FIMH 2003*, volume 2674 of *Lecture Notes in Computer Science*, pages 141–150. Springer, Berlin, Heidelberg, 2003.

[170] S. P. Raya and J. K. Udupa. Shape-based interpolation of multidimensional objects. *IEEE Transactions on Medical Imaging*, 9(1):32–42, 1990.

[171] D. Rey, G. Subsol, H. Delingette, and N. Ayache. Automatic detection and segmentation of evolving processes in 3D medical images: Application to multiple sclerosis. *Medical Image Analysis*, 6(2):163–179, 2002.

[172] J. P. Ridgway. Cardiovascular magnetic resonance physics for clinicians: part I. *Journal of Cardiovascular Magnetic Resonance*, 12:71, 2010.

[173] T. Rohlfing, R. Brandt, R. Menzel, D. B. Russakoff, and C. R. Maurer. Quo vadis, atlas-based segmentation? In J. S. Suri, D. L. Wilson, and S. Laxminarayan, editors, *Handbook of Biomedical Image Analysis: Volume III: Registration Models*, Topics in Biomedical Engineering International Book Series, pages 435–486. Springer, Boston, MA, 2005.

[174] D. Rueckert, L. I. Sonoda, C. Hayes, D. L. Hill, M. O. Leach, and D. J. Hawkes. Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Transactions on Medical Imaging*, 18(8):712–721, 1999.

[175] P. Scarborough, K. Wickramasinghe, P. Bhatnagar, and M. Rayner. Trends in coronary heart disease 1961–2011. British Heart Foundation, London, 2011.

[176] E. B. Schelbert, L.-Y. Hsu, S. A. Anderson, B. D. Mohanty, S. M. Karim, P. Kellman, A. H. Aletras, and A. E. Arai. Late Gadolinium-Enhancement Cardiac Magnetic Resonance Identifies Postinfarction Myocardial Fibrosis and the Border Zone at the Near Cellular Level in Ex Vivo Rat HeartClinical Perspective. *Circulation: Cardiovascular Imaging*, 3(6):743–752, 2010.

[177] J. A. Schnabel, D. Rueckert, M. Quist, J. M. Blackall, A. D. Castellano-Smith, T. Hartkens, G. P. Penney, W. A. Hall, H. Liu, C. L. Truwit, F. A. Gerritsen, D. L. G. Hill, and D. J. Hawkes. A generic framework for non-rigid registration based on non-uniform multi-level free-form deformations. In W. J. Niessen and M. A. Viergever, editors, *Medical Image Computing and Computer-Assisted Intervention, MICCAI 2001*, volume 2208 of *Lecture Notes in Computer Science*, pages 573–581. Springer, Berlin, Heidelberg, 2001.

[178] C. L. Schultz, R. Alfidi, A. D. Nelson, S. Y. Kopiwoda, and M. E. Clampitt. The effect of motion on two-dimensional Fourier transformation magnetic resonance images. *Radiology*, 152(1):117–121, 1984.

[179] B. N. Shah. Echocardiography in the era of multimodality cardiovascular imaging. *BioMed research international*, 2013, 2013.

[180] D. P. Shamonin, E. E. Bron, B. P. Lelieveldt, M. Smits, S. Klein, and M. Staring. Fast parallel image registration on CPU and GPU for diagnostic classification of Alzheimer's disease. *Frontiers in Neuroinformatics*, 7:50, 2014.

[181] W. Shi, J. Caballero, C. Ledig, X. Zhuang, W. Bai, K. Bhatia, A. M. S. M. de Marvao, T. Dawes, D. O'Regan, and D. Rueckert. Cardiac image super-resolution with global correspondence using multi-atlas patchmatch. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 9–16. Springer, 2013.

[182] W. Shi, M. Jantsch, P. Aljabar, L. Pizarro, W. Bai, H. Wang, D. O'Regan, X. Zhuang, and D. Rueckert. Temporal sparse free-form deformations. *Medical image analysis*, 17(7):779–789, 2013.

[183] W. Shi, X. Zhuang, H. Wang, S. Duckett, D. V. 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.

[184] W. Shi, X. Zhuang, H. Wang, S. Duckett, D. Oregan, P. Edwards, S. Ourselin, and D. Rueckert. Automatic segmentation of different pathologies from cardiac cine MRI using registration and multiple component EM estimation. In D. N. Metaxas and L. Axel, editors, *Functional Imaging and Modeling of the Heart, FIMH 2011*, volume 6666 of *Lecture Notes in Computer Science*, pages 163–170. Springer, Berlin, Heidelberg, 2011.

[185] 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 *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 332–340. Springer, 2017.

[186] P. J. Slomka, D. Fieno, A. Ramesh, V. Goyal, H. Nishina, L. E. 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.

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

[188] M. B. Srichai, R. P. Lim, S. Wong, and V. S. Lee. Cardiovascular applications of phase-contrast MRI. *American Journal of Roentgenology*, 192(3):662 – 675, 2009.

[189] H. Stark. An extension of the Hilbert transform product theorem. *Proceedings of the IEEE*, 59(9):1359–1360, 1971.

[190] C. Studholme, D. L. Hill, and D. J. Hawkes. Automated 3-D registration of MR and CT images of the head. *Medical Image Analysis*, 1(2):163–175, 1996.

[191] C. Studholme, D. L. Hill, and D. J. Hawkes. An overlap invariant entropy measure of 3D medical image alignment. *Pattern Recognition*, 32(1):71–86, 1999.

[192] M. G. S. J. Sutton and N. Sharpe. Left ventricular remodeling after myocardial infarction. *Circulation*, 101(25):2981–2988, 2000.

[193] 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.

[194] E. Tadmor and P. Athavale. Multiscale image representation using novel integro-differential equations. *Inverse Problems and Imaging*, 3(4):693–710, 2009.

[195] 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), 2013.

[196] Q. Tao, S. R. Piers, H. J. Lamb, and R. J. van der Geest. Automated left ventricle segmentation in late gadolinium-enhanced MRI for objective myocardial scar assessment. *Journal of Magnetic Resonance Imaging*, 42(2):390–399, 2015.

[197] L. Tautz, A. Hennemuth, M. Andersson, A. Seeger, H. Knutsson, and O. Friman. Phase-based non-rigid registration of myocardial perfusion MRI image sequences. In *Proceedings of the 2010 IEEE International Symposium on Biomedical Imaging: Nano to Macro*, pages 516–519. IEEE, 2010.

[198] L. Tautz, A. Hennemuth, and H.-O. Peitgen. Quadrature filter based motion analysis for 3D ultrasound sequences. In O. Camara, T. Mansi, M. Pop, K. Rhode, M. Sermesant, and A. Young, editors, *Statistical Atlases and Computational Models of the Heart. Imaging and Modelling Challenges, STACOM 2012*, volume 7746 of *Lecture Notes in Computer Science*, pages 169–177. Springer, Berlin, Heidelberg, 2013.

[199] P. M. Thompson, M. S. Mega, R. P. Woods, C. I. Zoumalan, C. J. Lindshield, R. E. Blanton, J. Moussai, C. J. Holmes, J. L. Cummings, and A. W. Toga. Cortical change in Alzheimer's disease detected with a disease-specific population-based brain atlas. *Cerebral Cortex*, 11(1):1–16, 2001.

[200] C. Tobon-Gomez, M. De Craene, K. Mcleod, L. Tautz, W. Shi, A. Hennemuth, A. Prakosa, H. Wang, G. Carr-White, S. Kapetanakis, et al. Benchmarking framework for myocardial tracking and deformation algorithms: An open access database. *Medical image analysis*, 17(6):632–648, 2013.

[201] N. Townsend, P. Bhatnagar, E. Wilkins, K. Wickramasinghe, and M. Rayner. Cardiovascular disease statistics 2015. British Heart Foundation, London, 2015.

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

[203] P. R. M. van Dijkman, J. Doornbos, A. de Roos, A. van der Laarse, S. Postema, N. A. A. Matheijssen, A. V. G. Bruschke, A. E. v. Voorthuisen, V. M. Cats, and E. E. van der Wall. Improved detection of acute myocardial infarction by magnetic resonance imaging using gadolinium-DTPA. *The International Journal of Cardiac Imaging*, 5(1):1–8, 1989.

[204] E. Van Reeth, I. W. Tham, C. H. Tan, and C. L. Poh. Super-resolution in magnetic resonance imaging: A review. *Concepts in Magnetic Resonance Part A*, 40(6):306–325, 2012.

[205] B. Villard, E. Zacur, E. Dall'Armellina, and V. Grau. Correction of slice misalignment in multi-breath-hold cardiac MRI scans. In *International Workshop on Statistical Atlases and Computational Models of the Heart*, pages 30–38. Springer, 2016.

[206] P. Viola and W. M. Wells, III. Alignment by maximization of mutual information. In *Proceedings of the 1995 5th International Conference on Computer Vision*, pages 16–23. IEEE, 1995.

[207] Visual Computing Lab ISTI - CNR. Meshlab. http://meshlab.sourceforge.net/.

[208] U. Vovk, F. Pernus, and B. Likar. A review of methods for correction of intensity inhomogeneity in MRI. *IEEE Transactions on Medical Imaging*, 26(3):405–421, 2007.

[209] 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(93), 2016.

[210] J. Weese, T. M. Buzug, C. Lorenz, and C. Fassnacht. An approach to 2D/3D registration of a vertebra in 2D x-ray fluoroscopies with 3D CT images. In J. Troccaz, E. Grimson, and R. Mösges, editors, *CVRMed-MRCAS'97*, volume 1205 of *Lecture Notes in Computer Science*, pages 119–128. Springer, Berlin, Heidelberg, 1997.

[211] D. Wei, Y. Sun, P. Chai, A. Low, and S. H. Ong. Myocardial segmentation of late gadolinium enhanced MR images by propagation of contours from cine MR images. In G. Fichtinger, A. Martel, and T. Peters, editors, *Medical Image Computing and Computer-Assisted Intervention, MICCAI 2011*, volume 6893 of *Lecture Notes in Computer Science*, pages 428–435. Springer, Berlin, Heidelberg, 2011.

[212] D. Wei, Y. Sun, S.-H. Ong, P. Chai, L. L. Teo, and A. F. Low. Three-dimensional segmentation of the left ventricle in late gadolinium enhanced MR images of chronic infarction combining long- and short-axis information. *Medical Image Analysis*, 17(6):685–697, 2013.

[213] J. Weickert. Coherence-enhancing diffusion filtering. *International Journal of Computer Vision*, 31(2):111–127, 1999.

[214] J. Weickert. Applications of nonlinear diffusion in image processing and computer vision. *Acta Mathematica Universitatis Comenianae. New Series*, 70(1):33–50, 2001.

[215] A. P. Wijnmaalen, R. J. van der Geest, C. F. van Huls van Taxis, H.-M. J. Siebelink, L. J. Kroft, J. J. Bax, J. H. Reiber, M. J. Schalij, and K. Zeppenfeld. Head-to-head comparison of contrast-enhanced magnetic resonance imaging and electroanatomical voltage mapping to assess post-infarct scar characteristics in patients with ventricular tachycardias: real-time image integration and reversed registration. *European Heart Journal*, 32(1):104–114, 2010.

[216] J. T. Willerson, F. Scales, A. Mukherjee, M. Platt, G. H. Templeton, G. S. Fink, and L. M. Buja. Abnormal myocardial fluid retention as an early manifestation of ischemic injury. *The American Journal of Pathology*, 87(1):159–188, 1977.

[217] D. E. Winchester, D. C. Wymer, R. Y. Shifrin, S. M. Kraft, and J. A. Hill. Responsible Use of Computed Tomography in the Evaluation of Coronary Artery Disease and Chest Pain. *Mayo Clinic Proceedings*, 85(4):358 – 364, 2010.

[218] R. P. Woods, J. C. Mazziotta, and S. R. Cherry. MRI-PET registration with automated algorithm. *Journal of Computer Assisted Tomography*, 17(4):536–546, 1993.

[219] M. Xavier, A. Lalande, P. M. Walker, F. Brunotte, and L. Legrand. An adapted optical flow algorithm for robust quantification of cardiac wall motion from standard cine-MR examinations. *IEEE Transactions on Information Technology in Biomedicine*, 16(5):859–868, 2012.

[220] R. S. Xu, P. Athavale, Y. Lu, P. Radau, and G. A. Wright. Myocardial segmentation in late-enhancement MR images via registration and propagation of cine contours. In *Proceedings of the 2013 IEEE 10th International Symposium on Biomedical Imaging (ISBI)*, pages 856–859. IEEE, 2013.

[221] J. Yang, J. Wright, T. S. Huang, and Y. Ma. Image super-resolution via sparse representation. *IEEE Transactions on Image Processing*, 19(11):2861–2873, 2010.

[222] E. A. Zerhouni, D. M. Parish, W. J. Rogers, A. Yang, and E. P. Shapiro. Human heart: tagging with MR imaging - a method for noninvasive assessment of myocardial motion. *Radiology*, 169(1):59–63, 1988.

[223] W. Zhang, J. A. Noble, and J. M. Brady. Real time 3-D ultrasound to MR cardiovascular image registration using a phase-based approach. In *Proceedings of the 2006 3rd IEEE International Symposium on Biomedical Imaging: Nano to Macro*, pages 666–669. IEEE, 2006.

[224] X. Zhen, X. Gu, H. Yan, L. Zhou, X. Jia, and S. B. Jiang. CT to cone-beam CT deformable registration with simultaneous intensity correction. *Physics in Medicine and Biology*, 57(21):6807–6826, 2012.

[225] X. Zhuang, K. S. Rhode, R. S. Razavi, D. J. Hawkes, and S. Ourselin. A registration-based propagation framework for automatic whole heart segmentation of cardiac MRI. *IEEE Transactions on Medical Imaging*, 29(9):1612–1625, 2010.

[226] B. Zitova and J. Flusser. Image registration methods: A survey. *Image and Vision Computing*, 21(11):977–1000, 2003.