Ph.D. Thesis

Towards automating cine DENSE MRI image analysis: segmentation, tissue tracking and strain computation

Bruce Shawn Spottiswoode

Submitted to the Faculty of Health Sciences at
the University of Cape Town in fulfillment of the requirements for the degree of
Doctor of Philosophy in Medicine in Biomedical Engineering
Cape Town, South Africa, July 2006

Declaration

Towards automating cine DENSE MRI image analysis: segmentation, tissue tracking and strain computation

I, BRUCE SHAWN SPOTTISWOODE, hereby declare that:

(a) the above thesis is my own unaided work both in concept and execution, and that apart from the normal guidance from my supervisor, I have received no assistance except as stated below:

(b) except as stated below, neither the substance nor any part of the above thesis has been submitted in the past, or is being, or is to be submitted for a degree at the University of Cape Town or any other university.

Calculating the strain directly from the trajectories is also described by Aaron Hess in "Calculating 3D intramyocardial strain tensors in a single slice of myocardium using MRI", M.Sc. Thesis, University of Cape Town, March 2006.

The thesis has been presented by me for examination for the degree of Doctor of Philosophy in Medicine in Biomedical Engineering.

Signed by candidate
Signature

30 November 2006
Date

Acknowledgements

I would like to express my sincere gratitude to the following people for their assistance and guidance:

University of Cape Town: A big thank you to my supervisor Dr Ernesta Meintjes for her support and encouragement in both avenues of academic and personal growth, and for her MRI expertise; to my second supervisor Professor Bongani Mayosi for his contagious enthusiasm and clinical insight; to Aaron Hess for his ideas and the countless stimulating conversations; to the radiographers at Groote Schuur for their flexibility and patience; and to the students and staff at UCT Biomedical Engineering who I had the pleasure of working with during my graduate studies. I would also like to thank the NRF, MRC and UCT Postgraduate Scholarships Office for funding my studies, travels and conferences.

University of Virginia: First and foremost Professor Fred Epstein for welcoming me into his research group, for his ideas, for his striking expertise in the field, and for his unwavering humour and inspiration; also a big thank you to Xiaodong Zhong for his patience, good nature, and outstanding programming skills; to Moriel Vandsburger for being so accommodating and helpful; to John Christopher for his expertise at operating the scanner; and to Dr Patrick Helm for his Matlab wizardry. I would also like to thank the departments of Biomedical Engineering and Radiology at the University of Virginia for welcoming and accommodating me during both visits, and for offering me the use of their excellent resources. My work at UVA was in part supported by NIH grants 5 RO1 EB 001763-03 and 1 RO1 EB 001763-01.

I am also sincerely grateful to Siemens Corporate Research USA for the experience gained in my summer internship with them, and to my supervisor Dr Christine Lorenz for her inspiration and support. Finally, I would like to thank my wonderful family for everything they have done for me.

Abstract

Over the past two decades, magnetic resonance imaging (MRI) has developed into a powerful imaging tool for the heart. Imaging cardiac morphology is now commonplace in clinical practice, and a plethora of quantitative techniques have also arisen on the research front. Myocardial tagging is an established quantitative cardiac MRI method that involves magnetically tagging the heart with a set of saturated bands, and monitoring the deformation of these bands as the heart contracts. Displacement encoding with stimulated echoes (DENSE) is a relatively new technique that encodes myocardial tissue displacement into the phase of the MRI image. Cine DENSE allows for rapid quantification of myocardial displacement at multiple cardiac phases through the majority of the cardiac cycle. The work presented here significantly improves and automates the processing of cine DENSE data.

For practical sensitivities to motion, relatively high displacement encoding frequencies are used and phase wrapping typically occurs. In order to obtain absolute measures of displacement, a two dimensional (2–D) quality-guided phase unwrapping algorithm was adapted to unwrap both spatially and temporally. A fully-automated algorithm and a faster semi-automated algorithm are proposed. The fully-automated phase unwrapping algorithm successfully unwrapped 767 out of 800 cine DENSE images (95.9%), and the semi-automated algorithm successfully unwrapped 786 out of these 800 images (98.3%). A method for computing the 2–D trajectories of discrete points in the myocardium as they move through the cardiac cycle is also introduced. The error in individual displacement measurements is reduced for each motion trajectory using temporal fitting with 5^{th} order Fourier basis functions. This improvement is in turn reflected in strain measurements, which are derived directly from the trajectories.

This 2–D tissue tracking method was validated both on a rigid rotating phantom and *in vivo*. End-systolic left ventricular (LV) circumferential shortening (Ecc) and radial thickening (Err) were measured using cine DENSE for 5 normal volunteers, and the values obtained are shown to agree with measures derived using myocardial tagging. The accuracy of this 2–D in-plane tracking algorithm for typical cardiac displacements measured using a rotating phantom is 0.24 ± 0.15 mm. Further measurements were made

to optimize the displacement encoding frequency and to estimate the baseline strain noise both for the phantom and *in vivo*. For the *in vivo* studies, the optimal displacement encoding frequency is in the region of 0.1 cycles/mm. For 2 scans of 17 second duration, the strain noise after temporal fitting was estimated to be $2.5 \pm 3.0\%$ at end-diastole, $3.1 \pm 3.1\%$ at end-systole, and $5.3 \pm 5.0\%$ at mid-diastole. The improvement in intra-myocardial strain measurements due to temporal fitting is apparent in strain-versus-time curves, strain histograms, and also in images identifying regions of dysfunctional myocardium in studies of patients with infarcts.

Slice following (SF) has previously been developed for myocardial tagging to account for tissue motion through the image plane. By incorporating SF into a cine DENSE sequence, two important implications arise: 1. the ability to obtain true 2–D and three dimensional (3–D) frame-to-frame motion trajectories of material points initially lying in a single plane, and 2. the ability to use these true motion trajectories to assist in accurately segmenting the myocardium at each cardiac phase. Algorithms are presented to realise these two objectives.

The SF cine DENSE 3–D tissue tracking algorithm was validated on a moving phantom, and a through-plane tracking accuracy of 0.48 ± 0.25 mm was achieved. Using the true 3–D motion trajectories, the effect that rotation of the myocardium through the image plane has on measures of strain was investigated. For 6 volunteers scanned in a mid-ventricular short-axis view, through-plane rotation caused the end-systolic Ecc to be reduced in magnitude by $3.91 \pm 0.43\%$, and the corresponding Err to be increased in magnitude by $6.01 \pm 1.07\%$. For LV motion-guided segmentation, epicardial and endocardial contours were manually drawn at a single cardiac phase, and the motion measured by SF cine DENSE was used to project the contours onto the remaining cardiac phases. Based on full sets of manually-drawn contours and an area measure of segmentation error, the accuracy of the segmentation algorithm was shown to be similar to inter-observer variability. In addition, a radial segmentation error metric was introduced for short-axis data, and the resulting mean errors for 6 SF and 6 non-SF cine DENSE datasets are 0.13 ± 0.12 and 0.21 ± 0.15 pixels, respectively. Although the segmentation error for non-SF cine DENSE is greater than for SF cine DENSE, motion-guided segmentation using non-SF cine DENSE typically yields acceptable results in mid-ventricular short-axis views of the heart.

For studies of LV function using cine DENSE, the implementation of these algorithms reduces the total user interaction to manually demarcating the myocardium on a single arbitrary frame. Regional strain-time curves and 2–D/3–D equations of motion of points in the myocardium can then automatically and rapidly be derived at an arbitrary spatial and/or temporal resolution.

The right ventricle (RV) has a complicated geometry, is thinner than the LV, and behaves in a more unruly manner than the LV. Methods for analysing RV motion using cine DENSE are introduced. Using a measure of midwall tangential strain (Ett) for both the LV and RV, interventricular dyssynchrony is investigated in both volunteers and patients with bundle branch blocks. Both left and right bundle branch blocks were successfully detected using the delay to onset of Ett, encouraging the use of this technique for prospectively identifying responders to cardiac resynchronisation therapy (CRT).

Harmonic phase (HARP) analysis is a Fourier filtering procedure that allows for quantification of material displacement and strain from tagged MRI images. A HARP tracking technique is introduced that rapidly produces frame-to-frame motion trajectories. A motion-guided tissue classification technique is also presented. The technique uses differences between the relatively ordered behaviour of trajectories in the myocardium and the chaotic trajectories in adjacent regions of noise (ventricular cavities and lungs). The tissue classification error was quantified relative to manual contours, and further validation was carried out using a simulated contracting LV.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Project purpose and scope

Heart disease is a leading cause of death in the Western world. In 2000, cardiovascular disease (CVD) caused 39% of deaths in the United States (AHA, 2003), and in 2001 CVD-related mortality in the United Kingdom was slightly higher than this (BHF, 2003). CVD is South Africa's second largest killer, being responsible for 19,8% of all deaths each year (HFSA, 2003).

Myocardial wall motion abnormalities occur in nearly all heart diseases. Wall motion imaging therefore plays an integral role in the diagnosis, prognosis, and clinical management of cardiac disease. Accordingly, quantitative myocardial motion tracking and strain analysis potentially has great clinical significance. Magnetic Resonance Imaging (MRI) is fast becoming an indispensable tool for the diagnosis of heart disease. MRI boasts a selection of techniques for quantifying intra-myocardial motion, including myocardial tagging (Zerhouni *et al.*, 1988; Axel and Dougherty, 1989b), velocity-encoded phase contrast (PC) imaging (van Dijk, 1984; Bryant *et al.*, 1984), and more recently, displacement-encoded imaging using stimulated echoes (DENSE) (Aletras *et al.*, 1999b).

Myocardial tagging yields accurate measures of displacement and strain. The main disadvantages of tagging are relatively low spatial resolution, as defined by the tag spacing, and the need for some degree of manual intervention for detection of tag lines during magnitude image analysis. These disadvantages have largely been eliminated with the

advent of Harmonic Phase (HARP) analysis (Osman and Prince, 2000a). Velocity-encoded phase contrast imaging measures instantaneous tissue velocity using the phase of the acquired signal. Velocity encoding achieves high spatial resolution because velocity is encoded into the phase of each pixel. The primary disadvantage of velocity encoding is that velocity measurement errors propagate during displacement and strain analysis. DENSE, which measures tissue displacement from a fixed encoding time using the phase of the acquired signal, offers both high spatial resolution and measurement of displacement rather than velocity. While initial DENSE sequences imaged the heart at a single cardiac phase, navigator-guided (Pai and Wen, 2003) and breath-hold (Kim *et al.*, 2004) cine DENSE sequences that trade SNR for temporal resolution have recently been developed.

A portion of this document describes the processing of cine HARP images, but the majority of the thesis focuses on improving cine DENSE image analysis. A new two dimensional (2–D) tracking algorithm is introduced for cine HARP. This is based on interpolated inter-frame displacement measurements, and allows for the rapid computation of frame-to-frame motion trajectories. A motion-based tissue segmentation technique is also proposed. If trajectories commencing from each pixel centre at the first cardiac phase are calculated, then coherent trajectories result for the myocardium and stationary tissue, and chaotic trajectories result for the cavities containing air or fluid. The nature of these trajectories can be used for tissue classification, and a number of discerning characteristics were explored to achieve this.

To date, the analysis of cine DENSE images has only been done on a frame-by-frame basis and has entailed conversion of phase to displacement, followed by the computation of strain using finite element techniques. The phase of an MRI image is inherently confined to lie between $-\pi$ and π radians, and aliasing (or phase wrapping) is inevitable given practical cine DENSE parameters. This document reports the first application of spatiotemporal phase unwrapping for cine DENSE images. Spatiotemporal phase unwrapping is required for absolute measures of displacement using cine DENSE.

An algorithm is introduced to track discrete material points in the myocardium from frame to frame using these absolute displacement measurements. The accuracy of this algorithm was validated using a moving phantom. If this tracking is commenced at each myocardial pixel centre at the first cardiac phase, then trajectories are obtained at approximately the resolution of the pixel size. Each trajectory is improved by fitting a

curve to each coordinate component as a function of time. Temporal fitting was done here using 5^{th} order Fourier basis functions, and the coefficients were rapidly derived using the discrete Fourier transform. Lagrangian strain was calculated directly from the trajectories, and the improvement offered by temporal fitting is evident both in the standard deviation of end-systolic strain values, and in curves representing strain versus time. The temporal fitting was also shown to be useful in extracting information in noisy cine DENSE images of patients with myocardial infarcts. Furthermore, random noise was added to both phantom and *in vivo* cine DENSE k-space data, and the effectiveness of the temporal fitting in reducing this noise was investigated.

A shortfall of phase contrast velocity encoding, myocardial tagging and DENSE is that a fixed plane is generally imaged, through which the heart moves during the cardiac cycle. Different portions of tissue can thus be present at each cardiac phase, which greatly complicates myocardial tissue tracking. This has been addressed in myocardial tagging with slice following (SF) (Fischer *et al.*, 1994), which uses a differential acquisition to track true myocardial motion.

The proposed cine DENSE tissue tracking algorithm can be directly extended to three dimensions (3–D) by applying cine DENSE to multiple adjacent image planes, and encoding each image plane in three orthogonal directions. This currently requires impractical imaging times for routine clinical use. By incorporating SF into a cine DENSE pulse sequence and applying displacement encoding in three orthogonal directions, it becomes possible to obtain true 3–D frame-to-frame motion trajectories from a single slice. This 3–D tissue tracking algorithm is introduced and was validated on a moving phantom. Volunteer SF cine DENSE data was then used to investigate the effect of through-plane rotation on myocardial strain measurements.

The most time consuming and labour intensive component of the cine DENSE image analysis is manually demarcating myocardial contours for each frame. If a large number of points in the myocardium are tracked using SF cine DENSE, then the 2–D in-plane trajectory components can be used to assist in segmenting the myocardium. Obtaining suitable trajectories is complicated by the presence of noisy displacement measurements in regions adjacent to the myocardium. A method is proposed to eliminate these noisy displacement measurements based on physiological motion parameters. The resulting trajectories can then be used to project a manually-defined portion of myocardium onto each cardiac phase. Contours were extracted using isocontours or active contour models.

The effectiveness of the method was gauged relative to manually-drawn contours based on a sum of false positive and false negative area measures, and also using radial distance measurements. This is the first reported segmentation algorithm for cine DENSE.

The cine DENSE image processing methods presented here allow for an accurate strain-time and displacement-time analysis of the human left ventricle (LV) and reduce the total manual interaction to the delineation of the myocardium at a single cardiac phase. These algorithms have all been implemented in a custom-written MATLAB graphical user interface which is described in Appendix A

Cardiac resynchronization therapy (CRT) has been shown to improve symptoms, quality of life, LV function, and mortality in patients with end stage heart failure (Bristow *et al.*, 2004). However, approximately 30% of patients that meet the current criteria for CRT do not respond positively to this costly and invasive therapy (Bristow *et al.*, 2004). Cine DENSE MRI methods have recently shown promise for quantifying LV mechanical dyssynchrony in a way that may be superior to tissue Doppler in terms of spatial coverage and the ability to assess circumferential strain rather that longitudinal velocities, strain rate, and strain (Patel *et al.*, 2006). The additional measurement of interventricular dyssynchrony between the right ventricle (RV) and LV further improves the ability to prospectively identify CRT responders (Penicka *et al.*, 2004).

The RV wall is notoriously difficult to image and analyse because it is thinner and behaves in a more unruly manner than the LV, and it has a complicated, asymmetric geometry. Only a limited number of RV studies have been carried out using myocardial tagging (Young *et al.*, 1996; Klein *et al.*, 1998; Fayad *et al.*, 1998; Haber *et al.*, 2000, 2005) and phase contrast velocity encoding (Kayser *et al.*, 2000). There are no reported RV studies using DENSE. Methods to study RV motion using cine DENSE are introduced here and these were used in conjunction with the LV analysis to investigate LV-RV interventricular dyssynchrony. These methods have been successful in identifying interventricular dyssynchrony in heart failure patients with left and right bundle branch blocks.

1.2 Document overview

The bulk of this document presents advances in cine DENSE image processing techniques. The work on cine HARP tissue tracking and motion-based tissue classification was done prior to the work on cine DENSE and is presented without significant phantom or patient validation. To facilitate the flow of the document, the cine HARP chapter is presented after the relevant background material has been covered in the cine DENSE chapters. A knowledge of MRI physics and image reconstruction is assumed throughout.

Chapter 2 gives an overview of cardiac anatomy and function, and the modalities available for imaging the heart. Quantitative cardiac MRI techniques capable of measuring myocardial strain are introduced, specifically phase contrast velocity encoding, myocardial tagging, HARP, and DENSE.

Chapter 3 introduces phase unwrapping algorithms for cine DENSE. Background concepts for phase unwrapping are presented and the requirements for the correct phase unwrapping of cine DENSE images are investigated. A fully-automated spatiotemporal phase unwrapping algorithm is introduced and was tested on cine DENSE data. The effectiveness of a faster semi-automated alternative is also investigated. Finally, the results of a study aimed at determining the optimal displacement encoding frequency based on maximising signal-to-noise-ratio (SNR) and phase unwrapping efficacy are presented.

Chapter 4 presents a 2–D tissue tracking algorithm for cine DENSE. This algorithm was validated on a moving phantom, and the tracking accuracy was quantified. Periodic curves are fitted to the motion trajectory coordinates as a function of time. Strain is calculated directly from the resulting motion trajectories and both the trajectories and the strain measurements are shown to be improved by the temporal fitting. The original cine DENSE displacement fields are also improved using interpolations from the fitted motion trajectories. Baseline strain noise measurements were made both on a phantom and *in vivo*.

Chapter 5 introduces slice following (SF) and discusses the integration of SF into the cine DENSE sequence. A 3–D tissue tracking algorithm for SF cine DENSE is introduced and results for a moving phantom are presented. This provides true 3–D motion trajectories of a single slice of tissue at roughly the resolution of the voxel size. The effect of through-

plane rotation of the myocardium on 2–D strain was quantified for short-axis views of the human heart.

Chapter 6 presents a 2–D tissue tracking algorithm that employs inter-frame cine HARP displacement measurements. The accuracy of the tracking method was quantified using the unique phase consistency properties of HARP images. Tracking points in the myocardium or stationary tissue results in coherent motion trajectories whereas tracking points in the noisy regions results in chaotic trajectories. A series of tissue classification algorithms based on the nature of these motion trajectories are introduced. This method for separating the myocardium from surrounding and enclosed noisy regions is compared with manually-defined regions.

Chapter 7 presents a myocardial tissue segmentation technique specific to cine DENSE images. A manually-defined region of myocardium is projected onto all other cardiac phases using the encoded displacement information. Contours are then extracted using isocontours or active contour models, and smoothed by fitting periodic functions. The segmentation method was applied to non-SF and SF cine DENSE data, and the resulting contours were compared to manual contours using both a sum of false positive and false negative area measures, and a tailored radial distance measurement.

Chapter 8 introduces techniques for analysing RV wall motion and deformation using cine DENSE. These were also applied to the LV, and used to assess LV-RV interventricular dyssynchrony in both patients and volunteers on the basis of time-to-onset of strain.

Chapter 9 summarises the conclusions from the preceding six chapters and discusses the relevance of this work in assisting diagnosis and understanding human heart function. Possible directions for future work are also explored.

The majority of this work has been submitted to or appears in the following peer-reviewed journal articles and conference proceedings:

- B. S. Spottiswoode, X. Zhong, C.H. Lorenz, B. M. Mayosi, E. M. Meintjes, and F. H. Epstein, "Slice Followed Cine DENSE MRI: True 3D Tissue Tracking and Motion-guided Segmentation of the Myocardium", *IEEE Transactions on Medical Imaging*, submitted post-review, November 2006.

- B.S. Spottiswoode, X. Zhong, C.M. Kramer, E.M. Meintjes, B.M. Mayosi, and F.H. Epstein, "Tracking Myocardial Motion from Cine DENSE Images using Spatiotemporal Phase Unwrapping and Temporal Fitting", *IEEE Transactions on Medical Imaging*, in press.

- B.S. Spottiswoode, "Myocardial contour detection in MRI tagged images using short-range correlations of displacement fields", *SAIEE Transactions*, vol. 96, no. 2, pp. 99-105, June 2005.

- B.S. Spottiswoode, X. Zhong, C.H. Lorenz, B.M. Mayosi, E.M. Meintjes, F.H. Epstein, "Motion-guided Segmentation of the Left Ventricle for Cine DENSE MRI", In: *Proc. 14^{th} ISMRM*, Seattle, USA, May 2006.

- B.S. Spottiswoode, X. Zhong, R.A. Patel, P.A. Helm, B.M. Mayosi, C.M. Kramer, E.M. Meintjes, and F.H. Epstein, "Measuring LV-RV Interventricular Dyssynchrony using Cine DENSE MRI", In: *Proc. 14^{th} ISMRM*, Seattle, USA, May 2006.

- B.S. Spottiswoode, X. Zhong, C.H. Lorenz, E.M. Meintjes, B.M. Mayosi, and F.H. Epstein, "Tracking 3–D Myocardial Motion using Slice-followed Cine DENSE", *Journal of Cardiovascular Magnetic Resonance*, vol. 8, no. 1, pp. 80-81, January 2006.

- B.S. Spottiswoode, X. Zhong, E.M. Meintjes, B.M. Mayosi, F.H. Epstein, "Phase Unwrapping for Absolute Cine-DENSE Myocardial Displacement Measurements using a 3D Guided Flood-fill Algorithm", In: *Proc. 13^{th} ISMRM*, Miami, USA, May 2005.

- B.S. Spottiswoode, X. Zhong, E.M. Meintjes, B.M. Mayosi, F.H. Epstein, "Improved Myocardial Tissue Tracking and Strain Accuracy in Cine-DENSE using Temporal Fitting", In: *Proc. 13^{th} ISMRM*, Miami, USA, May 2005.

Chapter 2

Background theory

2.1 Imaging the human heart

2.1.1 Cardiac anatomy and function

The heart is a pulsatile, four chamber pump with two atria and two ventricles. Its function is to pump deoxygenated blood to the lungs and oxygenated blood to the rest of the body. Figure 2.1 shows an anterior view of the heart in situe. The heart is bounded laterally by the lungs, anteriorly by the sternum and caudally by the diaphragm. The base of the heart is the cranial portion, wherefrom the majority of the great vessels arise. The blunt extremity of the heart formed by the left ventricle is known as the heart apex, which projects inferiorly, anteriorly and to the left of the left ventricle.

The heart is surrounded by a tough, fibrous bag called the pericardium. This prevents over distension of the heart and anchors the heart in position, being itself attached to the sternum and the diaphragm. Between the fibrous pericardium and the heart itself is the pericardial cavity, which may become filled with fluid under certain conditions. The heart muscle is termed myocardium, the exterior border of which is known as the epicardium, and the interior border (in contact with the blood) is known as the endocardium.

The atria serve as primer pumps for the ventricles. Deoxygenated blood enters the right atrium (RA) via the superior vena cava from above and the inferior vena cava from

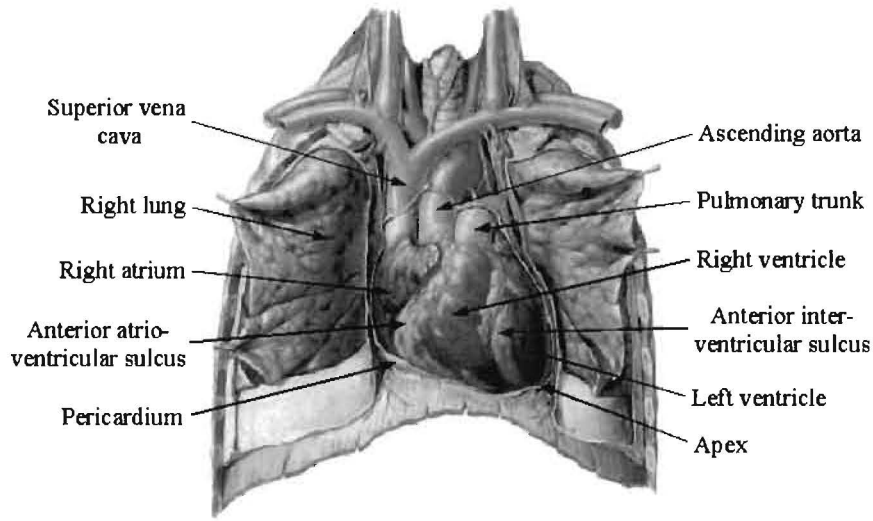

Figure 2.1: Anterior view of the heart in situe. The pericardium has been removed to expose the heart surface features. Image adapted from Netter (1998).

below. The RA pumps the blood through the tricuspid valve into the right ventricle (RV). Blood is then pumped by the RV via the pulmonary valve into the pulmonary trunk, which splits into the pulmonary arteries before entering the lungs. Upon leaving the lungs, the oxygenated blood enters the left atrium (LA) through the four pulmonary veins. The LA primes the left ventricle (LV), directing the blood through the mitral valve in the process. Finally, the contraction of the LV ejects the blood through the aortic valve and into the aorta.

The LV has the burden of pumping blood to the entire body and hence has relatively thick walls. The RV only pumps blood into the lungs and thus has far thinner walls than the LV. The interventricular septum is the common wall shared by the two ventricles, and is also made up of thick muscle because of its role in LV function. The portions of the ventricles excluding the septum are known as the ventricular free walls.

The period from the end of one heart contraction to the end of the next contraction is called the cardiac cycle. This is comprised of a period of relaxation called diastole, followed by a period of contraction called systole. The atria contract shortly before the ventricles and hence systole is subdivided into atrial and ventricular systole. As the atria contract, the tricuspid and mitral valve leaflets are initially pulled open by tendinous cords attached to muscular extensions of the ventricular wall known as papillary muscles. The papillary muscles also serve to prevent the atrioventricular valves from billowing into the atria when the ventricles contract. There are two large papillary muscles in

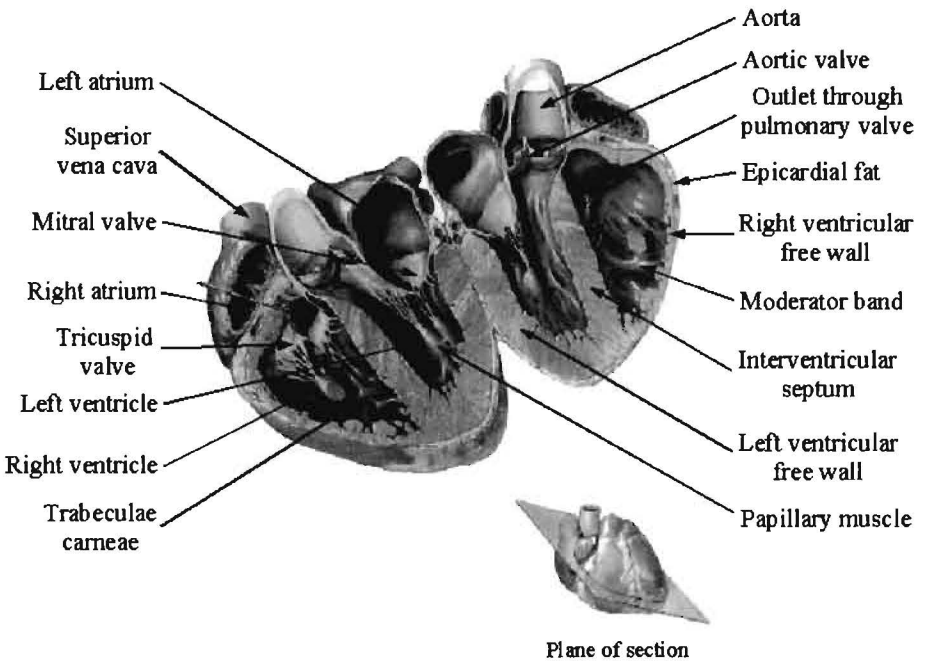

Figure 2.2: View of the heart sliced through its long-axis. Image adapted from Netter (1998).

each ventricle and these are often clearly visible in MRI images. The outflow tracts of both ventricles have a smooth interior, whereas the majority of the endocardial walls are rough, being comprised of prominent ridges of muscle known as trabeculae carneae. The moderator band is a small but prominent bridge of muscle in the right ventricle, and serves to conduct neural impulses. Figure 2.2 clearly depicts the tricuspid, mitral and aortic valves, as well as the papillary muscles, trabeculae carneae and the moderator band.

Immediately after ventricular contraction begins, there is a short period, known as isometric contraction, when all of the valves associated with the ventricles are closed. This occurs while the ventricles build up sufficient pressure to push the aortic and pulmonary valves open against the pressures in the aorta and pulmonary trunk. Isometric contraction implies that tension is increasing in the muscle but no shortening of the muscle fibres is occurring. This is not strictly true because there is apex-to-base shortening and circumferential elongation (Guyton, 1981).

Once the pressures in the LV and RV exceed the pressures in the aorta and pulmonary trunk, respectively, the semilunar valves open and the blood is ejected. About 60% of the blood is ejected during the first quarter of ventricular systole, and usually most of the remaining 40% during the next two quarters (Guyton, 1981). These three quarters

of ventricular systole are known as the period of ejection. Although very little blood flows out of the ventricles during the last quarter, the ventricle musculature remains contracted. This period is called protodiastole. Due to elastic recoil of the great arteries, the pressure in the ventricles during protodiastole actually falls below that within the great arteries.

The ventricular muscle fibres spiral around and down from the atrioventricular boundary to end either in the papillary muscles, or to spiral back up along a similar path. The effect of this is that when the ventricles contract they twist in an anticlockwise direction when viewed from the apex. It is this rotation that causes the heart to strike the anterior thoracic wall in a manner that we are all familiar with during exertion. In moving through the ventricular wall, muscle fiber direction undergoes a progressive change, with the fibers in the middle third of the wall being oriented circumferentially (Mirsky *et al.*, 1974).

The sinoatrial node is the natural pacemaker of the heart, and consists of a cluster of specialised cells situated in the wall of the RA. The electrical signal generated by the sinoatrial node moves down through the heart until it reaches the atrioventricular node, which is situated in the center of the heart between the atria and ventricles. After passing the atrioventricular node, the electrical signals travel to the ventricles along Purkinje fibers. There are three conducting pathways from the AV node to the tissue of the ventricles: the right bundle branch (RBB), the left anterior fascicle, and the left posterior fascicle. The combined left anterior and posterior fascicles are referred to as the left bundle branch (LBB).

2.1.2 Cardiac imaging overview

The three most commonly used imaging modalities for the diagnosis of heart disease are echocardiography, nuclear scintigraphy, and X-ray angiography. Positron electron tomography (PET) and single photon emission computed tomography (SPECT) are two advanced branches of nuclear medicine that are used for functional cardiac imaging. Computed tomographic (CT) scanners are also used for cardiac anatomical imaging. These latter three methods are, however, not routine techniques for the diagnosis of heart dysfunction and have thus only been given mention.

Echocardiography is typically the first-line noninvasive test to evaluate cardiac anatomy and function if a heart defect is suspected. It is based on analysing ultrasonic waves that have been reflected off the various blood-tissue and tissue-tissue interfaces that constitute the heart. By examining the Doppler shift of the ultrasonic waves due to flowing blood, normal and abnormal areas of blood flow can be quantified. Echocardiography is relatively inexpensive but it exhibits a lower image quality than MRI and CT. It is also limited by the fact that the ultrasonic waves penetrate poorly through bone and lung, hence restricting the use of echocardiography to a few specific viewing windows.

Nuclear scintigraphy involves infusing a radioactive compound, or tracer, into the patient. This compound collects in areas of the heart that receive a good blood supply. Gamma rays emitted by the compound are detected using a scintillation camera, and heart-wall movement and overall heart function can be evaluated with cardiac gating. Despite being an excellent tool for functional analysis, morphological detail is limited by the relatively low spatial resolution of the technique.

X-ray angiography has for decades served as the definitive technique for demonstrating the morphology of a variety of heart diseases. It is also the gold standard for defining coronary arterial anatomy and pathology. A catheter is inserted into a femoral vessel and guided up to the heart. An iodinated contrast medium is then injected, and its course is monitored by taking a rapid series of X-rays. X-ray angiography provides detailed morphological information, but it is invasive and typically only gives a two-dimensional representation of the process.

The ability of MRI to accurately discern soft tissues makes it well suited to studying a range of heart pathologies. In addition to examining soft tissue, MRI also has the ability to obtain images whose signals are proportional to blood flow in vessels, tissue oxygenation, tissue perfusion, mechanical contraction, and the concentration of metabolites in tissues.

Our understanding of intramyocardial motion and deformation is limited. The ability to monitor the detailed behaviour of the myocardium would better our knowledge of normal heart function, as well as assist in identifying abnormalities arising in cardiac disease. Assessment of intramyocardial motion was first analysed by implanting trackable materials into the myocardium. The earliest technique was based on radiopaque fiducial markers (Ingels *et al.*, 1975). These were very invasive and the method was

restricted by the paucity of the physical markers. A similar technique used ultrasonic crystals (Villarreal *et al.*, 1988), but had the same disadvantages. As these methods are too invasive to be performed repeatedly, research in this field has been limited. A less invasive technique, developed by Young *et al.* (1992), monitors the bifurcations of the coronary arteries using angiography. However, this method is also limited by the paucity of the features being tracked.

Noninvasive methods of quantifying myocardial motion based on electron beam PET, SPECT, and CT all rely on the movement of the epicardial and endocardial borders, and give limited insight into intramyocardial motion. The twisting of the heart during the cardiac cycle is particularly difficult to analyse using epicardial and endocardial borders. Both MRI and ultrasound allow for the measurement of intramyocardial motion. The available MRI methods will be discussed later in this Chapter. Ultrasound methods include tissue Doppler imaging (Sutherland *et al.*, 1994; Heimdal *et al.*, 1998; D'hooge *et al.*, 2000), which measures instantaneous myocardial velocity, and speckle tracking (Li *et al.*, 2006), which is based on tracking the correlated speckle pattern from frame to frame through the cardiac cycle.

2.1.3 Cardiac MRI

In cardiac MRI, electrocardiogram (ECG) triggering is done almost without exception. An ECG is a trace depicting the potential of the electrical currents involved in the heart's contraction. A normal ECG trace is shown in Figure 2.3(a). It is composed of a P wave, a QRS complex, and a T wave. The P wave is caused by electrical currents generated by the atria just before they contract, and the QRS complex is caused by currents generated by the ventricles just before they contract. The T wave is caused by currents generated when the ventricles recover from their contraction. The QRS complex is the most prominent feature in an ECG trace, and is used as the trigger for ECG gating. Image acquisition during cardiac MRI is a relatively slow process, and sampling is done over a number of heart beats. In order to minimise respiratory motion artifacts, the patient is often asked to hold his/her breath for each scanning sequence. The image acquisition is made more time-efficient by capturing a number of samples during each heart beat. This process is depicted in Figure 2.3(b), where the resulting samples from different phases of the cardiac cycle can be used in a multi-slice collection of images, or combined into a cine image. A number of fast imaging techniques have

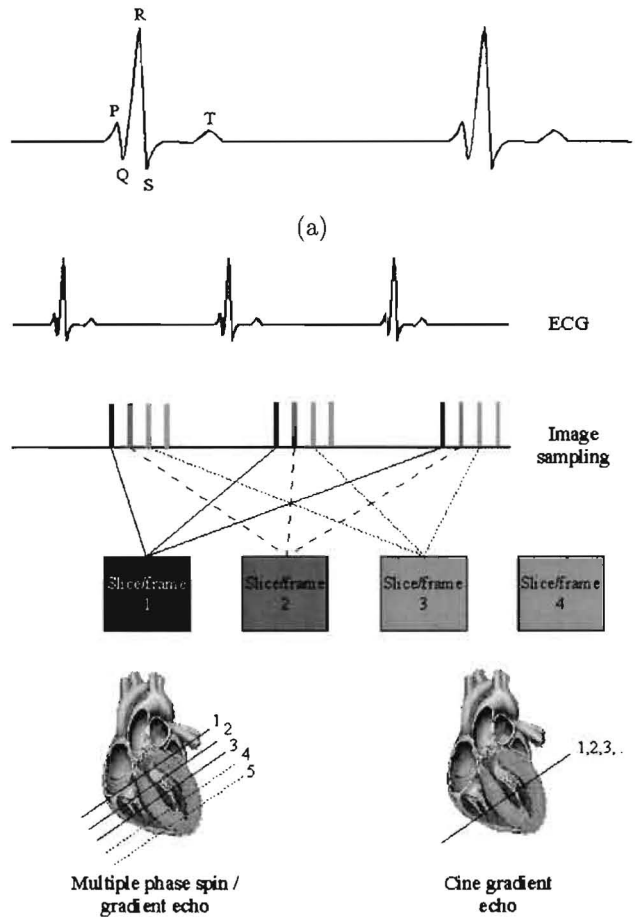

Figure 2.3: Electrocardiogram triggering in cardiac MRI. (a) A standard ECG trace depicting the most prominent electrical peaks. (b) Triggered gating showing the series of samples used to construct images of multiple slices or cine images.

been developed, including steady-state free precession (SSFP) and echo planar imaging (EPI). Details of these sequences can be found in many MRI texts, for example Haacke *et al.* (1999).

While MR imaging of the body often uses the conventional orthogonal imaging planes (i.e. transverse, sagittal and coronal), cardiac imaging is more frequently performed in planes orthogonal to the cardiac axes. The most common imaging planes are shown in Figure 2.4. The positions of the atrial cavities are labelled, but the resolution offered by the magnitude-reconstructed DENSE images is insufficient to clearly resolve the atrial walls. The long-axis of the left ventricle is defined by a line passing through the apex and the centre of the atrial valve (Dinsmore *et al.*, 1984). Planes parallel to this long-axis and aligned vertically or horizontally are known as vertical or horizontal LA planes

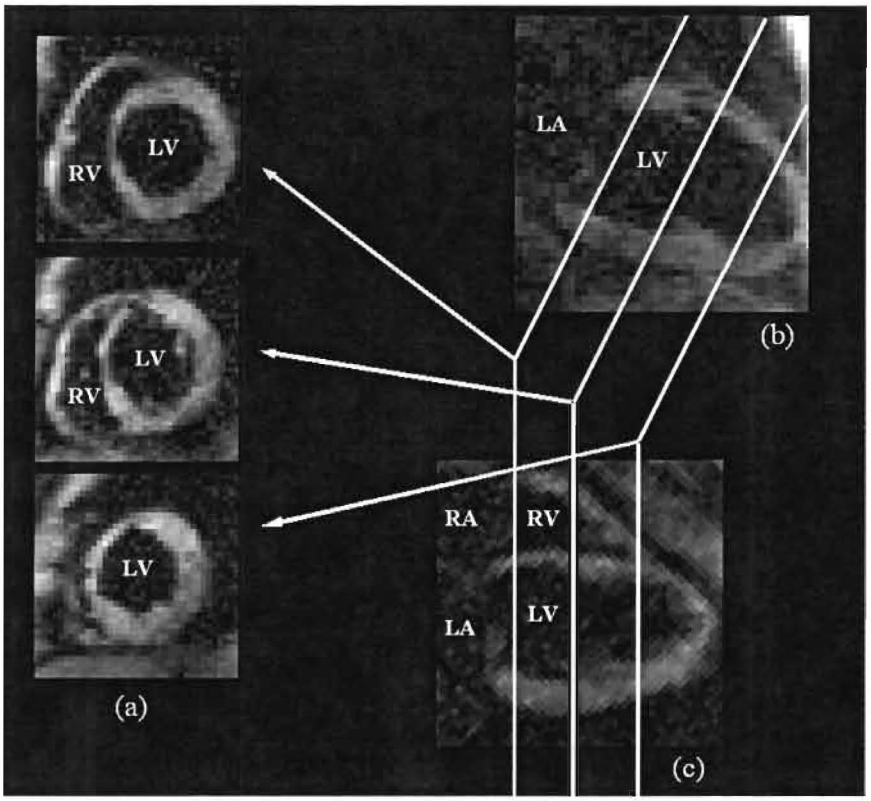

Figure 2.4: Illustration of magnitude-reconstructed DENSE images in three standard cardiac planes. (a) Short-axis view; (b) long-axis parallel to the septum; and (c) long-axis perpendicular to the septum. LV – left ventricle; RV – right ventricle; LA – left atrium; RA – right atrium.

respectively. The planes geometrically perpendicular to the long-axis of the left ventricle are known as short-axis (SA) planes. The apical two chamber view is the vertical long-axis plane which is parallel to the septum, but passes through the left ventricle, mitral valve and left atrium. The four chamber view is the horizontal long-axis plane which is perpendicular to the septum, but passes through both ventricles, both atrioventricular valves and both atria. The work presented in this manuscript was largely confined to mid-ventricular short-axis views of the heart. Major anatomical features seen in this image plane are shown in Figure 2.5.

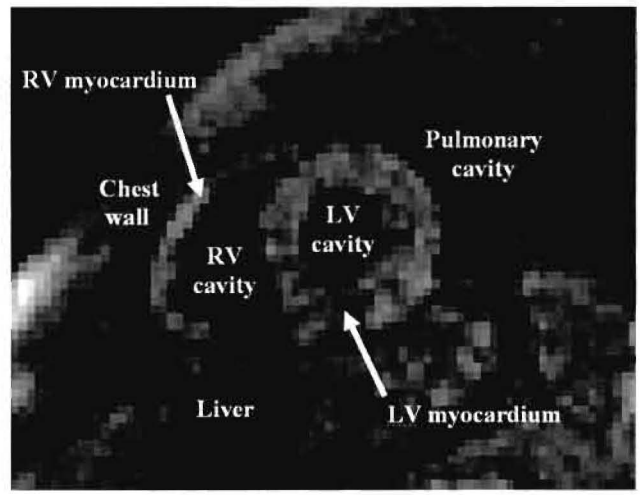

Figure 2.5: Short-axis DENSE magnitude-reconstructed overview image depicting major anatomical features.

2.2 Quantifying myocardial contractility using MRI

2.2.1 Phase contrast velocity encoding

Phase contrast (PC) MRI was first developed by van Dijk (1984) and, separately, by Bryant *et al.* (1984) to map the velocity of blood and myocardium. It was adapted for cine acquisitions by Nayler *et al.* (1986). Velocity encoding is based on the fact that the frequency of precession of a spin in a magnetic field gradient will change if it moves to another part of the gradient, as governed by the Larmor equation $f = B\gamma$, in which f is the frequency of precession, B is the magnetic field, and γ is the gyromagnetic ratio.

If a spin moves in a constant magnetic field gradient G with constant velocity v for time t, then an accumulated phase shift of $\Delta\phi = \gamma G v t^2$ results. This phase shift is directly proportional to the velocity of the moving particle. Bipolar gradients are typically used to cancel any phase shifts for stationary spins, thus ensuring that only moving spins acquire a net phase shift. To obtain velocities in different directions, gradients have to be applied in each direction. A map of velocities can then be extracted and strain-rate tensors can be calculated from the velocity field using finite element techniques.

Adjusting the magnetic field gradient effectively changes the range (window) of velocities that can be imaged. This window corresponds to a phase shift of 2π radians, and velocities outside the window will result in phase wrapping (or aliasing). The velocity

window thus needs to be set according to the range of expected myocardial velocities in the direction of interest. McVeigh (1996) states that the motion of the heart wall is of the order of 2.5 cm/s during systole and 10 cm/s during rapid filling.

Obtaining measures of strain (as opposed to strain-rate) requires that material points of myocardium be tracked through time. The strain tensor can be calculated directly from the motion trajectories (Zhu *et al.*, 1997), which can be obtained using one of several tracking algorithms (Constable *et al.*, 1994; Pelc *et al.*, 1995; Zhu *et al.*, 1996). Tissue tracking using velocity measurements results in an integration of measurement error. The tissue tracking algorithms attempt to correct this error by integrating forward and backwards in time (Constable *et al.*, 1994; Pelc *et al.*, 1995), and by using the periodicity of cardiac motion (Zhu *et al.*, 1996). The strain measurements can be smoothed further using a least squares fit of the trajectory data to a trajectory model (Zhu *et al.*, 1997). Despite these advances, a limited number of studies of myocardial motion have been conducted using velocity encoding (Masood *et al.*, 2000). This is in part due to the fact that velocity encoded images are slow to acquire and are prone to motion artifacts.

2.2.2 Myocardial tissue tagging

2.2.2.1 Theory

In MR myocardial tagging (Zerhouni *et al.*, 1988; Axel and Dougherty, 1989b), magnetically tagged points act as noninvasive fiducial markers and their motion can be quantitatively analysed like the motion of the invasively implanted markers mentioned previously. In addition to simple translation and rotation of the markers relative to each other, complex motions such as cardiac twist can be demonstrated (Lorenz *et al.*, 2000).

Myocardial tagging typically produces regular dark bands, or *tags*, of selective saturation. Since the saturation effect is a property of the protons being imaged, the saturation bands deform with the tissue. Specified regions of the myocardium are thus labelled and serve as markers during contraction. Figures 2.6(a) and 2.6(b) are examples of myocardial line tagging in a short-axis view. The tags are applied at end-diastole and are seen undeformed in Figure 2.6(a). Their distortion at end-systole is clear in Figure 2.6(b).

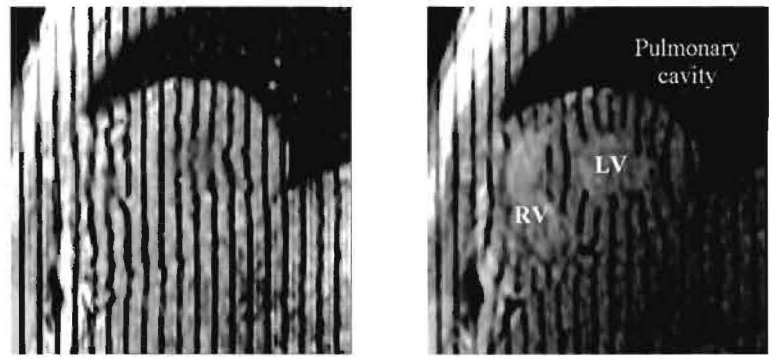

(a) (b)

Figure 2.6: Short-axis tagged images, (b) undeformed, at end-diastole, and (c) deformed, at end-systole.

The most common myocardial tagging sequence is the spatial modulation of magnetisation (SPAMM) sequence, which was proposed by Axel and Dougherty (1989b). The periodic spatial modulation of signal is achieved by using a sequence of two non-frequency selective RF pulses separated by a magnetic field gradient pulse. This is followed by a conventional imaging sequence. The direction of the applied gradient pulse determines the direction of the tag lines. The physics describing SPAMM myocardial tagging is presented in Appendix B. For 1–D tagging in the x direction, the longitudinal magnetisation M_z after applying the tag sequence is given by

$$M_z = M_{z0}(\cos^2\theta - \sin^2\theta \cos k_x x) \tag{2.1}$$

where M_{z0} is the longitudinal magnetisation prior to applying the tagging sequence, θ is the flip angle of the RF pulses about the z axis, and k_x corresponds to the tag spatial frequency. The applied gradient lies along the x direction, so that the resulting magnetisation is cosine modulated in this direction. More complicated tag profiles are possible using additional RF pulses (Axel and Dougherty, 1989a; Mosher and Smith, 1990), but only tagging with two RF pulses (1–1 SPAMM) will be considered in this manuscript.

Although the time dependence of the magnetisation is not shown in Equation 2.1, the tag contrast fades as a function of the $T1$ relaxation. Most commonly, the tags are applied as early as possible after detection of the QRS complex in order to ensure that motion is captured during systole. Zerhouni *et al.* (1988) noted that, in humans, tags

could be seen 60–450 ms after RF saturation, thus permitting sampling of the entire contractile phase of the cardiac cycle.

Until the nuclei in the tags relax, the tag lines are material properties of the tissue being investigated and hence move with it as it deforms. The displacement of the tag lines during the cardiac cycle indicates motion at right angles to the direction of the undeformed set of tag lines. Capturing motion in 2–D using orthogonal sets of line tags typically requires two acquisitions over two separate breath holds. The tagging sequence can be extended to produce a grid pattern by including a second modulation gradient orthogonal to the first. A grid arrangement allows one to resolve motion in 2–D from a single cine image acquisition, and also eliminates registration errors due to difference breath hold positions. However, grid tagging reduces the signal roughly twice as much as line tagging and in practice two separate sets of orthogonal line-tagged images are better suited to quantitative analysis. To identify motion in a third dimension, further image planes are required with components perpendicular to the original image plane. It is customary in MRI myocardial tagging to obtain two directions of motion from short-axis slices, and the third direction from long-axis slices.

SPAMM tagging suffers from the drawback that the tag lines are relatively poorly resolved. Tag contrast was improved by Fischer *et al.* (1993) using a technique called complimentary spatial modulation of magnetisation (CSPAMM), which separates the component of the magnetisation with the tagging information from the relaxed component. This is achieved by subtracting an image with a positive tagging grid from an image with a negative tagging grid, for which the modulating sinusoid was shifted by π radians. CSPAMM not only improves the contrast of the tags, but also prolongs the persistence of the tags, making motion assessment of the entire heart cycle possible. Fischer *et al.* (1993) further improved the contrast by optimising the flip angle of the RF pulses of the tagging sequence. The imaging excitation angle was adjusted for each heart phase to ensure that the magnetisation of the tag lines is equal for each image of the series. The drawback of the CSPAMM technique is that it requires two measurements, thus doubling the scan time of a SPAMM acquisition.

Tag spacings most commonly used range from 5 to 7 mm, although separations as low as 4 mm have been reported (Urayama *et al.*, 2000). Breath-hold sequences are typically used to eliminate respiratory motion effects over the 16–20 cardiac cycles required for data acquisition. Such breath-holds are feasible for cardiac patients even early after

acute myocardial infarction, provided that the patient is correctly oriented and cooperative, and that dyspnea due to severe pulmonary congestion or an accompanying disease such as severe obstructive airways disease is not present (Reichek, 1999).

2.2.2.2 Image analysis

Tagged cardiac images can be very informative when viewed in cine loops by an experienced observer. The strength of the technique, however, lies in the quantitative analysis methods that have been developed. Typically, the first step in such an analysis is marking the tag lines, or alternatively, the tag line intersections. The most common methods of semi-automatically or automatically identifying tag lines are presented in Appendix C. The interested reader is further referred to the review paper by Axel *et al.* (2005) for more detail on both the imaging and image analysis of myocardial tagging.

The second step typically involves deriving tissue displacements and/or strain measurements. More complicated motion parameters can also be derived, such as torsion, rotation velocity, and shear angle (Stuber *et al.*, 1998). Strain can be derived by triangulating the grid and using finite element methods on these triangles to identify deformation parameters. Figure 2.7 represents the deformation of four adjacent tag intersections. Myocardial motion and deformation are assumed to be homogenous within these triangles, that is, the tissue is assumed to be incompressible. The movement of two vertices of a triangular element can be used to calculate changes in magnitude and direction of the first (λ_1) and second (λ_2) principal strains within the element. Lengths of triangle sides before (Δp) and after (Δx) deformation are shown in Figure 2.7. A hypothetical unit circle at the centre of the equilateral triangle is seen to deform to an ellipse. Principal strains are aligned along major and minor axes of this ellipse, and the lengths of these axes give relative strains.

The motion of these elements can be characterised as a combination of rigid-body translation, given by the displacement of the triangle's centroid, and deformation, given by a tensor (F) that transforms Δp to Δx. Background material describing deformation analysis and strain tensors is presented in Appendix D.

The tensor F is independent of rigid body translation, but can be decomposed into its equivalent rigid body rotation and strain components. This is done using the polar decomposition theorem, which can be written in the form $F = R \cdot G$. Here R is an

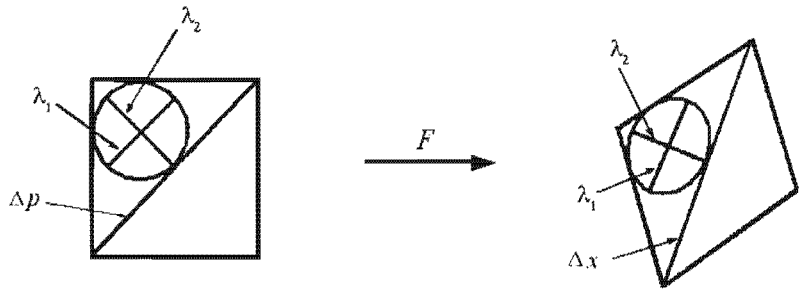

Figure 2.7: Finite element strain analysis depicted on a deforming square, where the corners of the square represent tag intersection points.

orthogonal tensor describing the rotation, and G, also known as Green's strain tensor, is a symmetric positive tensor describing the strain. Knowledge of the strain tensor at a point allows one to find the fractional change in length of an infinitesimal line segment oriented along any direction through that point. The eigenvalues of G represent the principal strains, while the eigenvectors represent the principal axes of the strain.

It is useful to fit the displacement information from the tagged images to a model of the heart. This not only smooths the displacement and strain measurements, but also acts as an interpolant and allows for the incorporation of data from multiple imaging slices. The reader interested in 3–D cardiac modelling is referred to the review paper by Frangi *et al.* (2001).

Kumar and Goldgof (1994) developed a mapping from one image to the next using a thin plate spline model as an interpolant. This eliminates the need for an explicit step that tracks the points and allows an arbitrary point in one image to be mapped to its corresponding point in the next image. In this way deformation parameters can be estimated for arbitrarily small regions. Kumar and Goldgof used equilateral triangles centred about chosen pixels and then used the thin plate spline mapping to compute the positions of the triangle corners in the next image. Once point correspondences are established, the motion and deformation parameters can be estimated using standard methods.

Another technique, developed by Young and Axel (1992), fits a 3–D deformable model to the motion of tagged points in two views (short and long axes). Young and Axel used a 16-element finite element model to develop a comprehensive kinematic model of the dynamic geometry of the left ventricle. This method has also been applied to tag stripes (Young *et al.*, 1995) and has been validated using gel phantoms (Kraitchman

et al., 1995). Park *et al.* (1996) developed a further volumetric model that describes complex motion using just a few intuitive motion parameters. This eliminates the need for extensive post-processing before the data is interpreted by a physician.

A similar method, known as the field-fitting approach, was devised by O'Dell *et al.* (1995). One-dimensional displacements are obtained from three independent sets of tag lines, one in the cardiac long-axis and two orthogonal sets in the short-axis. The deformed tag lines in each consecutive slice of the heart are treated as part of a surface that is described by a series expansion. The coefficients of the series expansion are found by using the coordinates of the undeformed and deformed tag lines. This series expansion describes the displacement field, from which the strain can be calculated.

2.2.3 Harmonic phase (HARP) analysis

The main disadvantages of tagging are relatively low spatial resolution as defined by the tag spacing, and the need for some degree of manual intervention for detection of tag lines during magnitude image analysis. These have been eliminated with the development of Harmonic Phase (HARP) analysis, which was first proposed by Osman and Prince (1998). The following description of HARP is largely adapted from Osman and Prince (2000a).

When viewing a tagged image in the Fourier domain, it is seen that the spectral information of the original untagged image is modulated onto a number of spectral peaks. An example of this is shown in Figures 2.8(a) and 2.8(b). These peaks are located at the tag frequency and multiples thereof. By filtering out all but one of the spectral peaks, and then performing an inverse 2–D Fourier transform, a complex image results. The magnitude of this complex image (Figure 2.8(a)) is a crude version of the original untagged image. The phase of this complex image, or the HARP image (Figure 2.8(b)), is intimately related to the tag lines. This phase, which is wrapped from tag line to tag line, is a material property of the tissue as it deforms.

2.2.3.1 HARP images

Let end-diastole be considered to be at time $t = 0$, and let every material point in the heart be labelled by its position $\mathbf{p} = (p_1, p_2)$ at this reference time. As the heart deforms, a material point moves from its reference point to a new spatial position $\mathbf{x}$ at time t. The reference map $\mathbf{p}(\mathbf{x}, t)$ characterizes this motion, where $\mathbf{p}(\mathbf{x}, 0) = \mathbf{p}$. Two dimensional vector notation will be used in this section since HARP analysis can also be performed using spectral peaks in 2–D grid tagging.

It is shown in Appendix E that a SPAMM tagged image is the sum of a number of complex harmonic images, each corresponding to a distinct spectral peak in k-space. The following unified description of the tagged MR image during motion is given by Osman and Prince (2000a)

$$I(\mathbf{p}(\mathbf{x}, t)) = \sum_{n=1}^{N} I_0(\mathbf{p}(\mathbf{x}, t)) C_n(t) e^{j\mathbf{k}_n^T \mathbf{p}(\mathbf{x}, t)} \tag{2.2}$$

were $I_0(\mathbf{p}(\mathbf{x}, t))$ is the image intensity in the absence of the tag pulse sequence, and the coefficients C_n account for the effect of tag fading. The number of spectral peaks N is 3 for 1–D 1–1 SPAMM and 9 for 2–D 1–1 SPAMM. Each of the N complex harmonic images corresponds to a distinct spectral peak identified by the frequency vector $\mathbf{k}_n$ ($\mathbf{k}_n^T$ denotes the transpose of $\mathbf{k}_n$).

The displacement field is defined as the transition from the reference configuration to the current configuration, that is, $\mathbf{u}(\mathbf{x}, t) = \mathbf{x} - \mathbf{p}(\mathbf{x}, t)$. Substituting this expression in Equation 2.2 for a single value of n (i.e. a single spectral peak) yields

$$I_n(\mathbf{p}(\mathbf{x}, t)) = I_0(\mathbf{p}(\mathbf{x}, t)) C_n(t) e^{j\mathbf{k}_n^T \mathbf{x}} e^{-j\mathbf{k}_n^T \mathbf{u}(\mathbf{x}, t)} \tag{2.3}$$

The spectral peaks in the Fourier domain arise because the sinusoids in the SPAMM tag pattern amplitude modulate the underlying image, acting as carriers that shift the underlying spectrum into various positions. From Equation 2.2, the phase $\phi_n(\mathbf{p}(\mathbf{x}, t))$ of a spectral peak $I_n(\mathbf{p}(\mathbf{x}, t))$ is given by $\mathbf{k}_n^T(\mathbf{p}(\mathbf{x}, t))$. This phase is inherently wrapped between $-\pi$ and π, and is described by the nonlinear wrapping function $\phi = mod(\phi^* + \pi, 2\pi) - \pi$, where ϕ^* is the unwrapped phase. Although the spectrum of each $I_n(\mathbf{p}(\mathbf{x}, t))$

is spread throughout the whole of k-space, the motion of the heart is small enough that most of its energy is concentrated around the spectral peak located at $\mathbf{k}_n$. Any rigid rotation of the tag pattern would rotate the position of the spectral peaks about the centre of k-space. Fortunately, the twisting of the heart during systole causes negligible angular motion of the spectral peaks.

A HARP image can be extracted by isolating a spectral peak using a 2–D bandpass filter with dimensions large enough to encompass the maximum spreading of the energy around the peak. Only one peak out of the symmetric pair is chosen to avoid phase cancellation. This estimate will have small errors due to interference from other spectral peaks, as well as due to high frequency information lost in the filtering process. Performing an inverse 2–D Fourier transform of this filtered image results in the complex image $I_n(\mathbf{p}(\mathbf{x}, t))$, the phase of which depicts the position of the tags. Figure 2.8 shows an estimate of the bandpass filter, as well as both magnitude and HARP images of this complex image. Both undeformed and deformed tagged images (Figures 2.8(a) and 2.8(e)) and their corresponding HARP images (Figures 2.8(d) and 2.8(f)) are shown. These images demonstrate that the phase of the tag pattern is a material tissue property and remains constant despite tag fading and intensity variations.

The spectral peaks are located at integer multiples of the fundamental tag frequency. Choosing the peak at the fundamental tag frequency results in a 2π phase wrap for each tag in the HARP image. Selecting peaks for the higher harmonics results in phase wraps that are integer multiples of 2π for each tag in the HARP image. The peak based at the fundamental frequency is chosen for the HARP image primarily because of its high energy content, but also to limit aliasing, which occurs when the motion between frames is greater than the tag spacing.

The HARP method can easily be extended to 2–D tagged images. Here, the phase images for the two tagging directions are determined independently using the first spectral peaks in each tag direction. This is demonstrated in Figure 2.9 for orthogonal 45 degree tagged images. The disadvantage of using 2–D tagged images is that less spectral information is captured using the k-space filter than with 1–D tagged images.

The shape and size of the spectral peak is a function of the tagging pulse sequence, the anatomy present in the image plane, and the motion of the heart. An elliptical bandpass filter with a Gaussian rolloff to eliminate Gibbs ringing has been proposed (Osman and

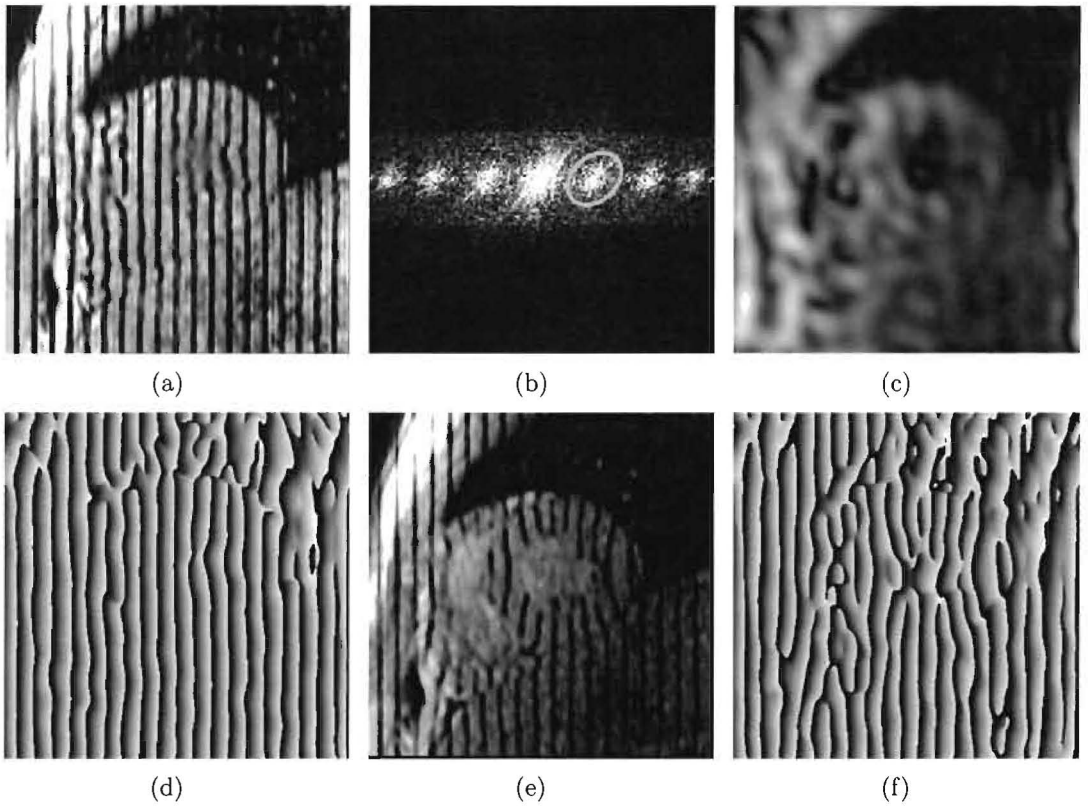

Figure 2.8: (a) A 1–D tagged magnitude image at end-diastole, (b) the corresponding k-space image and bandpass filter, (c) the magnitude image of the inverse 2–D Fourier transform of the filtered image, and (d) the corresponding phase image. (e) The 1–D tagged magnitude image at end-systole, and (f) the corresponding harmonic phase image, showing that the harmonic phase is a material property of the tissue.

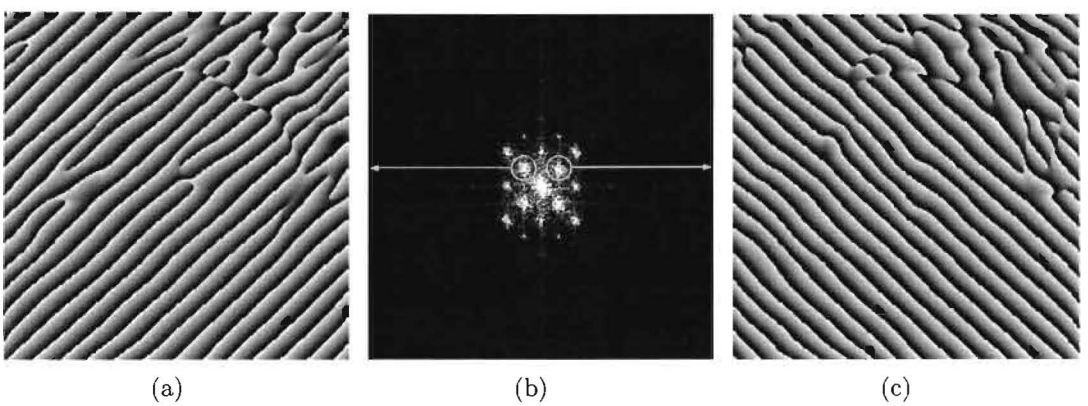

Figure 2.9: Filtering the spectral peaks in a 2–D tagged image.

Prince, 2000a). The use of bandpass filters to extract the spectral peaks suggests that the tracking resolution of HARP strain images is equal to the intrinsic Fourier resolution defined by the filter size. Parthasarathy *et al.* (2004) investigated the achievable dynamic range in terms of cardiac strain and twist. They found that the bandpass filter used

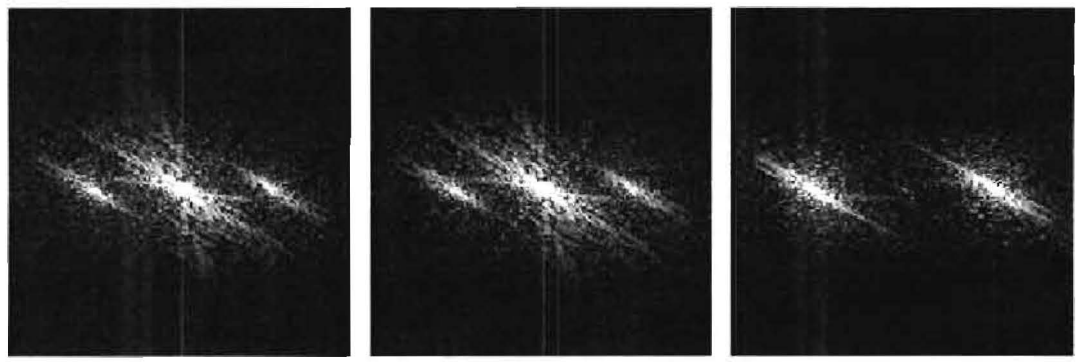

(a) (b) (c)

Figure 2.10: Effect of CSPAMM on k-space magnitude images. (a) k-space for 1-D 1–1 SPAMM, (b) k-space for 1–D 1–1 SPAMM where the sinusoid has been shifted by π radians, (c) CSPAMM k-space image obtained by subtracting (b) from (a). The DC peak is suppressed and the amplitudes of the harmonic peaks are doubled.

in a typical analysis sufficiently accommodated for the physiological limits of strain and twist in a human heart, but that the method is inadequate for dobutamine stress testing. In a further study, Parthasarathy and Prince (2004) investigated the strain step response of HARP, and concluded that the HARP resolution is almost equal to the Fourier resolution, which is of the order of the tag spacing.

Increasing the size of the bandpass filter seems like an obvious solution, but this introduces noise from adjacent spectral peaks. Using complimentary 1–1 SPAMM (CSPAMM) suppresses the DC spectral peak and allows half of k-space to be used for the HARP analysis as shown in Figure 2.10. The complex conjugate spectral peak can also be suppressed by using 3-point phase cycling, as described by Callot *et al.* (2003). This allows the whole of k-space to be used by eliminating the need for any filtering operation, and increases the spatial resolution to that of the pixel size.

HARP images can be used to generate synthetic tag lines by identifying constant-phase contours. Figure 2.11 depicts isocontours for $\phi = \pi$, corresponding to the tag minima. There is, in principle, no limit to the achievable density of tag lines that can be portrayed in this manner. Strain can be calculated directly from HARP images using the gradient of the locally unwrapped phase. This strain relationship is described in Appendix F.

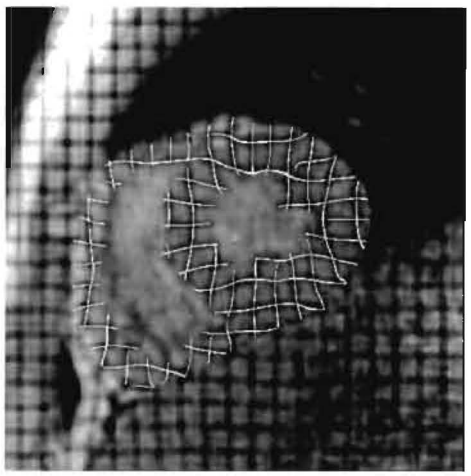

Figure 2.11: Harmonic phase isocontours corresponding to tag minima on a manually segmented image.

2.2.3.2 Measuring small motion from cine HARP images

Suppose that two orthogonal sets of 1–D tagged images are obtained, with tag frequencies $\mathbf{k}_1$ and $\mathbf{k}_2$ respectively. Consider a point $\mathbf{p}$ on the myocardium. The position of this point at time $t = 0$ is again given by the reference map $\mathbf{p}(\mathbf{x}, t)$. The phases ϕ_1 and ϕ_2 for the point $\mathbf{x}$ remain constant with time since the harmonic phase is a material property of the tissue. If it were possible to unwrap $\phi_1(\mathbf{x})$ and $\phi_2(\mathbf{x})$, then each point on the image would be defined by a unique set of phase coordinates, and tracking these points through time would be a straightforward process. There are unfortunately several technical challenges to unwrapping HARP phase images, some of which are mentioned in Section 3.

An alternative to unwrapping the HARP images involves subtracting HARP images of successive frames. Performing this operation, i.e. $\triangle\phi(\mathbf{x}, t_{f'}) = \phi(\mathbf{x}, t_{f+1}) - \phi(\mathbf{x}, t_f)$, where f is the frame number and f' is a frame difference, results in an image that is only wrapped in areas where motion has occurred between frames. Figure 2.12 demonstrates the result of this process. Successive frames of HARP images for vertical tag lines are shown in Figures 2.12(b) and 2.12(a), and for horizontal tag lines in Figures 2.12(f) and 2.12(e). If successive HARP images are subtracted, the wrapping shown in Figures 2.12(c) and 2.12(g) results. Unwrapping such images simply involves using 1–D line integrals that add 2π where the phase difference $\triangle\phi < -\pi$, and subtract 2π where $\triangle\phi > \pi$. These unwrapped images are defined as $\triangle\phi_1^*(\mathbf{x}, t_{f'})$ and $\triangle\phi_2^*(\mathbf{x}, t_{f'})$, respectively. In Figure 2.12(d), light image intensities indicate motion to the left, and dark

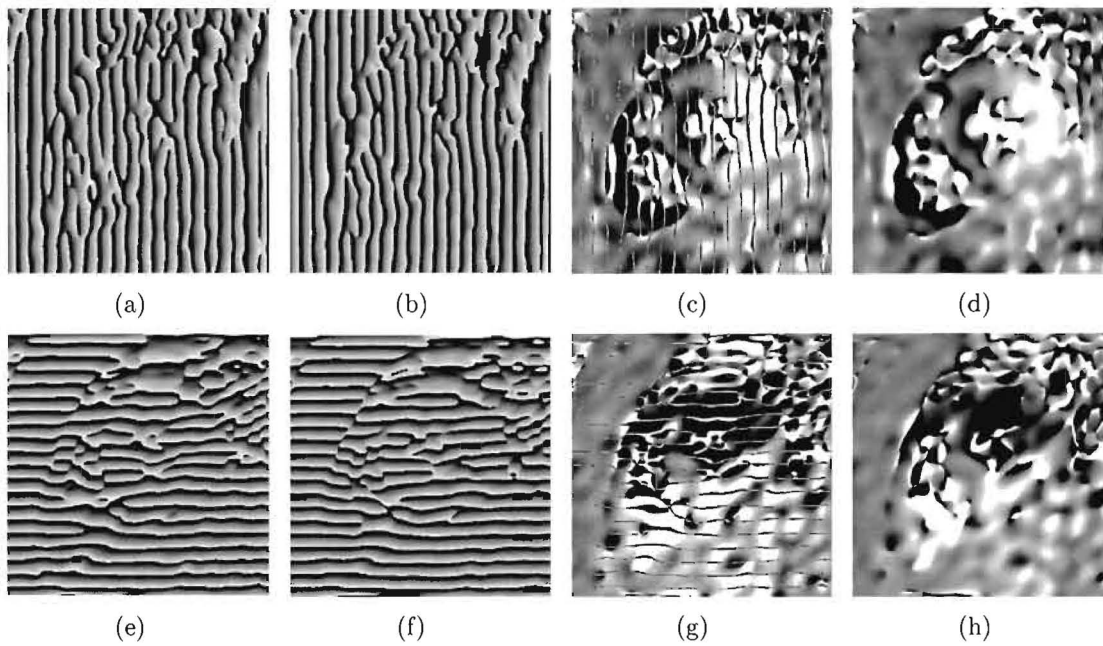

Figure 2.12: Subtraction of HARP images to derive components of displacement. (a) HARP image of frame $f+1$ of a sequence with vertical tag lines; (b) HARP image of frame f of a sequence with vertical tag lines; (c) Wrapped horizontal displacement field obtained by subtracting image (b) from image (a); (d) unwrapped horizontal displacement. Here light intensities imply motion to the left and dark intensities imply motion to the right. (e) HARP image of frame $f+1$ of a sequence with horizontal tag lines; (f) HARP image of frame f of a sequence with horizontal tag lines; (g) Wrapped vertical displacement field obtained by subtracting image (f) from image (e); (h) unwrapped vertical displacement. Here light intensities imply upward motion and dark intensities imply downward motion.

image intensities indicate motion to the right. Similarly, in Figure 2.12(h) light image intensities indicate upward motion and dark image intensities indicate downward motion. The radial contractile action of the left ventricle is clearly evident from these two images.

With reference to the phase term in Equation 2.2, the 2–D displacement $\mathbf{u} = [u_1 \; u_2]^T$ can be calculated according to

$$\mathbf{u}(\mathbf{x},t) = \kappa W^{-1} \begin{bmatrix} \triangle\phi_1^*(\mathbf{x},t) \\ \triangle\phi_2^*(\mathbf{x},t) \end{bmatrix} \tag{2.4}$$

where $W = diag(\mathbf{k}_1, \mathbf{k}_2)$ and κ is a constant scaling factor that relates pixels to millimetres. The displacement field corresponding to the phase components in Figures 2.12(d) and 2.12(h) is shown in Figure 2.13.

Figure 2.13: The HARP inter-frame displacement vector field corresponding to Figures 2.12(d) and 2.12(h).

A limitation of this method of calculating displacement fields is that aliasing will occur if the movement of the myocardium between frames is greater than half the tag spacing. Thus Equation 2.4 only holds if $|\mathbf{k}_n \mathbf{u}(\mathbf{x}, t)| < \pi$, $n = 1, 2$. Using a larger tag spacing is an obvious solution to this problem, but this also has shortfalls. Osman and Prince (2000a) simulated the HARP analysis using a mathematical phantom, and found that the error in the displacement estimate increased almost monotonically as the tag size was increased. This occurs because the spectral peaks in the Fourier domain move closer together as the tag size is increased, thus introducing noise from adjacent spectral peaks, which in turn corrupts the displacement estimate. Note that the measures of Eulerian strain for HARP described in Appendix F are not affected by the limitation $|\mathbf{k}_n \mathbf{u}(\mathbf{x}, t)| < \pi$, $n = 1, 2$.

2.2.4 Displacement encoding with stimulated echoes (DENSE)

DENSE encodes displacement into the phase of the MRI image. The DENSE pulse sequence is shown schematically in Figure 2.14. The sequence commences by position

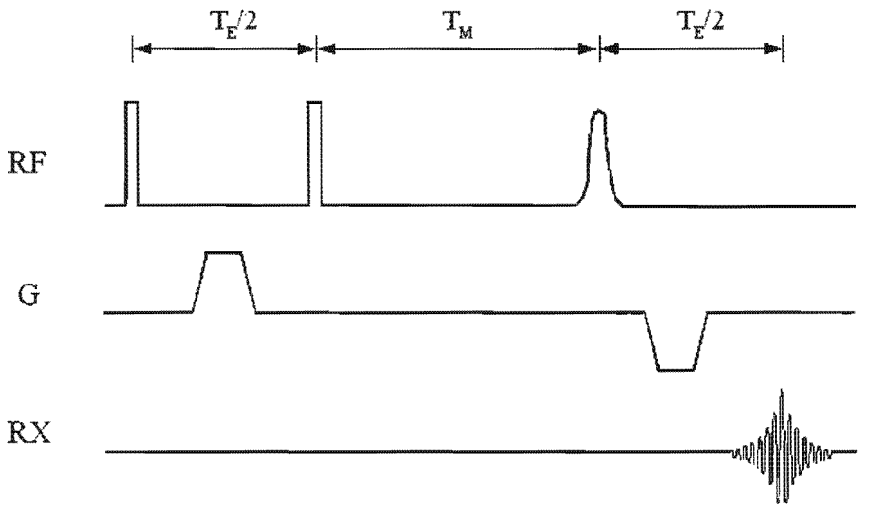

Figure 2.14: Schematic timing diagram of the DENSE pulse sequence. RF - radio frequency pulses; G - magnetic field gradient; RX - receiver coil; T_E - echo time; T_M - mixing time. The two gradient lobes have equal area and opposite signs.

encoding the longitudinal magnetisation using 1–1 SPAMM. After the SPAMM kernel, displacement is accrued as a phase shift during a long mixing time T_M, after which an RF pulse recalls the longitudinal magnetisation into the transverse plane. Here, another gradient lobe (with the same amplitude but opposite polarity as the first) rewinds the phase dispersion caused by the first lobe. Spins that have remained stationary during the mixing time thus have a zero net phase, and spins that have moved during the mixing time accumulate a phase proportional to their displacement.

Cine DENSE allows for displacement encoding in any direction, but 2 or 3 orthogonally encoded directions are typically applied. The following description of 1–D cine DENSE has been adapted from Kim *et al.* (2004). Let spatial position in the imaging slice be described by $\mathbf{x} = (x_1,\ x_2)$ and spatial frequency in the corresponding k-space be described by $\mathbf{k} = (k_1,\ k_2)$. Notation has been simplified throughout this document by using the indices x_1 and x_2 to denote both position and direction. It is also assumed that displacement encoding is always parallel to one of the image axes.

If the flip angles of the two RF pulses in the SPAMM sequence are both 90 degrees, then Equation 2.1 shows that the longitudinal magnetisation for encoding in the x_1 direction is given by

$$M_z(x_1) = M_{z0}\cos(k_e x_1) \tag{2.5}$$

where M_z is the longitudinal magnetisation, M_{z0} is the longitudinal magnetisation prior to application of the displacement encoding pulses, and k_e represents the displacement encoding frequency and describes the phase sensitivity to displacement. The minus sign has been dropped for notational convenience. The magnitude and k-space images of a stationary phantom immediately after the SPAMM kernel are shown in Figures 2.15(e) and 2.15(a), respectively. This corresponds to a normal 1–1 SPAMM tagged image. The effects of $T1$ relaxation at subsequent points in the cardiac cycle can be seen by incorporating the M_z component of the solution to the Bloch equations (Appendix G). This yields the following relationship

$$M_z(x_1, t) = [M_{z0}\cos(k_e x_1) - M_0]e^{-t/T1} + M_0 \tag{2.6}$$

where M_0 is the thermal equilibrium value of the longitudinal magnetisation. The unencoding gradient pulse is applied just before readout, and RF pulses are played out at several points in the cardiac cycle to create the following transverse magnetisation

$$M_{xy}(x_1, t) = \left[[M_{z0}\cos(k_e x_1) - M_0]\, e^{-t/T1} + M_0\right] \cdot e^{-jk_e(x_1+u_1)}\sin(\alpha) \tag{2.7}$$

where α is the flip angle of the RF pulse that recalls longitudinal magnetisation and u_1 is the tissue displacement that occurred in the x_1 direction between the position encoding and the current time t. By applying the relation $\cos(k_e x_1) = (e^{jk_e x_1} + e^{-jk_e x_1})/2$ and rearranging terms, we obtain

$$\begin{aligned} M_{xy}(x_1, t) = &\frac{M_{z0}}{2}\sin(\alpha)e^{-t/T1}e^{-jk_e u_1} \\ &+ \frac{M_{z0}}{2}\sin(\alpha)e^{-t/T1}e^{-jk_e(u_1+2x_1)} \\ &+ M_0\sin(\alpha)(1 - e^{-t/T1})e^{-jk_e(u_1+x_1)} \end{aligned} \tag{2.8}$$

The three terms of Equation 2.8 appear as three distinct echoes in k-space and are named (in order of presentation) the stimulated echo, complex conjugate echo, and $T1$ echo. These echoes are centred in k-space at $\mathbf{k} = (0, 0)$, $(2k_e, 0)$ and $(k_e, 0)$, respectively,

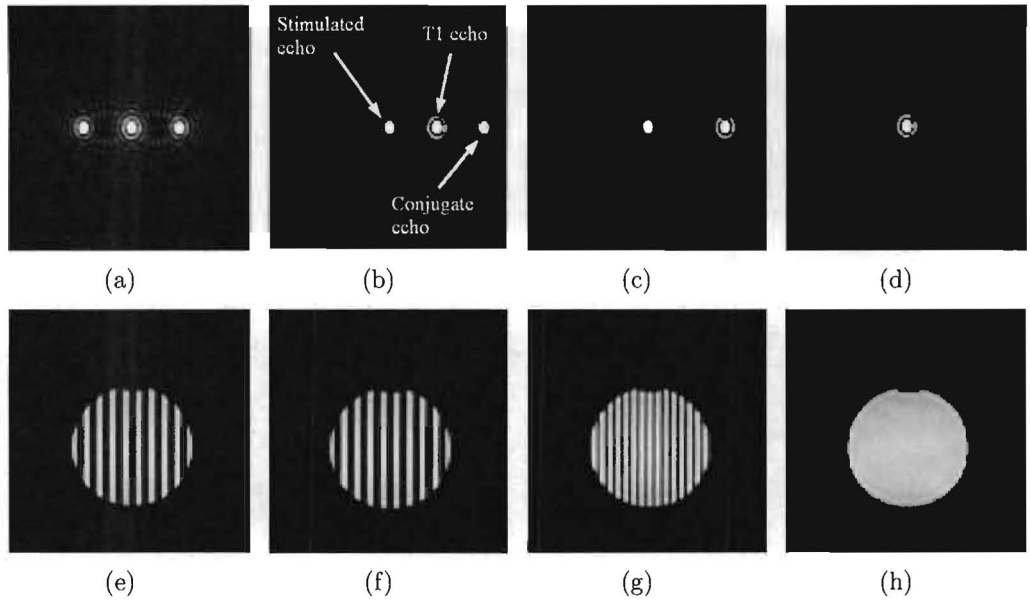

Figure 2.15: Stationary phantom k-space and image space portraying the suppression of the unwanted echoes. (a) and (e) represent k-space and magnitude images after the SPAMM kernel, respectively. (b) and (f) portray k-space and magnitude images after the stimulated echo has been centred by the unencoding gradient. (c) and (g) show the k-space and magnitude images where the displacement encoding frequency k_e is set high enough that the conjugate echo shifts out of the sampled k-space. (d) and (h) show k-space and magnitude images after suppressing the $T1$ echo using CSPAMM. Images courtesy of Professor F. H. Epstein.

and are shown in Figure 2.15(b), with the corresponding magnitude image shown in Figure 2.15(f).

The stimulated echo is the desired DENSE echo, whose phase is directly proportional to the tissue displacement u_1. The stimulated echo has an inherently low SNR, as is evident in the $M_{z0}/2$ factor in Equation 2.8. The other two echoes produce undesirable tag-like artifacts, and can be suppressed using inversion recovery (Aletras and Arai, 1999), through-plane dephasing (Aletras *et al.*, 1999a; Zhong *et al.*, 2005), or various degrees of phase cycling (Kim *et al.*, 2004; Callot *et al.*, 2003; Epstein and Gilson, 2004). Suppression using inversion recovery can only be done at a single time point, and is unsuitable for cine DENSE. The through-plane dephasing suppresses all echoes in a manner proportional to their distance from DC in k-space, but at the expense of SNR. Phase cycling can be applied to selectively isolate any of the echoes (Tsao and Laurent, 2005), simultaneously improving the SNR at the expense of lengthy acquisition times. The tradeoffs between phase cycling and through-plane dephasing for suppressing unwanted DENSE echoes were explored by Bennett *et al.* (2006).

For the data presented in this document, the displacement encoding frequency was set high enough to shift the conjugate echo out of the sampled k-space, as shown in Figures 2.15(c) and 2.15(g). CSPAMM was then used to suppress the $T1$ echo (Kim *et al.*, 2004), as well as to improve the SNR by signal averaging. The tradeoff is that two separate acquisitions are required for CSPAMM. The final DENSE k-space and magnitude images for a stationary phantom are shown in Figure 2.15(d) and 2.15(h).

A series of reference images are acquired for this CSPAMM configuration to compensate for phase shifts caused by background magnetic field inhomogeneities. Using cine DENSE to measure displacement in 2–D thus entails six acquisitions, and can be performed in 19 heartbeats with a 60 ms temporal resolution and 2.7 $\times$ 2.7 mm spatial resolution (Kim *et al.*, 2004). Since the background phase varies slowly compared to the image spatial resolution, low resolution phase reference images are adequate. This was implemented by Zhong *et al.* (2006b), where the imaging time was reduced to 12 heartbeats. The SNR of magnitude-reconstructed cine DENSE images decays with time because of the $T1$ dependence of the stimulated echo, as well as artifacts arising due to growth of the $T1$ echo. As a result, cine DENSE can only reliably capture about two thirds of the cardiac cycle.

Figure 2.16(a) shows the cine DENSE magnitude images at three instants in time. Displacement was encoded in a vertical direction. Figure 2.16(b) shows the corresponding phase-reconstructed images where the phase of each pixel is proportional to tissue displacement, and where white represents π radians and black represents $-\pi$ radians. Since only the final phase accrual of the spins is measured, MRI phase is inherently confined to the range $-\pi < \phi < \pi$. The true phase is thus wrapped to lie within this range. This wrapping is evident in the regions of 12 o'clock and 5 o'clock in the end-systolic phase image in Figure 2.16(b). Using manually defined LV contours and a variation of Itoh's 1–D phase unwrapping method (Itoh, 1980), the unwrapped images in Figure 2.16(c) were obtained. This phase unwrapping method is not suited to automation, as it is very sensitive to noise, and it has to be initiated on a pixel with known absolute phase (typically a pixel that has endured no phase wrapping).

The first term of Equation 2.8 shows that this phase is simply related to displacement in the x_1 direction by the displacement encoding frequency k_e. Vector combination of these unwrapped phase images with a corresponding set of images encoded in the x_2 (horizontal) direction yields the 2–D displacement fields shown in Figure 2.16(d).

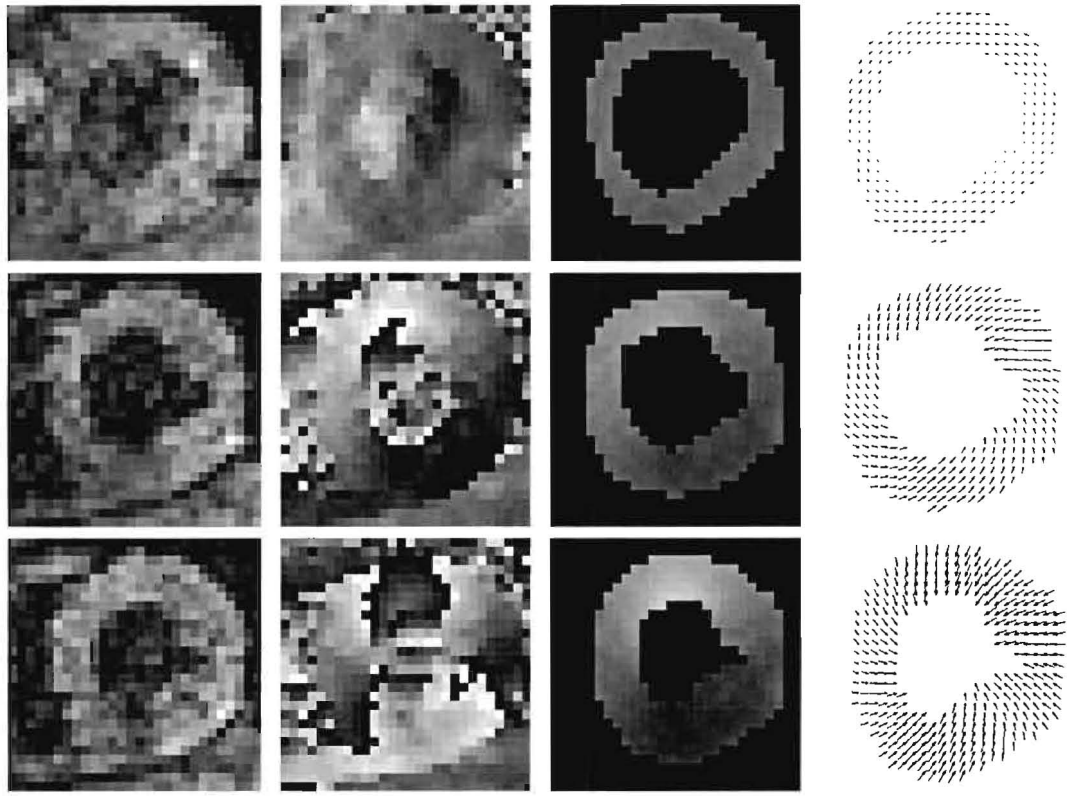

Figure 2.16: (a) Cine DENSE magnitude-reconstructed images at early-systole (top), mid-systole (middle) and end-systole (bottom). (b) Corresponding phase-reconstructed DENSE images encoded for motion in the vertical direction. Field inhomogeneities have already been corrected for using the reference images. (c) Corresponding contoured and unwrapped DENSE phase images. Light pixels indicate downwards motion and dark pixels indicate upwards motion. (d) Corresponding DENSE displacement fields derived from unwrapped phase images encoded in both x_1 and x_2 directions.

While initial DENSE sequences imaged the heart at a single cardiac phase, navigator-guided (Pai and Wen, 2003) and breath-hold (Kim *et al.*, 2004) cine DENSE sequences that trade SNR for temporal resolution have recently been developed. To date, the analysis of DENSE images has occurred on a frame-by-frame basis, and predominantly includes phase unwrapping, conversion of phase to displacement, and deriving strain directly from the displacement fields. The computation of Lagrangian strain is achieved by Kim *et al.* using isoparametric formulation with quadrilateral elements (Moaveni, 1999). It should be noted that localized phase unwrapping is sufficient for this method of quantifying strain, which uses regions of four neighbouring displacement vectors.

2.2.5 Comparison of methods

Phase contrast (PC) velocity encoding measures instantaneous velocity, whereas DENSE measures displacement relative to a fixed time. Measuring displacement from tagged MRI images yields similar results to DENSE, but instantaneous inter-frame displacement measurements are also sometimes made using tagging (e.g. Section 2.2.3.2). The DENSE displacement measurements, which portray the motion of a deformed state into an undeformed grid, are inherently in a Eulerian frame of reference, whereas the translation of an undeformed tagging grid into a deformed state implies a Lagrangian frame of reference. Nevertheless, the strain derivation described for tagging in Section 2.2.2.2 also applies to DENSE data.

Phase contrast MRI offers a higher spatial resolution than both tagging and DENSE methods. A high temporal resolution is also inherent in the PC MRI acquisitions. These virtues are compensated for by imaging time, which is longer for PC than both tagging and DENSE. The PC technique is further complicated by the fact that sampling a rate of change of displacement can adversely affect the true instantaneous displacement estimate, which is required to calculate strain. Although PC does not suffer from tag fading due to $T1$ recovery, it is more susceptible to motion artifacts. These include respiratory artifacts, and phase shifts that result from random motions of ventricular blood. These phase shifts can be largely eliminated by presaturating the surrounding blood, and the respiratory artifacts can be significantly reduced using respiratory navigator control (Masood, 2003).

Phase contrast, DENSE and HARP techniques all suffer from phase distortion due to B_0 inhomogeneities, eddy currents and Maxwell terms. Phase contrast MRI is typically based on gradient echoes whereas DENSE is based on stimulated echoes, so that the SNR of DENSE is inherently less than velocity encoding. HARP only uses one out of the two conjugate echoes, and thus exhibits an effective SNR similar to DENSE. Finally, only a limited number of clinical studies have been carried out using PC myocardial velocity mapping, and it has not received as much attention as myocardial tagging (Masood *et al.*, 2000).

Despite the existence of dedicated software packages for tag analysis using active contours (Axel *et al.*, 1993; Guttman *et al.*, 1994), the analysis is labour intensive and time consuming. A major benefit of the HARP method of tagged image analysis is the short

processing times and high spatial resolution if echo suppression is used. The processing times for strain measurements using DENSE are also fast and automated. HARP evolved as a postprocessing method for myocardial tagging and DENSE evolved from stimulated echo imaging. However, these two techniques are arguably identical, differing mainly in the framework for motion and strain analysis (Kuijer *et al.*, 2006). For example, the HARP phase maps provide a direct measure of Eulerian strain, whereas the DENSE phase maps yield Lagrangian strain.

For a given acquisition time, the spatial resolution achievable with DENSE is less than phase contrast MRI but greater than tagging methods. The spatial resolution for DENSE and HARP is comparable if all of the unwanted echoes are suppressed using phase cycling. Spatial resolution in DENSE and HARP can be traded for temporal resolution by sampling only the portion of k-space enclosing the required spectral peak. This has been implemented by Sampath *et al.* (2001) for HARP analysis, and has been termed FastHARP.

Table 2.1 summarises the advantages and disadvantages of these MRI methods for myocardial motion analysis.

Table 2.1: Comparison of myocardial tagging, phase contrast velocity encoding, HARP and DENSE MRI imaging.

	Advantages	Disadvantages
Tagging	Large displacements can be measured over long periods. Images suitable for qualitative analysis. Large number of clinical case studies have been done. 2–D displacement can be measured from a single series.	Spatial resolution limited to tag spacing. 3–D analysis is laborious and time consuming. Contrast degradation occurs due to tag fading. Manual myocardial segmentation required. Poor contrast between blood and endocardial borders.
Phase contrast	High spatial resolution. High temporal resolution. Processing software exists for all major manufacturers. Straightforward postprocessing. Measures velocity over the entire cardiac cycle.	Lengthy acquisition times. Manual myocardial segmentation required. Poor contrast between blood and endocardial borders. Phase distortion due to eddy currents and Maxwell terms. Measuring the derivative of displacement may give a poor displacement estimate. Prone to motion artifacts and may require respiratory navigator control.
DENSE and HARP	Large displacements can be measured over long periods. Good contrast between blood and endocardium (DENSE). Straightforward postprocessing. Resolution greater than tagging resolution.	Low SNR associated with the selected echo. Few clinical case studies have been done. Manual myocardial segmentation required. Contrast degradation due to regrowth of the $T1$ echo.

Chapter 3

Absolute cine DENSE displacement measurements using spatiotemporal phase unwrapping

3.1 Introduction

It is possible to set the displacement encoding frequency k_e to be small enough to ensure that no phase wrapping occurs for physiological limits of cardiac motion. Although lowering the value of k_e increases the SNR (Epstein and Gilson, 2004), it also has an associated decrease in sensitivity to displacement. Experience has shown that practical values of k_e makes phase wrapping inevitable (Aletras *et al.*, 1999b; Kim *et al.*, 2004).

Two dimensional phase unwrapping has its roots in synthetic aperture radar and optics (Ghiglia and Pritt, 1998), but has also found a number of applications in MRI. These include magnetic field mapping (Jezzard and Balaban, 1995), chemical shift mapping (Glover and Schneider, 1991), and phase contrast velocity encoding (Yang *et al.*, 1996). Two dimensional phase unwrapping has been used previously for processing DENSE images (Aletras *et al.*, 1999b), but with predefined contours to isolate the myocardium. This Chapter presents the first reported spatiotemporal phase unwrapping application for cine DENSE images, and the algorithm is robust enough to operate successfully without predefined myocardial contours.

3.2 Theory

Consider a pixel with 2–D indices i and j, phase ϕ_{ij} and a wrapping operator W where $W(\phi_{ij}) = mod(\phi_{ij} + \pi, 2\pi) - \pi$. The operator mod is the congruence modulus. This can also be represented as $W(\phi_{ij}) = \phi_{ij} + 2k_{ij}\pi$, where k_{ij} is an integer chosen so that $W(\phi_{ij}) \in (-\pi, \pi]$. The phase unwrapping problem involves solving k_{ij} for each pixel. Phase unwrapping can be done along a particular path by integrating the differences, $\Delta\phi$, of the locally unwrapped phases. This is equivalent to integrating the phase gradient $\nabla\phi$, and is embodied in the line integral

$$\phi(\mathbf{r}) = \int_C \nabla\phi(\mathbf{r}) \cdot d\mathbf{r} + \phi(\mathbf{r}_0) \tag{3.1}$$

where C is any path in the image domain connecting the points $\mathbf{r}_0$ and $\mathbf{r}$. Note that $\Delta\phi$ is constrained to lie between $-\pi$ and π to avoid aliasing and hence ensure correct unwrapping. This unwrapping process can be implemented in 2–D by using a region-growing algorithm. Although phase unwrapping is conceptually straightforward, it can in practice be difficult to implement. Errors introduced due to image noise are inclined to propagate with devastating results. Fortunately, there are ways of identifying problematic features in phase images.

For a phase surface where discontinuities are solely a result of the wrapping process, the integral of the phase gradient over any closed loop is zero. Figure 3.1(a) shows a simple 2–D phase surface, where black represents $-\pi$ radians and white represents π radians. The phase surface in Figure 3.1(b) contains a single phase wrap. In both cases the integral of the phase differences over any closed loop in the image is zero.

In cine DENSE images, 2–D phase unwrapping is complicated by image noise, phase shear and the fact that only a few pixels span the myocardial walls. Phase shear occurs where regions of adjacent tissue slide across each other in opposite directions. The integration of the phase gradient over a path encompassing a 2–D domain of interest thus often yields nonzero values, implying that the unwrapped phase is dependent on the path followed. These inconsistencies can be ascribed to phenomena known as phase residues. Two examples of phase residues are shown in Figures 3.1(c) and 3.1(d). Here the line integral of the phase gradient around the centre of the image accumulates a phase of $\pm 2\pi$ radians. It can thus be concluded that a residue is contained within a

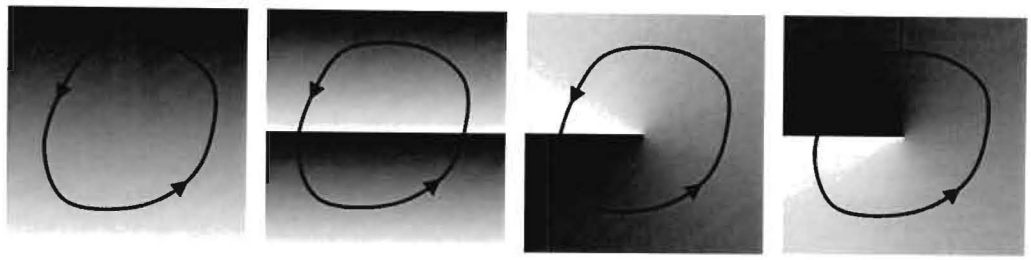

(a) (b) (c) (d)

Figure 3.1: Simple closed line integrals on (a) a non-wrapped phase surface, (b) a phase surface containing a single phase wrap, (c) a positive phase residue, and (d) a negative phase residue. Black represents $-\pi$ radians and white represents π radians.

closed path if the accrued phase is $\pm 2\pi$ radians. The sign of this phase denotes the *charge* of the residue and the sum is calculated in a counter clockwise direction around the loop by convention. Note that it is possible, in more complicated residues, to accrue a phase that is an integer multiple of 2π radians. These scenarios are unlikely, and it is generally assumed that only residues of the form shown in Figure 3.1 are encountered.

Extending this supposition to include a number of residues within the closed path yields the residue theorem, which is given by

$$\oint \nabla\phi(\mathbf{r}) \cdot d\mathbf{r} = 2\pi \times (sum\ of\ enclosed\ residue\ charges) \tag{3.2}$$

where $\mathbf{r}$ is the length of the integration path and $\nabla\phi$ is the gradient of ϕ. This states that the closed-path integral around a number of phase residues will equal some integer multiple of 2π radians. When the residue charges are balanced within a region, the line integral around that region is zero for any path taken. Consistent phase unwrapping is thus possible if and only if all integration paths do not encircle unbalanced residue charges.

Residues are detected in practice by integrating around each 2×2 block in the image. By convention, even though the residue in fact exists between the pixels, the upper left pixel is chosen to designate the residue's position. Figure 3.2(a) shows a HARP image, and the corresponding residues are shown in Figure 3.2(b). The white points depict positive residues and the black points depict negative residues. The two arrows illustrate a residue dipole, which is a positive and negative residue pair occurring as a result of a discontinuous region of phase wrapping. Figure 3.2(c) shows the streaking effect of the

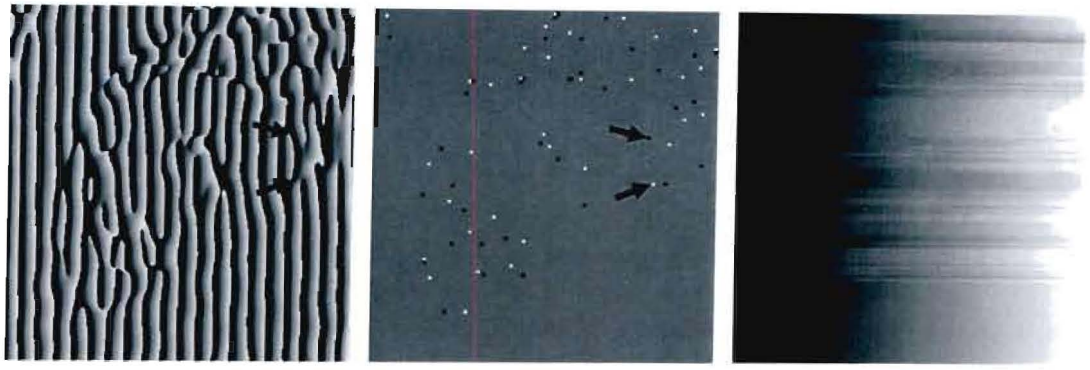

(a) (b) (c)

Figure 3.2: (a) Wrapped phase image. (b) Corresponding phase residues where white dots correspond to positive residues and the black dots correspond to negative residues. The arrows illustrate a residue dipole. (c) Itoh's one dimensional phase unwrapping technique produces streaks that emanate from the residues and continue in the direction of the integration paths.

residues when Itoh's 1–D phase unwrapping technique (Itoh, 1980) is applied. Itoh's technique unwraps phase by integrating the locally unwrapped phase in 1–D.

There are two main methods for 2–D phase unwrapping: path following and minimum L^p-norm methods (Ghiglia and Pritt, 1998). Path following methods involve unwrapping the phase along a path that depends on some measure of phase consistency. Minimum L^p-norm methods apply a more global approach, and seek to minimize the integral of the differences, in an L^p-norm sense, between the gradients of the solution, and those of the locally unwrapped original phase values.

The minimum L^p-norm methods are typically more robust than path following methods, but are also more computationally expensive. In the interest of efficiently unwrapping through both space and time (see Section 3.3), only the path following methods were investigated. Path following methods lend themselves well to cardiac images, where the myocardium is a continuum and generally provides a reliable path of integration. Two common 2–D path following methods are Goldstein's branch cut algorithm, and the phase quality-guided algorithm.

3.2.1 Goldstein's branch cut algorithm

A branch cut is a barrier that the phase unwrapping path of integration may not cross. Branch cuts are typically placed between residues, or between residues and the image

border. Figure 3.3(a) depicts a wrapped phase image with a positive (+) and a negative (-) residue. Paths 1 and 3 from point A to B (each crossing from $-\pi$ to π five times) result in the same phase value at B, whereas path 2 (crossing from $-\pi$ to π four times) yields a different result. If a branch cut is placed between the two residues, as shown in Figure 3.3(b), then any path of integration between points A and B will yield the same phase value. The act of connecting residues of opposite polarity is known as residue charge balancing. Alternatively, if branch cuts are placed such that they connect the two residues to the edge of the image (Figure 3.3(c)), then the integration also becomes path independent.

Although these two choices of branch cut placement result in path-independent phase unwrapping, they give different results for the estimated phase in the upper half of the image. The choice of branch cuts is often not obvious, and some criteria must be established to guide the branch cut placements. The branch cuts must serve not only to balance the residue charges, but must also be placed such that important regions of the wrapped phase images are not isolated and hence made inaccessible to the unwrapping paths of integration. An obvious criterion for connecting residues is to use branch cuts that are as short as possible. This is suitable for some applications, but is generally impractical due to the computational power required to examine all possible choices of cuts in order to find the shortest. Typical cropped DENSE phase-reconstructed images contain several hundred residues.

Many techniques have been proposed to optimise the branch cut placement. The Goldstein method is the most commonly used, and has proven to be both effective and computationally efficient. The method "clumps" multiple residue pairs in a manner that attempts to minimise the sum of the branch cut lengths as described below.

A set of matrices depicting all residue positions and charges is defined. For each of the unbalanced residues, the following algorithm is followed: A 3×3 pixel box is centred around the residue and its perimeter is searched for other residues. If a residue is found, a branch cut is placed between the residues, even if they have the same polarity. If the residues have opposite polarities, they are designated as "balanced" by the branch cut, and the search starts again on the next unbalanced residue in the list. If, however, the residues have the same polarity, the search within the 3×3 box continues. Whenever a new residue is encountered, it is connected by a branch cut to the residue at the centre of the box. This connection is performed even if the new residue has previously been

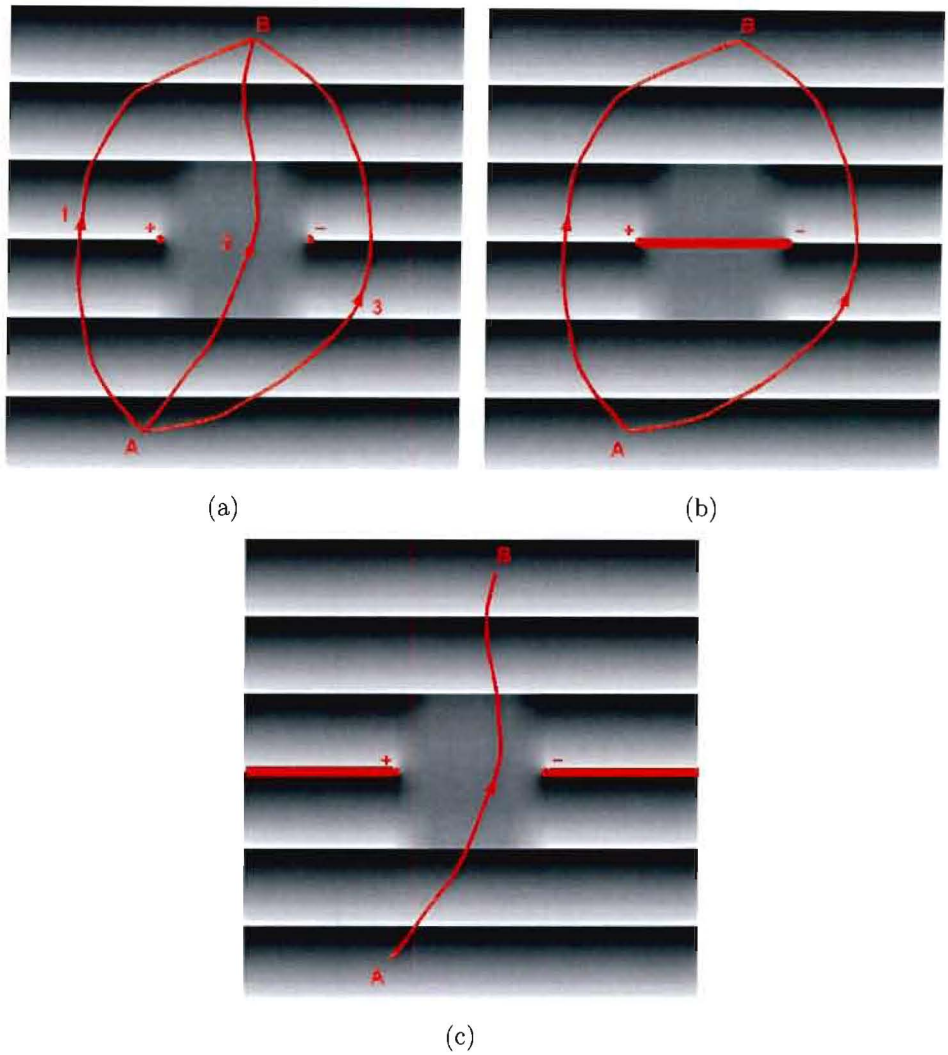

Figure 3.3: Path dependent phase unwrapping in the presence of two residues of opposite polarity. (a) Starting at point A, paths 1 and 3 will result in the same phase value at point B, whereas the phase value predicted using path 2 will be different. (b) If a branch cut is placed between the two residues, then any possible integration path between A and B will yield an identical phase result. (c) Alternatively, if the two residues are connected to the image border by branch cuts, integration between A and B along any path will also yield consistent phase results.

connected to some other residue. If the newly connected residue has not been connected to another residue in another cluster, then its polarity (-1 or 1) is added to the sum of the polarities of the other residues in the box. If, however, the residue has already been connected to other residues, its polarity is not added since it has already been accounted for in another residue cluster.

If the cumulative charge is nonzero when the search of the 3×3 box ends, then the box is centred at each of the other residues (within the initial box) in turn and searched.

Any residues found are again connected to the initial central residue. If the cumulative charge is still nonzero, then the box is enlarged to 5×5 pixels, and is centred around each of the residues in turn, starting with the initial residue. This process of enlarging and repositioning the box continues until either the net charge is balanced, the box touches the image border, or the box reaches a maximum pre-defined size. For all of these cases, all residues acted upon are designated as balanced, and the search continues from the beginning on the next unbalanced residue. If the box touches the image border, then a branch cut is placed to the border. This branch cut thus prevents any path integral from encircling the residues.

The path integration takes the form of a "flood-fill" algorithm, which acts in analogy to a flood of liquid filling the rooms of a house, where the branch cuts are the various walls of the house. The algorithm keeps track of a list of pixels which adjoin (or border) the unwrapped pixels. A starting pixel is selected and its phase is stored as the unwrapped value in a solution matrix. The four neighbouring pixels are then placed in an "adjoin" matrix, which is a list of wrapped pixels that neighbour unwrapped pixels. These four neighbours are unwrapped, the unwrapped values are placed in the solution matrix, and the adjoin matrix is updated. The algorithm then proceeds by selecting a pixel in the adjoin list, unwrapping it, adding it to the solution matrix, removing it from the adjoin list, and adding its wrapped neighbours to the adjoin list. Unwrapping may not use the branch cut pixels. Eventually, the adjoin list becomes empty, which means that either all the pixels have been unwrapped, or there is a region of the image that is isolated by the branch cut. After all non-branch cut pixels have been unwrapped, the branch cut pixels that border unwrapped pixels can be unwrapped. This can be done because, strictly speaking, the branch cuts lie between the pixels and not on them.

Figure 3.4 gives an example of Goldstein's phase unwrapping algorithm applied to a DENSE phase-reconstructed image. Here the image size is 40×40 pixels, and the maximum box size is 5×5 pixels. Note the phase discontinuity depicted by the arrow in the right ventricle in Figure 3.4(d). This effect is likely if a branch cut spans the width of the ventricular wall, as is the case here (arrow in Figure 3.4(c)). If the maximum box size is reduced, then fewer residue clusters are likely to be balanced, and vertical and horizontal unwrapping artifacts akin to those shown in Figure 3.2(c) result. On the other hand, enlarging the maximum box size increases the chance of isolating important portions of the image, and of creating branch cuts that span the widths of the ventricular walls.

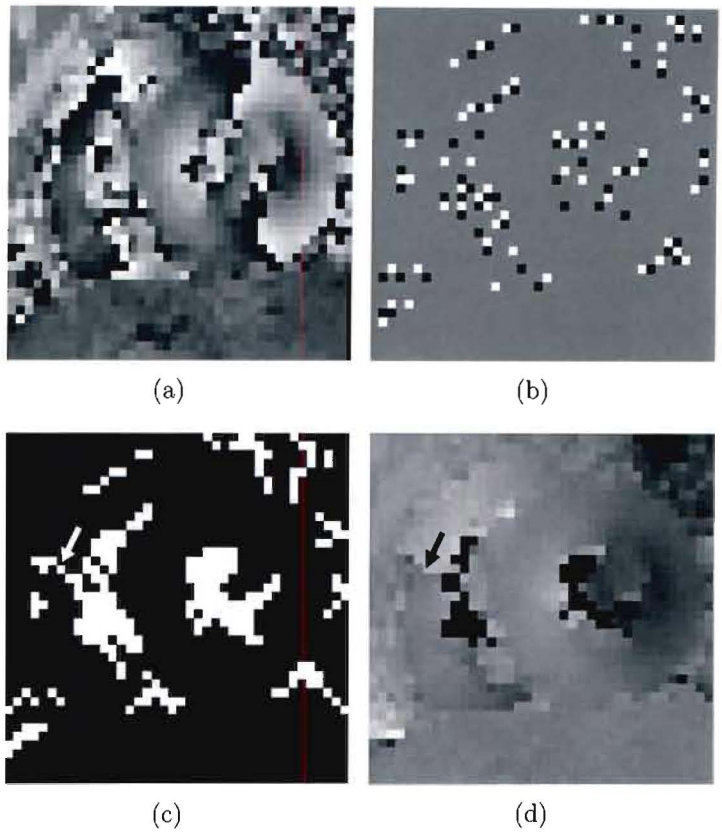

Figure 3.4: Performance of Goldstein's phase unwrapping method on DENSE images. (a) Wrapped DENSE phase image depicting motion in the horizontal direction, (b) corresponding phase residues, (c) corresponding branch cuts, and (d) unwrapped phase. The arrow depicts an erroneous phase discontinuity in the right ventricle. The black pixels depict regions that were isolated by the branch cuts.

3.2.2 Quality-guided path following

The quality-guided path following phase unwrapping method uses a measure of phase "quality" to guide the path of unwrapping (Ghiglia and Pritt, 1998). The approach does not identify residues, and simply assumes that the path of integration will not encircle any unbalanced residues. Although this assumption seems theoretically unsound, the method is surprisingly robust in practice.

A good measure of phase quality is a root-mean-square measure of the variances, within a local $n \times n$ pixel region, of the partial derivatives in the x_1 and x_2 directions (Ghiglia and Pritt, 1998). This is given by

$$\rho_{pq} = \frac{\sqrt{\sum(\Delta_{i,j}^{x_1} - \overline{\Delta}_{p,q}^{x_1})^2} + \sqrt{\sum(\Delta_{i,j}^{x_2} - \overline{\Delta}_{p,q}^{x_2})^2}}{n^2} \tag{3.3}$$

where for each sum the indexes (i, j) range over the $n \times n$ window centred at the pixel (p, q). The terms $\Delta_{i,j}^{x_1}$ and $\Delta_{i,j}^{x_2}$ are the partial derivatives of the locally unwrapped phase in the x_1 and x_2 directions, respectively, and the terms $\overline{\Delta}_{p,q}^{x_1}$ and $\overline{\Delta}_{p,q}^{x_2}$ are the averages of these partial derivatives in the $n \times n$ windows.

Since only a few pixels span the myocardial walls in DENSE images, only the variances with respect to the four nearest neighbours were calculated. As with Goldstein's method, the unwrapping path is guided by a flood-fill algorithm, but this variation floods into regions of high quality phase before those with lower quality. The quality-guided flood-fill algorithm can be summarized as follows:

1. A measure of phase quality for each pixel is calculated using Equation 3.3.
2. A starting point in the region of interest is selected. This is stored in the solution matrix.
3. The four pixels adjacent to the starting point are placed in an adjoin matrix, which keeps track of wrapped pixels with adjacent unwrapped pixels.
4. The pixel in the adjoin matrix with the highest phase quality is selected and unwrapped using an adjacent unwrapped pixel. This pixel is removed from the adjoin matrix and added to the solution matrix.
5. The new wrapped nearest neighbours are included in the adjoin matrix.
6. Steps 4 and 5 are repeated until the adjoin matrix is empty.

An example demonstrating the algorithm on a DENSE phase-reconstructed image is shown in Figure 3.5. Figure 3.5(a) is the same wrapped DENSE phase image as shown in Figure 3.4(a), and Figure 3.5(b) is the corresponding phase quality map obtained using Equation 3.3. Figure 3.5(c) shows the unwrapped phase image, and Figures 3.5(d)–(g) capture a few stages of the flood-fill procedure, where the white pixels represent pixels in the adjoin matrix. As with Goldstein's method, a phase unwrapping error occurs in the right ventricle (arrow in Figure 3.5(c)). This is partly due to poor phase quality in this region, and partly to phase shear between the right ventricle and the adjacent liver.

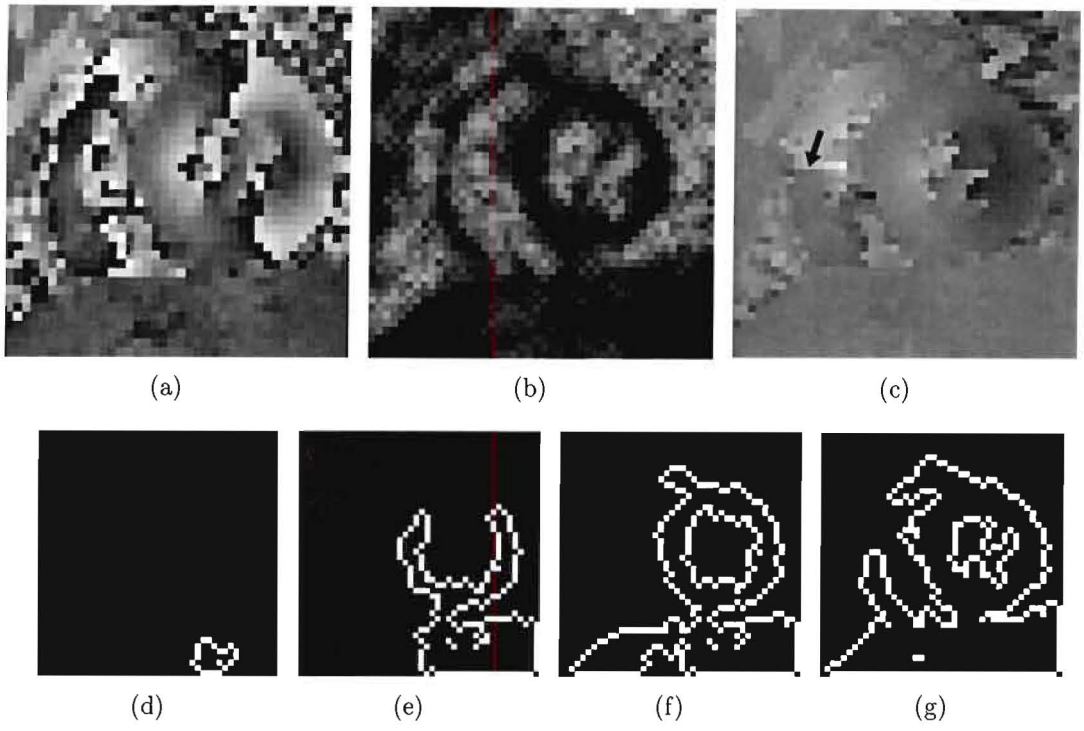

Figure 3.5: Performance of the quality guided path following phase unwrapping method on a DENSE image. (a) Wrapped DENSE phase-reconstructed image depicting motion in the horizontal direction. (b) Phase quality map where dark pixels represent high phase quality. (c) Unwrapped phase image using quality-guided path following. (d)-(g) Progression of the guided flood-fill operation using the phase quality map in (b). The calculation intervals are arbitrary, and the white pixels represent pixels in the adjoin matrix.

3.3 Methods

Obtaining true measures of displacement using cine DENSE requires an absolute measure of unwrapped phase. Disregarding unwrapping errors due to noise, the problem of obtaining absolute phase involves correctly estimating the phase of the starting seed point for the unwrapping algorithm ($\mathbf{r}_0$ in Equation 3.1). Unless the phase unwrapping integration commences at a point with a correct or known phase, the entire unwrapped image will be offset by the amount of phase error in the initial estimate. In the case where the seed point is placed on the myocardium, or another region of high quality phase, this error will be an integer multiple of 2π. The selected starting points in Figures 3.4 and 3.5 were manually chosen somewhere on the liver, since it lies adjacent to the LV and often provides a path of consistent phase into the LV. This is however not a robust solution since phase shear commonly exists between the liver and the LV, and phase wrapping often occurs in the inferior LV.

A method of ensuring that absolute phase reference seed points are present on each frame of a cine DENSE series was developed using spatiotemporal phase unwrapping. Very little cardiac motion takes place between the displacement encoding and the acquisition of the first frame of cine DENSE data. It is therefore characteristic that no phase wrapping will have occurred on the first frame, so that all phase measurements in the myocardium on this frame will be absolute. A measurement of absolute phase for each frame can thus be obtained by selecting any seed point on the myocardium on the first frame, and unwrapping the phase through space and time along any path restricted to the myocardium.

Goldstein's method is ill-suited to DENSE data since only a few pixels span the width of the myocardium, making it likely that branch cuts will be placed transmurally and prevent portions of the myocardium from being unwrapped. Goldstein's method was applied to 95 DENSE phase-reconstructed images. For maximum box sizes of 3×3, 5×5 and 7×7 pixels, the LV was correctly unwrapped for 70%, 81%, and 78% of images, respectively. Although a 3–D extension of the branch cut method does exist (Huntley, 2001) and has been successfully applied to phase contrast velocity encoded images (Salfity *et al.*, 2005), the results from Goldstein's 2–D implementation on DENSE images discouraged the use of this method for spatiotemporal phase unwrapping. Initial results with the 2–D quality-guided path following method were more promising, making this the method of choice.

3.3.1 Fully-automated spatiotemporal phase unwrapping algorithm

Cine DENSE phase images were unwrapped by extending the 2–D quality guided flood-fill algorithm to 3–D (one temporal and two spatial dimensions) and by commencing the unwrapping on all pixels at the first cardiac phase. The extension of the phase quality map to a third dimension (x_3 or time) is straightforward, and is given by

$$\rho_{pqr} = \frac{\sqrt{\sum(\Delta^{x_1}_{i,j,k} - \overline{\Delta}^{x_1}_{p,q,r})^2} + \sqrt{\sum(\Delta^{x_2}_{i,j,k} - \overline{\Delta}^{x_2}_{p,q,r})^2} + \sqrt{\sum(\Delta^{x_3}_{i,j,k} - \overline{\Delta}^{x_3}_{p,q,r})^2}}{n^3} \tag{3.4}$$

where for each sum the indexes (i, j, k) range over the $n \times n \times n$ window centred at the pixel (p, q, r). The terms $\Delta^{x_1}_{i,j,k}$, $\Delta^{x_2}_{i,j,k}$ and $\Delta^{x_3}_{i,j,k}$ are the partial derivatives of the locally unwrapped phase in the x_1, x_2 and x_3 directions, respectively. The terms $\overline{\Delta}^{x_1}_{p,q,r}$, $\overline{\Delta}^{x_2}_{p,q,r}$ and $\overline{\Delta}^{x_3}_{p,q,r}$ are the averages of these partial derivatives in the $n \times n \times n$ windows. The 3–D flood-fill algorithm is identical to the 2–D algorithm described earlier, except that the solution and adjoin matrices are 3–D, and six rather than four adjacent pixels are examined. Although this method is completely automated, the processing times can be prohibitively long. This is largely due to the number of pixels in the adjoin matrix, that is, the search space for the highest phase quality pixel that lies adjacent to an unwrapped pixel.

3.3.2 Combination of semi-automated spatiotemporal and two dimensional spatial phase unwrapping algorithm

Although the effectiveness of the fully-automated spatiotemporal phase unwrapping method was investigated in this study, a more computationally efficient semi-automated approach is also presented. The proposed algorithm is as follows:

1. Calculate 2–D phase quality maps for all frames.
2. Place a seed point in the myocardial mid-wall region on the first frame. This is the only manual interaction required.
3. Create correctly unwrapped seed points on the rest of the frames by rapidly threading an unwrapping path through time, allowing sufficient spatial unwrapping to follow the myocardium as it moves. This is achieved using 3–D quality-guided phase unwrapping, and by limiting the spatial unwrapping to a specified radius per frame. As soon as an unwrapped pixel appears on a particular frame, all unwrapping for all previous frames is stopped.
4. Once unwrapped pixels are present on all frames, 2–D quality-guided phase unwrapping is performed for each frame independently, commencing at the correctly unwrapped reference pixels.

The 3–D phase unwrapping was achieved here by stacking 2–D phase quality maps (Equation 3.3) for a cine series on top of each other. This is reasonable because each

frame of the cine series is an independent acquisition. No improvement in the performance of the algorithm was found if full 3–D phase quality maps were used. If myocardial contours have been defined, the seed point placement can be automated by selecting any mid-wall myocardial pixel with a high corresponding phase quality. The radius size in Step 3 was set based on the maximum expected amount of cardiac displacement per frame. It should be large enough to allow the 3–D unwrapping to migrate spatially with the myocardium between frames, but small enough to limit the processing time and to restrict the 3–D unwrapping to relatively high values of phase quality. A radius of 2 pixels was used here.

3.3.3 Quantifying phase unwrapping effectiveness

A useful way of estimating where phase unwrapping errors have occurred involves combining the phase quality map with a discontinuity map (Ghiglia and Pritt, 1998). The discontinuity map specifies the locations of pixels in the unwrapped image that differ from a neighbour by more than $|\pi|$ radians. Unwrapping errors are likely to have occurred within regions isolated by the combined map.

Both the fully- and semi-automated phase unwrapping algorithms were applied to 16 sets of cine DENSE images, each with 25 frames, k_e=0.1 - 0.15, and displacement encoding applied in the x_1 and x_2 directions. The data represents a population of patients with infarcts and normal volunteers. No myocardial contours were used, and the images were cropped to about twice the end-diastolic epicardial diameter. The combined phase quality and discontinuity maps were used to aid a visual inspection of unwrapping errors in the LV. If a single error was encountered, then the image was deemed incorrectly unwrapped.

3.3.4 Optimising k_e

Equation 2.8 shows that phase sensitivity to displacement of the stimulated echo is dependent on the displacement encoding frequency k_e. Increasing k_e increases the phase sensitivity to displacement, but at the expense of decreased SNR and increased phase aliasing. A tradeoff thus exists in the interest of finding a maximum practical phase sensitivity to displacement. The effects of varying k_e on SNR and phase unwrapping

were investigated by acquiring mid-ventricular short-axis data from a single volunteer with k_e = 0.02, 0.06, 0.10, 0.14, 0.18, 0.22, 0.26 and 0.30 cycles/mm. The scans were performed on a 1.5 T Siemens Avanto MRI scanner (Siemens Medical Systems, Erlangen, Germany) using a four-channel body phased array RF coil and the following imaging parameters: FOV = 360 × 248 mm; echo train length (ETL) = 9; effective TE = 9.82 ms; TR = 20 ms; matrix size = 128 × 75 pixels; pixel size = 2.81 × 2.81 mm; slice thickness = 8 mm; and cardiac phases = 33. A bottom-up flyback interleaved k-space trajectory was used to remove ghosting artifacts (Kim *et al.*, 2003).

In order to remove the effects of the conjugate echo, which moves into the sampled k-space for low values of k_e, the CANSEL phase cycling method (Epstein and Gilson, 2004) was applied[1]. This increased the breath hold time, and scans for the two encoding directions were acquired in two breath holds of 22 heart beats each. The effectiveness of both fully- and semi-automated phase unwrapping methods with manually-drawn contours was investigated with reference to k_e. The SNR in the myocardium as a function of k_e and time was also investigated. The mean myocardial signal per frame was defined using manually defined contours, and the noise level was defined as the standard deviation of a sample of background noise.

3.4 Results

The fully-automated method correctly unwrapped 767 out of the 800 images (95.9%), and the semi-automated method correctly unwrapped 786 out of the 800 images analysed (98.3%). The semi-automated method was 5 - 9 times faster than the fully-automated method.

Magnitude-reconstructed end systolic images for a single volunteer with varying values of k_e are shown in Figure 3.6. The SNR becomes significantly lower in the inferior LV wall and the majority of the RV as k_e is increased. This signal loss is attributed to intravoxel dephasing (Fischer *et al.*, 1995), and is observed in DENSE because incomplete intravoxel rephasing occurs when tissue deforms. This effect is more pronounced for higher values of k_e. The presence of signal in the blood in Figure 3.6(a) is because k_e is too low for even the motion of blood to cause significant intravoxel dephasing during

[1]This was implemented by Xiaodong Zhong in the Siemens IDEA and ICE environments.

the displacement encoding period. The signal in the blood for $k_e = 0.02$ cycles/mm only washes out during diastole.

The phase unwrapping errors as a function of k_e for both semi- and fully-automated phase unwrapping methods are shown in Figure 3.7. The percentage of incorrectly unwrapped frames increases dramatically for $k_e > 0.1$ cycles/mm, and largely occur due to the regions of low SNR associated with intravoxel dephasing.

Mean myocardial SNR as a function of time is shown in Figure 3.8 for the various values of k_e. As expected, intravoxel dephasing causes the SNR to decrease as k_e is increased. The reason for the SNR spike at frame 3, particularly for the higher values of k_e, is not clear. Given that frame 3 corresponds to early-systole and that this effect is a function of k_e (and hence displacement), an association with myocardial twist or higher order motion such as acceleration could be speculated. The effects of intravoxel dephasing decrease as the myocardium relaxes during diastole. The corresponding SNR increase is greater than the $T1$-related signal decay, explaining the small SNR peak around frame 13.

3.5 Discussion

A branch cut algorithm more suited to DENSE images than Goldstein's method is presented in Buckland *et al.* (1995). Here residue dipoles are joined using branch cuts, and the total length of these branch cuts is minimized in a least squares sense. Unfortunately, it is not straightforward to extend this algorithm to 3–D. If the manually delineated LV borders are treated in the same way as the image borders for Goldstein's method, then branch cuts can be connected to these borders to balance any residual charge. For the 96 DENSE images investigated using Goldstein's method, 100% were correctly unwrapped if predefined borders were used.

For 800 DENSE images without predefined myocardial contours, the proposed fully-automated spatiotemporal quality-guided phase unwrapping method was 95.9% successful, and the semi-automated method was 98.3% successful. Based on efficacy and computation speed, the semi-automated method is used for all remaining studies in this manuscript. The increased number of unwrapping errors using the fully-automated method stems mostly from isolated unwrapping errors propagating through time. Step 3

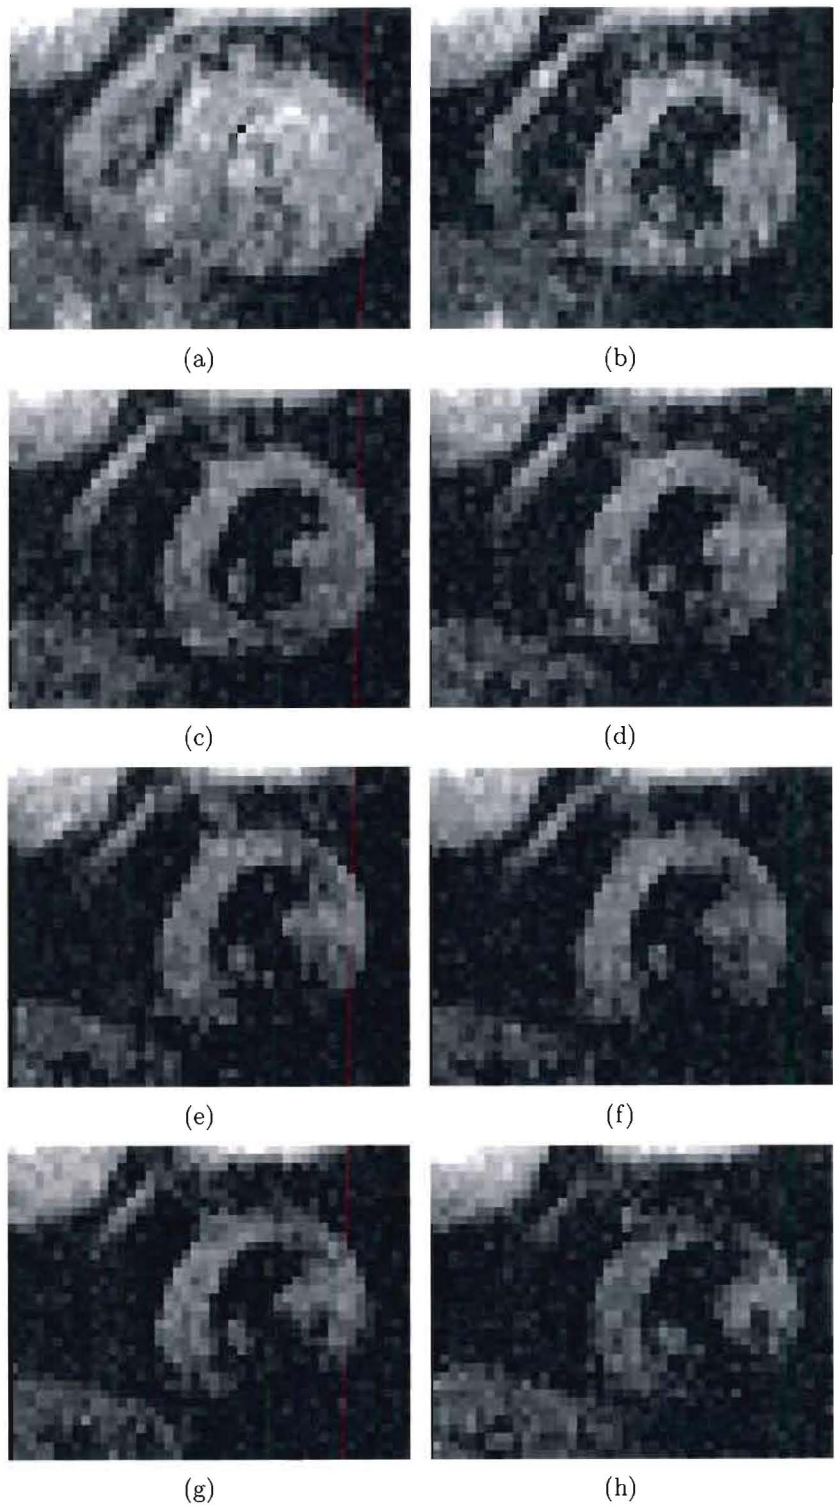

Figure 3.6: End systolic cine DENSE magnitude-reconstructed images for a healthy volunteer with k_e values of (a) 0.02, (b) 0.06, (c) 0.10, (d) 0.14, (e) 0.18, (f) 0.22, (g) 0.26, and (h) 0.30 cycles/mm. Decreased signal intensity due to intravoxel dephasing becomes apparent in the RV and inferior LV as k_e is increased.

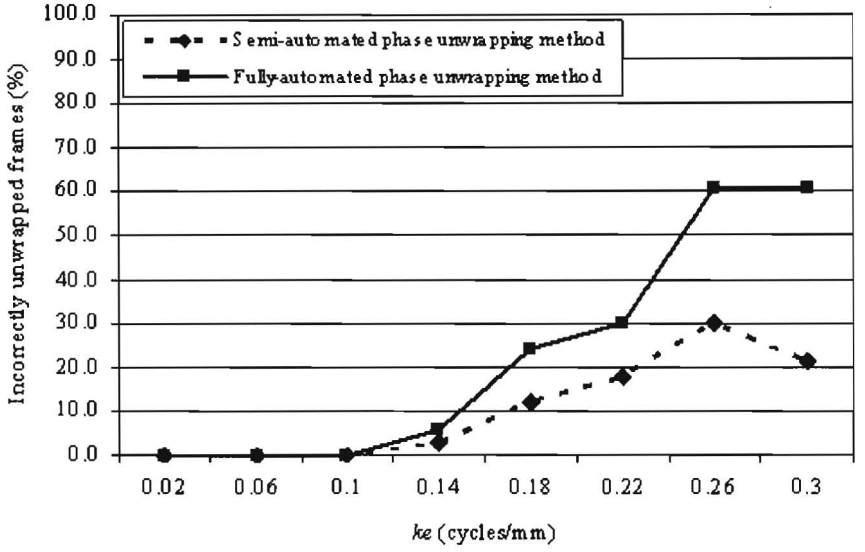

Figure 3.7: Phase unwrapping errors as a function of k_e for both semi- and fully-automated phase unwrapping methods. The majority of the incorrectly unwrapped frames for the high values of k_e are due to regions of low SNR that result from intravoxel dephasing.

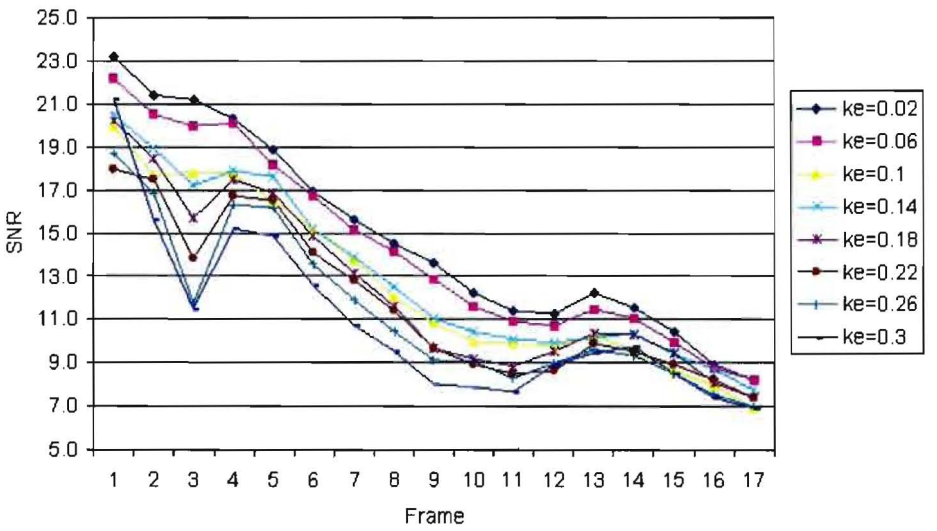

Figure 3.8: Mean myocardial SNR as a function of time for various values of k_e.

of the semi-automated algorithm limits the unwrapping through time to regions of high phase quality, thus generally ensuring that phase unwrapping errors are confined to individual frames. The radius parameter is dependent on both the spatial and temporal resolution, and should be adjusted accordingly if the imaging parameters differ greatly from those used in this study.

The phase unwrapping algorithms are more robust for lower values of k_e, as is shown in Figure 3.7. The displacement encoding frequency for the cine DENSE sequence should be set as high as possible for large phase sensitivity to displacement, but low enough to ensure acceptable phase unwrapping efficacy and a relatively uniform SNR in the LV. Figures 3.6 and 3.7 show that encoding frequencies in the region of 0.1 cycles/mm provide optimal results. Indeed, when limited to images acquired with $k_e = 0.1$ cycles/mm, the fully-automated phase unwrapping method was successful for 97.8% and the semi-automated method was successful in 98.8% of images.

A 3–D phase unwrapping method for MRI images is presented in Cusack and Papadakis (2002). The method is similar to that presented here in that unwrapping is also guided using a measure of phase quality, which is either based on the magnitude image, or on a 3–D smoothed map of phase inconsistencies. Such smoothing is undesirable in the cine DENSE application because of the limited number of pixels spanning the walls of the myocardium. More recently, a fast N-dimensional phase unwrapping algorithm has been proposed (Jenkinson, 2003). The method uses a region merging approach to optimize a cost function that penalizes differences across boundaries. This method could be applied to the cine DENSE data, but a mask would be beneficial because of the high ratio of noisy phase data to good phase data. Either the magnitude images or the phase quality maps could be used as a mask. An alternative mask for DENSE phase unwrapping has been proposed by Sinele *et al.* (2006), which involves removing low SNR pixels that lie adjacent to phase residues, and also removing pixels where displacement shear effects are present.

A spatiotemporal phase unwrapping method has been proposed for phase contrast velocity encoding (Yang *et al.*, 1996). The method is similar to the proposed semi-automated method in that it effectively implements temporal and spatial unwrapping in separate steps. Yang *et al.* use a variation of optical flow on the magnitude images for spatiotemporal unwrapping, and remaining inconsistencies are removed by using a 2–D minimum L^p-norm method to unwrap the phase separately for each frame. This method is ill-suited to cine DENSE because the optical flow spatiotemporal guiding is complicated by the $T1$-related signal decay.

The phase unwrapping method proposed here does not necessitate myocardial contours. This has been implemented with the aim of using the tracking to assist in segmenting the myocardium in cine DENSE images (Chapter 7). Myocardial contours are typically

drawn manually. The presence of myocardial contours greatly simplifies the phase unwrapping problem, improving both its speed and robustness. It is worth noting that spatiotemporal phase unwrapping could also be applied to HARP images, but the higher phase gradient in HARP imposes more stringent temporal unwrapping constraints than for cine DENSE. In fact, a recent study on phase unwrapping for HARP based on this work shows promising results (Rutz *et al.*, 2006).

The Nyquist sampling theorem states that the sampling has to be at least twice the highest frequency present in the pattern. The consequence of this for correct phase unwrapping is that the true phase difference between adjacent pixels on one frame or between the same pixel on adjacent frames needs to lie in the range $-\pi < \Delta\phi < \pi$ radians. This imposes both spatial and temporal phase unwrapping constraints, which are defined by the limits of cardiac strain and velocity, respectively. Considering the Nyquist sampling theorem and the phase of the stimulated echo in Equation 2.8, this motion constraint is $k_e \Delta x < 0.5$ cycles, where k_e is in units of cycles/mm. Thus for k_e = 0.1 and 0.15 cycles/mm, Δx = 5.0 and 3.3 mm, respectively. This places no spatial phase unwrapping constraints for practical ranges of pixel size and cardiac strain. The temporal constraints are however more stringent. For a maximum myocardial velocity of 100 mm/s, a minimum temporal resolution of 50 and 34 ms is required for k_e = 0.1 and 0.15 cycles/mm, respectively. Up to now, no temporal phase unwrapping errors have been encountered using the semi-automated method, even with k_e = 0.1 cycles/mm and a 60 ms temporal resolution, where the velocity limit is 83 mm/s. A high temporal resolution is nevertheless recommended, and the seed point for the semi-automated phase unwrapping algorithm should be placed in the interventricular septum, where the velocity is typically less than in the LV free wall.

Although phase quality maps and discontinuity maps were used here to manually detect phase unwrapping errors, there are currently no satisfactory automated methods for error detection. An operator should be aware that it is possible for phase unwrapping errors to exist, and the unwrapped phase images should be given a brief visual scrutiny. Most phase unwrapping errors are immediately obvious, even for the inexperienced observer. For a clinical setting, a phase unwrapping confidence factor could be introduced based on the number of discontinuities lying in regions of high quality phase. Temporal phase unwrapping errors could automatically be detected by looking for unreasonably large DC phase offsets between adjacent frames. Phase unwrapping errors occur in regions of very low SNR or ambiguous phase, and are often a result of patient motion.

The phase unwrapping errors should thus ideally be detected online so that the patient can be immediately re-scanned if necessary.

At least two possible improvements can be made to the phase unwrapping algorithms presented here. Firstly, the magnitude images can be used in addition to the phase quality map to guide the unwrapping path, as in Cusack and Papadakis (2002). Secondly, zero filling of k-space can be implemented to reduce the effective pixel size. This does not increase the information contained in the data, but serves to increase the number of pixels spanning the walls of the heart. Zero filling would be particularly beneficial where the myocardium is only a pixel or two thick, as in the case of the RV.

Chapter 4

Two dimensional frame-to-frame tissue tracking from cine DENSE displacement fields

4.1 Introduction

Both myocardial tagging and velocity encoded phase contrast have been used previously to track the motion of myocardial points through time. Since thinner tag lines tend to disappear within a shorter time interval, the spacing of tag lines typically employed in clinical studies is 6 mm or greater (Urayama *et al.*, 2000). Spatial resolution is a limiting factor with myocardial tagging, and only a few points spanning the width of the myocardium can be reliably tracked. Several phantom validations have been reported using myocardial tagging. In a study using a rigid rotating phantom, McVeigh and Zerhouni (1991) reported a 2–D tag tracking accuracy of 0.1 - 0.2 mm. In a further study by Moore *et al.* (1994), myocardial tagging applied to a 1–D deformable phantom was compared to photographs of surface markings. With the tags detected using a Fourier filter, the authors reported errors in estimates of tag position in the range of 0.090 - 0.185 mm. By simply scaling these values by $\sqrt{2}$, it can be inferred that 2–D frame-to-frame tracking errors will roughly correspond to 0.127 - 0.262 mm. Altar and McVeigh (1994) used a stationary phantom and a least squares estimation algorithm for validation and reported a 1–D tracking accuracy of less than 0.1 mm. This roughly

corresponds to a 2–D tracking error of less than 0.141 mm. A further deformable phantom study by Young *et al.* (1995) using active contour models to identify tag lines found the 2–D tracking accuracy to be 0.27 mm. Optical flow methods have also been used to track tag lines. In a study by Dougherty *et al.* (1999), rotational motion was simulated using stationary phantom data and the 2–D tracking error was found to decrease from 0.16 mm to 1.18 mm during 1 second. For the HARP tracking algorithm proposed by Osman *et al.* (1999), the tracking accuracy is estimated to be similar to the 0.1 mm reported by Guttman *et al.* (1994), where active contour models were used. A preliminary deformable phantom study by Osman *et al.* (1999) suggests that FastHARP is capable of tracking material points with a 2–D accuracy of 0.22 mm. In a study by Urayama *et al.* (2000) using a simulated high resolution tagging grid of 4 mm, and active contour models for tag localization, the 2–D tracking accuracy was estimated to be 0.47 ± 0.17 mm.

For phase contrast velocity encoding, the forward-backward tracking technique (Pelc *et al.*, 1995) has been used to quantify tracking accuracy on a rigid rotating phantom (Lingamneni *et al.*, 1995). For a 1×1 mm region of interest the authors reported a 2–D tracking error of 2 mm for a trajectory that is 69 mm long, having travelled at 125 mm/sec. Drangova *et al.* (1998) reported 2–D tracking accuracies of 0.7 mm for a 2.2×2.2 mm region of interest using the compensated Fourier integration technique (Zhu *et al.*, 1996) on a deformable phantom. Using "forward-backward tracking with forward and reverse integration averaged" on a gel phantom that was both rotated and translated, typical errors greater than 1 mm have been reported (Constable *et al.*, 1994).

Processing of cine DENSE images has in the past been done on a single frame basis, and strain-time relationships have been derived with respect to a fixed spatial frame of reference. The application of spatiotemporal phase unwrapping to cine DENSE phase-reconstructed images yields absolute measures of displacement. These can be used to derive frame-to-frame displacements, and in turn discrete tissue motion trajectories.

This Chapter proposes a tissue tracking algorithm for cine DENSE. These motion trajectories are improved by fitting curves that are a function of time for each ordinate direction. Trajectories are obtained at roughly a pixel resolution, and strain is calculated directly from these trajectories. The tracking accuracy was validated on a rotating phantom for typical ranges of cardiac velocity and displacement, and the technique was applied to both healthy volunteers and patients. Finally, strain noise measurements were

made to verify the effectiveness of the cine DENSE technique for measuring myocardial strain.

4.2 Methods

4.2.1 Imaging protocol

All cine DENSE scans were performed on 1.5 T scanners (Sonata or Avanto, Siemens Medical Systems, Erlangen, Germany). All subjects provided informed consent and were studied in accordance with research protocols approved by the Human Investigations Committee at the University of Virginia.

A four-channel body phased array RF coil was used for all subjects. Two-dimensional EPI cine DENSE acquisitions (Kim *et al.*, 2004) were acquired in a short-axis, mid-ventricular imaging plane. Imaging parameters included FOV = 360 × 236 mm; echo train length (ETL) = 3; effective TE = 6 ms; TR = 12 ms; matrix size = 128 × 84 pixels; pixel size = 2.81 × 2.81 mm; slice thickness = 8 mm; cardiac phases = 14; and displacement encoding frequency k_e = 0.1 - 0.15 cycles/mm. A bottom-up flyback interleaved k-space trajectory was used to remove ghosting artifacts (Kim *et al.*, 2003). The same sequence was used for a rotating phantom, which was mounted in the head coil of the scanner and rotated about the longitudinal axis of the magnet at a frequency of 0.52 Hz. The phantom (described in Section 4.2.7) was imaged in a transverse imaging plane with the same imaging parameters except FOV = 280×280 mm with pixel size = 2.19×2.19 mm. Through-plane dephasing (Aletras *et al.*, 1999a; Zhong *et al.*, 2005) was also applied to the phantom to further suppress the gradient echo of relaxed magnetisation and the conjugate stimulated echo. Scans for the two encoding directions were acquired in two breath holds of 17 heart beats each.

Rate two view sharing was used to increase the number of frames from 14 to 27, thus increasing the temporal resolution from 40 ms to 20 ms. View sharing (Foo *et al.*, 1995) constructs intermediate temporal frames by interleaving lines of k-space acquired later in time from the previous frame, and lines of k-space acquired earlier in time from the following frame in a cine acquisition. For example, for rate two view sharing, the latter

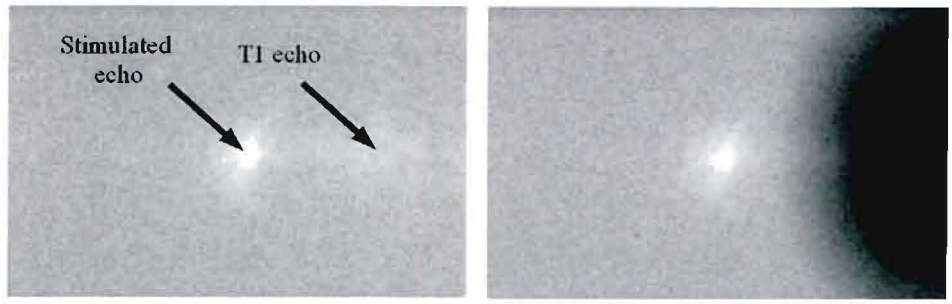

(a) (b)

Figure 4.1: Logarithmic k-space magnitude images showing the effect of the elliptical Fermi filter. (a) k-space image after the CSPAMM subtraction, and at mid-diastole, where the $T1$ echo has had time to grow. (b) k-space image after application of the filter.

half of the previous frame's k-space is interleaved with the former half of the following frame's k-space. Both non view shared and view shared sets of images were analysed.

Images were reconstructed online and analysed off-line using MATLAB (The Mathworks Inc, Natick, MA). Data sets for encoding in the x_1 and x_2 directions were co-registered using 2–D normalized cross-correlation (Gonzalez and Woods, 1992) of the DENSE magnitude-reconstructed images. Good correlations were achieved using frames at mid-systole, where the contrast between blood and myocardium is high. This high contrast occurs because of a relatively high SNR and because intravoxel dephasing has destroyed the signal in the blood pools. Residual signal from the gradient echo of relaxed magnetisation was eliminated by applying an elliptical Fermi k-space filter. For displacement encoding in the x_1 direction, the filter is centred at the edge of k-space $(k_{1_{max}}, 0)$, and takes the form

$$F_{x_1}(k_1, k_2) = \frac{1}{1 + e^{-(S_{x_1}(k_1,k_2)-1)/\sigma}} \tag{4.1}$$

where

$$S_{x_1}(k_1, k_2) = \sqrt{\frac{(k_1 - k_{1_{max}})^2}{r^2} + \frac{k_2^2}{\lambda r^2}} \tag{4.2}$$

Here σ represents a decay rate and λ is a shape parameter. The radius r is set such that the half-power level lies half way between the centre of k-space and the $T1$ echo, that is $F_{x_1}(1.5k_e, 0) = 0.5$. The setting of σ and λ were not rigorously optimised, but were done based on a visual interpretation of artifacts in image space and the shape of the echoes in k-space, respectively. A value of $\sigma = 0.05$ was used here, and no Gibbs ringing was noted. The elliptical shape parameter λ was set to 0.7. An example of the effect of this k-space filter is shown in Figure 4.1.

4.2.2 Cine DENSE tissue tracking algorithm

In DENSE, the displacement of each voxel is encoded as a phase shift that has accrued since the time that the displacement encoding was initiated (t_0). Consider a cine DENSE series of reference-corrected phase images encoded in x_1 and x_2 directions, and acquired at times $t_f = t_1, \ldots, t_F$, where F is the number of frames in the series. The time between t_0 and t_1 is 3 ms, which is much shorter than subsequent time intervals $t_2 - t_1, \ldots, t_F - t_{F-1}$ that are of the order of 60 ms. The vectors in a DENSE displacement field at any time t_f all reference the material position of the myocardium at time t_0. These displacement fields allow for a straightforward and independent calculation of strain for each frame. However, their use as indicators of myocardial motion is limited since they provide a net measure of displacement from t_0, but not measurements of frame-to-frame motion. A method of computing the trajectories of discrete elements of the myocardium as they move through the cardiac cycle was developed.

The vectors in a 2–D DENSE displacement field have heads pointing to the centres of the pixels at t_f, and tails originating from the positions of these material points at t_0. The trajectories could be initiated at t_0, but t_1 was chosen for the convenience of using fixed time intervals, and to initiate the tracking at the pixel centres of the first frame. It is thus assumed that the motion in the 3 ms interval between t_0 and t_1 is negligible.

The proposed tracking method is demonstrated for a single frame-to-frame trajectory in Figure 4.2. Figure 4.2(a) shows a DENSE displacement field of the LV at end-systole, and Figure 4.2(b) is an isolated portion of this image. The displacements shown are overlaid on the end-systolic DENSE magnitude-reconstructed image. All solid vectors have heads at pixel centres. Consider any point in the myocardium at t_1. The tail of the purple dotted vector depicts an arbitrarily chosen starting point at t_1. For each frame, the 3 vector tails closest to this starting position were identified. The red vectors correspond to the 3 vectors whose tails, at this time frame, lie closest to the chosen point. Two dimensional distance-weighted linear interpolation with these vectors was used to obtain an estimate of the displacement vector stemming from the selected starting point (purple dotted vector in Figure 4.2(b)) for this frame. By subtracting interpolated vectors of successive frames from each other, the frame-to-frame trajectory of the material point is obtained. This is demonstrated in Figure 4.2(c)

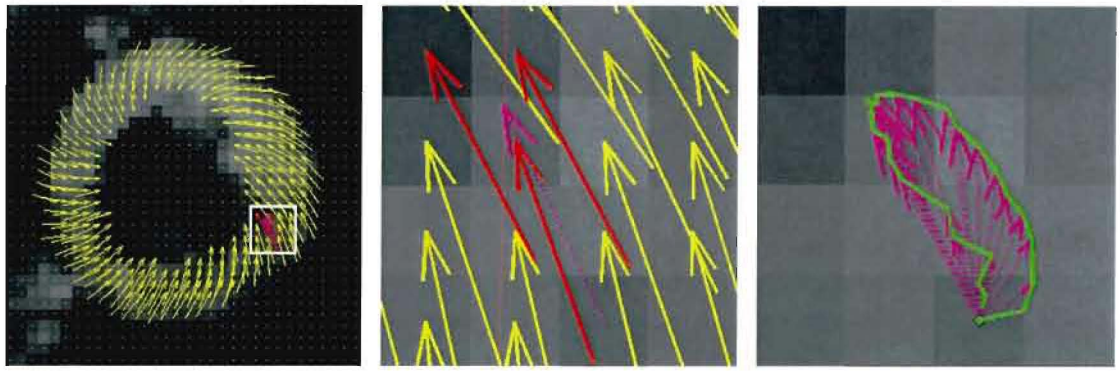

(a) (b) (c)

Figure 4.2: Cine DENSE tissue tracking. (a) End-systolic DENSE displacement field. (b) Magnified portion of this displacement field. The tail of the dotted purple vector corresponds to the chosen material point to be tracked from t_1. The three red vectors are the vectors at this frame used for the interpolation, and the dotted purple vector is the interpolated vector. (c) A frame-to-frame trajectory is obtained by subtracting the interpolated vectors of successive frames.

4.2.3 Temporal fitting

Successive non view shared cine DENSE displacement measurements are independent, providing a framework suited to temporal fitting. The accuracy of the estimate of a point's position can thus be improved based on its estimated positions at preceding and/or following cardiac phases. This was implemented here by fitting curves to each of the spatial coordinate components as a function of time.

As previously mentioned in Young (1991), the periodic behaviour of cardiac motion naturally suggests that a periodic descriptor should be used for the fit. The fitting was done using Fourier basis functions. For motion in the x_1 direction,

$$f_{x_1}(t) = \frac{1}{2}a_0 + \sum_{m=1}^{\infty} a_m \cos(m\omega t) + \sum_{m=1}^{\infty} b_m \sin(m\omega t) \tag{4.3}$$

where a_m and b_m are the Fourier series coefficients, $\omega = 2\pi/T$ is the angular frequency, and T is the cardiac period. A similar expression holds for x_2. It has been shown in an analysis of the spectral content of cardiac motion, that the energy of motion is largely contained within the first 5 harmonics of such a Fourier series (Young, 1991). This corresponds to 11 Fourier series coefficients. These coefficients can be obtained rapidly by performing a discrete Fourier transform of the data for each of the spatial

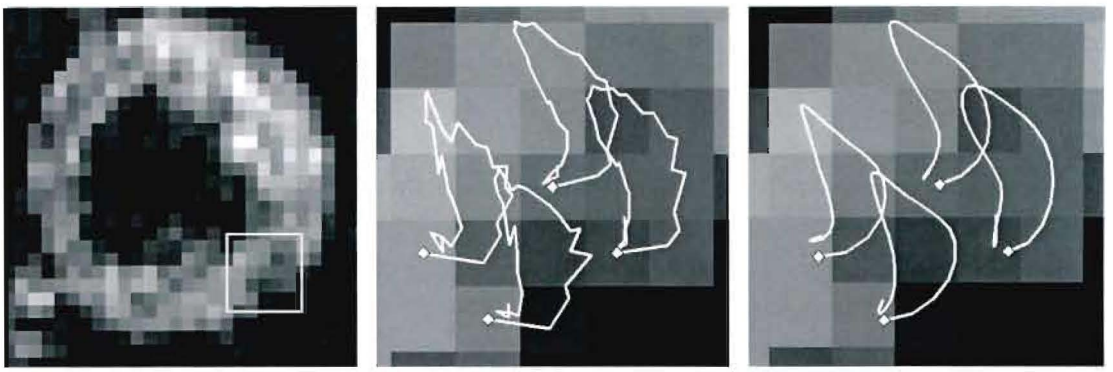

(a) (b) (c)

Figure 4.3: Improving motion trajectories by fitting 5^{th} order Fourier basis functions in x_1 and x_2 directions for a 27 frame 2–D cine DENSE data set (view shared). (a) Overview image, (b) selected raw trajectories, and (c) corresponding fitted trajectories.

coordinates, and relating the results to the complex description of the Fourier series coefficients

$$a_m + jb_m = \frac{2}{T}\int_x^{x+T} f_{x_1}(t)e^{jm\omega t}dt \tag{4.4}$$

Since only about 2/3 of the cardiac cycle is imaged and uniform time intervals are required for Equation 4.4, the position of a material point was linearly interpolated from the end point back to the starting point for the fit. The fitted information for this unsampled portion of the cardiac cycle was discarded. Note that this filling process can be avoided by fitting using least squares to estimate the Fourier coefficients, but with a considerable increase in processing time. For the *in vivo* cases where motion trajectories form almost completely closed loops, no significant difference was found between coefficients derived by fitting using least squares and the Fourier method.

Given that two coefficients are obtained from each frequency component, at least 11 temporal samples are required to capture the first 5 harmonics. More than 11 cardiac phases are thus required if any improvement is to be offered by a 5^{th} order fit.

Figure 4.3 shows a few such trajectories before and after curve fitting. Much of the behaviour of the trajectories seen in Figure 4.3(b) is discontinuous and inconsistent with physiological motion. The fitted trajectories in Figure 4.3(c) give a more realistic portrayal of cardiac motion. Figure 4.4 depicts trajectories for a healthy volunteer. These trajectories commence at all myocardial pixel centres at t_1.

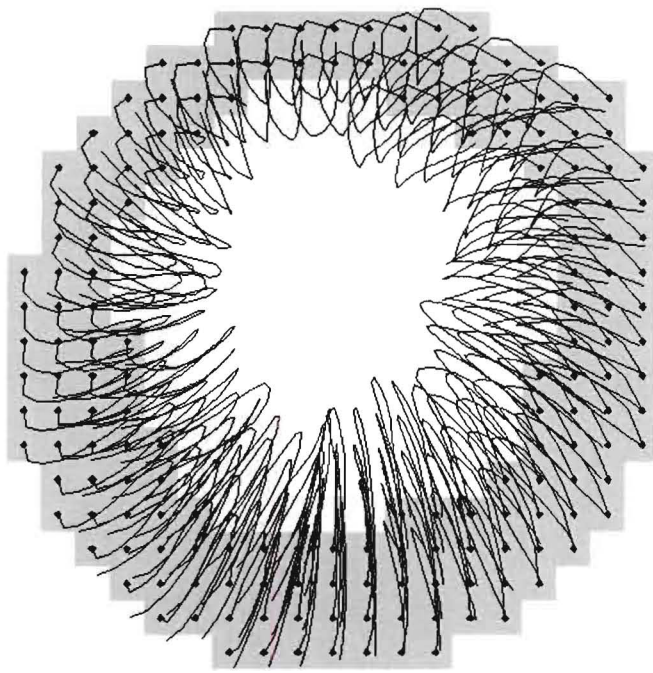

Figure 4.4: Trajectories for a healthy volunteer, commencing at all myocardial pixel centres at t_1 and traversing 75% of the cardiac cycle. The septum is centred at 9 o'clock in the image.

4.2.4 Reconstructing DENSE displacement fields

If all of the myocardial pixel centres at time t_1 are tracked and temporally fitted, then the resulting accumulation of temporal information can be used to improve the original cine DENSE displacement fields. This is done by performing a 2–D distance-weighted linear back-interpolation from the fitted trajectories to recreate vectors with heads positioned at the centres of all myocardial pixels at all cardiac phases. For each cardiac phase, the 3 closest trajectory positions to each pixel centre are found, and vectors are defined that reference these trajectory positions to time t_1. Two dimensional distance-weighted linear interpolation is then done using sets of these 3 vectors to estimate vectors with heads positioned at the centres of all myocardial pixels. The reconstructed DENSE displacement fields yield more uniform strain maps than the original DENSE displacement fields.

4.2.5 Calculating strain from the trajectories

The standard method of calculating strain-time curves is based on circumferential cardiac segments that are geometrically fixed with respect to a non-cardiac frame of reference. Cardiac twist and asymmetrical contraction cause portions of the myocardium to move into adjacent segments. Obtaining strain-time relationships using the trajectories eliminates this problem and temporal fitting yields more realistic stain-time curves. Furthermore, since temporal fitting provides equations of motion, the displacement and strain can effectively be interpolated to any temporal resolution. If all myocardial pixel centres at time t_1 are tracked, then strain can be calculated at a spatial resolution approximately equal to the pixel size[1].

If X_1 and X_2 are unit basis vectors aligned at t_0 with x_1 and x_2, respectively, then the 2–D deformation gradient tensor F is given by

$$\mathbf{F} = \begin{bmatrix} \frac{\partial X_1}{\partial x_1} & \frac{\partial X_1}{\partial x_2} \\ \frac{\partial X_2}{\partial x_1} & \frac{\partial X_2}{\partial x_2} \end{bmatrix} \tag{4.5}$$

The 2–D Lagrangian finite strain tensor $\mathbf{E}$ can then be calculated using $\mathbf{E} = \frac{1}{2}(\mathbf{F}^T\mathbf{F} - \mathbf{I})$, where I is the identity matrix, and the superscript T denotes the transpose operation. Consider a single trajectory position $C = (x_1, x_2)$ at any cardiac phase. Let $L = (x_1 - 1,\ x_2)$, $R = (x_1 + 1,\ x_2)$, $A = (x_1,\ x_2 + 1)$ and $B = (x_1,\ x_2 - 1)$ correspond to the trajectory positions to the left, right, above and below position C, respectively. The tensor $\mathbf{F}$ can be calculated using any 3 adjacent trajectories, corresponding to the configurations ACL, LCB, BCR or RCA. All four of these configurations will be present for the majority of the pixels in the myocardium. Fewer configurations will be present for pixels bordering the epicardium or endocardium. The average tensor $\mathbf{F}$ was calculated for all configurations that are present, and was used to estimate $\mathbf{E}$ for the selected trajectory position. For short-axis views, 2–D Lagrangian strain in the LV is typically decomposed into circumferential and radial components, defined at the first cardiac phase. Circumferential strain (Ecc) is the component of $\mathbf{E}$ taken in a direction that circumnavigates the LV centroid. Radial strain (Err) is the component of $\mathbf{E}$ perpendicular to Ecc.

[1]The algorithm described here was implemented in MATLAB by Aaron Hess.

4.2.6 Measuring strain noise

Measurements were made to estimate the noise in strain data computed from cine DENSE images. Such measurements have been made previously for single phase DENSE (Aletras *et al.*, 1999a), where instead of acquiring reference, x_1, and x_2 encoded images, all three images were acquired with the same displacement encoding gradients. The strain noise was then defined as the mean myocardial strain in these images.

For the cine DENSE technique, the pulse sequence was modified to acquire two identical pairs of CSPAMM images, where all acquisitions used identical displacement encoding gradients. This was done for encoding in both x_1 and x_2 directions, and the resulting value of the mean first principal strain was used as an indication of strain noise. The principal strains were calculated from the Lagrangian strain tensors using eigenvector decomposition (as shown schematically in Figure 2.7).

4.2.7 Phantom validation

A rigid rotating phantom was used to quantify tracking accuracy. The phantom comprises a plexiglass cylinder filled with agarose gel, and was rotated about the longitudinal axis of the magnet at a frequency of 0.52 Hz. It was driven by a brush servo motor (model 3528, Cleveland Motion Control) via a 30:1 gear box (Boston Gear USA). The cine DENSE sequence was gated using a notched disc attached to the drive shaft of the phantom, and a photodiode/photoreceptor circuit (EE-SX1106, Omron, Schaumburg, IL).

The centre of rotation of the phantom was estimated from the DENSE reference images, since they inherently have a higher SNR than the DENSE magnitude images. A k-means clustering algorithm (Gonzalez and Woods, 1992) was used to separate the phantom from the background. The mean centroid of the phantom for frames 2 to 6 was designated as the centre of rotation. This point provided the reference for simulated trajectories, which were compared to the measured trajectories to give an indication of tracking accuracy. The seed point for the semi-automated phase unwrapping algorithm was placed near the centre of the phantom to avoid temporal phase unwrapping errors due to the large displacements occurring near the periphery of the phantom. Baseline strain noise measurements were also calculated.

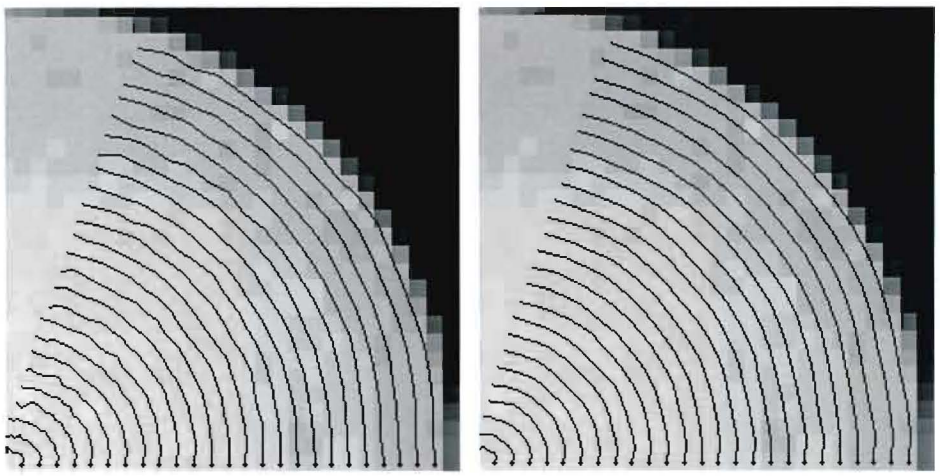

(a) (b)

Figure 4.5: Tracking of selected pixel centres on the rotating phantom. (a) Raw trajectories, and (b) least squares fitted trajectories using 5^{th} order Fourier basis functions.

4.3 Results

4.3.1 Phantom tracking

Raw and fitted trajectories obtained for selected starting pixel centres on the rotating phantom are shown in Figure 4.5. The raw trajectories portray a decrease in accuracy with time. This is due to decay of the stimulated echo as a result of $T1$ relaxation, as well as the decrease in 2–D interpolation accuracy that accompanies increasing DENSE displacement vector lengths. Since a small portion of the rotational cycle was imaged, the calculation of the Fourier coefficients for the fitted trajectories was done using least squares, and not the discrete Fourier transform approach described above. The fitted trajectories provide a clear improvement, particularly for the latter frames.

The unfitted rotating phantom trajectories were compared to their theoretical counterparts, yielding a tracking accuracy of 0.11 ± 0.07 pixels (0.24 ± 0.15 mm). These figures were derived using physiological velocity and displacement limits of 0-200 mm/sec and 0-20 mm, which correspond to a maximum radius of 23 pixels from the centre of the phantom and a maximum rotation angle of 0.33 radians (≈ 8 frames), respectively. No significant improvement was evident for the fitted trajectories because the measured accuracy approaches the mechanical tolerance of the rotating phantom assembly. The improvement became more pronounced when noise was added to the k-space data of

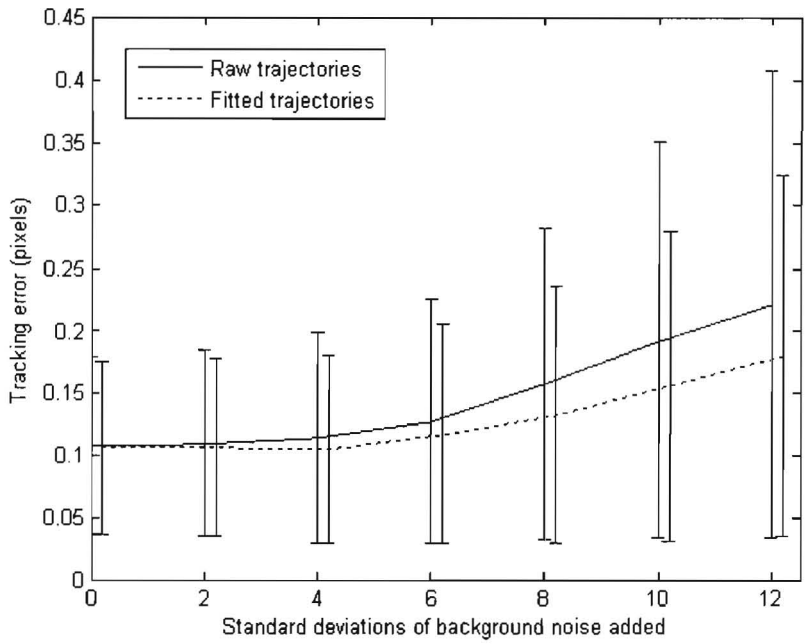

Figure 4.6: Raw and fitted tracking error plotted as a function of the number of standard deviations of background noise added to the DENSE and reference data. Temporal fitting was done using least squares because only a small portion of the rotational cycle was imaged. Error bars are ± 1 standard deviation tracking error.

the phantom. Multiples of the standard deviation of a background noise sample (σ_n) were used to scale the standard deviation of normally distributed, zero-mean noise. This noise was added separately to the real and imaginary components of the k-space data. Figure 4.6 demonstrates this, where tracking error is less for fitted than raw data as the amount of added noise is increased.

4.3.2 Strain analysis

Five healthy volunteers and two patients with myocardial infarcts were imaged in a mid ventricular short axis view using cine DENSE. The effects of tissue tracking and temporal fitting on end-systolic circumferential (Ecc) and radial (Err) strain were investigated.

Figures 4.7(a) and 4.7(b) show an example of an original cine DENSE displacement field and the corresponding Ecc map for a normal volunteer at end-systole. Note that end-systole occurs roughly at $t_F/2$, where F is the total number of frames. Note again that F corresponds to about two thirds of the cardiac cycle. The strain here was derived

using the finite element method described in Moaveni (1999). If all pixel centres at t_1 are tracked and temporal fitting is done using 5^{th} order Fourier basis functions, and if the original displacement fields are reconstructed using 2–D back-interpolation, then the displacement field and Ecc map in Figures 4.7(c) and 4.7(d) result, respectively. Temporal fitting has removed much of the noise in the reconstructed maps of Ecc, while still maintaining the transmural gradient of higher Ecc near the endocardium to a lower Ecc near the epicardium. Although these images demonstrate the improvement offered by temporal fitting, the focus of the work presented in this thesis is on the temporal or 'cine' aspect of cine DENSE analysis, and all strain calculations hereafter were derived directly from the trajectories.

Table 4.1: Mean peak radial and circumferential strain calculated from cine DENSE images of five healthy volunteers. The strain was derived using the trajectories and was calculated both before and after temporal fitting, and both with and without rate two view sharing. All fitting was done using 5^{th} order Fourier basis functions.

	No view sharing (14 frames)	View sharing (27 frames)
Mean original Ecc	-0.18 ± 0.07	-0.17 ± 0.08
Mean fitted Ecc	-0.18 ± 0.06	-0.17 ± 0.06
Mean original Err	0.30 ± 0.23	0.30 ± 0.22
Mean fitted Err	0.28 ± 0.21	0.25 ± 0.19

For the normal volunteers strain was calculated from the trajectories both before and after temporal fitting. The results for non view shared and view shared sets of data are presented in Table 4.1. Figure 4.8 shows end-systolic Ecc maps of a patient with an infarct centred at seven o'clock in the images. The strains in Figures 4.8(a) and 4.8(b) were derived from raw and fitted trajectories, respectively, for non view shared data. To further accentuate the ability of the temporal fitting to extract useful strain information, noise was added to the k-space data of this patient as described previously. Figure 4.9 shows the results for 9 standard deviations (σ_n) of noise added to both the cine DENSE and reference data. The infarct, which is nearly indistinguishable in Figure 4.9(a), becomes clearly apparent in Figure 4.9(b) after temporal fitting. Adding more than $9\sigma_n$ of noise renders the myocardium almost indistinguishable in the DENSE magnitude images, and phase unwrapping errors become apparent.

The combined Ecc histograms of the two infarct patients are shown before and after temporal fitting in Figure 4.10. Here temporal fitting reduces the standard deviation of Ecc from 0.50 to 0.11. The regions of infarct with contractile dysfunction are, in these examples, evident as a distinct peak in Figure 4.10(b) (arrow).

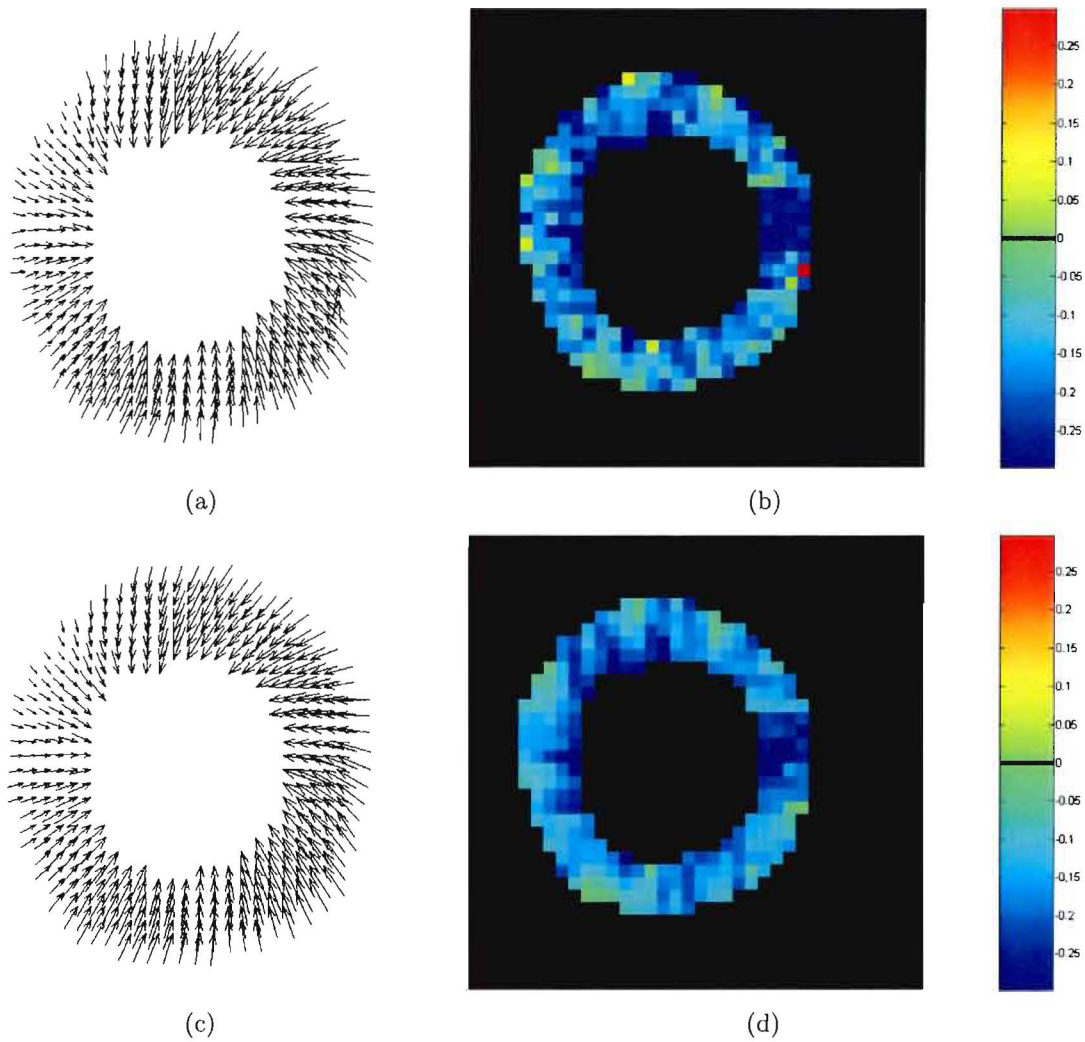

Figure 4.7: (a) End-systolic cine DENSE displacement field of a normal volunteer, and (b) corresponding map of Ecc. (c) Reconstructed displacement field, and (d) corresponding map of Ecc. Trajectories were fitted using 5^{th} order Fourier basis functions. The septum is centred at 9 o'clock in these images.

Strain-time curves for typical cardiac segments were obtained by plotting the mean strain values of trajectories originating within these segments as a function of time. Figure 4.11 shows a few Ecc and Err strain-time curves before and after temporal fitting for a normal volunteer, and Figure 4.12 shows the corresponding curves for a patient with an anteroseptal infarct. Both Ecc and Err are significantly impaired in the segment containing the infarct. The curves for strain obtained from the fitted trajectories look more realistic than those from the original trajectories. Similar improvements were noted in the strain-time curves from the second infarct patient and for all other volunteers.

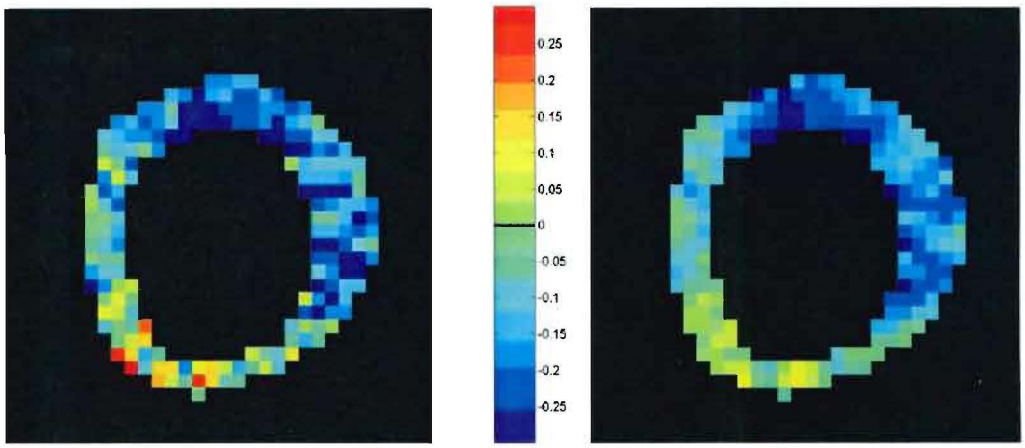

(a) (b)

Figure 4.8: Ecc at end-systole for a patient with an infarct centred at seven o'clock, (a) Ecc before fitting, and (b) Ecc after fitting 5^{th} order Fourier basis functions. The septum is centred at 9 o'clock in these images.

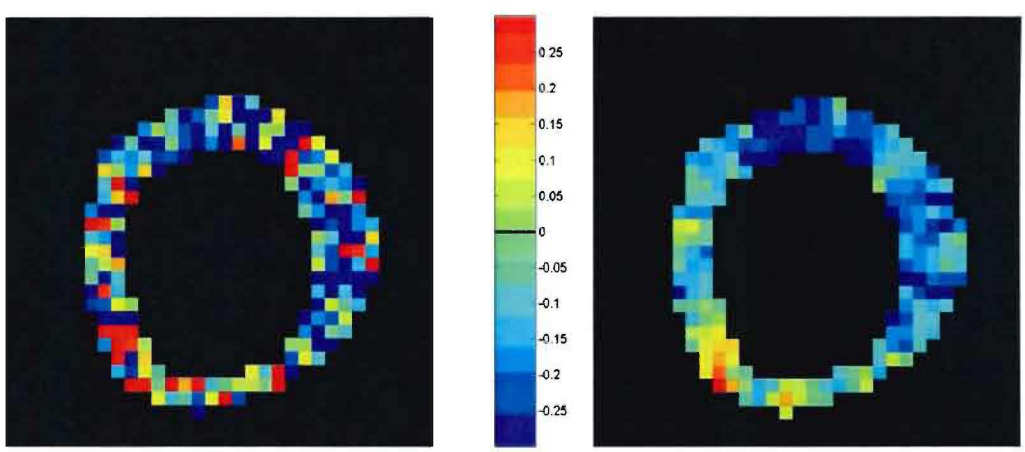

(a) (b)

Figure 4.9: Effect of temporal fitting on noisy data. (a) Ecc at end-systole, obtained after adding $9\sigma_n$ of noise to the k-space data. (b) Ecc, obtained after adding $9\sigma_n$ of noise to the k-space data and after fitting 5^{th} order Fourier basis functions. The septum is centred at 9 o'clock in these images.

The principal strain noise for the phantom was measured to be $1.3 \pm 1.6\%$ at t_0, increasing to $1.6 \pm 1.9\%$ 800 ms later due to growth of the $T1$ echo. Similar phantom studies using single phase DENSE measured a strain noise level of 0.5% (Aletras *et al.*, 1999a). Adjusting the strain noise level for comparison using simply the square root of the acquisition times (34 heart beats here versus 24 heart beats in Aletras *et al.* (1999a)) reduces this strain noise level to 0.42%. The SNR of single phase DENSE is inherently greater than cine DENSE, and a higher strain noise is thus expected for cine DENSE.

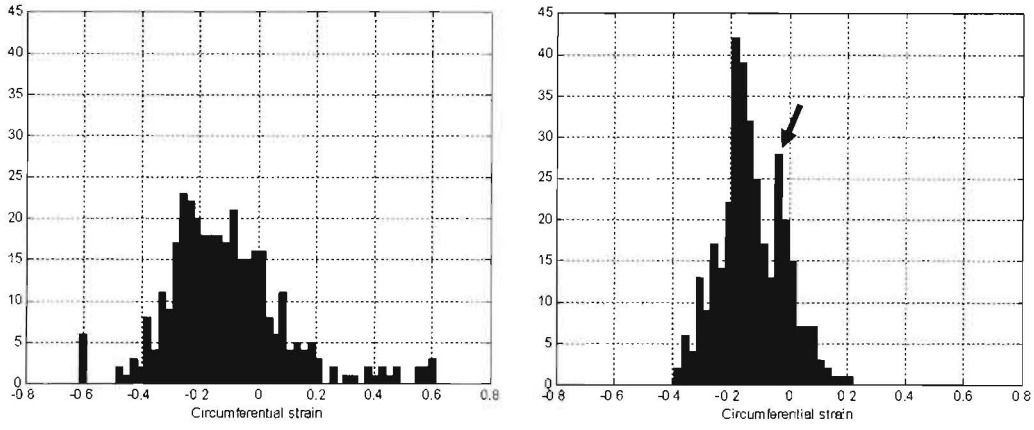

(a) (b)

Figure 4.10: Histogram depicting Ecc at end-systole for two infarct patients using (a) original data, and (b) a 5^{th} order Fourier fit. A peak depicting the low strain region of the infarcts is evident in (b) (arrow).

Using the raw trajectories, the first principal strain noise *in vivo* was measured to be 4.3±5.1% at end-diastole, 4.5±5.1% at end-systole, and 9.3±9.1% in mid-diastole. After temporal fitting, the strain noise was reduced to 2.5 ± 3.0% at end-diastole, 3.1 ± 3.1% at end-systole, and 5.3 ± 5.0% in mid-diastole. The reported in vivo mean strain using single phase DENSE was 2.8 ± 2.5% for both contraction and dilation (Aletras *et al.*, 1999a). Adjusting this strain noise level for comparison reduces the mean to 2.35%, which is again less than was achieved for the cine DENSE data.

4.4 Discussion

Phantom studies with myocardial tagging reported 2-D tracking accuracies of the order of 0.10 - 0.27 mm, and similar studies using velocity encoding reported accuracies of 0.7 - 1.0 mm. The rotating phantom cine DENSE tracking accuracy of 0.24 mm is comparable to those obtained using myocardial tagging, and superior to those obtained using phase contrast velocity encoding. The image spatial resolution offered by cine DENSE is, however, greater than that achievable using myocardial tagging. Tracking errors in velocity-encoded techniques accumulate as the velocity fields are integrated, whereas displacement measurements for both tagging and DENSE are independent of time. Velocity-encoded images are, however, advantageous since they do not suffer from a decay in SNR with time. Factors contributing to errors in the cine DENSE

tracking include: decaying SNR with time; 2–D interpolation errors, particularly for large displacements; through-plane motion; and translation of the start and end positions of the trajectories due to periodicity constraints of the Fourier fit. Imperfect phase correction also contributes to the error. The phase reference images compensate for phase variations caused by, among other things, main magnetic field inhomogeneities and eddy currents.

Adding artificial noise to the cine DENSE data to investigate the effect of temporal fitting is useful because of the anticipated variation in the quality of images that might realistically be acquired from the MR system. It is common to use a number of different scanning protocols for different clinical purposes where SNR is traded for higher spatial resolution, higher temporal resolution, and/or reduced scan time. For example, for some protocols a voxel size of $3 \times 3 \times 8$ mm is used, and for other protocols a voxel size of $1 \times 1 \times 8$ mm is used, representing a 9-fold difference in SNR for this single example. By adding noise to real data, insight into the accuracy of the proposed methods as a function of SNR is gained and it can be better understood how practical adjustments to the scanning parameters may affect the accuracy of the displacement estimates using both raw cine DENSE data and temporal fitting. Only the effects of temporal fitting on random noise have been investigated thus far. Regions with non-random noise are expected to benefit less from temporal fitting. More specifically, transient errors due to ghosting and motion artifacts are expected to be reduced by temporal fitting, but errors that are present on more than a few frames of the cine series would not be improved by the technique.

For normal subjects, the mean peak mid-ventricular Ecc predicted using conventional processing of tagged images has been measured to be -0.19 ± 0.04 (Moore *et al.*, 2000). The Ecc strain values calculated using the non view shared raw and fitted trajectories are -0.18 ± 0.07 and -0.18 ± 0.06, respectively. These are in agreement with the results from tagging, and the temporal fitting serves to reduce the variability of the cine DENSE strain measurements as demonstrated by the strain-time curves in Figures 4.11 and 4.12. The fact that the fitted strain is more regular than the unfitted strain is, however, not proof that the results are correct. The corresponding Err measurements are 0.30 ± 0.23 and 0.28 ± 0.21, respectively. These are lower than the corresponding Err values of 0.42 ± 0.21 reported in Moore *et al.* (2000) and the mean value of 0.38 reported in Bogaert and Rademakers (2001), but greater than the values reported by Osman *et al.* (1999) and Tustison *et al.* (2003). A small reduction in Err calculated

using the trajectories is expected due to the effect of through-plane motion on the tracking process. Cine DENSE images are acquired in a fixed imaging plane, through which the heart moves during the cardiac cycle. For the short-axis view where the heart has a torus shape, the larger base moves into the image plane during systole, and then returns during diastole. This through-plane motion means that vector tails to be used for interpolation are displaced radially from the trajectory starting points at t_0. The consequence is that the same 3 nearest vector tails can be used for different radially-separated trajectory starting points. This results in a decreased Err, with the effect being greatest in the sub-endocardium. Since the closest vector tails will generally be found in a radial direction when through-plane motion is present, the effect of the through-plane motion on the tracking has a less pronounced effect on Ecc.

The values of Ecc and Err derived using the view shared cine DENSE data are similar to those derived using non view shared data. Although view sharing doubles the number of points that can influence the harmonics of the 5^{th} order Fourier basis functions, these points do not deviate significantly from the curves obtained from the non view shared data.

Given that the first principal strain is generally oriented in the radial direction in short-axis images, measures of Err taken together with the strain noise measurements give an indication of the strain-to-noise ratio. For typical values of cardiac Err, the strain noise measurements show that the methods presented here yield useful strain data.

The average processing times for 25 frames of a typical contoured 2–D cine DENSE data set on an Intel 1.4 GHz Pentium M with 512 MB RAM are: phase unwrapping = 13.1 seconds; tracking and temporal fitting = 8.3 seconds; calculating strain from the trajectories = 14.3 seconds. Thus, excluding manual segmentation, the analysis of a single slice of 2–D cine DENSE data in MATLAB takes less than a minute.

In summary, the following approximations and assumptions were made when processing the cine DENSE data: 1. The use of temporal fitting implicitly constrains cardiac behaviour to an underlying model of motion; 2. although view sharing provides a first level of temporal smoothing it does not improve the true temporal resolution; 3. the Fermi filter effectively zero-fills a portion of the k-space, thus removing some high frequency spectral information and slightly reducing the true spatial resolution; and 4.

any 2–D tracking method using tagging or DENSE implicitly assumes that the through-plane motion is uniform throughout the image plane.

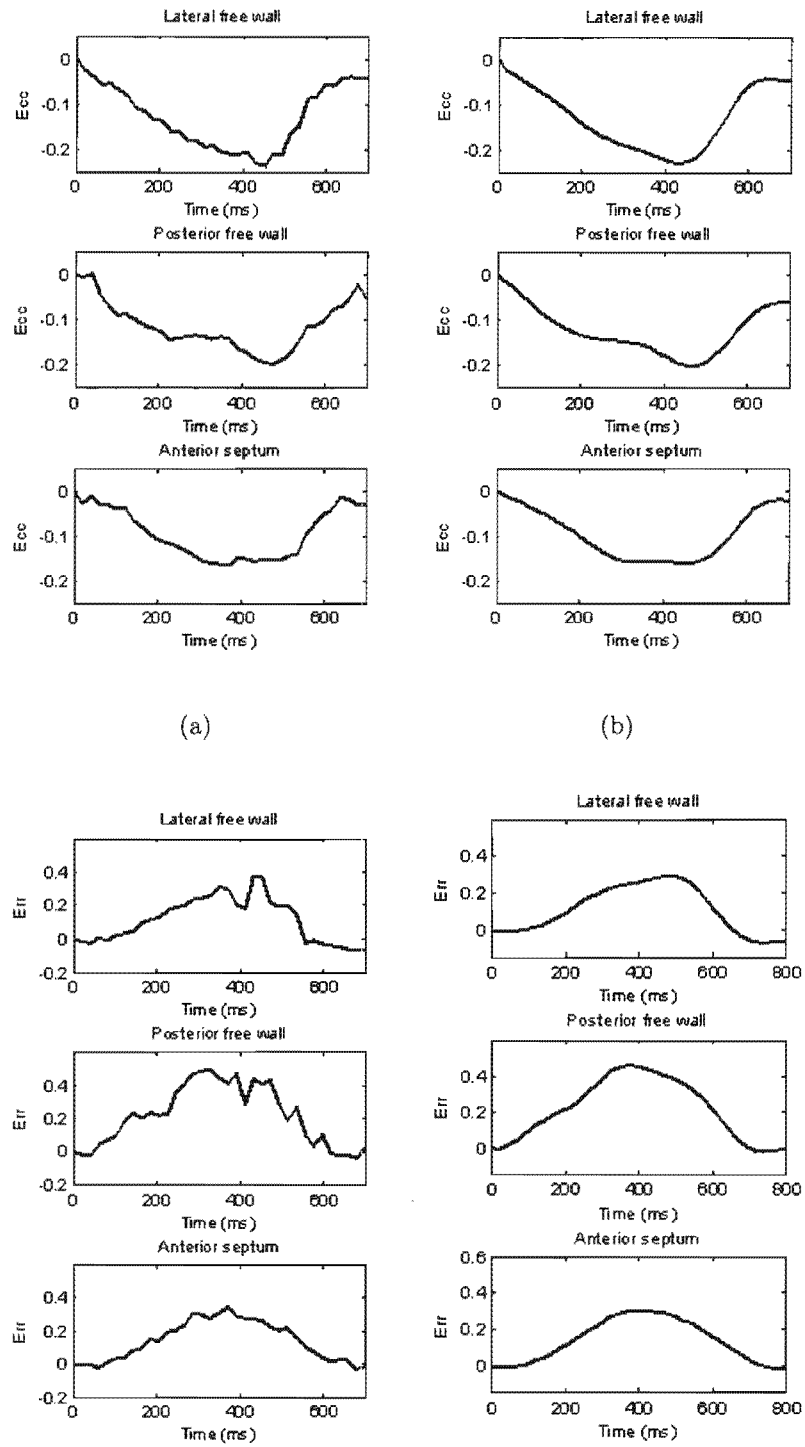

Figure 4.11: Strain versus time for three LV cardiac segments of a normal volunteer. (a) Ecc from unfitted trajectories, (b) Ecc from fitted trajectories, (c) Err from unfitted trajectories, and (d) Err from fitted trajectories.

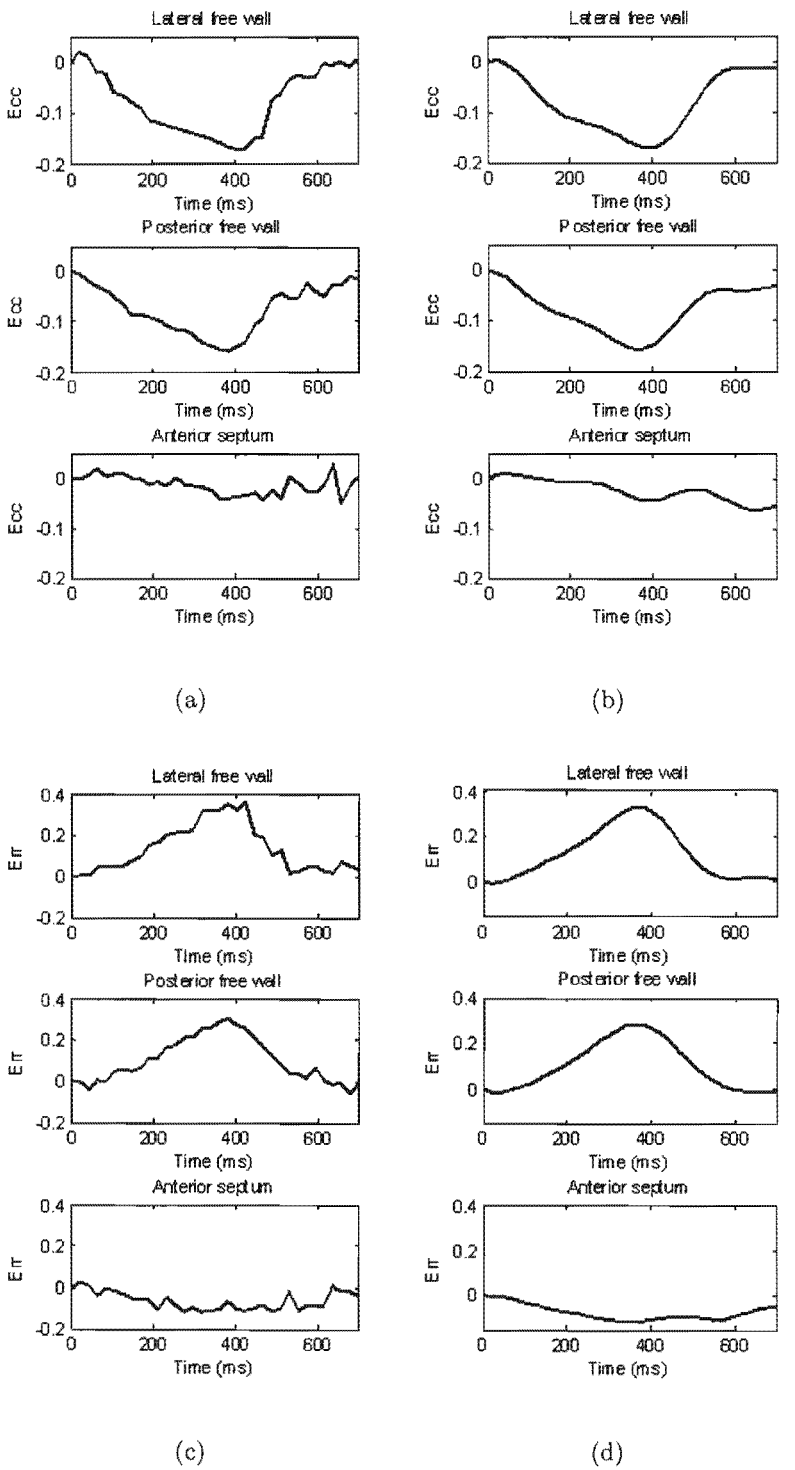

Figure 4.12: Strain versus time for three LV cardiac segments for a patient with an anteroseptal infarct. (a) Ecc from unfitted trajectories, (b) Ecc from fitted trajectories, (c) Err from unfitted trajectories, and (d) Err from fitted trajectories.

Chapter 5

True two/three dimensional tissue tracking using slice followed cine DENSE

5.1 Introduction

A shortfall of the quantitative MRI techniques described in Chapter 2 is that a fixed plane is typically imaged, through which the heart moves during the cardiac cycle. Different portions of tissue can thus be present at each cardiac phase, greatly complicating myocardial tissue tracking methods. Slice following (SF) (Fischer *et al.*, 1994; Stuber *et al.*, 1999) is a variant of CSPAMM myocardial tagging that takes through-plane motion into account. By incorporating SF into a cine DENSE sequence, true 2–D and 3–D frame-to-frame motion trajectories of material points initially lying in a single plane can be calculated at an arbitrary spatial resolution.

5.1.1 Slice following

CSPAMM serves to separate the component of the magnetisation with the tagging information from the component resulting from $T1$ relaxation. This is achieved by subtracting a tagged image from a second tagged image where the modulating sinusoid

is shifted by π radians. Through improving the contrast of the tags, CSPAMM also prolongs their persistence.

In both tagging and DENSE, the entire excitable volume of tissue is typically position encoded. Slice following makes use of spatially selective excitation and a CSPAMM subtraction to account for through-plane motion. The SPAMM kernel is comprised of two non-selective 90 degree RF pulses separated by a gradient lobe applied in the direction of the desired encoding. In SF, a thin slice is selected by applying a through-plane gradient with the first 90 degree RF pulse. The imaged slice is chosen to be large enough to accommodate the motion of the encoded tissue. By applying the CSPAMM subtraction, the unencoded tissue is cancelled out and only the signal from the encoded tissue remains. Figure 5.1 demonstrates this concept. Slice following offers a lower SNR than conventional tagging because (i) fresh spins are not moving into the imaged slice, and (ii) errors due to noise and patient motion are amplified by the CSPAMM subtraction of the thick imaged slice. Note that the nonselective application of the second 90 degree RF pulse also serves to presaturate the non-tagged tissue, thus improving the SNR for the early heart phases.

To compensate for tag fading, the imaging excitation angle is varied with time such that the signal intensity is constant for each cardiac phase. This ramped flip angle series, described in Stuber *et al.* (1999), is given by

$$\alpha_{f-1} = \arctan\left(\sin(\alpha_f)e^{-\Delta t/T1}\right) \tag{5.1}$$

where $\alpha_1, \ldots, \alpha_F$ are the RF excitation angles, Δt is the temporal resolution, and f is the frame number. Numerical simulation for single breath hold CSPAMM predicts that a consistent signal intensity is ensured if $\alpha_F = 21$ degrees (Stuber *et al.*, 1999), where F corresponds to the last cardiac phase in the series.

The longitudinal motion of the heart is greatest at the base, and decreases as one moves towards the apex. In order to minimise the noise introduced by image subtraction of non-tagged tissue in SF, it is desirable to use a thin image slice. For a tagged slice thickness of 8 mm, Stuber *et al.* suggest imaged slice thicknesses for short-axis views of 30 mm at the base, 25 mm at the mid-ventricle, and 20 mm at the apex of the heart. The recommended offset of the centre of the tagged slice relative to the centre of the

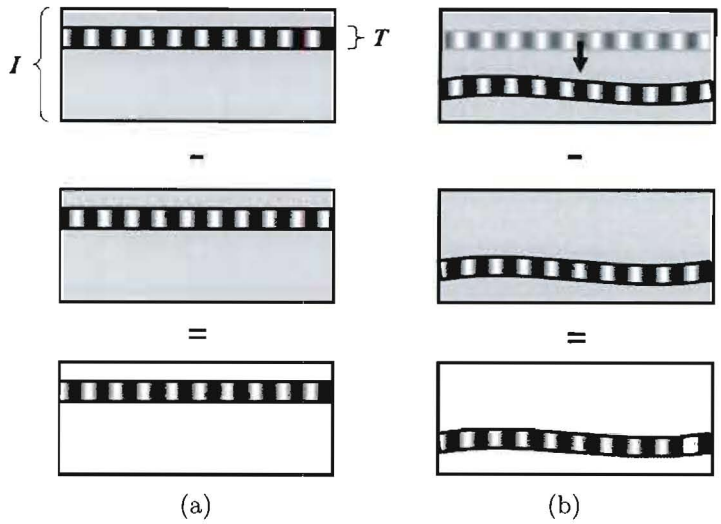

Figure 5.1: Schematic description of slice following. T = tagged slice; I = imaged slice. CSPAMM subtraction removes all untagged tissue, leaving only the tagged slice. (a) Before tissue deformation, and (b) after deformation.

imaged volume is 12 mm at the base, 6 mm at the mid-ventricle, and 0 mm at the apex. The tagged slice is offset towards the base in short-axis images.

5.1.2 Three dimensional MRI cardiac tissue tracking

Myocardial tracking in 3–D using tagged data typically involves acquiring tagged images in both short-axis and long-axis views of the heart. Spatial resolution is then not only limited by the tag spacing, but also by the intersections of the short-axis and long-axis image planes. The 3–D motion of any point in the myocardium is typically inferred using various degrees of model-based and non model-based interpolation (Young and Axel, 1992; O'Dell *et al.*, 1995; Young *et al.*, 1995; Park *et al.*, 1996; Kuijer *et al.*, 2000; Chandrashekara *et al.*, 2002; Deng and Denney, Jr., 2004). A comprehensive review of the modelling techniques developed for 3–D cardiac image analysis is given in Frangi *et al.* (2001). Three dimensional tracking using HARP has been reported in a study using short-axis and long-axis tagged image planes with a simple material mesh model (Pan *et al.*, 2003). In further work, Pan *et al.* (2005) proposed a true 3–D tracking algorithm using SF with HARP, but the spatial resolution is still limited to the intersection points of short-axis and long-axis image planes. Recently, true 3–D tracking of points within a single slice has been demonstrated using HARP with SF CSPAMM tagging to measure in-plane motion, and zHARP encoding to measure through-plane motion

(Abd-Elmoniem *et al.*, 2006). The zHARP technique (Abd-Elmoniem *et al.*, 2005) is a modification of SF-CSPAMM, where encoding gradients add a phase proportional to the through-plane motion to every material point in the acquired slice. Data acquisition for zHARP requires no extra time over a conventional CSPAMM acquisition.

Phase contrast velocity encoding is advantageous because it offers a higher spatial resolution than myocardial tagging. It also allows velocity to be encoded in the through-plane direction. Its accuracy for tissue tracking is, however, limited by the fact that velocity, and not displacement, is measured. Several tracking algorithms have been developed for velocity encoding (Pelc *et al.*, 1995; Zhu *et al.*, 1996; Constable *et al.*, 1994), all of which can be readily generalized to 3–D. The common drawback of these velocity encoding tracking methods is that errors propagate with time. Also, obtaining clinically acceptable data acquisition strategies for 3–D tissue tracking using velocity encoding is a challenging task (Zhu *et al.*, 1997). True 3–D tissue tracking using velocity encoding requires a large volume of the myocardium to be imaged over multiple parallel slices, all of which must be velocity encoded in three orthogonal directions (Zhu and Pelc, 1999). It is worth noting at this point a study that combines in-plane tagging and through-plane velocity encoding into a single sequence (Kuijer *et al.*, 1999). Kuijer *et al.* also describe an algorithm that uses this sequence on two adjacent imaging planes to track 3–D myocardial motion. Although this approach is faster than 3–D tracking using velocity encoding, it also suffers from the adverse effects of through-plane motion, and the accumulation of errors in through-plane tracking.

Whereas velocity encoding measures instantaneous tissue velocity, DENSE measures tissue displacement relative to a fixed encoding time. Cine DENSE thus provides independent displacement measurements at each cardiac phase. A 2–D myocardial tissue tracking algorithm for cine DENSE was proposed in Chapter 4. If cine DENSE images encoded in the through-plane direction are acquired in addition to those encoded in the two in-plane directions, the resulting 3–D displacement fields have vector heads all lying on the image plane, with the corresponding vector tails projecting out of the image plane to their positions at the time of displacement encoding (t_0). This is illustrated in Figures 5.2(a) and 5.2(b). Obtaining accurate 3–D trajectories from this 3–D cine DENSE data is not a matter of simply combining the in-plane trajectories with the measured through-plane displacement. This is because through-plane measurements for each voxel correspond to points arising from varying distances out of the image plane. Since the position of the vectors at t_0 is of primary interest for frame-to-frame tracking,

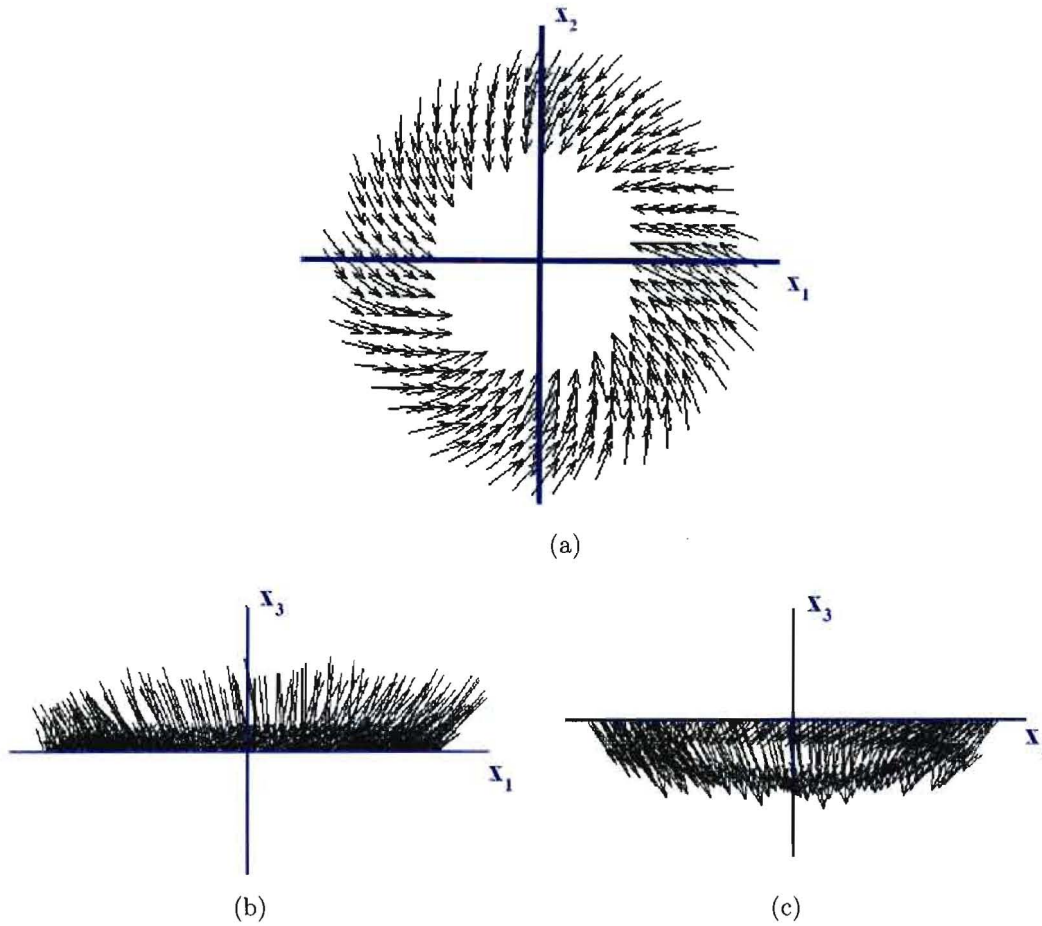

Figure 5.2: Three dimensional DENSE displacement fields for the LV at end-systole. (a) Plan view, (b) side view: conventional DENSE, and (c) side view: slice followed DENSE.

the direct extension of the 2–D tracking method to 3–D would require acquiring multiple adjacent slices of cine DENSE data, each with encoding in three orthogonal directions.

A tracking algorithm for approximating 3–D motion using conventional cine DENSE in a single imaging plane is now proposed. Let the position of a voxel in 3–D space be $\mathbf{x} = (x_1, x_2, x_3)$, where x_1 and x_2 are the in-plane components and x_3 is the through-plane component. Let t_{es} be the time at end-systole, and let d_{es} be the through-plane component of the cine DENSE displacement measurement at this time. The motion of a discrete point can be tracked in 3–D as follows (refer to Figure 5.3):

1. Select the starting point to be at end-systole, which typically corresponds to the largest measure of through-plane displacement. This forms the only true cine DENSE measurement of the 3–D position along the trajectory. The through-

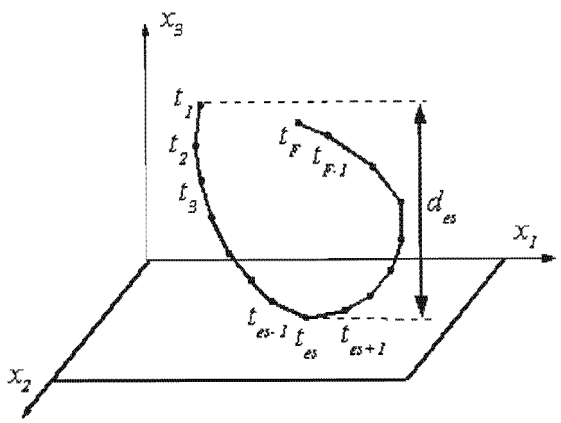

Figure 5.3: Schematic representation of a proposed 3–D tracking procedure using a single slice of conventional cine DENSE encoded in 3 orthogonal directions.

plane cine DENSE measurements for all other cardiac phases will refer to material points out of the image plane.

2. The through-plane displacement for points between t_{es} and t_1 can be estimated by working backwards in time using differences in displacement between successive frames, i.e. $t_{es-1} - t_{es}$, $t_{es-2} - t_{es-1}$, $\ldots$, $t_2 - t_1$.

3. At t_1, the point of interest should be displaced a distance d_{es} from the image plane in the x_3 direction. If this is not the case, normalize the estimates of the through-plane differences for the frames between t_1 and t_{es} to correct this.

4. The through-plane displacements for motion between t_{es} and t_F are calculated by working forwards in time using differences in displacement between successive frames, i.e. $t_{es+1} - t_{es}$, $t_{es+2} - t_{es+1}$, $\ldots$, $t_F - t_{F-1}$. If it were possible to capture the whole of diastole, then d_{es} could again be used to normalize these estimates.

The x_3 component at each time point is derived using the 2–D distance weighted linear interpolation described in Section 4.2.2. The distance d_{es} places a bounding limit on the accumulated through-plane tracking error. Note that the selection of t_{es} is not crucial, although the errors in systole will be minimised if d_{es} is as large as possible. This 3–D tracking method provides an approximation to 3–D motion, and suffers from similar drawbacks to velocity encoding tracking algorithms, namely that errors accumulate with each estimate of the difference in the through-plane displacement, and that true through-plane motion is not taken into account. Furthermore, it is difficult to obtain useful strain measurements from a set of trajectories produced using this method.

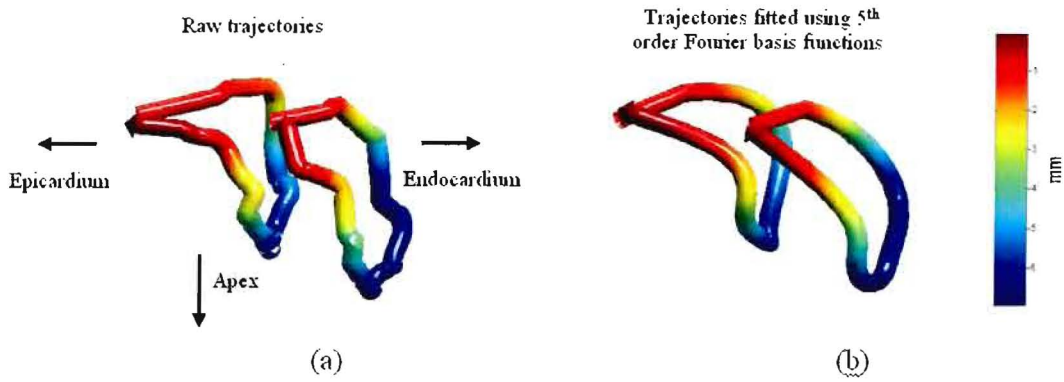

Figure 5.4: Two isolated 3–D trajectories in the septum, (a) raw, and (b) fitted using 5^{th} order Fourier basis functions. The trajectories are colour-coded by through-plane displacement (in mm).

5.2 Methods

5.2.1 True three dimensional tissue tracking using cine DENSE

The shortfalls of the 3–D cine DENSE tracking methods mentioned above can be overcome by incorporating SF into the cine DENSE sequence. SF cine DENSE simplifies the 3–D tracking process, and allows for true 2–D and 3–D tissue tracking without the need to scan adjacent slices. As shown in Figure 5.2(c), the SF DENSE 3–D displacement vectors are represented with tails all lying on the same image plane. For 3–D tracking, the search for vector tails closest to any chosen in-plane starting point at t_1 is thus reduced to 2–D. The through-plane position of a tracked point at each cardiac phase can simply be obtained using distance-weighted linear interpolation as described previously. If temporal fitting is applied in all three encoding directions, trajectories such as those shown in Figure 5.4 result. Further smoothing is required because of the low SNR associated with SF. This was done by applying a five pixel 2–D median filter to the displacements prior to tissue tracking and temporal fitting. It is worth noting here that 2–D tissue tracking using SF cine DENSE yields true 2–D motion trajectories.

5.2.2 Phantom validation

A rigid rotating phantom was used to quantify the 3–D tracking accuracy. The phantom assembly was the same as described in Section 4.2.7. The phantom was imaged in a

coronal plane, thus constraining motion to largely lie in the through-plane and one in-plane direction (see Figure 5.6(a)). Nevertheless, the analysis was carried out using data acquired with displacement encoded in three orthogonal directions. Simulated trajectories were again compared to measured trajectories. Image reconstruction was done online, and the centre of rotation of the phantom was estimated as follows.

The cine DENSE magnitude-reconstructed images of the rotating phantom eroded from the periphery towards the centre of rotation as the encoded slice moves out of the imaged slice. The offset between the centres of the encoded slice and the imaged slice thus introduces an asymmetry in the cine DENSE images of the rotating phantom. For this reason, estimating the centre of rotation can not be done using centre of mass calculations with magnitude images, or by using the mean displacement measurements. Let x, y, and z correspond to the orthogonal scanner coordinates, where z is aligned with the bore of the scanner. The phantom motion is largely constrained to lie in the (x, y) plane. The centre of rotation was estimated by reducing the rotating plane to a line by averaging the coordinates of the trajectories in the z direction. For all F frames, the gradients $m_1, \ldots, m_F$, and $x = 0$ intercepts $c_1, \ldots, c_F$ were found using least squares line fitting. The centre of rotation coordinate (x_r, y_r) was then be found by solving

$$\begin{bmatrix} -m_1 & 1 \\ \vdots & \vdots \\ -m_F & 1 \end{bmatrix} \begin{bmatrix} x_r \\ y_r \end{bmatrix} = \begin{bmatrix} c_1 \\ \vdots \\ c_F \end{bmatrix} \tag{5.2}$$

The phantom has a radius of 160 mm, and only a portion of the rotational cycle can be imaged before the encoded slice moves out of the imaged slice. For the imaging parameters used here, and a typical maximum cardiac velocity and displacement of 20 cm/sec and 2 cm, respectively, only 10 out of the 25 acquired frames could be used before the encoded slice moved out of the imaged slice. The displacement data was smoothed using a five pixel 2–D median spatial filter, and no temporal fitting is applied because so few frames are present and such a small portion of the cycle is captured.

5.2.3 Effect on planar strain

Through-plane motion is typically ignored when calculating 2–D strain using myocardial tagging, PC velocity encoding, and DENSE. However, through-plane tissue rotation introduces errors in 2–D strain measurements. Let x_1 and x_2 lie in the image plane and let x_3 correspond to the through-plane direction as before. If X_1, X_2, and X_3 are unit basis vectors aligned at t_0 with x_1, x_2, and x_3, respectively, then the 3–D deformation gradient tensor $\mathbf{F}$ is given by

$$\mathbf{F} = \begin{bmatrix} \frac{\partial X_1}{\partial x_1} & \frac{\partial X_1}{\partial x_2} & \frac{\partial X_1}{\partial x_3} \\ \frac{\partial X_2}{\partial x_1} & \frac{\partial X_2}{\partial x_2} & \frac{\partial X_2}{\partial x_3} \\ \frac{\partial X_3}{\partial x_1} & \frac{\partial X_3}{\partial x_2} & \frac{\partial X_3}{\partial x_3} \end{bmatrix} \tag{5.3}$$

The six components of Equation 5.3 that are associated with x_1 and x_2 correspond to the 2–D components of $\mathbf{F}$. These comprise the 2–D gradient deformation tensor (Equation 4.5), and two X_3 components that correspond to through-plane rotation. All six components can be calculated from the 3–D trajectories. Calculating the three components associated with x_3 require additional imaging planes, and these were set to zero here. The Lagrangian finite strain tensor $\mathbf{E}$ was then derived using $\mathbf{E} = \frac{1}{2}(\mathbf{F}^T\mathbf{F} - \mathbf{I})$ as before. The effect of through-plane rotation on both Ecc and Err was investigated for mid-ventricular short-axis views of the heart.

5.2.4 Imaging protocol

The cine DENSE scans were performed on 1.5 T Siemens Sonata and Avanto scanners (Siemens Medical Systems, Erlangen, Germany). Six normal volunteers were scanned. All subjects provided informed consent and were studied in accordance with research protocols approved by the Human Investigations Committee at the University of Virginia. Both SF and conventional cine DENSE acquisitions were acquired for each volunteer, with the number of cardiac phases varying between 15 and 22.

A four-channel body phased array RF coil was used and the subjects were imaged in a short-axis mid-ventricular view. The slice following was incorporated into the ECG-

gated 2–D EPI cine DENSE sequence[1] described in Kim *et al.* (2004). The excitation flip angles were varied to compensate for $T1$-related decay of the stimulated echo as described in Section 5.1.1. Imaging parameters included FOV = 360 × 236 mm; echo train length (ETL) = 9; effective TE = 11 ms; TR = 22 ms; matrix size = 128×94 pixels; pixel size = 2.81 × 2.81 mm; encoded slice thickness = 8 mm; imaged slice thickness = 24 mm (offset 12 mm towards the apex and 4 mm towards the base of the heart); final flip angle = 21 degrees; and displacement encoding frequency k_e = 0.1 cycles/mm. Scans for the three encoding directions were acquired in three breath holds of 14 heart beats each.

The same sequence was used for the rotating phantom, which was mounted in the head coil of the scanner and rotated about the longitudinal axis of the magnet at a frequency of 0.52 Hz. The phantom was imaged in a coronal plane with the same imaging parameters except: FOV = 280 × 280 mm; ETL = 3; effective TE = 6 ms; TR = 12 ms; pixel size = 2.19 × 2.19 mm; and displacement encoding frequency k_e = 0.15 cycles/mm.

5.3 Results

Four cardiac phases of a SF cine DENSE acquisition for a mid-ventricular short-axis slice of a volunteer are shown in Figure 5.5 in two orthogonal views. A five pixel 2–D median filter has been applied to reduce noise in each of the displacement components and the trajectories have been fitted with 5^{th} order Fourier basis functions. Each point in Figure 5.5(a) represents the centre of a voxel at end-diastole, and 5.5(b), 5.5(c) and 5.5(d) portray the positions of these points along their trajectories at different cardiac phases. The results are similar for the other data sets. LV radial thickening and circumferential shortening is comparable to that seen in tagged images. Also, the RV through-plane motion is greater than LV through-plane motion as expected. In a myocardial tagging study by Haber *et al.* (2005), through-plane motion in a mid-ventricular short-axis view is shown to be around 10.8 mm for the RV free wall, 7.2 mm for the septum, and 7.5 mm for the LV free wall. Note that the septal through-plane displacement portrayed in Figure 5.4 is in agreement with the 7.2 mm reported by Haber *et al.* The mean 3–D tracking time per frame was 0.61 seconds on a 2.8 GHz Intel Celeron with 256 MB of RAM.

[1]This was implemented by Xiaodong Zhong in the Siemens IDEA environment.

Figure 5.6(a) shows a schematic overview of the rotating phantom with the imaging slice orientation, and Figures 5.6(b), 5.6(c) and 5.6(d) show three frames of the tracked rotating phantom data. As in Figure 5.5, each point in Figure 5.6(b) represents the centre of a voxel at t_1, and Figure 5.6(c) and 5.6(d) portray the positions of these points along their trajectories at later frames. The motion is relatively uniform, and the grid maintains its structure even at the late frames. Two dimensional in-plane tracking accuracy using conventional cine DENSE was shown in Chapter 4 to be 0.24 ± 0.15 mm. The corresponding through-plane tracking accuracy using SF cine DENSE is 0.48 ± 0.25 mm, which is lower than the in-plane accuracy because the voxel dimensions are not uniform and the SNR is lower for SF cine DENSE. This tracking error was derived by comparing the through-plane distance between the theoretical and estimated trajectories for all points and at all frames. No spatial median filtering or temporal fitting was applied to obtain this measurement. It is worth noting at this point that small errors will also be introduced because the cine DENSE phase reference images are not slice followed.

An example of an Err versus time curve and corresponding Ecc versus time curve is shown in Figure 5.7. The solid lines depict strain calculated from true 2–D trajectories, and the dotted lines depict the corresponding strain calculated from the true 3–D trajectories. The difference between strain values calculated using 2–D and 3–D trajectories represents the error in strain due to through-plane rotation. For the six volunteers imaged, the end-systolic Ecc calculated using 3–D trajectories is reduced in magnitude by $3.91 \pm 0.43\%$ compared to Ecc calculated using 2–D trajectories. The corresponding end-systolic Err is increased in magnitude by $6.01 \pm 1.07\%$.

The measured SNR for SF cine DENSE with varying flip angle compensation is on average 1.7 times less than that of conventional cine DENSE without varying the flip angle. This ratio is highest at t_0, and is seen to decrease with time due to the $T1$-related decay in SNR for conventional cine DENSE. The measured SNR for SF cine DENSE with varying flip angle compensation is constant throughout the cardiac cycle.

5.4 Discussion

The major discerning parameters affecting the image quality in SF cine DENSE relative to conventional cine DENSE are the encoded and imaged slice thicknesses, and the variable flip angle settings. Both of these have been optimized by Stuber *et al.* (1999). The SNR is lower for SF cine DENSE than conventional cine DENSE, largely because previously unexcited protons that move into the imaging slice during the cardiac cycle do not contribute to the signal intensity in SF cine DENSE. The ratio of the encoded slice thickness to the imaged slice thickness is proportional to SNR and should be kept as large as possible. The performance of the phase unwrapping algorithm is unaffected by the lower SNR images, thus indicating that SF cine DENSE in a mid-ventricular short-axis view of the heart is practical and suitable for routine use. The tissue tracking accuracy and phase unwrapping efficacy as a function of SNR has been investigated in Chapter 4, and the results suggest that SF cine DENSE will be reliable for all of the encoded to imaged slice thickness ratios presented in Stuber *et al.* (1999).

By increasing the amount of phase cycling in the cine DENSE acquisition, the SNR can be improved while also suppressing the unwanted echoes. Since the SF technique is based on image subtraction, it can most easily be applied to N-point phase cycling where N is even, such as the CANSEL method where $N = 4$ (Epstein and Gilson, 2004). Increasing the amount of phase cycling has an associated increase in scan time, but this could be regulated using respiratory belt gating, navigator diaphragm tracking, or the repetitive, coached breathing pattern described in Fischer *et al.* (1994).

The effect of through-plane rotation in a mid-ventricular short-axis view is to increase the magnitude of Err and decrease the magnitude of Ecc. This effect has also been noted recently in a study using zHARP (Abd-Elmoniem *et al.*, 2006). It can be explained by considering two arbitrary adjacent trajectory starting points in the myocardium at t_1, and the corresponding position of these points at any later time t_f. The 3–D trajectories typically introduce a different through-plane displacement component to each point at t_f, and will thus result in a greater separation between the points than if only the in-plane 2–D trajectories were used. The result of this greater separation between the points is to increase positive values of strain (Err) and decrease negative values of strain (Ecc).

Tracking and temporal fitting of true 3–D myocardial motion using SF cine DENSE provides 3–D equations of motion at roughly a voxel resolution. This allows true myocardial planar strain to be calculated, but also paves the way for further measures of myocardial kinematics. For example, Figure 5.8 shows a kinesia map which portrays the total distance covered by each tracked point on a mid-ventricular view of a normal volunteer. In this case the septum moves a smaller total distance than the LV free wall, and the RV free wall moves the largest total distance. Further examples are pixel-wise representations of the time-to-onset of strain, the time-to-peak strain, and the area under the strain-time curves.

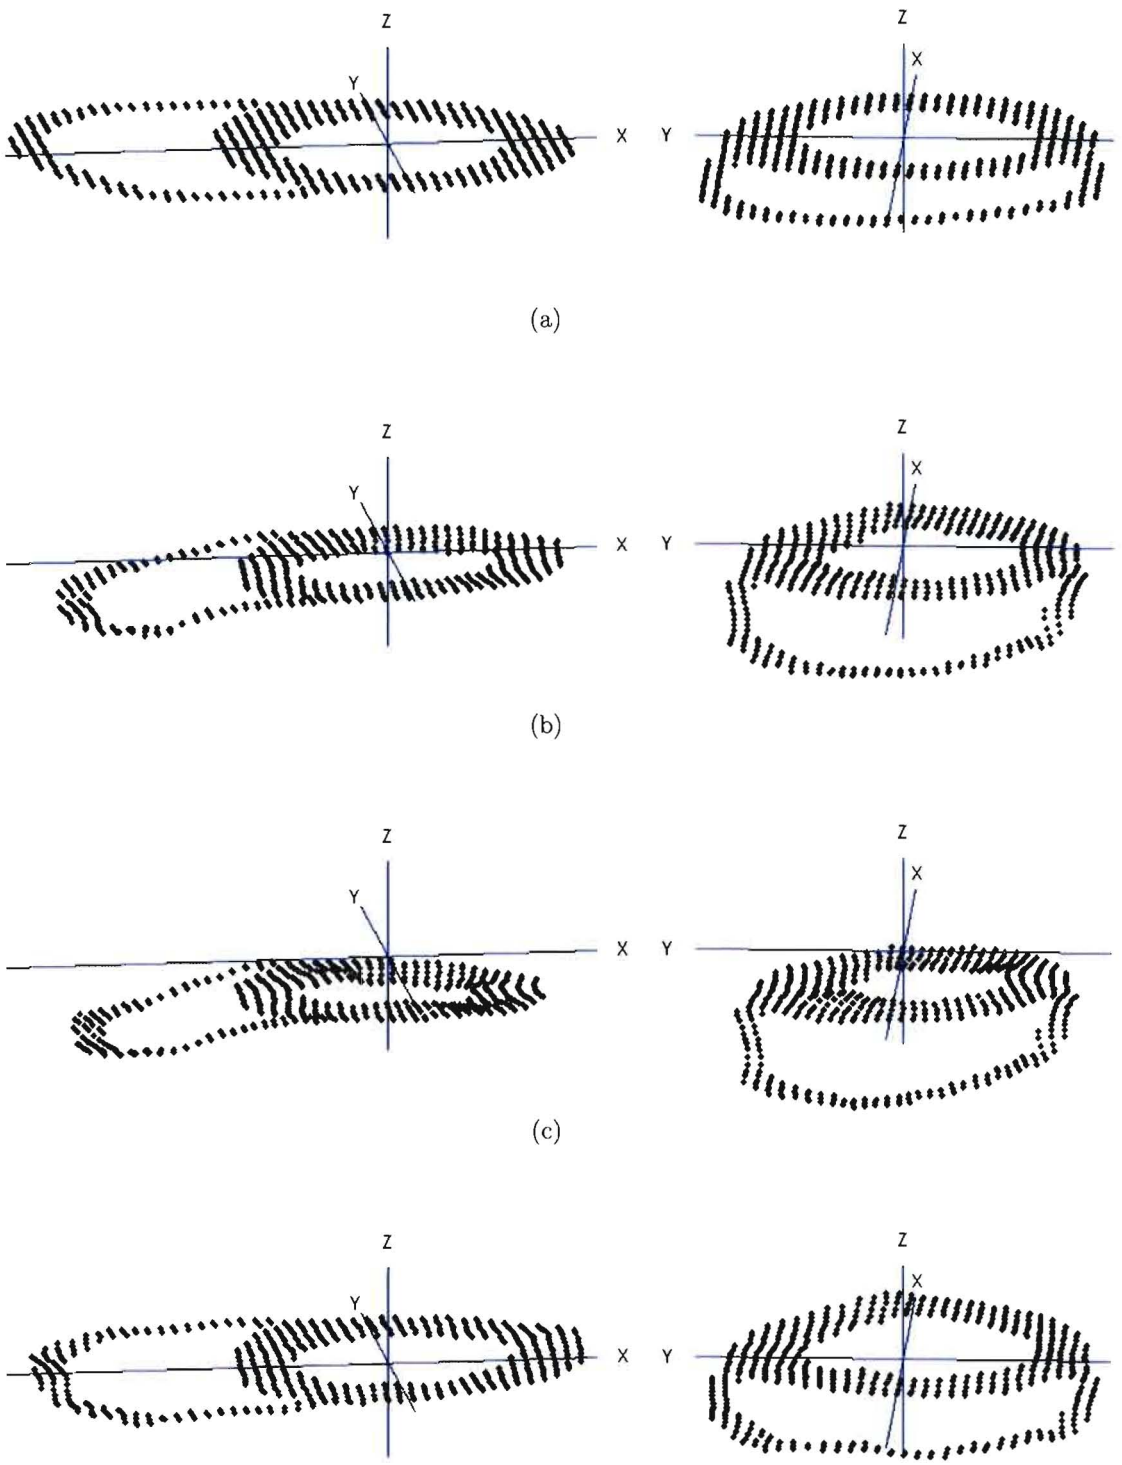

Figure 5.5: Three dimensional SF cine DENSE trajectory positions of a mid-ventricular short-axis slice of a healthy volunteer at (a) end-diastole, (b) mid-systole, (c) end-systole, and (d) late-diastole.

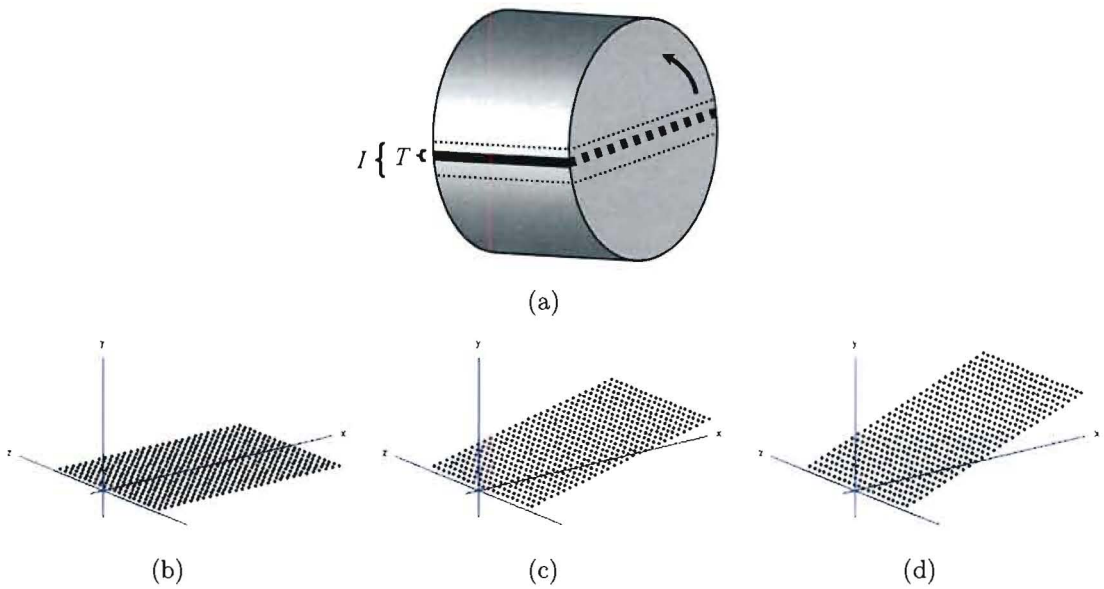

Figure 5.6: Three dimensional tracking of the rotating phantom. (a) Overview image depicting the imaging plane where T = tagged slice, and I = imaged slice. Trajectory positions at (b) frame 1, (c) frame 5, and (d) frame 9. The component of motion in the z direction (parallel to the bore of the magnet) is small, but it is still used for the tracking.

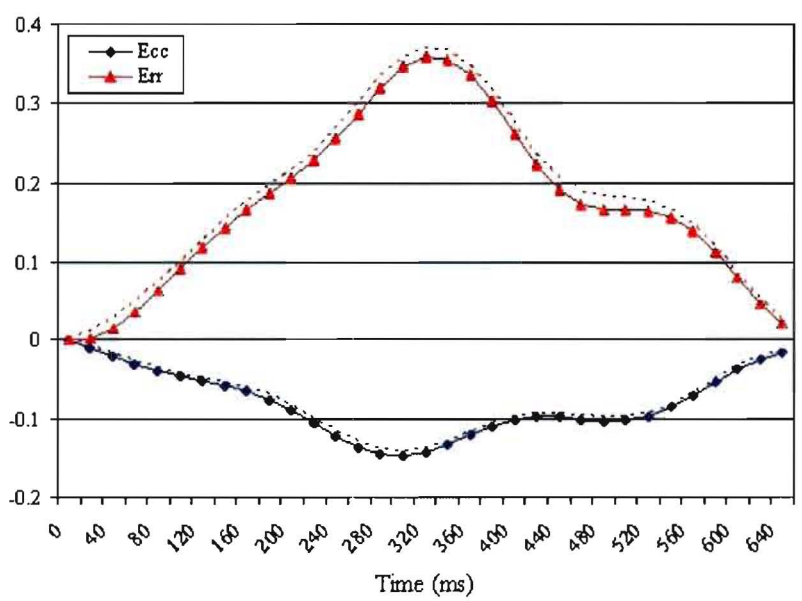

Figure 5.7: Example mid-ventricular short-axis Lagrangian radial (Err) and corresponding circumferential (Ecc) strain-time curves depict the effect of through-plane rotation on the 2–D planar strain measurements. The solid lines depict strain derived from 2–D trajectories, and the dotted lines depict the corresponding strain derived from 3–D trajectories.

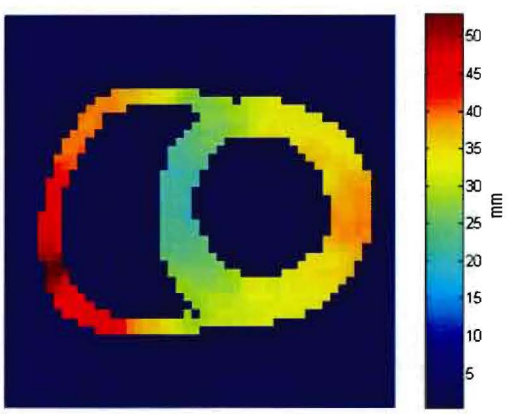

Figure 5.8: Map of myocardial kinesia portraying the total true 3–D trajectory length for each tracked point on a mid-ventricular short-axis view of a normal volunteer. In this case the septum moves a smaller total distance than the LV free wall. The RV free wall moves the largest total distance.

Chapter 6

Fast two dimensional tissue tracking and motion-based tissue classification using cine HARP

6.1 Introduction

Manual segmentation of the myocardial borders in MRI tagged images is a subjective and time consuming process, and accounts for approximately 80% of the analysis time (Montillo *et al.*, 2002). Myocardial contours are required for regional strain-time analysis, as well as for computing clinical indicators such as ejection fraction and wall thickening. The lack of a fast, automated method of segmenting tagged images is still a major factor limiting the clinical use of myocardial tagging.

As can be seen in the tagged images in Figure 2.6(a) and 2.6(b), the low intensity troughs of the tags complicate the use of intensity-based segmentation methods such as edge detection and active contour models. The appearance of tags in the blood for the first few frames further complicates the image analysis. For example, in Figure 2.6(a) it is not even possible to manually define the endocardial borders.

Earlier semi-automated techniques for defining myocardial boundaries on tagged images were based on user-guided active geometry (Kumar and Goldgof, 1994; Kraitchman *et al.*, 1995; Young *et al.*, 1995). Osman *et al.* (1999) simply made use of thresholded

harmonic amplitude images to separate foreground and background. Further automatic methods that actually define myocardial borders on tagged images include the use of 3–D greyscale morphological opening for the endocardium and active contour models for the epicardium (Montillo *et al.*, 2002); and the use of a 4–D deformable model (Montillo *et al.*, 2003). A comprehensive overview of boundary surface extraction from tagged images is given in Axel *et al.* (2005). Although the automated methods provide an accurate segmentation, they typically have two limitations: 1. they are unable to define the endocardium for the first few frames where tags are present in the blood; and 2. the methods rely on tagged image information from a number of imaging planes, and don't reliably work for a single image plane.

This Chapter introduces a novel tissue classification technique that is based on the behaviour of frame-to-frame motion trajectories. Myocardial trajectories typically exhibit relatively ordered behaviour, whereas 'noise' trajectories located in the enclosed blood or surrounding pulmonary cavity behave in a more chaotic fashion. A fast 2–D tracking algorithm is presented for cine HARP enabling the rapid quantification of interpolated motion trajectories. A number of motion-based tissue classification methods are proposed to discern between coherent and noisy trajectories. The accuracy of the method was estimated based on manually drawn contours, and a theoretical phantom simulation was implemented to determine the feasibility of the proposed tissue classification technique. This work was done prior to the work on cine DENSE, which shows greater promise for routine clinical use. As a result, comparably little validation was carried out on *in vivo* HARP data.

6.2 Methods

6.2.1 Imaging protocol

Even though the original HARP analysis allows one to measure the phase in the HARP images at a pixel resolution, the true spatial resolution of tracked material points in the myocardium is limited. The use of bandpass filters to extract the spectral peaks suggests that the tracking resolution of HARP images is equal to the intrinsic Fourier resolution defined by the filter size. The tracking resolution can be improved by using CSPAMM, which allows one to isolate half of k-space in the HARP filtering process.

This is, however, accompanied by the disadvantage that CSPAMM takes twice as long to acquire as conventional SPAMM.

Three volunteers were studied using a 1.5 T Siemens Symphony Quantum scanner (Siemens Medical Systems, Erlangen, Germany). The subjects provided informed consent and were scanned with the steady-state free precession (SSFP) CSPAMM tagging sequence described in Zwanenburg *et al.* (2003). A four-channel body phased array RF coil was used and the images were acquired in a mid-ventricular short-axis imaging plane. Imaging parameters included FOV = 300×300 mm; TE = 2.3 ms; TR = 4.7 ms; matrix size = 256×77; image size = 256×256 pixels; slice thickness = 6 mm; cardiac phases = 14; and tag spacing 9 mm.

Images were analysed off-line using MATLAB (The Mathworks Inc, Natick, MA).

6.2.2 Fast two dimensional tissue tracking for cine HARP

The HARP displacement fields for single frame differences, such as the example shown in Figure 2.13, provide limited information about myocardial motion. Osman *et al.* (1999) describe a method for tracking phase coordinates from frame to frame in cine HARP images. They make use of a Newton-Raphson iterative procedure to locate the closest set of identical harmonic phase coordinates in successive frames of a tagged series. Although this method is accurate, it becomes computationally expensive if a vast number of points are to be tracked.

The cine HARP displacement fields for all frame differences have vector tails lying on a unform grid. A fast and straightforward means of tracking points from frame to frame using cine HARP images can be achieved by adjoining displacement field vectors from end to end. In this study, tracking is commenced at all pixel centres on the first frame (t_1), but it could be implemented at any spatial resolution. The positions of the displacement vector heads at any later frame t_f form the starting points for displacement vectors at t_{f+1}. Vectors emanating from these starting points are obtained by 2–D distance-weighted linear interpolation using the 3 vectors at t_{f+1} with tails closest to each starting point. Note that this is different to the 2–D cine DENSE tracking proposed in Chapter 4, where the starting coordinates for the 2–D linear interpolation remain fixed through time.

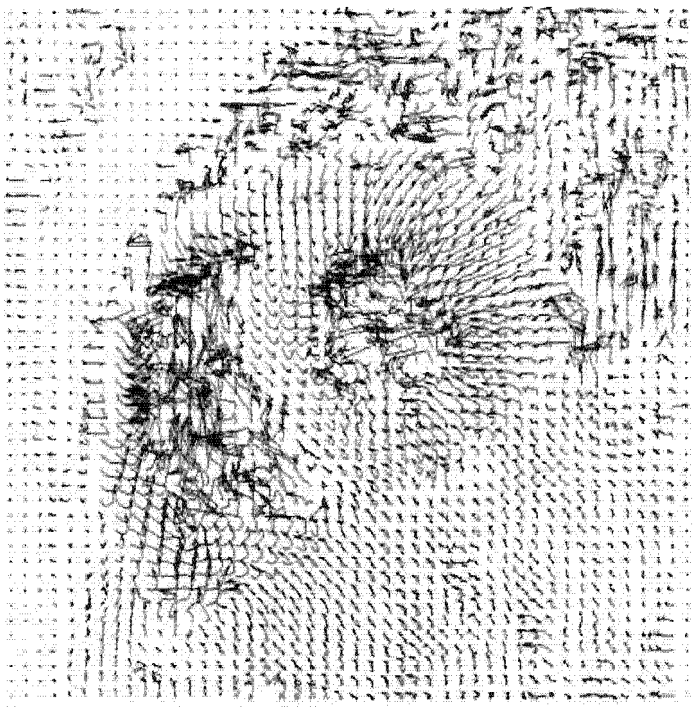

Figure 6.1: Two dimensional cine HARP motion trajectories from end-diastole to end-systole.

A Lagrangian coordinate system is preferable for describing motion trajectories. If all of the reference pixels $\mathbf{p}$ lie on a uniform grid at time $t_1 = 0$, then they can be considered to lie on the deformed grid $\mathbf{x}(\mathbf{p}, t_f)$ at time t_f. The displacement field vectors for each frame difference $t_{f'}$ are given by $\mathbf{u}(\mathbf{p}, t_{f'}) = \mathbf{x}(\mathbf{p}, t_{f+1}) - \mathbf{x}(\mathbf{p}, t_f)$. Let $\hat{\mathbf{u}}(\mathbf{p}, t_{f'})$ correspond to the vector field after 2–D linear interpolation, where the vector tails lie on the deformed grid position $\mathbf{x}(\mathbf{p}, t_f)$ at time t_f. Note that $\hat{\mathbf{u}}(\mathbf{p}, t_{1'}) = \mathbf{u}(\mathbf{p}, t_{1'})$. The motion trajectories are shown by plotting $[\hat{\mathbf{u}}(\mathbf{p}, t_{1'}), \ldots, \hat{\mathbf{u}}(\mathbf{p}, t_{(F-1)'})]$ on the same axes, where F is the last frame of the series. Figure 6.1 shows an example set of trajectories for a tagged series from end-diastole to end-systole.

6.2.3 Motion-based tissue classification

A striking feature in Figure 6.1 is the chaotic behaviour of the trajectories in the pulmonary cavity and blood pools. In the pulmonary cavity, this is because no tags can be placed in spaces containing air and hence any readings of harmonic phase can be classified as noise. The tags initially placed in the blood pools are washed away as blood is ejected from the ventricles, resulting in similar nonsensical harmonic phase values for these areas. It is proposed here that the chaotic behaviour of the vectors in the lungs

and blood pools, versus the more ordered behaviour of the vectors in the myocardium, be used as a means of differentiating between these regions.

The trajectories are represented here as two sets of three dimensional matrices containing the spatiotemporal coordinates of $\mathbf{x}(\mathbf{p}, t_f)$ for all frames $f = 1, \ldots, F$. One matrix thus contains coordinates for the x_1 direction, and the other contains coordinates for the x_2 direction. The two sets of 2–D matrix entries corresponding to t_1 thus simply describe the initial grid point coordinates. The following sections describe a number of methods devised to segment tagged images by operating on these matrices. By virtue of this trajectory data structure, the tissue classification results are most easily reflected at t_1, on the initial grid.

6.2.3.1 Spatial derivatives

For each frame, motion vectors behaving in a manner similar to their neighbours will have low spatial derivatives, whereas disordered behaviour will result in higher values. Such a tissue classification can be obtained by calculating

$$M(x_k(\mathbf{p}, t_1)) \approx \sum_{f'=1}^{F-1} \left[\left(\frac{\partial \hat{u}_k(\mathbf{p}, t_{f'})}{\partial x_1} \right)^2 + \left(\frac{\partial \hat{u}_k(\mathbf{p}, t_{f'})}{\partial x_2} \right)^2 \right] \tag{6.1}$$

for $k = 1, 2$, and then summing $M(x_1(\mathbf{p}, t_1))$ and $M(x_2(\mathbf{p}, t_1))$. Large values of M should correspond to the lungs and blood pools and smaller values of M should correspond to myocardium and stationary tissue.

6.2.3.2 Changes in vector components with time

Noting changes in vector components along a single trajectory with time will give an indication of the eccentricity of its behaviour as it moves within the continuum. Such a mask can be derived as follows

$$M(x_k(\mathbf{p}, t_1)) \approx \sum_{f'=1}^{F-2} \left[\hat{u}_k(\mathbf{p}, t_{f'+1}) - \hat{u}_k(\mathbf{p}, t_{f'}) \right]^2 \tag{6.2}$$

for $k = 1, 2$, and then summing $M(x_1(\mathbf{p}, t_1))$ and $M(x_2(\mathbf{p}, t_1))$ as before. Here one is inherently limited to $F - 2$ operations. A moving point with large changes in trajectory

angle and/or displacement between frames will thus result in a large value of M, and should correspond to points in the lung and blood pools. Vectors maintaining a similar magnitude and direction between frames will yield a small value of M and should be more characteristic of the myocardium and stationary tissue.

6.2.3.3 Nearest neighbours: distance

This approach identifies short range correlations in the vector field by measuring the deviation of distances from the mean distance between a point and its eight nearest neighbours. If the neighbours remain equidistant from the point with time, then the deviation is zero. If the tissue is deformed, then the deviation increases in proportion to the amount of deformation. The more ordered behaviour of the vectors within the myocardium implies that they maintain the starting uniform grid shape better than the random vectors. These nearest neighbour operations are performed for each frame in a series and the results are squared and summed as for the previous methods.

6.2.3.4 Nearest neighbours: angle

This is similar to the previous method, except that the vector angle $\arctan(\hat{u}_1(\mathbf{p},t)/\hat{u}_2(\mathbf{p},t))$ between the point of interest and its nearest neighbours is used instead of the distance $|\hat{\mathbf{u}}(\mathbf{p},t)|$. The rationale behind this is analogous to the previous method.

6.2.3.5 Two dimensional bandpass elliptical k-space filter

The rapid variation in both magnitude and angle of trajectories in the lung and blood pools relative to the myocardium and stationary tissue imply that a difference exists in the high frequency components in the Fourier domain. Performing a highpass filtering operation on the spectral information should thus isolate image regions corresponding to the lungs and blood pools. However, the presence of other spectral peaks can introduce unwanted artifacts, and a k-space bandpass filter is thus used to capture only a portion of this high frequency spectral data.

The filtering operation is implemented as follows: Two sets of x_1 and x_2 trajectory matrices are obtained, one for the optimal HARP elliptical filter size $\hat{\mathbf{u}}_{inner}(\mathbf{p},t)$, and

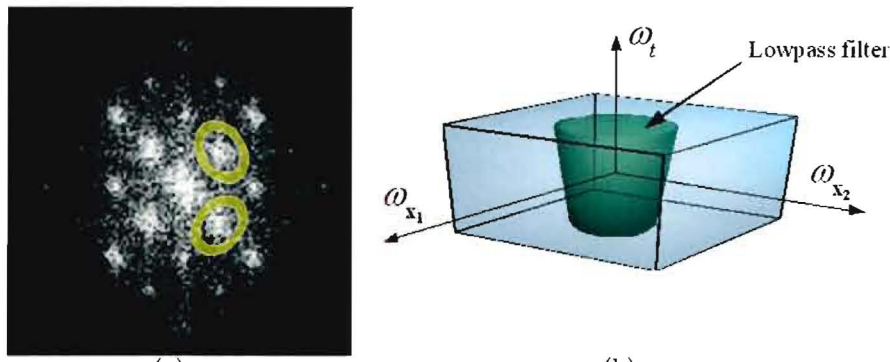

(a) (b)

Figure 6.2: (a) Bandpass elliptical filters in k-space for grid tagging. (b) Segmentation using a cylindrical spatiotemporal highpass Fourier filter. The 3–D Fourier space is represented as a cube with frequency axes ω_{x_1}, ω_{x_2} and ω_t. The low frequency spectral information is removed using a cylinder centred at $\omega_{x_1} = \omega_{x_2} = 0$.

one for a larger (or smaller) elliptical filter $\hat{\mathbf{u}}_{outer}(\mathbf{p}, t)$. Optimised elliptical filters for HARP are presented by Osman and Prince (2000a) and Osman and Prince (2000b). The trajectory components of the two data sets are subtracted from each other, and spatial derivatives are performed on the result as done in Equation 6.1. This filtering is demonstrated for grid tagging in Figure 6.2(a), where information is isolated in elliptical bands for the two orthogonal tag directions.

6.2.3.6 Three dimensional highpass cylindrical Fourier filter

This method attempts to incorporate aspects of both the spatial and temporal evolution of the trajectories in a single filtering operation. Performing a 3–D Fourier transform on both x_1 and x_2 motion component matrices of the trajectory field results in two matrices with frequency coordinates ω_{x_1} and ω_{x_2} corresponding to the two in-plane coordinates, and a third frequency coordinate ω_t corresponding to time. This frequency data is represented as a box in Figure 6.2(b). The spectral information describing the motion of all pixels in the image is centred about $\omega_{x_1} = \omega_{x_2} = \omega_t = 0$.

A 3–D highpass filter would involve removing a specified volume about $\omega_{x_1} = \omega_{x_2} = \omega_t = 0$. The number of frames F is typically less than the number of pixels in the x_1 and x_2 directions. When looking at the spectra for constant values of ω_t, it is seen that the spectral spread does not decrease significantly for high positive and negative values of ω_t. A 3–D cylindrical highpass filter is thus appropriate to remove the low

frequency components of the motion. This is represented graphically in Figure 6.2(b). An exponential roll-off is applied to reduce the Gibbs ringing that is introduced by the filtering process. If d is the distance from $\omega_{x_1} = \omega_{x_2} = 0$, r is the filter radius, and s is the image size, then the filter can be described by

$$F(d) = \begin{cases} 0, & d \leq r \\ 1 - e^{-(d-r)/(s\sigma)}, & d > r \end{cases} \tag{6.3}$$

The constant σ (unitless) defines the rate of decay of the exponential. The higher σ, the less the ringing, but the poorer the contrast-to-noise ratio (CNR) between the two regions being segmented. Using $\sigma = 0.05$ was found to suppress the ringing sufficiently.

The filtered image is then transformed back to the spatial domain using a 3–D inverse Fourier transform, and a mask is obtained using spatial derivatives as described previously.

6.2.4 Segmenting the remaining frames

The segmentation of the remaining frames of the series was achieved by following the trajectories that were classified in the first frame. The image derived using one of the classification methods just introduced is first lowpass filtered using a 3×3 pixel Gaussian kernel with a standard deviation of 1 pixel. A threshold that depends on the classification method is then applied to dichotomise the image, thus defining a binary mask of the myocardium and stationary tissue at the first frame.

Each trajectory point at each remaining cardiac phase is then represented by a 2–D Gaussian function with an integrated intensity of 1 pixel and a spread of 1.25 pixels. Thresholding the resulting images at a value of 0.5 gives the position of the myocardium and stationary tissue at all other cardiac phases.

6.2.5 Error analysis and parameter optimisation

The accuracy of this cine HARP tissue tracking method was gauged by comparing the phase coordinates of points at t_1 with those at their corresponding tracked positions at

t_F. The slope of the harmonic phase image at t_F in the directions of x_1 and x_2 was then used to give an estimate of the tracking error for each point.

A means of evaluating this motion-based tissue classification was necessary to quantify the accuracy of the methods relative to true contours. In quantifying the segmentation of medical images, it is customary to compare the computed boundaries with expert-drawn contours. The process of obtaining expert-drawn contours is complicated by the presence of tag lines in the blood pools at the first few frames of a series. The third frame of the tagged series was thus used to validate the segmentation, as the tag lines had then typically washed out of the blood pools.

To reduce the errors inherent in manual segmentation, specific regions were identified based on the presence of clearly defined borders. Regions of the myocardium as well as noisy regions adjacent to the myocardium (blood and lungs) were selected. Figure 6.3 shows the six manually delineated areas used here, where three are for the myocardium and three are for the non-myocardium (lungs and blood pools). Areas corresponding to stationary tissue could not be included, as they are classified in the same category as the myocardium.

The areas for the myocardium were then combined, as were the areas for the non-myocardium. Udupa *et al.* (2002) suggest the following measures for the assessment of segmentation accuracy

$$False\ negative\ volume\ fraction = \frac{|T_m - E_m|}{|T_m|} \tag{6.4}$$

$$False\ positive\ volume\ fraction = \frac{|T_o - E_o|}{|T_o|} \tag{6.5}$$

Here E corresponds to the segmentation estimate by the method being investigated and T corresponds to the true segmentation by an expert. The subscript m refers to the combined myocardium areas (Figure 6.3(a)), and the subscript o refers to combined non-myocardium areas (Figure 6.3(b)). The contours were delineated by the author.

The false negative volume fraction ($FNVF$) corresponds to the fraction of myocardium missed by the segmentation. The false positive volume fraction ($FPVF$) corresponds to the amount of tissue falsely identified as myocardium. It is desirable that values of both $FNVF$ and $FPVF$ are low. For these initial results, these measures were given equal weighting and the final measure of segmentation error used was simply $FNVF+FPVF$.

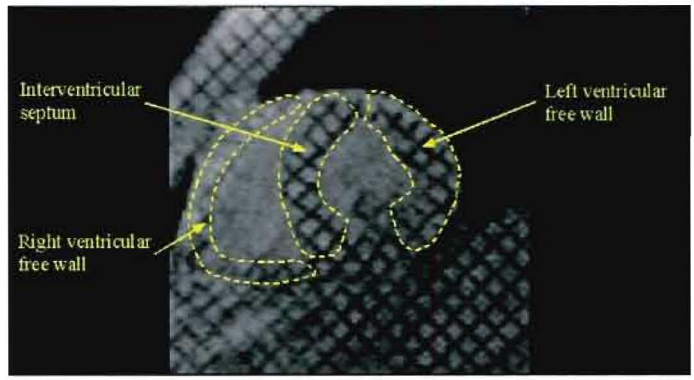

(a)

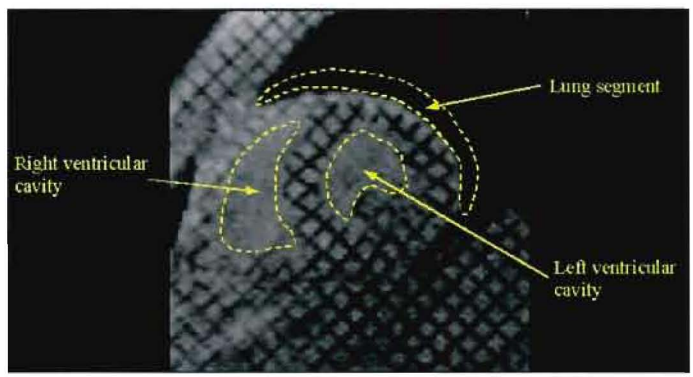

(b)

Figure 6.3: Practical regions of interest for (a) the myocardium, and (b) regions adjacent to the myocardium.

The various segmentation methods needed to be optimised before any comparison could take place. The optimisation involved varying a number of parameters that could affect the shape of the final mask. These included the number of frames used for the masking procedure, the threshold level used to dichotomise the binary mask, the ratio of inner to outer elliptical filter sizes (in the case of the 2–D bandpass elliptical filter), and the radius of the cylinder for the 3–D highpass cylindrical filter method. Although using a large number of frames results in a greater contrast between the myocardium and non-myocardium, the errors introduced due to noise and tag fading accumulate. The segmentation method should preferably use as few frames as possible.

6.2.6 Theoretical phantom study

A phantom simulation was conducted to examine the relationship between tagged image pixel size and achievable segmentation accuracy. The phantom mimics the LV in a short-

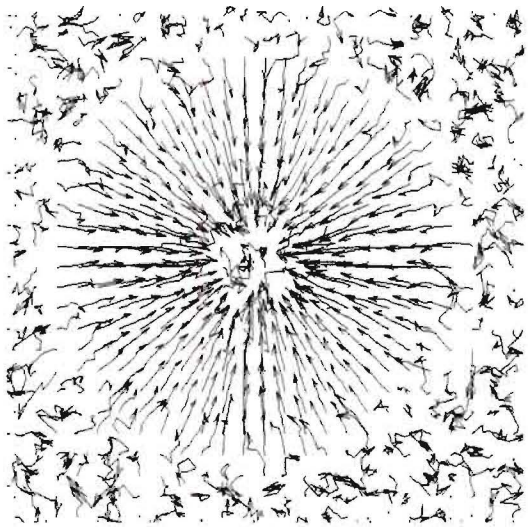

Figure 6.4: Phantom simulation describing a contracting torus.

axis view and takes the form of a contracting torus. Figure 6.4 shows the simulated phantom trajectories. The inner surface of the torus contracts more rapidly than the outer surface, thus exhibiting the wall thickening associated with normal ventricular contraction.

A number of noise components were introduced. These are described by their standard deviations and are implemented by scaling the standard deviation of simulated Gaussian noise. The first of these noise components is a background noise, which was applied to each vector component within the trajectory structures. An estimate of this was obtained by examining the *in vivo* x_1 and x_2 components of a relatively stationary portion of tissue. The subcutaneous tissue in the chest wall was used here. The distribution could be well approximated by a Gaussian envelope, and the resulting worst-case standard deviation (σ) of vector component lengths is 0.24 pixels.

The second noise component describes the vectors in the lungs and blood pools. Noise samples taken in these regions reveal Gaussian distributions with $\sigma = 2.84$ pixels. A third noise component was introduced to simulate the effects of tag fading. Groot-Koerkamp *et al.* (1999) describe the decay of CNR between tag peaks and troughs for a typical tagged data set. In their data, the CNR decays from 0.96 to 0.59 in 0.36 seconds. Assuming that the contractile motion takes place within 0.36 seconds, this gradient was used to estimate the effective noise introduced by tag fading.

Ventricular volume normally varies between 160 ml at end-diastole, to 80 ml at end-systole. This contraction ratio was modelled by reducing the cavity area by a factor of two, i.e. reducing the radius of the endocardium by $\sqrt{2}$. The velocity of contraction of the epicardial border was scaled to yield the typical value of 30% myocardial thickening at end-systole.

Figure 6.4 represents the resulting phantom for a 40×40 pixel image with the parameters introduced above. A typical LV end-diastolic epicardial radius of 50 mm thus results in a pixel size of 1.25×1.25 mm. The variables for segmentation using this phantom data were the image resolution, the threshold level, and the number of frames of displacement information. The same measure of error, $FNVF + FPVF$, was used for this study, and the third frame was analysed to maintain consistency with the previous method.

6.3 Results

Figure 6.5 shows the images that result when applying the six proposed methods of discerning between noisy and coherent trajectories. For Figure 6.5(e), the inner ellipse for the bandpass filter was set to 0.9 times the optimal size proposed in Osman and Prince (2000a). The Gibbs ringing artifacts are still evident near the borders of Figure 6.5(f).

The first image in Figure 6.6 shows the smoothed and thresholded binary mask corresponding to Figure 6.5(a). The remaining three images are three successive frames that were obtained by following the mask using the motion trajectories.

The tracking error for one of the data sets is depicted in Figure 6.7. The mean tracking error for all three data sets is 4.48 ± 5.03 pixels for the entire image and 1.36 ± 2.23 pixels for the myocardium. Masks based on spatial derivatives were used to define the myocardium here.

Table 6.1 shows the optimal settings for the various classification methods. Note that the threshold is described on a per-frame basis, since the contrast between the myocardium and non-myocardium increases in proportion with the number of frames used. The processing times are for determining the mask of the first frame of a series. The image size is 256×256 pixels and the processing was done on a 1.6 GHz Pentium 4 with 256 MB RAM. Surprisingly, the elliptical bandpass filter method gives the most accurate

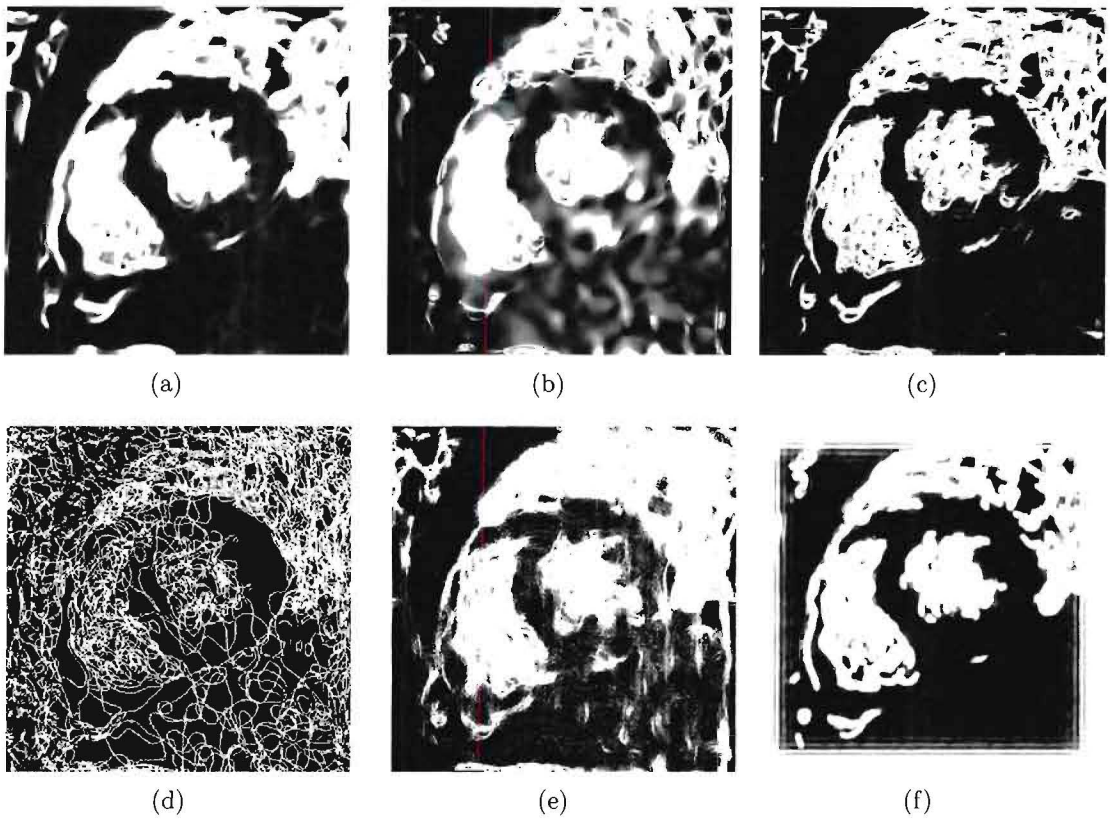

Figure 6.5: Masks of the myocardium and stationary tissue based on: (a) spatial derivatives; (b) changes in vector components with time; (c) eight nearest neighbour distance deviations; (d) eight nearest neighbour angle variations; (e) bandpass elliptical k-space filter, using an inner elliptical filter 0.9 times the optimal size; and (f) 3–D cylindrical Fourier filter.

Figure 6.6: Masks based on spatial derivatives followed from t_1 using the motion trajectories. Four cardiac phases during systole are portrayed.

result. However, the processing time for this method is double that of the other methods since the trajectories also have to be calculated for the inner elliptical filter. The second most accurate result corresponds to the spatial derivative method. The values of error based on $FNVF$ and $FPVF$ all roughly lie between 10 and 30%. This is inadequate for diagnostic purposes, where aspects such as wall thickening would be measured.

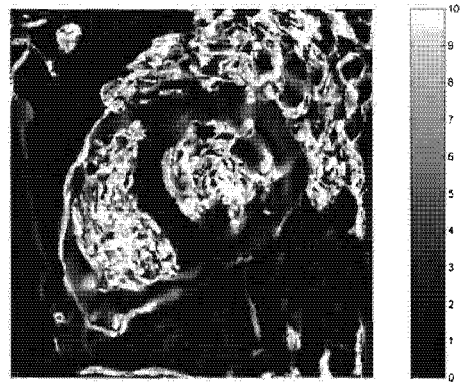

Figure 6.7: Example of the tracking error (in pixels), represented on the first frame of a tagged series.

Table 6.1: Optimal number of frames and threshold level per frame for the various motion-based tissue classification techniques. The area error measurements are percentile and time is measured in seconds. F_1 is the 2–D bandpass filter ratio and F_2 is the 3–D cylindrical highpass filter radius as a percentage of image size. *This time excludes the computation of a second cumulative displacement field data set for the 2–D bandpass filter.

Method	Frames	Thresh.	$FNVF$	$FPVF$	Error	Time	F_1	F_2
Spatial derivatives	5	0.10	21.1	12.0	33.1	6.34	n/a	n/a
Temporal variation	3	6.10	17.8	26.4	44.2	2.69	n/a	n/a
Neighbours: distance	5	0.08	21.6	13.0	34.6	4.94	n/a	n/a
Neighbours: angle	7	0.15	26.9	31.7	58.6	3.56	n/a	n/a
2–D bandpass filter	7	0.02	16.0	11.6	27.6	8.56*	0.55	n/a
3–D highpass filter	7	0.08	22.7	15.0	37.7	45.47	n/a	5.38

For the theoretical phantom, three frames were found to be optimal for the segmentation. This is not surprising since the detrimental effects of the noise on the segmentation accumulates with the number of frames used. This figure would likely increase if tag-washout were included in the model. Figure 6.8 shows the phantom segmentation error as a function of the threshold level and the pixel size, with the number of frames used fixed at three.

6.4 Discussion

A number of differences become apparent upon closer inspection of the segmented masks in Figure 6.5. Observe, for example, the differences in shape and position of the papillary muscle at 2 o'clock on the left ventricle. The method that uses only spatial derivatives could in principle be improved by performing spatial derivatives in any number of radially oriented directions. A shortfall of this method is that is doesn't take into account

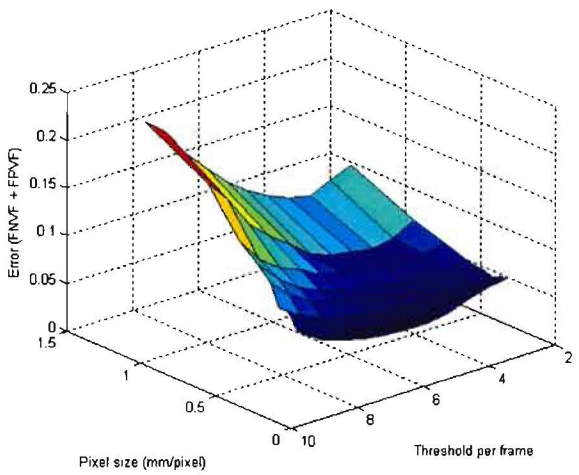

Figure 6.8: Simulated phantom segmentation error as a function of threshold level and pixel size.

the changes in individual vector components with time. If the motion of the vectors in the lungs and blood pools is random, then not only should there be differences in the behaviour of adjacent vectors, but also in the temporal behaviour of vectors in individual trajectories. Similarly, the disadvantage of using changes in vector components with time is that each material point is treated without regard to its surroundings. It seems advantageous to use aspects of both spatial and temporal discretion, although combining the two methods is a challenging task. A downfall of using the temporal behaviour of individual vectors is evident in Figure 6.5(b), where an incorrectly segmented portion of tissue can be seen at 6 o'clock on the left ventricle. This can be explained with reference to Figure 6.1, where the vectors within this portion of the myocardium display a tight circular motion.

Using the eight nearest neighbour distance operation also yields a clear result. A problem with the nearest neighbour operations are that portions of the myocardium that undergo considerable deformation could be incorrectly assigned. The corresponding angle operation yields strange artifacts, for which no explanation has been found. The method that uses the 2–D elliptical bandpass filter is expected to produce an erroneous estimate because of the small amount of spectral information used. Finally, even though the method that uses the 3–D cylindrical highpass filter produces a greater error than some of the other methods, it offers a solution to integrating both spatial and temporal aspects of vector discretion, and it also has no effect on the HARP filter.

As expected, the phantom segmentation error decreases as the pixel size is decreased. For the 1.25×1.25 mm simulated pixel size, the total segmentation error is 11.9%. This is considerably less than the figures depicted in Table 6.1. A portion of this discrepancy is due to errors in the manual segmentation for the *in vivo* studies, which is a difficult task given the paucity of the tag lines. The greatest contributing factor to this error is, however, the limited image resolution and the tracking resolution of HARP. For a pixel size of 0.21×0.21 mm, the error predicted by the phantom only reduces to 3.9%. The scanning time increases as the square of the increase in resolution, so halving the pixel size increases the acquisition time fourfold. It is clear from these results that this motion-based tissue classification technique is not accurate enough for clinical use.

The analysis described herein only applies to 2–D, or apparent motion. The effects of through-plane motion also introduces errors into this segmentation process. Figure 6.9 shows an example of mask borders superimposed onto tagged images for four systolic cardiac phases. The effect of through-plane motion is made evident in these images by the appearance of an additional tag line in the RV, as shown by the arrows in Figures 6.9(c) and 6.9(d). Note again that the endocardial border is completely invisible to the human eye for the first frame, when the tag lines appear in the blood pools.

Tagged images are typically preprocessed before applying any intensity-based segmentation methods. Techniques have been developed to compensate for RF field inhomogeneities, and for tag fading due to T_1–relaxation. The RF nonlinearities can be accommodated for using the scale-based method suggested by Zhuge *et al.* (2002), and the tag fading can be reduced using the histogram-based method of contrast manipulation introduced by Groot-Koerkamp *et al.* (1999). Both of these methods were implemented and their effect on the segmentation was investigated. The combined effect of both preprocessing steps reduced the error by a fraction of a percent. This is not surprising since HARP is a phase-based method, and the HARP filtering operation takes place in the frequency domain.

Figure 6.10 shows the trajectories emanating from a mask of the myocardium and stationary tissue. This image is useful to give an indication of wall motion abnormalities. Note that no user interaction is required to obtain this information from the tagged images.

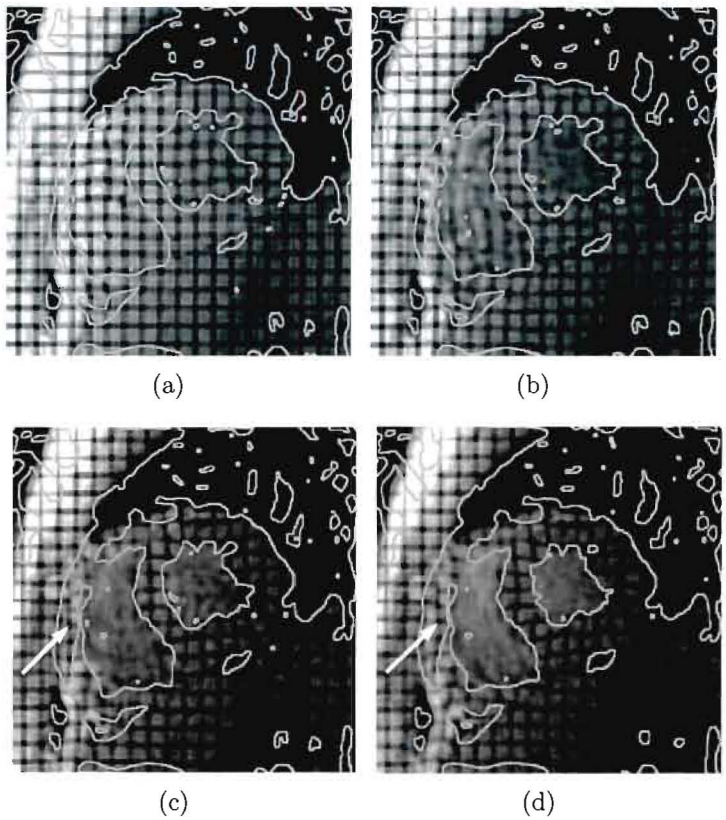

Figure 6.9: Myocardium mask overlaid onto a summation of vertical and horizontal tagged images. (a) Frame 1, (b) frame 3, (c) frame 5, and (d) frame 7. The arrows indicate an additional tag line that has arisen because of through-plane motion.

The accuracy of the segmentation was investigated only for a single frame of the tagged series. Tag fading has the effect of degrading the segmentation accuracy with time. Considering that the accuracy has been deemed inadequate for the third frame of a series, the segmentation accuracy of subsequent frames was not even considered.

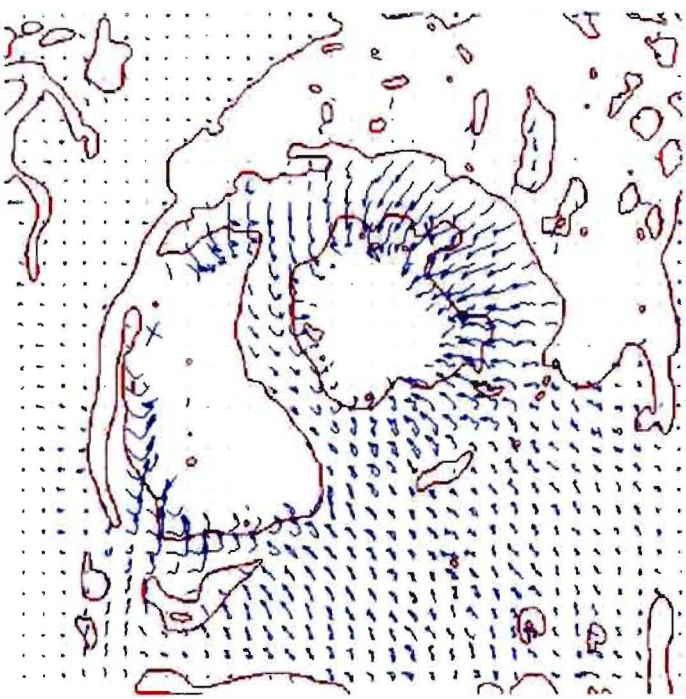

Figure 6.10: Motion trajectories shown emanating from a mask based on spatial derivatives.

Chapter 7

Motion-guided segmentation using cine DENSE

7.1 Introduction

Defining epicardial and endocardial contours is an integral step in quantifying regional cardiac wall motion for all of the MRI methods introduced in Chapter 2. In cine DENSE, these contours are currently manually delineated for all cardiac phases. This is a laborious process, and currently forms the most time-consuming component of the cine DENSE image analysis.

Myocardial contour detection techniques based on image intensity are ill-suited to cine DENSE data for a number of reasons: 1. the image resolution is relatively low compared to other MRI cine image acquisition techniques; 2. there is a $T1$-related decay in SNR with time; 3. boundaries between the myocardium and adjacent tissue (e.g. the liver) are often indiscernible; and 4. high signal is present in the encoded blood of the first few frames, before it is washed out of the image plane.

As the heart deforms, a material point moves from its reference point $\mathbf{p} = (p_1, p_2, p_3)$ at time t_0 to a new spatial position $\mathbf{x}$ at time frame t_f. The reference map $\mathbf{p}(\mathbf{x}, t_f)$ characterizes this motion, where $\mathbf{p}(\mathbf{x}, t_0) = \mathbf{p}$. Note that this is an Eulerian frame of reference, whereas the trajectories in Chapter 6 were described in a Lagrangian frame of reference. The 2–D DENSE displacement field is given by $\mathbf{u}(\mathbf{x}, t_f) = \mathbf{p}(\mathbf{x}, t_f) - \mathbf{p}$.

It is proposed here that the encoded motion in cine DENSE images be used to guide the segmentation process. This is achieved by using the myocardial motion trajectories to project a portion of myocardium that has been manually defined on a single frame onto all other cardiac phases. The starting points of the cine DENSE displacement vectors for every pixel in the manually contoured region define the position of this region at t_0. The initial contour(s) can thus be drawn on any cardiac phase. Tracking these starting points through time identifies the position of this myocardium at all other frames. Note that this is the first reported segmentation algorithm for cine DENSE. The only related work found in the literature was by Wong *et al.* (2002), who presented a myocardial segmentation technique where velocity fields from phase contrast velocity encoding are combined with image intensity-based techniques to propagate contours from frame to frame. Although Wong *et al.* (2002) only presented results from a single study, they showed that the velocity constrained approach performs better than intensity information alone.

7.2 Methods

7.2.1 Tracking refinement

If spatiotemporal phase unwrapping is performed without pre-defined contours then regions of noisy phase produce near-randomly scaled and oriented displacement vectors. An example of an un-contoured DENSE displacement field for the region portrayed in Figure 7.1(a) is shown in Figure 7.1(b). As was shown for HARP in Chapter 6, the phase data in the blood and pulmonary cavity is comprised of noise, and the appearance of these vectors is random. The smoothly varying motion in the LV is clearly contracting towards the centre of the image. The aim here is to extract only vectors in the myocardium from such displacement fields given a set of predefined myocardial starting points at t_0. Note that only LV segmentation is discussed here because the phase unwrapping algorithm presented in Chapter 3 is not robust enough to reliably unwrap the thin RV wall without predefined contours.

Obtaining suitable motion trajectories from these noisy displacement fields requires a series of refinements to the motion tracking algorithm described in Chapter 4, where the 3 nearest vector tails were used to define vectors for 2–D interpolation. If a manually-

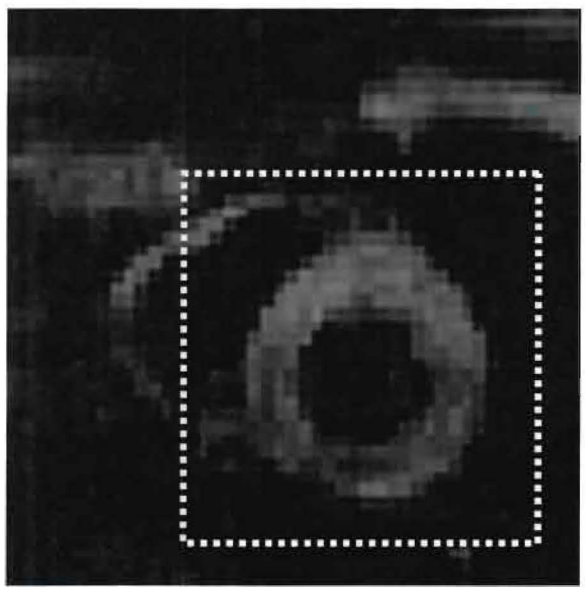

(a)

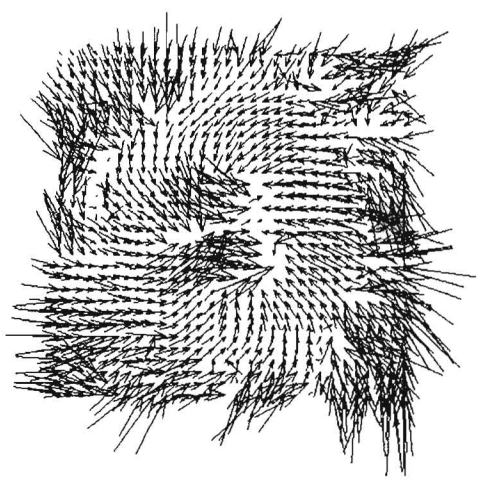

(b)

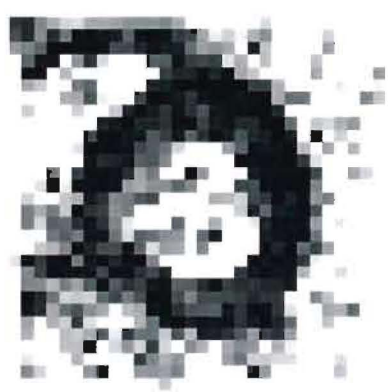

(c)

Figure 7.1: (a) DENSE magnitude-reconstructed overview image, (b) corresponding DENSE displacement field derived by unwrapping phase without pre-defined myocardial contours, and (c) the corresponding deformation mask.

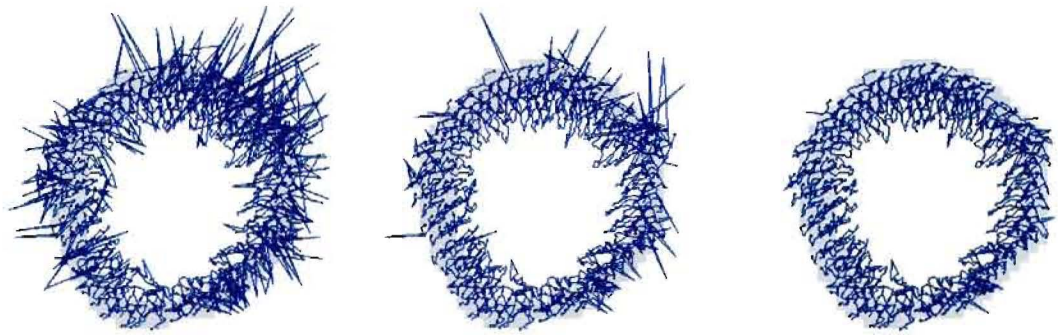

(a) (b) (c)

Figure 7.2: (a) Frame-to-frame trajectories obtained by tracking using the 3 closest vector tails, (b) frame-to-frame trajectories after applying the deformation mask, and (c) frame-to-frame trajectories after applying both the deformation mask and the heuristic vector magnitude and angle deviation criteria.

defined set of myocardial points at t_0 are tracked using this algorithm without applying predefined myocardial contours prior to phase unwrapping, then frame-to-frame trajectories such as those shown in Figure 7.2(a) result. Many of these trajectories are incorrect because noisy vectors from the lungs and blood pools intrude into the myocardium and are inadvertently being used for the 2–D interpolations.

These trajectories can be improved by applying the following three steps:

1. Remove pixels with a low SNR from the manually-contoured region of myocardium before tracking. This is done with reference to the magnitude images, where all pixels with magnitude less than 2 standard deviations below the mean are ignored. This limits the number of incorrectly labelled points at t_0, which would cause errors on every frame.

2. Remove noisy displacement vectors using a modulus deformation mask, which is obtained for each frame t_f by combining orthogonal spatial derivatives of the displacement fields as follows:

$$M(\mathbf{x}, t_f) = \frac{1}{2} \sum_{k=1,2} \left[\left(\frac{\partial u_k(\mathbf{x}, t_f)}{\partial x_1} \right)^2 + \left(\frac{\partial u_k(\mathbf{x}, t_f)}{\partial x_2} \right)^2 \right]^{1/2} \tag{7.1}$$

   The deformation mask corresponding to Figure 7.1(b) is shown in Figure 7.1(c). By setting a threshold that excludes all deformation values greater than those found in the human heart, the majority of the noisy displacement vectors are

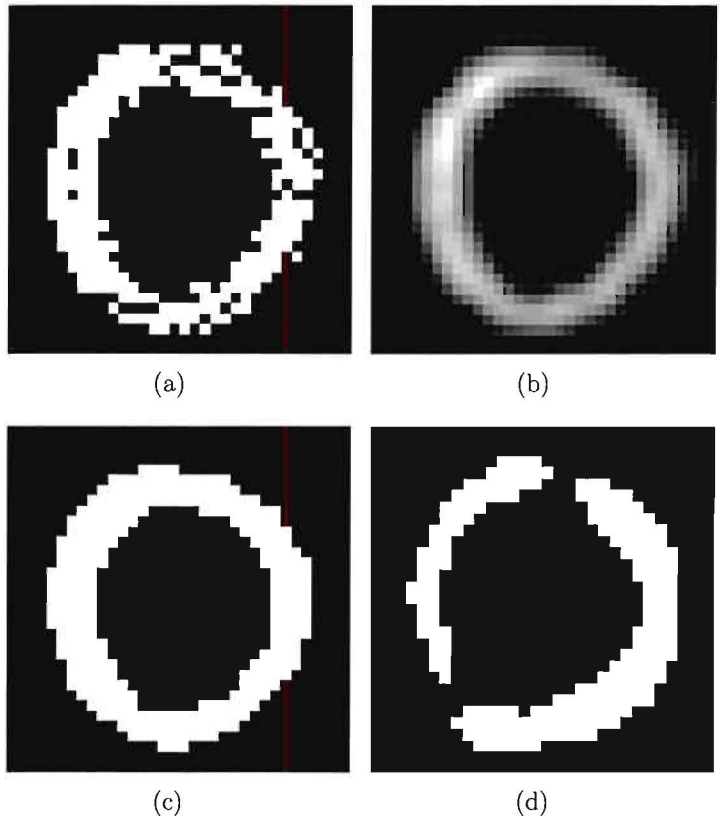

Figure 7.3: (a) Myocardial mask for a single cardiac phase obtained by rounding each trajectory position at that frame to the nearest pixel. (b) Myocardial image obtained by representing each trajectory position at the current frame by a 2–D Gaussian function with an integrate intensity of unity and a standard deviation of 1.25 pixels. (c) Applying a threshold of 0.5 to the smooth myocardial image gives a binary mask of the myocardium. (d) Transmural discontinuities appearing due to a thin wall, conservative contours, or poor data.

For E_{int} , an elastic energy was implemented to minimize $\partial v(s)/\partial ds$ and force points along $v(s)$ to be equidistant. For E_{con}, a radial ballooning force for the endocardial contour was applied, and a radial shrinking force for the epicardial contour was applied. A force was also implemented to limit the temporal evolution of the snake, providing a linear resistance to change from the contour defined at the preceding or succeeding cardiac phase. Finally, for E_{image}, based on the binary myocardial mask, an energy was implemented to oppose both the elastic energy and the radial energy. For each point $m = 1, \ldots, M$ along the arc length s, and for each direction x_k, where $k = 1, 2$, the discrete representation of these forces is given by:

$$F^i_{k\ elastic}(s_m, t_f) = w_{elastic} \times \left[\frac{x^i_k(s_{m+1}, t_f) + x^i_k(s_{m-1}, t_f)}{2} - x^i_k(s_m, t_n) \right] \tag{7.4}$$

$$F^i_{k\ radial}(s_m, t_f) = w_{radial} \times r^i_k(s_m, t_f) \tag{7.5}$$

$$\begin{aligned} F^i_{k\ binary}(s_m, t_f) =& w_{binary} \times \left[-F^i_{k\ elastic} - F^i_{k\ radial}\right] \\ & \times B\left(x^i_1(s_m, t_f), x^i_2(s_m, t_f)\right) \end{aligned} \tag{7.6}$$

$$\begin{aligned} F^i_{k\ temporal}(s_m, t_f) =& w_{temporal} \times \left[x^i_k(s_m, t_f) - x^N_k(s_m, t_{f\pm1}\right] \\ & \times \left[1 - B\left(x^i_1(s_m, t_f), x^i_2(s_m, t_f)\right)\right] \end{aligned} \tag{7.7}$$

where $w_{elastic}$, w_{radial}, w_{binary}, and $w_{temporal}$ are user-defined weights; i is the iteration number; N is the total number of iterations; f is the frame number; B is the binary myocardial mask; and the radial component $r^i_k(s_m, t_f)$ is defined as the component in the x_k direction of a unit vector oriented along the line from $(x^i_1(s_m, t_f), x^i_2(s_m, t_f))$ to $(\bar{x}^i_1(s_m, t_f), \bar{x}^i_2(s_m, t_f))$, where $\bar{x}^i_k(s_m, t_f) = \frac{1}{M}\sum_{m=1}^{M} x^i_k(s_m, t_f)$.

For all points m, the position of the snake is adjusted iteratively in the directions x_k, where $k = 1, 2$, as follows:

$$\begin{aligned} x^{i+1}_k(s_m, t_f) =& x^i_k(s_m, t_f) + F^i_{k\ elastic}(s_m, t_f) + F^i_{k\ radial}(s_m, t_f) \\ & + F^i_{k\ binary}(s_m, t_f) + F^i_{k\ temporal}(s_m, t_f) \end{aligned} \tag{7.8}$$

Before initiating the active contour on a preceding or proceeding frame, each node in the epicardial snake is moved radially outwards by 1 pixel, and each node in the endocardial snake is moved radially inwards by 1 pixel. This is done in the directions of the unit vector components $r^i_k(s_m, t_f)$. For the cine DENSE pixel size and temporal resolution used here, and the range of heart sizes studied, suitable parameters for the epicardium are: $M = 50$, $N = 60$, $w_{elastic} = 0.2$, $w_{radial} = -0.03$, $w_{binary} = 1.1$, and $w_{temporal} = 0.05$. Suitable parameters for the endocardium are $M = 40$, $N = 60$, $w_{elastic} = 0.2$, $w_{radial} = 0.05$, $w_{binary} = 1.1$, and $w_{temporal} = 0.05$. Note that w_{radial} is less in magnitude for the epicardial snake than the endocardial snake because the elastic force has a shrinking effect on a closed contour. The weight w_{binary} is chosen to be slightly greater than unity to ensure that the snake nodes do not get "stuck" inside

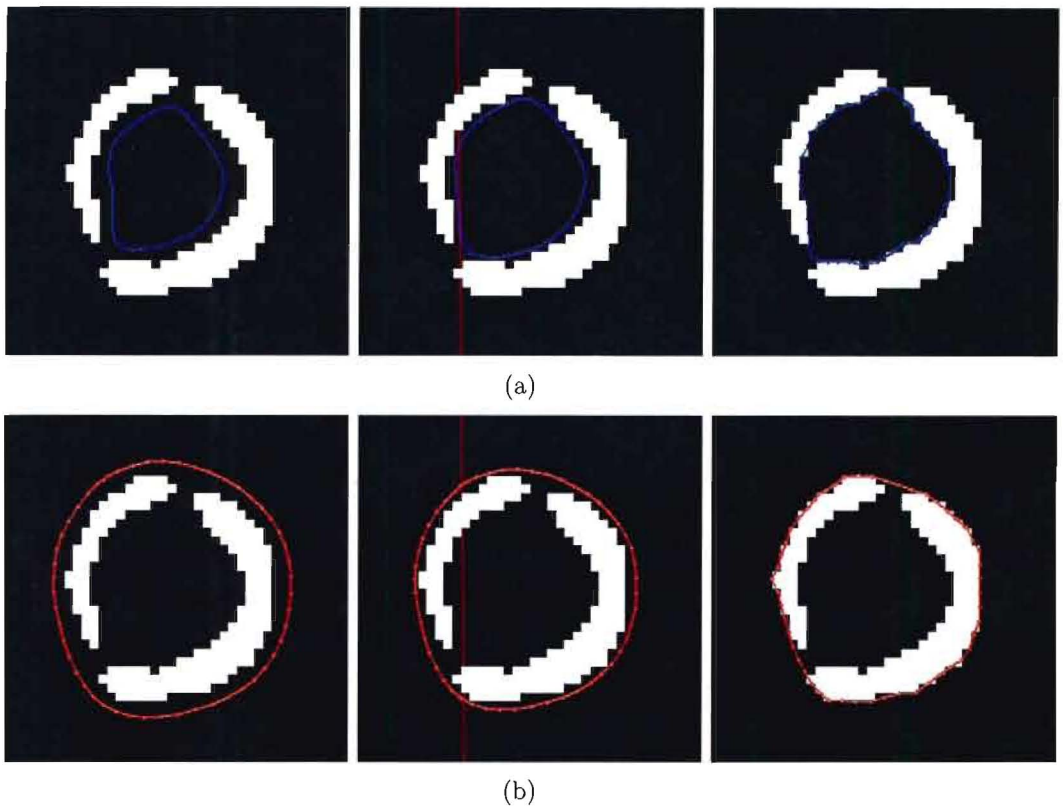

(a)

(b)

Figure 7.4: Iterations 1, 11 and 41 for (a) a ballooning endocardial active contour, and (b) a shrinking epicardial active contour.

the binary mask of the myocardium. The $(1 - B)$ term in $F_{temporal}$ ensures that the temporal force has no effect on a node that lies within the myocardium, thus making is difficult for nodes to be pulled through the myocardial mask. Figure 7.4 shows the action of these active contour models on the binary image in Figure 7.3(d).

Finally, the resulting contours (either isocontours or snakes) are smoothed by fitting periodic Fourier basis functions to each of the spatial directions x_1 and x_2 as a function of arc length s. This is given by

$$f_{x_1,x_2}(s) = \frac{1}{2}a_0 + \sum_{m=1}^{M} a_m \cos(m\omega s) + \sum_{m=1}^{M} b_m \sin(m\omega s) \tag{7.9}$$

where a_m and b_m are the Fourier series coefficients, $\omega = 2\pi/S$ is the spatial angular frequency, S is the total arc length of the contour, and M is the order of the fit. Fitting was done using least squares, and 4^{th} order functions, requiring 9 coefficients, were found to be adequate to describe the short-axis contour shapes encountered in this study. The

Fourier coefficients for a particular frame were used as starting coefficients for fitting the contours of the preceding and/or succeeding frame.

7.2.3 Defining segmentation accuracy

The algorithm was tested on 6 SF and 6 conventional mid-ventricular short-axis cine DENSE data sets. For 4 data sets, the automatic contours were compared to manually-drawn contours from two independent observers using the sum of false positive and false negative area measures. For each observer, the manual contours at mid-systole were used to initiate the motion-guided segmentation process. The false negative area measure identifies the percentage of myocardium missed by the algorithm. If C_M is the region enclosed by the manually-drawn contour and C_A is the region enclosed by the algorithm's contour, then the false negative error is given by $(C_M - (C_M \cap C_A))/C_M$, where $\cap$ denotes intersection. The false positive area measure, which identifies the percentage of tissue falsely identified by the algorithm, is given by $(C_A - (C_M \cap C_A))/C_M$. Inter-observer variability was calculated using the same method. Note that these measures of segmentation accuracy provide identical results to the method presented in Section 6.2.5.

The conventional cine DENSE data were acquired to investigate the effects of through-plane motion on the motion-guided segmentation. Figure 7.5 shows a four-chamber view of the heart, at early- and end-systole, with tags portraying longitudinal motion. A fixed mid-ventricular short-axis plane is shown as a dotted line, as is the motion of discrete points in the RV and LV free walls relative to this plane. The larger-diameter base can be seen to move into the image plane during systole. In conventional cine DENSE, base-to-apex motion in a short-axis view will thus cause the motion-guided contours to appear shrunken compared to imaged myocardium during systole and inflated during diastole. To minimize this error, the manual contours were drawn at mid-systole in conventional cine DENSE data. Any frame can be used in SF data.

To quantify the effect of through-plane motion and to give a further measure of segmentation accuracy, manual distance measurements were made between the motion-guided contour and the manually-defined epicardium on the LV free wall. As shown in Figure 7.6, these measurements were made radially in four places, along lines stemming from the centroid of the epicardial contour, and separated by an angle of $\pi/6$ radians. This portion of the epicardium was chosen because the border is clearly visible through-

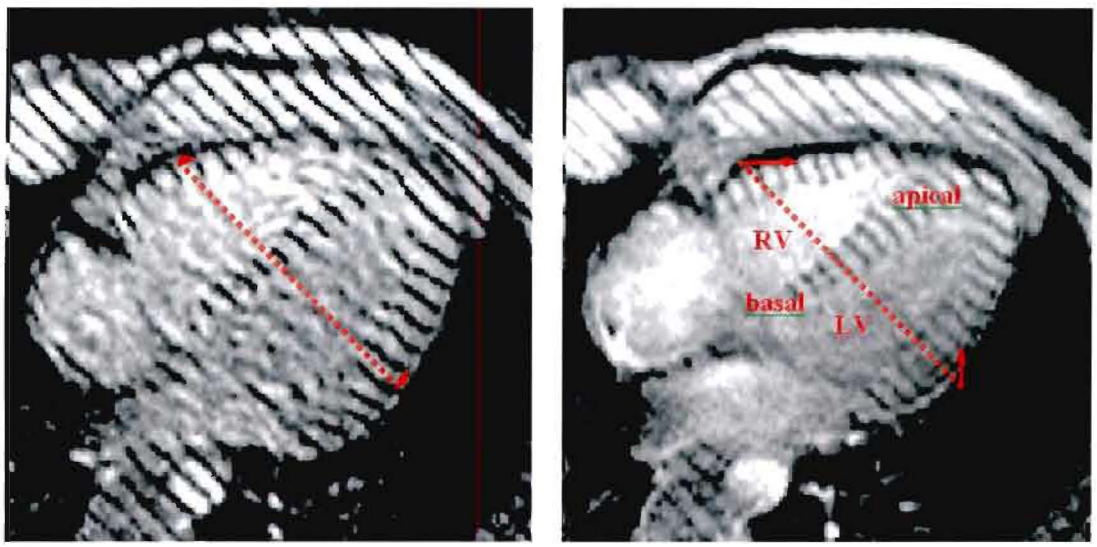

(a) (b)

Figure 7.5: Longitudinal motion depicted on tagged images at (a) early-systole, and (b) end-systole. A fixed mid-ventricular short-axis slice is shown as a dotted line in both images, and the position of discrete points in the LV free wall and RV free wall are shown relative to this imaging plane.

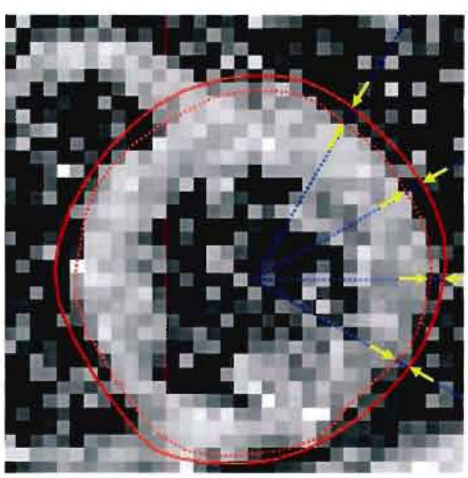

Figure 7.6: Short-axis view of the heart showing the four radial spokes along which differences between manual and automated contours were measured. Dotted contour = manual; solid contour = automated. The arrows depict the radial errors that were averaged for each frame. The error in the automated contour in this example has been exaggerated for demonstration purposes.

out the cardiac cycle, and because the LV through-plane motion is greatest in the free wall. The endocardial border was not investigated here because of the presence of papillary muscles and the high signal intensity in the blood during the first few frames. A mean error in pixels was then calculated for each cardiac phase.

7.3 Results

Typical mid-systolic SF cine DENSE images of magnitude, phase, and an inverted phase quality map, are shown in Figure 7.7(a), 7.7(b) and 7.7(c), respectively. The manually-drawn set of contours are shown on each image. Figure 7.7(d), 7.7(e) and 7.7(f) depict the corresponding motion-guided contours on inverted phase quality maps for a few cardiac phases. To demonstrate the flexibility of this method, a papillary muscle has also been contoured using the same algorithm (dotted lines). Note that no evidence of through-plane motion is present in these images. Figure 7.8 shows the effect of through-plane motion when the segmentation algorithm is applied to conventional cine DENSE data. Figure 7.8(a) and 7.8(b) depict the apparent swelling and shrinking of the contours that accompanies diastole and systole, respectively. The effect was amplified here by drawing the manual contour at end-systole for Figure 7.8(a), and at late-diastole for Figure 7.8(b). As expected, the error is greater for the LV free wall than for the septum.

Based on the mean of false positive and false negative area measures with manually-defined contours for four of the datasets, the segmentation error between the two observers is $18.12 \pm 5.98\%$, between observer A and the algorithm is $12.70 \pm 4.77\%$, and between observer B and the algorithm is $11.80 \pm 4.12\%$. This demonstrates that the accuracy of the algorithm relative to manually-drawn contours can be superior to inter-observer variability. The mean and standard deviation of the radial segmentation error for the 6 SF and 6 non-SF data sets is 0.13 ± 0.12 and 0.21 ± 0.15 pixels, respectively. Note that these error values are specific to short-axis mid-ventricular cine DENSE data. Incorrect phase unwrapping caused the algorithm to fail in 3 out of the total 414 frames (0.72%). Based on visual interpretation, a radial segmentation error greater than half a pixel was deemed unacceptable. Using this threshold, the segmentation reliability for SF and conventional cine DENSE of a short-axis mid-ventricular view is 98.2% and 94.8%, respectively. The mean segmentation processing time per frame was 2.39 seconds on a 2.8 GHz Intel Celeron with 256 MB of RAM.

To further demonstrate the flexibility of the method, the algorithm was applied to a set of long-axis images. Figure 7.9 shows the results of this process, where Figure 7.9(c) corresponds to the manually drawn contour, and the remaining images portray the mapping of this contour onto preceding and successive frames. Conventional cine DENSE was used here, and the higher SNR allows these results to be best portrayed on magnitude-

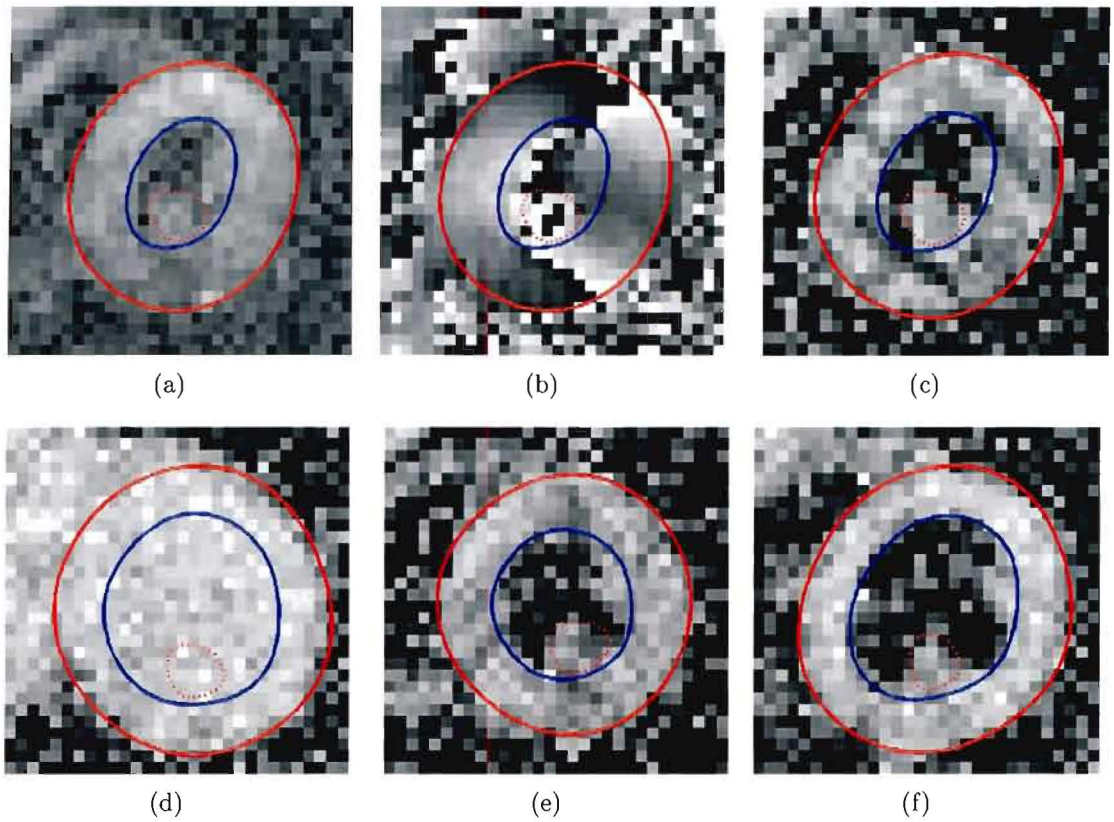

Figure 7.7: (a) Magnitude, (b) phase, and (c) inverted phase quality map of a short-axis view at end-systole depicting manually drawn contours. Corresponding motion-guided contours depicted on inverted phase quality maps for (d) end-diastole, (e) early-systole, and (f) mid-diastole. The dotted line depicts a papillary muscle that has been tracked and contoured using the same algorithm.

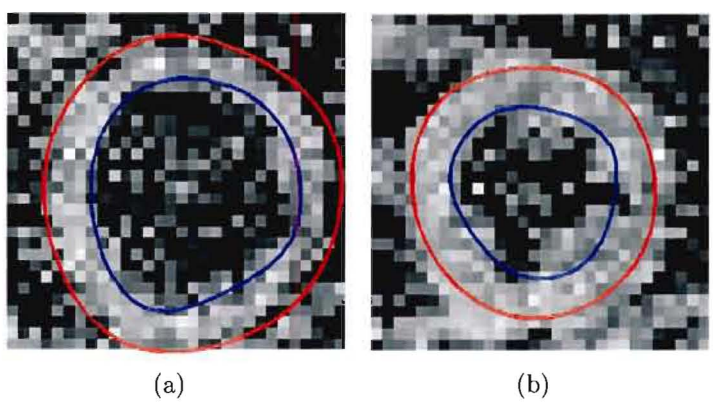

Figure 7.8: Effect of through-plane motion on motion-guided segmentation in conventional cine DENSE data of a mid-ventricular short-axis view of a normal volunteer. Contours are depicted on inverted phase quality maps. (a) Motion-guided contour at late-diastole, where the manually defined contour was drawn at end-systole. (b) Motion-guided contour at end-systole, where the manually defined contour was drawn at late-diastole.

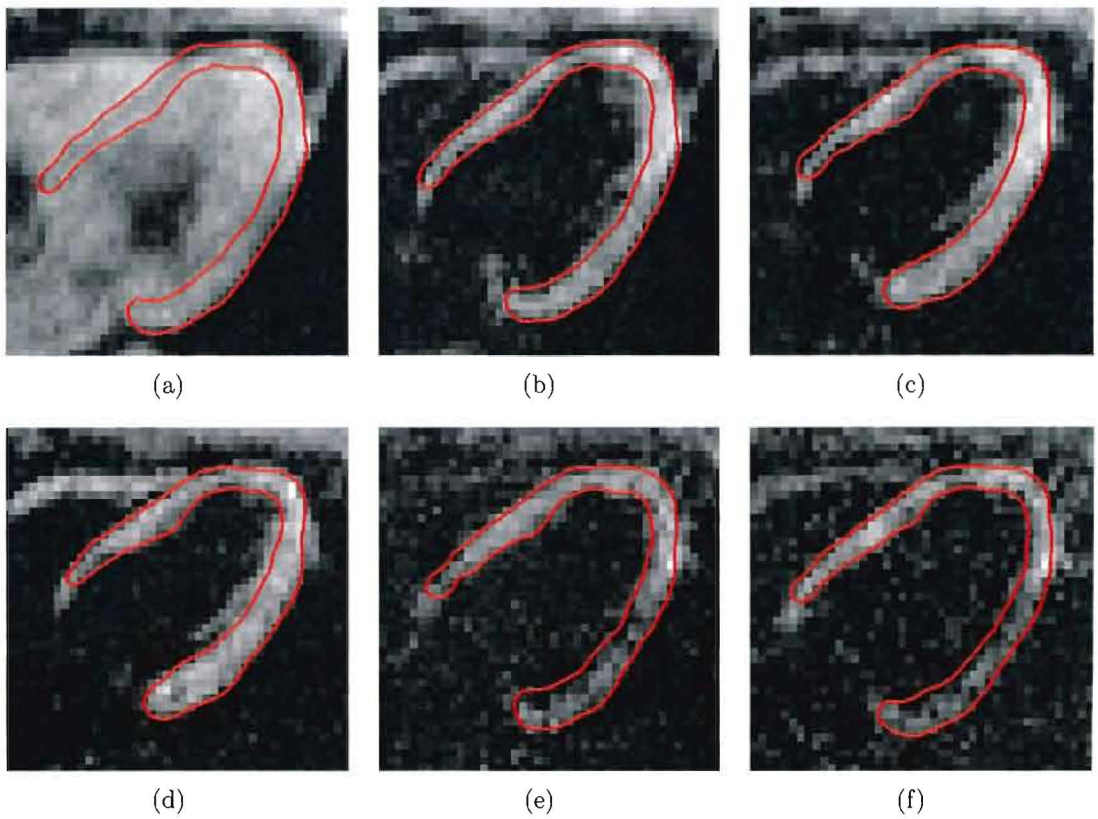

Figure 7.9: Motion-guided segmentation applied to long-axis conventional cine DENSE data. (a) End-diastole, (b) early-systole, (c) mid-systole, (d) end-systole, (e) mid-diastole, and (f) late-diastole. The manual contour was drawn at mid-systole, and all contours are portrayed on magnitude-reconstructed cine DENSE images.

reconstructed images. Due to the shape of the LV in the long-axis view, no periodic Fourier functions were used to smooth the contours.

7.4 Discussion

There are a number of unique advantages to the motion-guided segmentation method described here: 1. the parameters used to discern between useful and noisy vectors are all based on practical physiological limits; 2. the method ensures that only "high quality" displacement data is used for analysis; 3. contours are calculated for the first few cardiac phases, where it is difficult to visually distinguish blood from myocardium; 4. whereas many segmentation methods are specific to a particular view of the heart, this motion-based segmentation method is independent of the imaging plane; and 5. the

method is independent of the shape of the tissue delineated, and any arbitrary portions of the myocardium can be tracked through time.

The active contour models and periodic Fourier fits are all geared towards short-axis contours, where the LV is shaped like a torus. If the segmentation method is used for long-axis images, or arbitrary regions of myocardium, then the periodic Fourier fits or this active contour model implementation would not generally be suitable. Fitting using periodic splines is a suitable alternative.

Given that the segmentation accuracy is comparable to inter-observer variability, it can be argued that the algorithm will sometimes produce more accurate contours than a trained user. In particular, whereas it may be difficult for a user to consistently identify the same portion of myocardium for each cardiac phase, the segmentation algorithm inherently only includes myocardial vectors that contribute to the tissue tracking.

A limiting factor for this motion-guided segmentation technique is the robustness of the phase unwrapping algorithm. The segmentation method is equally appropriate for the RV, but the phase unwrapping algorithm used here can not reliably unwrap the phase in the RV at the current pixel size. Suitable alternative phase unwrapping algorithms were discussed in Chapter 3.

The use of motion trajectories to guide a set of contours through time could also be applied to myocardial tagging, although the spatial resolution may prove limiting. The HARP tissue classification technique described in Chapter 6 is advantageous because it requires no user interaction, but the results are not suitable for clinical use. More acceptable contours could be obtained by using an initial manually-defined region of myocardium as done in this Chapter, and refining the HARP tracking algorithm to identify and ignore noisy displacement vectors.

Chapter 8

Investigating LV-RV interventricular dyssynchrony using cine DENSE

8.1 Introduction

A bundle branch block (BBB) of the heart's electrical pathways is a delay or blockage of the conduction of electrical impulse to the ventricles. A BBB may be related to degeneration with age, or it may be associated with a specific pathology such as myocardial infarction. BBB's impair the synchrony of contraction of the ventricles, and are evident as a widening of the QRS complex. The resulting dyssynchrony reduces cardiac systolic function, increases oxygen consumption, and may be a source of arrhythmia (Kass, 2002).

Cardiac resynchronisation therapy (CRT) involves implanting pacemakers with pacing leads directed to sites affected by BBB's. CRT has been shown to improve symptoms, quality of life, LV function, and mortality in patients with end stage heart failure (Bristow *et al.*, 2004). Current criteria for CRT include New York Heart Association (NYHA) class III or IV symptoms (patients only comfortable at rest), LV ejection fraction (LVEF) less than 35% (normally greater than 55%), and a QRS duration greater than 130 msecs (normally 60 - 100 ms). However, when these criteria are applied, as many as 30% of patients are non-responders (Bristow *et al.*, 2004).

For healthy myocardium it is characteristic that mechanical activation follows electrical activation. Strain-time relationships give a measure of mechanical activation, and modalities such as tissue Doppler echocardiography and quantitative cardiac MRI can be applied to assess mechanical dyssynchrony. Recent tissue Doppler echocardiographic studies have shown that measurement of mechanical dyssynchrony improves the ability to prospectively identify CRT responders compared to QRS duration and morphology (Penicka *et al.*, 2004). Moreover, while intraventricular LV dyssynchrony is superior to QRS duration, the additional measurement of interventricular dyssynchrony between the LV and RV further improves the ability to prospectively identify CRT responders (Penicka *et al.*, 2004).

Mechanical LV dyssynchrony has been identified with myocardial tagging (McVeigh *et al.*, 1998; Prinzen *et al.*, 1999; Curry *et al.*, 2002; Leclercq *et al.*, 2002; Kerckhoffs *et al.*, 2005; Helm *et al.*, 2005; Zwanenburg *et al.*, 2005), but the resolution of the analysis was limited to the tag spacing and no results for the RV were reported. The use of HARP to measure mechanical dyssynchrony has also been reported by Spragg *et al.* (2005), and Lardo *et al.* (2005) identify strain-encoded (SENC) MRI (Osman *et al.*, 2001) as a further promising technique. A database of normal peak radial and longitudinal LV velocities, as well the time to these peak velocities has recently been published for phase contrast velocity encoding MRI of the human heart (Petersen *et al.*, 2006), but limited work has been reported that quantifies mechanical dyssynchrony using this technique (Delfino *et al.*, 2006). Normal human benchmark data also exists for myocardial tagging, where onset to time of circumferential shortening, and time-to-peak circumferential shortening is presented (Zwanenburg *et al.*, 2004).

Myocardial tagging, cine HARP and cine DENSE may be superior to tissue Doppler because they can measure circumferential shortening, which has been shown to be a better predictor of mechanical dyssynchrony than longitudinal shortening (Helm *et al.*, 2005). The high spatial resolution offered by cine DENSE makes it an attractive modality for measuring interventricular dyssynchrony. Patel *et al.* (2006) described the application of cine DENSE for measuring LV mechanical dyssynchrony. Patel *et al.* also identified regions of myocardial scarring using Gadolinium (Gd) delayed-enhancement, and the combined information offered by the two techniques was used to identify regions of healthy myocardium that could benefit from CRT. Other preliminary work on measuring mechanical dyssynchrony with cine DENSE is presented in Wen *et al.* (2006a), where normal volunteers are imaged using a respiratory-gated cine DENSE sequence.

In this Chapter, the cine DENSE methods used to image the LV are extended to measure strain in the RV. Both volunteer and patient data were acquired, and interventricular dyssynchrony in heart failure patients with left and right bundle branch block (LBBB and RBBB, respectively) was quantified. The majority of this work is geared towards the analysis of RV function using cine DENSE, and this forms the first reported study where RBBB delays are detected using quantitative MRI. Apart from predicting response to CRT, the ability to assess RV function is valuable in many cardiac and cardiopulmonary diseases such as pulmonary hypertension, congenital heart disease, and ischaemic heart disease.

The RV is notoriously difficult to image because of its thin walls, complicated geometry, and unruly behaviour. The normal thickness of the LV and interventricular septum at end-diastole is 9 – 11 mm, whereas the RV only measures 2 – 3 mm (Pettigrew, 1992). Also, the RV is heavily trabeculated and may not offer as clearly defined margins as typically seen in the LV. The analysis can be further complicated by the presence of the moderator band, which appears in certain imaging planes. The assumption of 2–D motion is reasonable for short-axis views of the LV (Teceläo *et al.*, 2006), but the through-plane motion of the RV should not be ignored in these views. The large through-plane RV motion is clearly demonstrated in Figure 5.5, and is quantified in Haber *et al.* (2005).

Measuring RV motion using myocardial tagging is limited by the number of tag lines present on the wall. A number of studies have been done to describe both normal and diseased RV motion using multiple slices of tagged images (Young *et al.*, 1996; Klein *et al.*, 1998; Fayad *et al.*, 1998; Haber *et al.*, 2000, 2005). Although these studies provided detailed information about RV behaviour, the acquisition and processing times were too lengthy for the methods to merit routine clinical use. In particular, contours and points to be tracked were marked manually by Young *et al.* (1996); Klein *et al.* (1998); Fayad *et al.* (1998), and semi-automatically using snakes by Haber *et al.* (2000, 2005).

Phase contrast velocity encoding applied in three orthogonal directions has also been used to study RV motion in a single mid-ventricular short-axis imaging plane (Kayser *et al.*, 2000). Kayser *et al.* studied normal volunteers and patients with acute myocardial wall infarction. Contours were drawn manually, and the remaining processing was fast and automated. Although through-plane velocity was measured, this technique does not take the through-plane RV motion into account.

There are three main challenges for RV analysis using cine DENSE: phase unwrapping; through-plane motion; and obtaining a suitable measure of strain. These are addressed in the following sections.

8.2 Methods

8.2.1 Phase unwrapping

The RV is typically 2 - 3 mm thick, and spans only 1 - 3 pixels in cine DENSE images. The spatiotemporal phase unwrapping algorithms presented in Chapter 3 require a path of high-quality phase throughout the myocardium. The phase quality measure used is based on phase uniformity in a local region, and thin portions of the RV will not necessarily have an associated high phase quality. In preliminary tests using 2–D quality-guided phase unwrapping without predefined contours and with seed points manually placed in the LV, the RV was correctly unwrapped in 252 out of 318 images (79.9%). This discourages the use of the motion-guided segmentation algorithm for contouring the RV. If predefined LV and RV contours are used, then 315 out of the 318 images were correctly unwrapped (99.1%). The use of manual RV contours thus leads to a significant improvement in the reliability of phase unwrapping in the RV.

The direct application of the semi-automated phase unwrapping method described in Section 3.3.2 with predefined LV and RV contours works adequately for the cases considered here. Temporal unwrapping was limited here to regions where myocardial masks from adjacent frames overlap. Phase unwrapping errors were identified using discontinuity maps, which specify the locations of pixels in the unwrapped image that differ from a neighbour by more than $|\pi|$ radians. Two examples of RV phase unwrapping errors identified in this manner are shown in Figure 8.1, where the red dots depict discontinuities.

8.2.2 Midwall tracking and tangential strain

The large oblique through-plane motion of the RV in short-axis views implies that a realistic study of RV behaviour requires a true 3–D description of its motion. Tissue

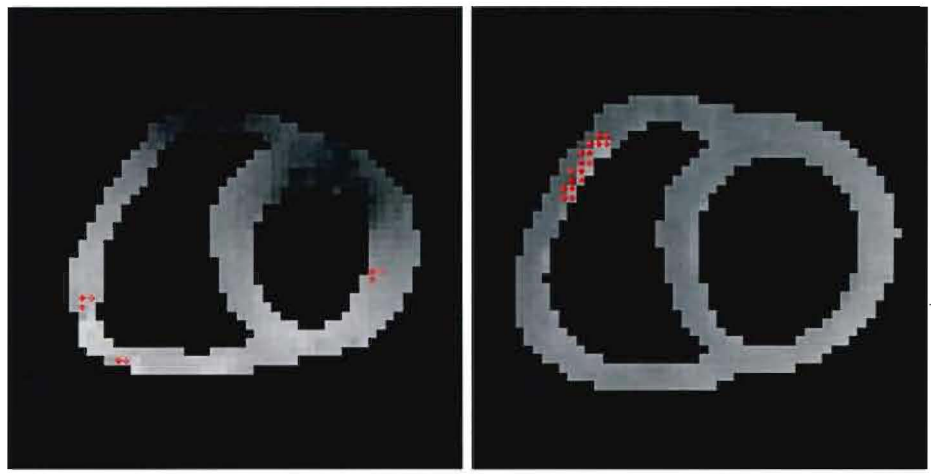

Figure 8.1: Two examples of phase discontinuities being used to identify phase unwrapping errors.

tracking using slice followed (SF) cine DENSE provides a fast and practical way of describing this motion.

Circumferential strain (Ecc) has previously been shown to be a sensitive metric of intraventricular LV dyssynchrony (Helm *et al.*, 2005). Ecc is a reasonable measure of strain for the annulus-shaped LV in the short-axis view. The shape of the RV is more complex and there are typically too few transmural pixels to calculate the 2–D strain tensor. It is therefore proposed that a measure of Lagrangian strain taken tangential to the midwall (Ett) be used. Figure 8.2(a) shows the DENSE displacement vectors for both ventricles of a normal volunteer at end-systole, and Figure 8.2(b) shows the corresponding 2–D linearly interpolated DENSE displacement vectors spaced a pixel distance apart along the LV and RV midwalls. Once again, the heads of the vectors refer to the position of the material points at the current time (t_f), and the tails refer to the position at the time of displacement encoding (t_0). Ett at point p_n can be obtained using the two adjacent vectors along the midwall. If l_{t_0}, the undeformed length, is the distance measured by cine DENSE between the tails of the vectors at p_{n-1} and p_{n+1}, and l_{t_f}, the deformed length, is the distance between the heads of these two vectors at time t_f, then Ett at p_n is given by

$$Ett = \frac{1}{2}\left(\left(\frac{l_{t_f}}{l_{t_0}}\right)^2 - 1\right) \tag{8.1}$$

The shape of the RV varies considerably through the cardiac cycle, and rigid sectors are not suitable for strain-time analysis. If points initially spaced a pixel distance apart

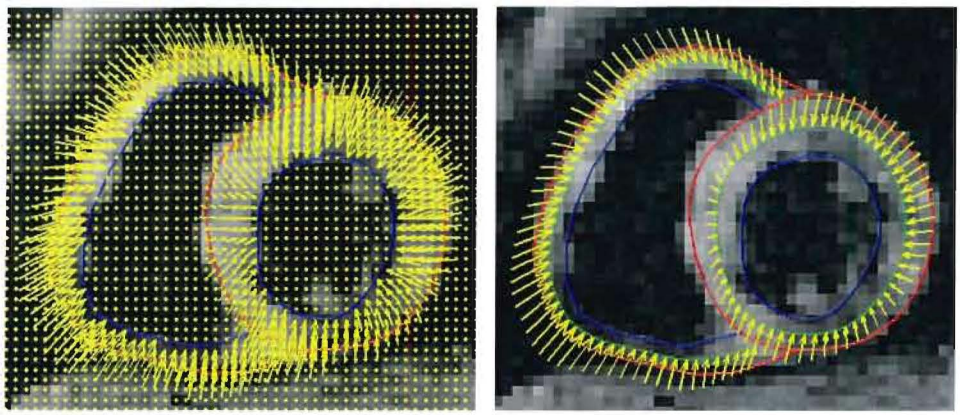

(a) (b)

Figure 8.2: (a) DENSE displacement vectors of both the LV and RV, and (b) interpolated displacement vectors at pixel-spaced points along the midwall.

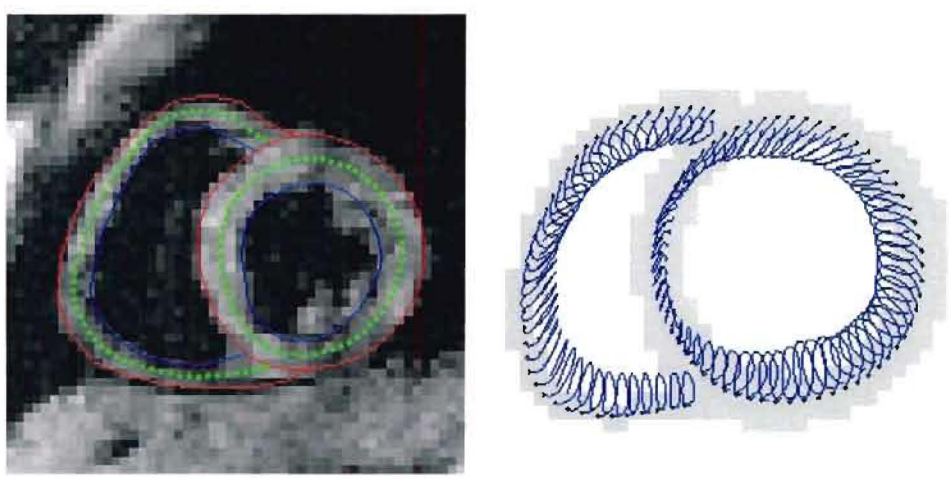

(a) (b)

Figure 8.3: (a) Pixel-spaced points defined along the arc length of the midwall contour at t_0, and (b) 2–D motion trajectories of these midwall starting points.

along the midwall at the first cardiac phase are tracked through the cardiac cycle as described in Chapter 4, then sectors defined along the midwall move with the heart as it deforms. For this reason and for the benefits of temporal fitting, trajectories were used for the computation of Ett. Trajectories were improved by temporal fitting with 5^{th} order Fourier basis functions as before. An example of the pixel-spaced starting points along the midwall is shown in Figure 8.3(a), and the resulting temporally fitted trajectories are shown in Figure 8.3(b).

The LV midwall was divided into six standard segments. The anterior two-thirds of the RV midwall was classified as the RV free wall, and the remaining third was classified as

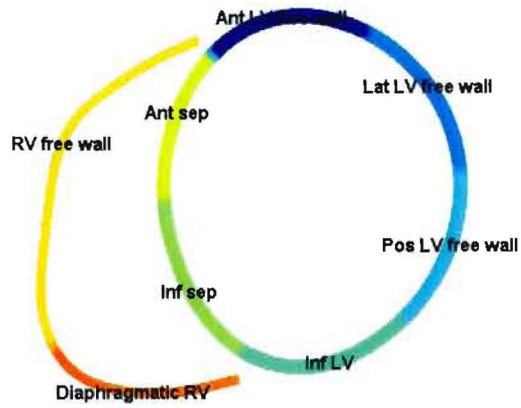

Figure 8.4: Myocardial segments for interventricular strain-time analysis.

the diaphragmatic RV. These divisions are shown in Figure 8.4. Finally, the Ett data were smoothed by sliding a 5 pixel mean filter along the midwall data. This is similar to the anatomically guided strain filtering described by Wen (2005).

8.2.3 Strain confidence

The measures of Ett for the LV are typically obtained from regions of higher SNR than the measures of Ett for the RV. The thickness of the RV is of the order of the cine DENSE in-plane pixel size, and less than the slice thickness. Partial volume effects can thus significantly reduce the SNR in voxels encompassing the RV. Incomplete intravoxel rephasing in simulated echo imaging also leads to signal loss when tissue deforms (Fischer *et al.*, 1995). This latter effect is particularly evident for high values of k_e, as seen in Figure 3.6.

A measure of strain confidence was derived to limit interpretation errors associated with low SNR. This also serves to alert one to the presence of phase unwrapping errors or misalignment of images from the two breath holds. Since it is the phase of the MRI image that is of interest, the phase quality maps were used for this measure. As described in Section 4.2.2, 2–D tissue tracking involves 2–D distance-weighted linear interpolation using the three closest vector tails to a starting point at t_0. For each encoding direction, a measure of confidence for each tracked point at each frame is derived by 2–D distance-weighted linear interpolation using the phase quality measures associated with each of the three pixels used for the tracking. The confidence measures for the two encoding

directions are then averaged and normalised to give a final measure of strain confidence. These can be displayed in conjunction with the stain-time curves.

8.2.4 Imaging protocol

The cine DENSE scans were performed on a 1.5 T Siemens Avanto scanner (Siemens Medical Systems, Erlangen, Germany). Two normal volunteers, two patients with LBBB, and two patients with RBBB were imaged. All human subjects provided informed consent and were studied in accordance with research protocols approved by the Human Investigations Committee at the University of Virginia. Patients were imaged without SF as part of a larger study, and the volunteers were imaged with SF.

A four-channel body phased array RF coil was used and the subjects were imaged in a short-axis mid-ventricular view. Imaging parameters included FOV = 360 × 252 mm; echo train length (ETL) = 9; TE = 11 ms; TR = 20 ms; slice thickness = 8 mm; pixel size = 2.81 × 2.81 mm; temporal resolution = 20 ms (rate 2 view sharing); and displacement encoding frequency = 0.1 cycles/mm. Scans for the two encoding directions were acquired in two breath holds of 14 heart beats each.

The LV contours were derived using the segmentation technique described in Chapter 7 and the RV contours were drawn manually. Image reconstruction was performed online, and image analysis was done offline using MATLAB (The Mathworks Inc, Natick, MA). The custom written MATLAB GUI for bi-ventricular cine DENSE image analysis is described in Appendix A.

8.3 Results

Figures 8.5 and 8.6 depict interpolated midwall cine DENSE displacement fields for a normal volunteer at various cardiac phases. The corresponding measures of Ett before and after applying the 5 pixel mean anatomical filter are also shown. Figure 8.7 shows the corresponding Ett maps derived using the temporally fitted trajectories to calculate strain, and after applying anatomical smoothing. An initial cardiac twist in both ventricles is evident in both the interpolated DENSE displacement fields (Figure 8.5(b)),

and the trajectories (Figure 8.7(b)). This effect has previously been noted in Klein *et al.* (1998).

Measures of normalised strain confidence for the designated cardiac segments are shown in Figure 8.8(a). The decay in confidence with time, associated with the $T1$-related decay in SNR, is evident in all of the segments, and a low measure of confidence is evident for both RV segments. The corresponding strain-time curves for this volunteer are presented in Figure 8.8(b), with the normalised strain confidence represented as a colour bar.

Midwall Ett strain-time curves for two segments in a normal volunteer are presented in Figure 8.9(a), depicting synchronous onset of contraction for both ventricles. Corresponding strain-time curves for a patient with a LBBB are shown in Figure 8.9(b), demonstrating a delayed onset of strain of 80 ms in the posterior LV free wall.

Figure 8.10(a) shows the Gd-enhanced image of a patient with a RBBB and an anteroseptal infarct. The corresponding end-systolic Ett map is shown in Figure 8.10(b), where the infarct is clearly evident as a region of near-zero strain. The Ett versus time curves are shown in Figure 8.10(c). A 40 - 60 ms delay in onset of strain is evident in the RV free wall compared to the LV, and no significant strain is present in the infarcted anteroseptal region. This data suggests that this patient would benefit from pacing the RV, but that pacing the anteroseptal region would yield no improvement.

The results for the remaining subjects are similar, and the Ett versus time plots for all segments are presented in Appendix H. A Gd delayed-enhanced image of a mid-ventricular short-axis view from another study is shown in Figure 8.11(a). The infarct is clearly evident in the anterior septum and part of the anterior RV. The infarct also appears as a slightly positive strain in the corresponding map of end-systolic Ett in Figure 8.11(b). This positive strain is due to tethering effects with the adjacent healthy myocardium, and these results further support the use of Ett for RV analysis.

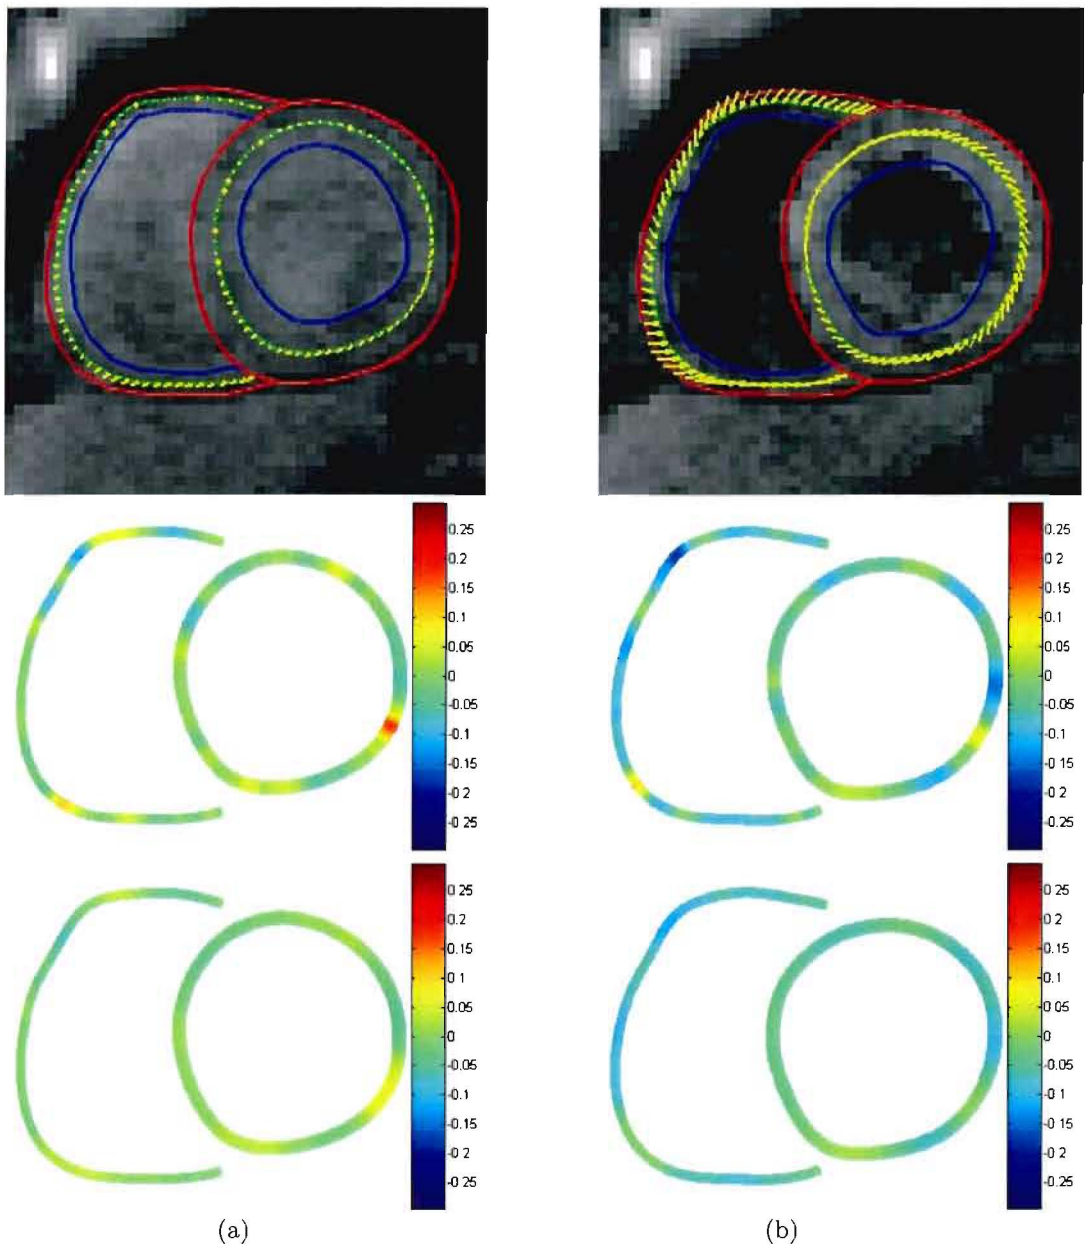

Figure 8.5: Interpolated midwall displacements and corresponding raw and anatomically smoothed maps of Ett of a normal volunteer at (a) end-diastole, and (b) early-systole.

8.4 Discussion

A cine DENSE technique suitable for imaging both the LV and RV has been developed. Initial results in volunteers and heart failure patients suggest that this technique can be used to assess RV strain and as such interventricular dyssynchrony.

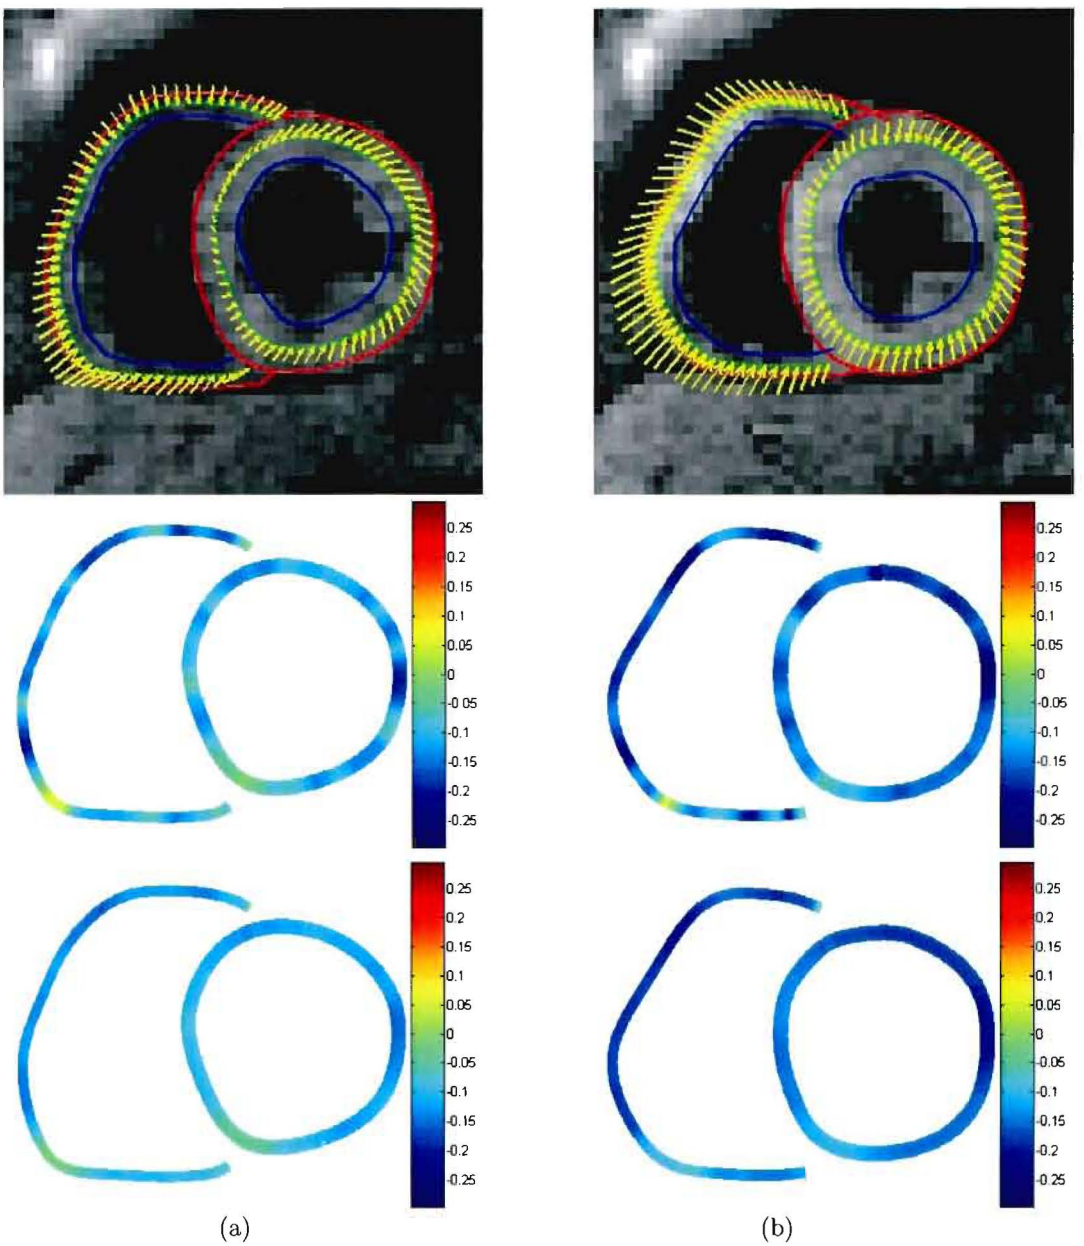

Figure 8.6: Interpolated midwall displacements and corresponding raw and anatomically smoothed maps of Ett of a normal volunteer at (a) mid-systole, and (b) end-systole

Higher spatial resolution would simplify the phase unwrapping and tissue tracking algorithms. For cases where a more robust phase unwrapping algorithm is required, the fully automated quality-guided algorithm described in Section 3.3.1 could be used, where unwrapping is again confined to lie within the predefined contours. The unwrapping can commence at all myocardial pixels on the first cardiac phase. This method successfully unwrapped 316 out of the 318 test images used in this study (99.4%).

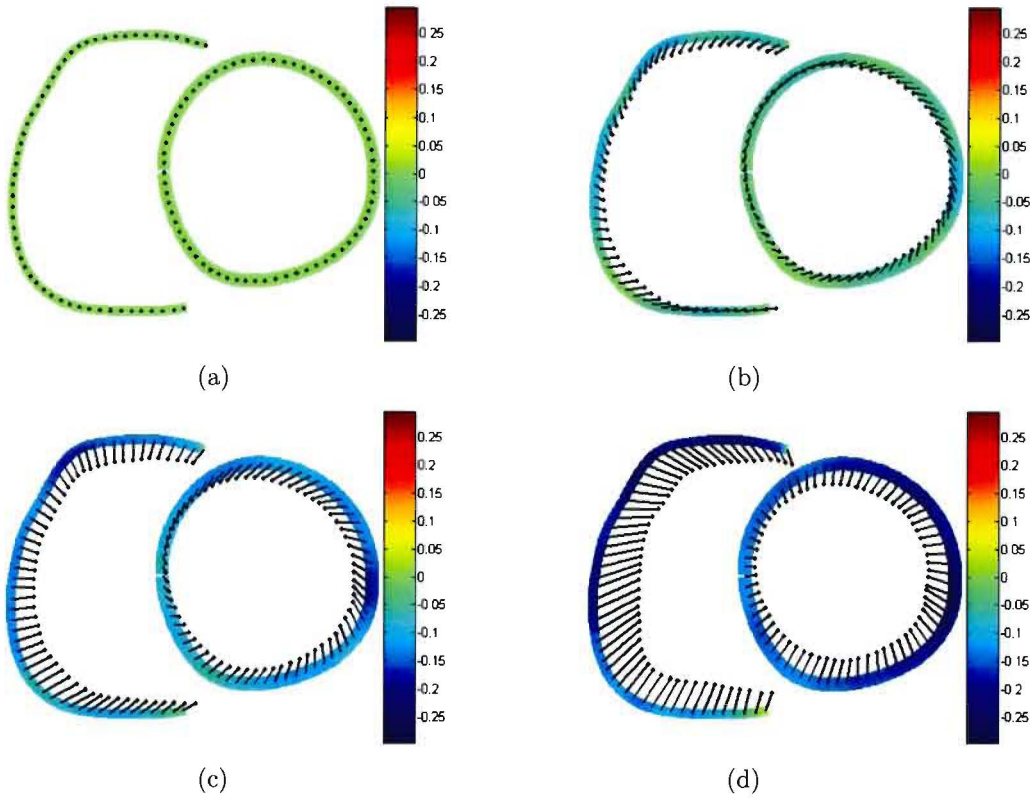

Figure 8.7: Anatomically smoothed Ett and displacement fields obtained from temporally fitted trajectories of a normal volunteer at (a) end-diastole, (b) early-systole, (c) mid-systole, and (d) end-systole.

Although view sharing and the equations of motion from temporal fitting can be used to improve the temporal resolution, the true temporal resolution of the cine DENSE sequence used here remains 40 ms. Improving this temporal resolution is the focus of future work. Respiratory gating can be used to increase the spatial and/or temporal resolution, but with a significant increase in acquisition time. The respiratory-gated EPI cine DENSE sequence used by Wen *et al.* (2006a) provides a spatial resolution of $1.4 \times 2.4 \times 7\ mm$, a true temporal resolution of 20 ms, and requires a scan time of 3.5 - 4.8 min/slice.

Although the measures of strain confidence provide a useful indicator of the reliability of the data, a way of automatically identifying unacceptable SNR and/or phase unwrapping errors in the RV is still desirable. The low SNR associated with SF together with the large amount of RV through-plane dephasing still poses challenges for routine RV analysis using cine DENSE. In some cases, it may be necessary to reduce the value of displacement encoding frequency k_e to limit the intravoxel dephasing seen in Figure 3.6.

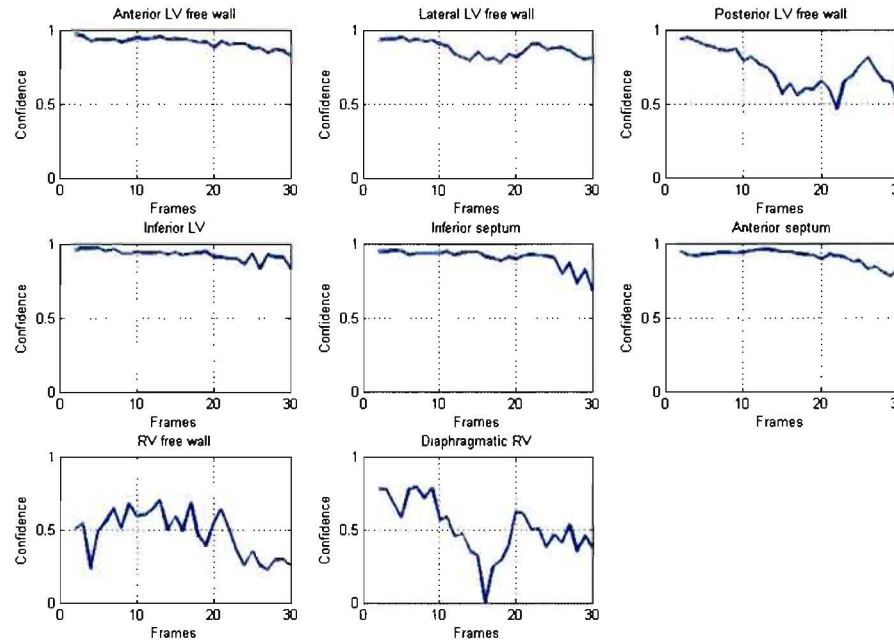

(a)

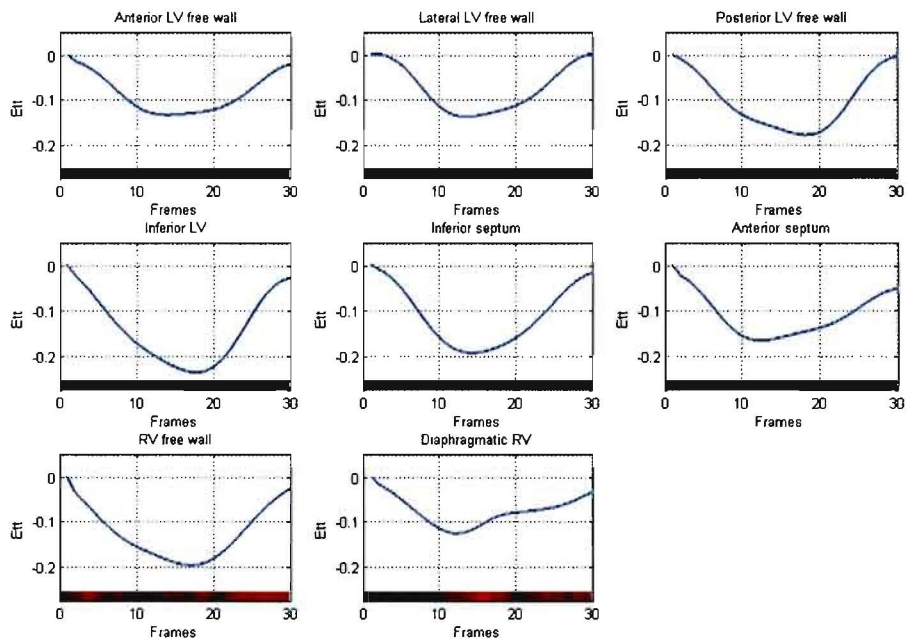

(b)

Figure 8.8: Normalised strain confidence plotted (a) as graphs, and (b) on colourbars with the Ett versus time curves. This is data from a normal volunteer.

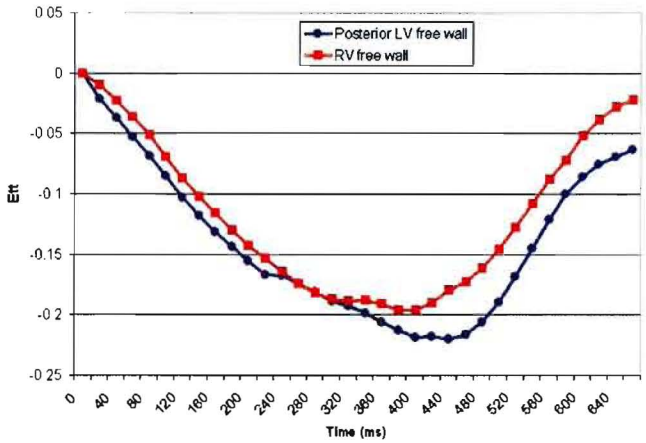

(a)

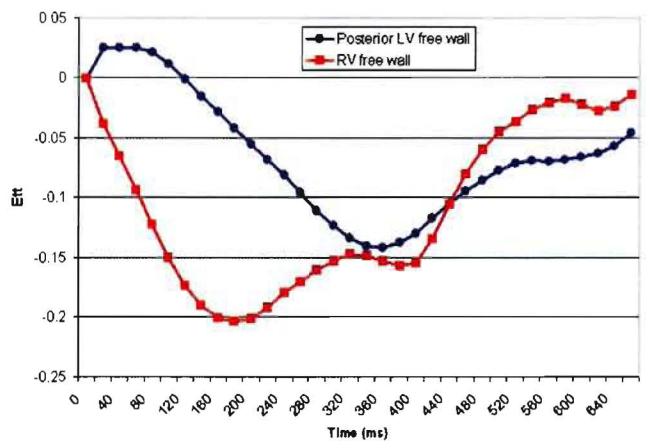

(b)

Figure 8.9: (a) Mid-ventricular short-axis Ett versus time for a normal volunteer, and (b) a patient with a LBBB.

Only the time-to-onset of strain was investigated here as a metric for identifying mechanical dyssynchrony. CRT only influences the onset times directly, but time-to-peak strain is often used to select patients for CRT. The relationship between these two metrics has been investigated by Zwanenburg *et al.* (2005). Zwanenburg *et al.* found the relation between peak time and onset time of shortening to be strongest in nonischaemic patients, and that this relationship was most consistent when the time to first peak, instead of time to maximum peak, is used. A variety of other methods of detecting dyssynchrony using strain measurements are summarised in Lardo *et al.* (2005), where both spatial and temporal aspects of strain distribution were investigated. All of these methods could be applied to cine DENSE analysis of LV dyssynchrony, but the low

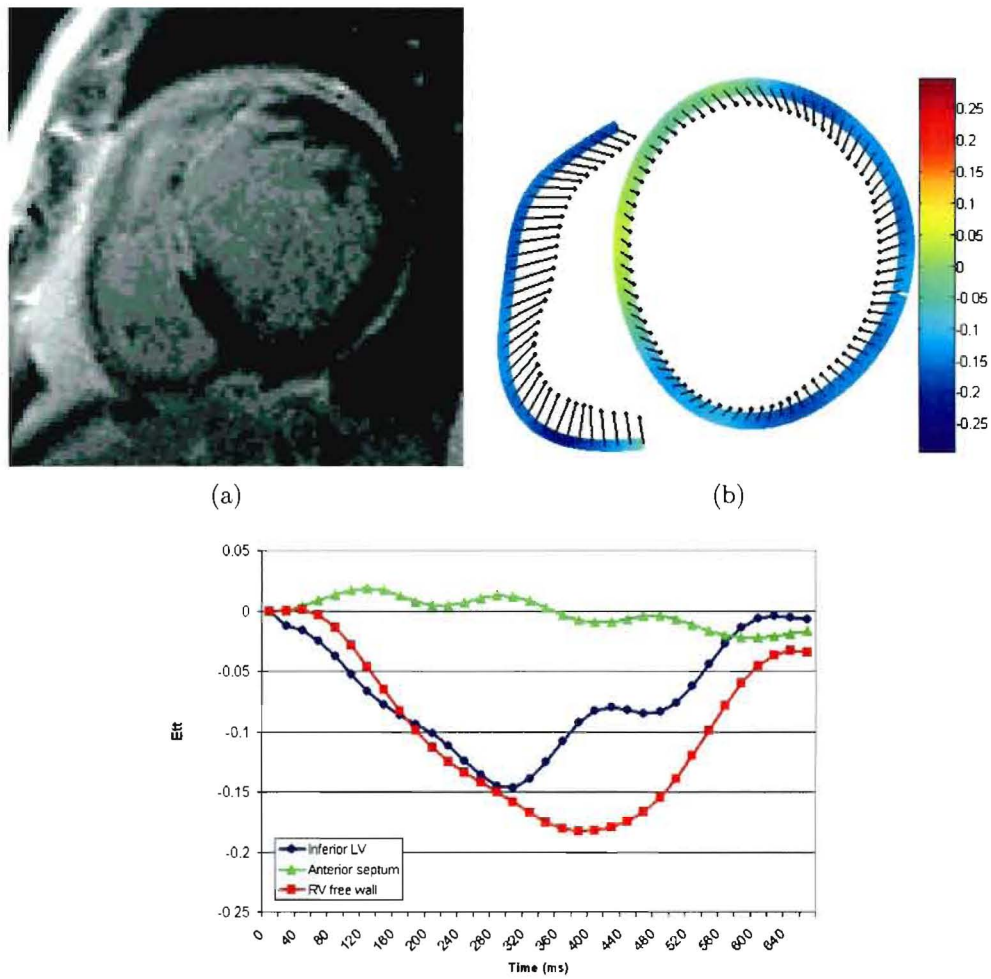

Figure 8.10: (a) Mid-ventricular short-axis Gd-enhanced image of a patient with a RBBB and an anteroseptal infarct, (b) corresponding map of end-systolic Ett, and (c) corresponding Ett versus time curves.

spatial resolution and the lack of spatial periodicity of the RV limits their use for RV analysis.

In conclusion, when used in conjunction with delayed contrast-enhanced MRI and cine DENSE measurements of intraventricular LV dyssynchrony, the additional measurement of interventricular RV-LV dyssynchrony may improve the ability to prospectively identify CRT responders.

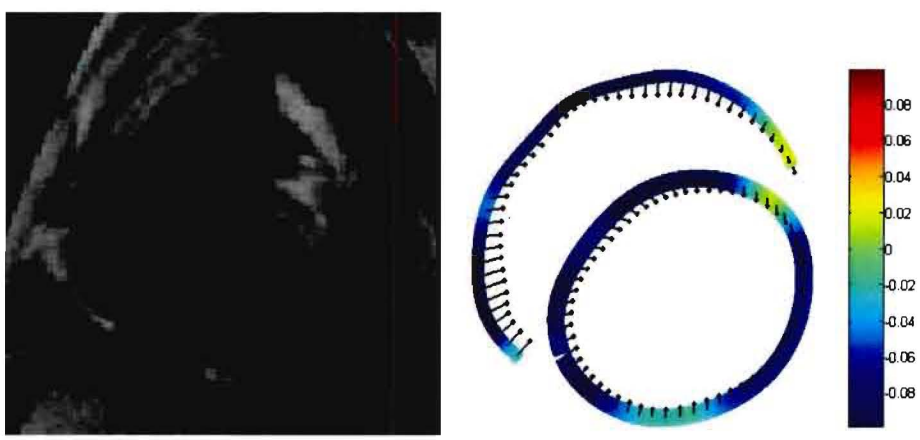

(a) (b)

Figure 8.11: (a) Mid-ventricular short-axis Gd-enhanced image depicting an infarct due to a left anterior descending (LAD) coronary artery occlusion. (b) End-systolic map of Ett and displacement, where the infarct appears as a region of positive or near-zero strain.

Chapter 9

Conclusions and future directions

Reliable automated and semi-automated spatiotemporal phase unwrapping methods have been developed to obtain absolute phase measurements for cine DENSE. The semi-automated method is faster than the fully-automated method, and requires the user to only specify a single seed point in the myocardium. Based on these absolute displacement measurements, a 2–D technique for tracking the motion of discrete material points within the myocardium was presented.

The 2–D tissue tracking was improved by fitting curves to the temporal evolution of the motion trajectories. This tracking was done at roughly a pixel resolution and strain was calculated directly from the trajectories. Using sets of 5^{th} order Fourier basis functions to describe the trajectories reduces the variability in the strain measurements, while still retaining sufficient information to distinguish pathology. Equations of motion were obtained for each pixel tracked, thus offering a wealth of information for studies of myocardial motion.

The improvement in Ecc due to temporal fitting is evident visually on strain images, in strain-time curves, and in a comparison of histograms. Also, the values of Ecc closely match those predicted using conventional myocardial tagging methods. There is more variability in the literature regarding normal values of Err, and the results presented here for Err are within the range of previously reported values. The accuracy of the tissue tracking was quantified using a rotating phantom, but the effects on strain would benefit from further validation using a deformable phantom. For physiological cardiac velocities and displacements, the 2–D in-plane tracking error is 0.24 ± 0.15 mm.

The optimal displacement encoding frequency for studies of the human heart using the cine DENSE sequence described in Kim *et al.* (2004) was shown to be in the region of 0.1 cycles/mm. For 2 scans of 17 second duration, the *in vivo* strain noise after temporal fitting was estimated to be $2.5 \pm 3.0\%$ at end-diastole, $3.1 \pm 3.1\%$ at end-systole, and $5.3 \pm 5.0\%$ in mid-diastole. Given that the first principal strain is generally oriented in the radial direction in short-axis images, measures of Err taken together with these strain noise measurements give an indication of the strain-to-noise ratio. For typical values of cardiac Err, the strain noise measurements show that the methods presented here yield useful strain data.

Three-dimensional tissue tracking using conventional cine DENSE involves scanning multiple image slices and performing a large amount of post processing. Slice following was implemented in the cine DENSE sequence and a true 3–D myocardial motion tracking algorithm using data from a single imaging plane was introduced. Spatial median filtering and 3–D temporal fitting were shown to improve the trajectories without any noticeable adverse effects on the deformation. The through-plane tracking accuracy using SF cine DENSE on a rotating phantom with typical ranges of cardiac displacement and velocity was shown here to be 0.48 ± 0.25 mm. The effects of through-plane rotation on 2–D strain in a mid-ventricular short-axis view of the heart was also quantified. For six normal volunteers, through-plane rotation caused the end-systolic Ecc to be reduced in magnitude by $3.91 \pm 0.43\%$, and the corresponding Err to be increased in magnitude by $6.01 \pm 1.07\%$.

A versatile and effective segmentation method has been developed for cine DENSE, where displacement encoded motion in the myocardium is used to guide a set of contours through time. The total user interaction in cine DENSE image analysis is thus reduced to the manual demarcation of the myocardium on a single frame. Based on manually-drawn contours, and an area measure of segmentation error, the accuracy of the algorithm was shown to be similar to inter-observer variability. A radial segmentation error metric was introduced for short-axis data, and the resulting error for 6 SF and 6 non-SF data sets is 0.13 ± 0.12 and 0.21 ± 0.15 pixels, respectively. Although the segmentation error for conventional cine DENSE is greater than for SF cine DENSE, motion-guided segmentation using conventional cine DENSE yields acceptable results in a mid-ventricular short-axis view.

A processing technique suitable for analysing RV wall motion using cine DENSE has also been developed. Manual segmentation of the RV borders is required for robust phase unwrapping, and SF cine DENSE is recommended to compensate for the large oblique through-plane motion of the RV in short-axis views. Circumferential shortening for the RV was defined as the 1–D strain orientated tangential to the RV midwall. Two dimensional tissue tracking of RV midwall points was performed and this tangential strain (Ett) was calculated directly from the trajectories. Temporal fitting was applied to the trajectories as before, and the resulting strain was smoothed using an anatomically-guided mean filter. By also applying these techniques to the LV, the interaction between the two ventricles could be assessed. In particular, interventricular dyssynchrony was studied here using regional time-to-onset of Ett for both ventricles. Initial results in volunteers and heart failure patients suggest that this technique can be used to assess interventricular dyssynchrony. When used in conjunction with delayed contrast-enhanced MRI and cine DENSE measurements of intraventricular LV dyssynchrony, the additional measurement of interventricular LV-RV dyssynchrony may improve the ability to prospectively identify CRT responders.

A higher cine DENSE spatial resolution would simplify the phase unwrapping and tissue tracking algorithms, and also allow for RV motion-guided segmentation, and the measurement of RV radial thickening. Low regions of SNR in the RV due to partial volume effects and intravoxel dephasing are commonplace, and can lead to unusable data or phase unwrapping errors. Although a measure of data confidence has been introduced here, there is still room for further integrated measures of SNR and phase unwrapping efficacy. A further problem for RV analysis using the breath hold cine DENSE sequence described in Kim *et al.* (2004) is spatial mismatch due to variations in breath hold position. These effects could potentially be reduced using image morphing and matching with affine transformations.

A cine HARP tissue tracking algorithm and methods of motion-based tissue classification of tagged images were also introduced. The tissue tracking method is based on inter-frame displacement measurements and suffers from the same main drawback as phase contrast velocity encoding tissue tracking methods, namely the integration of tracking error. Unfortunately, the $T1$-related decay in SNR with time means that motion periodicity characteristics can not be used to improve these cine HARP trajectories.

The proposed cine HARP myocardial tissue classification method is based on differences in behaviour between motion trajectories in tissue and those in regions of noise. This provides a means of classifying myocardium and stationary tissue without necessitating any manual user intervention. The cine HARP tracking process and trajectory fields could also be adapted to guide a manually defined set of contours in a similar manner to the cine DENSE motion-guided segmentation method. This cine HARP tissue classification method is capable of delineating the endocardium even in the first frame of a tagged series, where tag lines are present in the blood. Based on comparisons with manually drawn regions, and taking computation time into account, optimal tissue classification is achieved using temporally accumulated spatial derivatives of the motion trajectories. The values of false-positive and false-negative area error measures are each in the region of 10%, implying that this tissue classification technique is inadequate for diagnostic purposes, where aspects such as wall thickening would be measured. Contracting torus phantom simulations were conducted to investigate the relationship between pixel size and classification error. The errors predicted using the phantom were lower than those measured in practice. This is partially due to through-plane motion and errors inherent in the cine HARP tracking process, but also to inaccuracies in the manual delineation of myocardial contours.

Note that the cine DENSE phase unwrapping, tissue tracking and motion-guided segmentation can be implemented for any view of the heart. Although view sharing can be used to increase the perceived temporal resolution, it has been shown here to offer little improvement to trajectory behaviour and strain-time curves. The 2–D and 3–D tissue tracking algorithms can be implemented at any spatial resolution, and the temporal fitting results in 2–D or 3–D equations of motion at any temporal resolution.

Despite the fact that stimulated echoes have an inherently low SNR, cine DENSE has proven to have both a high tracking accuracy and a high spatial resolution. Furthermore, the post processing is reliable, computationally efficient and relatively user-independent. Although the cine DENSE method still needs validation on an appreciable patient base, it holds great promise as a tool for automated, detailed analysis of myocardial motion.

Future work on the cine DENSE sequence will involve reducing the scan times, increasing the spatial and temporal resolution, and/or improving image SNR. Other recent applications for cine DENSE include studies of skeletal muscle (Zhong *et al.*, 2006a) and studies of arterial walls (Wen *et al.*, 2006b). Cine DENSE has also been combined

with through-plane strain encoding (Osman *et al.*, 2001) to measure the full 3-D strain tensor at a voxel resolution (Hess *et al.*, 2006). Some work has already been done to incorporate stages of the cine DENSE processing directly onto the scanner, but considerable work still remains to portray DENSE strain maps online.

In the event that volumetric cine DENSE imaging becomes available, a number of these processing techniques can be directly extended to embrace the additional spatial dimension. The 3–D phase unwrapping can be easily extended to 4–D (three spatial and one temporal dimension). True 3–D tissue tracking using volumetric data is a straightforward extension of the 2–D method introduced for conventional cine DENSE. Slice following isn't necessary for volumetric data. Using the encoded motion to project the position of a predefined volume of myocardium through time is also a direct extension of the ideas implemented in the 2–D case. In fact, the large amount of spatiotemporal information available for volumetric segmentation opens new doors for a completely automated approach, whereby surface models and/or motion-based tissue clustering could also be used. Implementing these algorithms in reasonable processing times will form a significant part of the challenge.

Appendix A

CineDENSEAnalysis: MATLAB Graphical User Interface

A MATLAB graphical user interface (GUI) titled *CineDENSEAnalysis* was created to process cine DENSE data. *CineDENSEAnalysis* was coded by the author, with contributions from Xiaodong Zhong, Aaron Hess, Wesley Gilson, and Daniel Kim.

CineDENSEAnalysis for the LV

Figure A.1 shows the GUI for LV analysis. The majority of the cine DENSE algorithms described in this thesis are implemented as functions in the GUI. Both magnitude- and phase-reconstructed images are portrayed in the GUI, and the user can freely navigate through the cardiac phases and between encoding directions. A status box is also present to give the user an indication of the progress of the algorithms. The *Save workspace* and *Load workspace* pull-down menus save and load all data and variables in the GUI workspace, respectively, thus removing the need to re-process Cine DENSE data. In addition, clicking on the *Restore Original Images* button at any stage during the analysis returns the workspace to the state just after loading the images. Image analysis using the GUI is divided into four sections, 1. *Load data*, 2. *Contours*, 3. *Analysis*, and 4. *Plot and Export.*

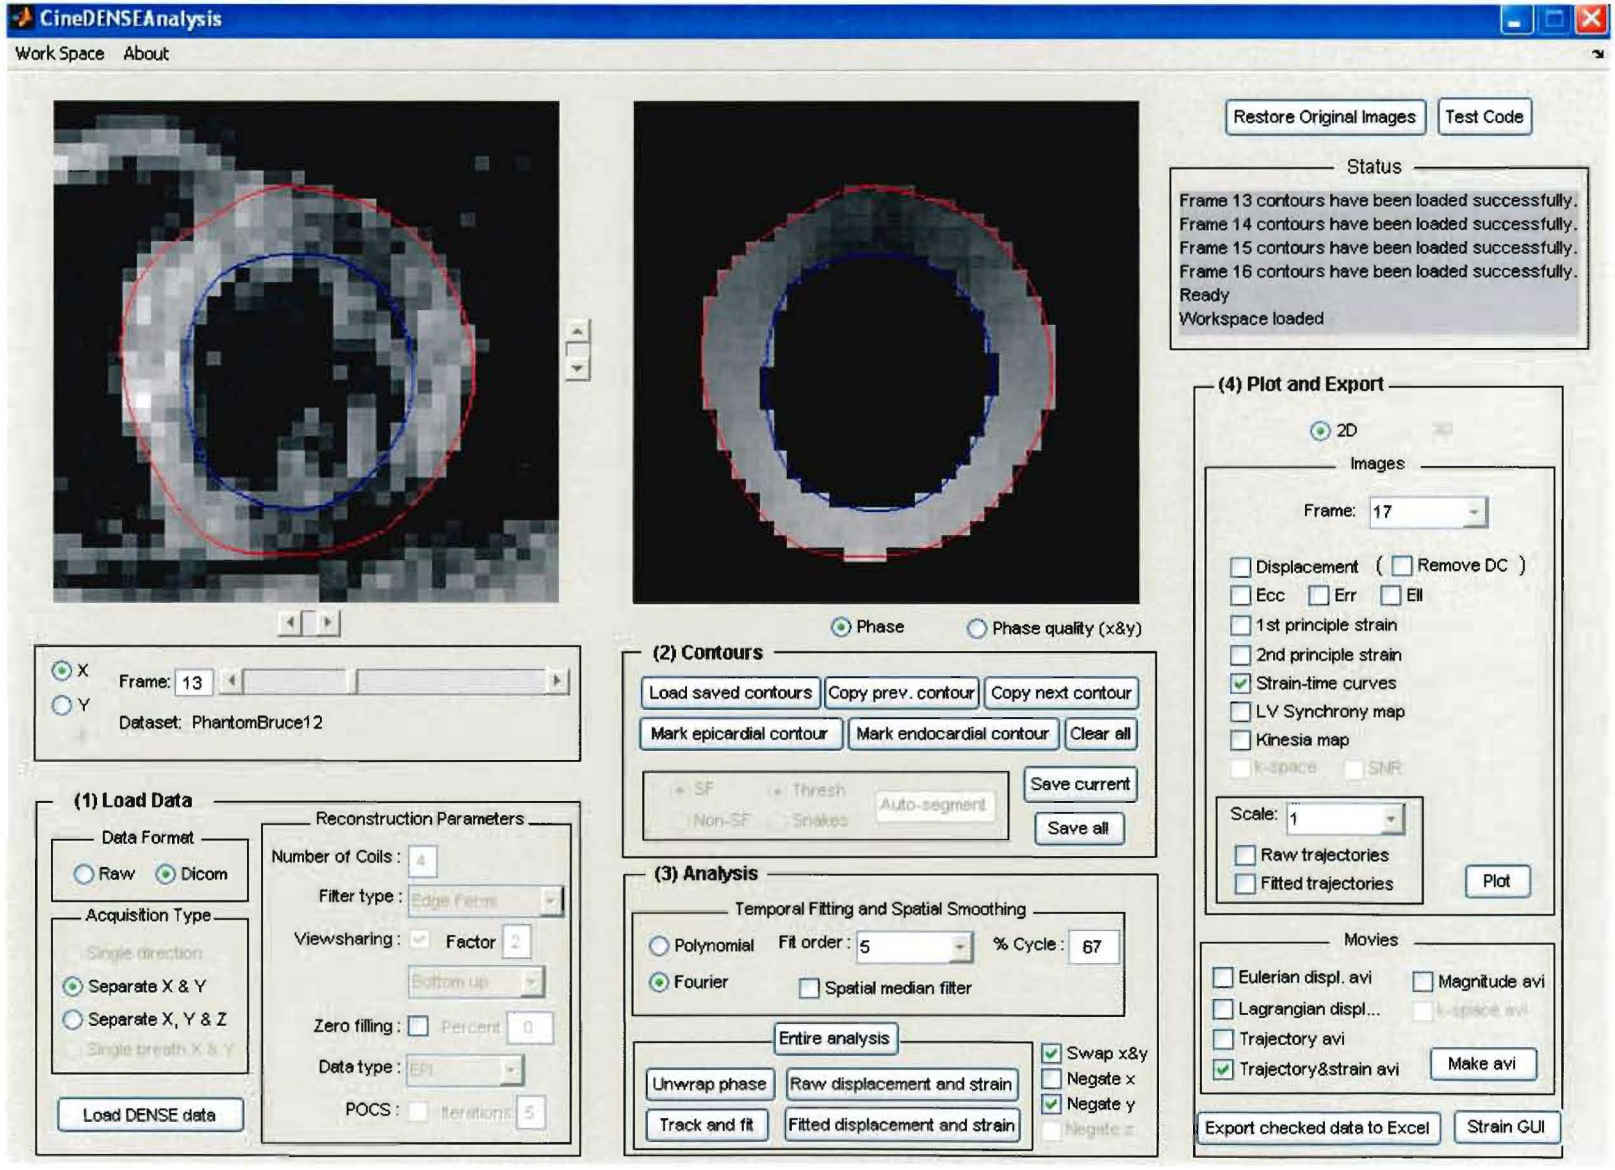

Figure A.1: Custom written *CineDENSEAnalysis* MATLAB graphical user interface for processing LV cine DENSE data.

Load Data

Radio buttons allow the user to select between loading raw k-space data files or Dicom files. A further set of radio buttons define whether data encoded in 1–D, 2–D, or 3–D should be loaded. Data can also be loaded for two cine DENSE encoding strategies: 1. a phase reference and a CSPAMM set of images for each encoding direction; or 2. a phase reference and two encoding directions in a single breath hold. For the latter method, through-plane dephasing (Zhong *et al.*, 2005) is typically applied at the expense of SNR to suppress the $T1$ and conjugate echoes.

If raw k-space data is loaded, a number of reconstruction parameters can be set: the number of RF coils; a selection of k-space filters for suppressing the $T1$-echo; zero filling of the k-space periphery to interpolate to a smaller pixel size; varying rates of view sharing for bottom-up or flyback EPI k-space trajectories (as described in Foo *et al.* (1995)); and projection onto convex sets (POCS) phase correction (Xu and Haacke, 2001) to accommodate for partial k-space sampling.

Load DENSE data button. When pressed, the user is asked to input a number of image reconstruction parameters specific to the type of data being loaded. For Dicom images, this includes series numbers, the displacement encoding frequency, and the number of frames to be loaded. For raw data, in addition to the number of frames and the displacement encoding frequency, a list of parameters for reconstruction are also included: readout field of view (FOV); percent phase resolution; percent of FOV in the phase encoding direction; EPI factor; and segments (scan line numbers per phase in one heart beat). The latter two parameters are only important if view sharing is being implemented.

Once the images are loaded and reconstructed, the user is asked to manually crop a region of interest (ROI). Normalized cross correlation (Gonzalez and Woods, 1992) is then used to align corresponding images of the two or three encoding directions. Slider bars adjacent to the axes have also been included to manually shift the images in the event of imperfect alignment.

Contours

Contours are typically drawn on magnitude-reconstructed DENSE images. An option also exists for viewing the combined phase quality maps instead of the phase maps. If the *Phase quality (x&y)* radio button is selected, then contours must be defined on these images and not the magnitude-reconstructed images. The phase quality maps from all encoding directions are combined in such a way that overlapping myocardium from the different breath holds can be visualised. Note that the phase quality maps are inverted to more closely resemble the magnitude-reconstructed images. If an existing contour is being re-drawn, the previous version is displayed as a dotted line for reference.

Mark epicardial contour button. This allows the user to manually define the epicardial contour on the current frame. The contour is smoothed using a periodic spline.

Mark endocardial contour button. This allows the user to manually define the endocardial contour on the current frame. The contour is smoothed using a periodic spline.

Save current button. Saves the contour points for the current frame in a text file in the patient folder. Both epicardial and endocardial contours are required. Note that if raw data was loaded, then different filename conventions exist for non-view shared and view shared frames.

Load saved contours button. Loads all contours associated with the current data set.

Save all button. Saves all marked contours for the current data set.

Clear all button. Clears all loaded or marked contours.

Copy previous contour button. Copies the epicardial and endocardial contours from the previous cardiac phase. These can be used as a reference for drawing the current set of contours.

Copy next contour button. Copies the epicardial and endocardial contours from the next cardiac phase. These can be used as a reference for drawing the current set of contours.

Auto-segment button. This initiates the motion-guided segmentation algorithm described in Chapter 7. If the segmentation algorithm is run with epicardial and en-

docardial contours defined on the displayed cardiac phase, then the algorithm defines motion-guided epicardial and endocardial contours for all other cardiac phases. Defining these contours at mid-systole minimises the segmentation error due to through-plane motion. If the myocardial wall is thin, then active contour models can be selected using the *Snakes* radio button. However, using isocontours on binary masks, as selected with the *Thresh* radio button, is typically sufficient and more accurate. If the segmentation algorithm is run without predefined contours, then the user is asked to manually define a region of the myocardium using a single contour on the current frame. This is useful for tracking and identifying select portions of the myocardium. Note that this button only becomes active after phase unwrapping is done without predefined contours and after displacement is calculated.

Analysis

This includes the main data processing features.

Unwrap phase button. This runs the semi-automated spatiotemporal phase unwrapping algorithm described in Chapter 3. If no predefined contours exist, then the phase quality map of the first cardiac phase is displayed and the user is requested to place a manual seed point in the myocardium. If contours exist, then only myocardial pixels are phase unwrapped and the seed point is defined as the pixel with highest phase quality lying within the contours at the first cardiac phase. The unwrapped phase images replace the wrapped phase images in the display after this algorithm has been run.

Raw Displacement and Strain button. This calculates the displacement fields for all cardiac phases. If the *Spatial Median Filter* box is checked, then a 5 pixel spatial median filter (centre pixel plus its four nearest neighbours) is applied to all components of all displacement fields. If contours exist, then the median filter for each pixel is applied using only the neighbouring pixels that lie within the contours. First ($E1$) and second ($E2$) principal strains are also calculated for all cardiac phases, as are circumferential (Ecc) and radial (Err) strains. Ecc and Err are derived relative to the centroid of the epicardial border at the first cardiac phase. These strain computations are described in Gilson (2004). If no contours exist then only the displacement fields are calculated.

Radio buttons are also available to flip the displacement vector directions into their correct orientation. This is done because it is not straightforward to determine the encoding direction relative to the image orientation defined in the Dicom headers. Note that changes in vector orientation only occur after the *Raw Displacement and Strain* button has been pressed.

Track and Fit button. If 2–D cine DENSE data is loaded, this applies the 2–D cine DENSE tracking algorithm described in Chapter 4. If 3–D slice followed cine DENSE data is loaded, the 3–D tracking algorithm described in Chapter 5 is applied. The trajectories are initiated at each myocardial pixel centre on the first cardiac phase. Temporal fitting is applied to each trajectory as described in Chapter 4. The radio buttons select whether polynomial functions or Fourier basis functions should be fitted, and the order of the fit can be selected from the pull-down menu. The percentage of the cardiac cycle captured can also be set for fitting with Fourier basis functions. This figure is used to gauge the number of interpolation points required to connect the ending and starting points of the trajectories, and is required to fulfill the periodicity constraints of the discrete Fourier transform.

Fitted displacement and strain button. This reconstructs the displacement fields for each frame using 2–D back interpolation from the fitted trajectories. This process is described in Chapter 4. $E1$, $E2$, Ecc, and Err are then calculated from the reconstructed displacement fields.

Entire analysis button. This runs the processing steps associated with the *Unwrap phase*, *Raw Displacement and Strain*, *Track and Fit*, and *Fitted displacement and strain* buttons.

Plot and Export

Plot button. This plots the selected images for the frame chosen using the pull-down menu.

Displacement checkbox. Plots the 2–D or 3–D displacement field for the selected frame. If the *Remove DC* checkbox is selected, then the mean displacement component is subtracted from each respective displacement component.

Ecc checkbox. Plots the Ecc map for the selected frame.

Err checkbox. Plots the Err map for the selected frame.

1st principal strain checkbox. Plots the $E1$ map for the selected frame. Eigenvector directions are also shown.

2nd principal strain checkbox. Plots the $E2$ map for the selected frame. Eigenvector directions are also shown.

Strain-time curves checkbox. Plots the $E1$, $E2$, Ecc, and Err strain versus time curves for 6 circumferential LV segments. The user is asked to manually mark the anterior RV insertion point on a magnitude-reconstructed image to provide a reference for the LV segments. If no tracking has been done, the strain-time relationships are calculated using fixed segments and the original strain calculations. If tracking and temporal fitting has been done, then the strain-time relationships are calculated from the fitted trajectories directly. If tracking, temporal, fitting and reconstructing the original displacement fields have been done, then the strain-time relationships are calculated using fixed segments and strain measures derived from the reconstructed displacements.

Kinesia map checkbox. This represents on a colour image the total 2–D/3–D distance covered by each tracked point. Tissue tracking is a prerequisite.

k-space map checkbox. This displays the k-space of all encoding directions at the selected frame, and is only available if raw k-space data was loaded.

SNR checkbox. This plots a graph of SNR as a function of time for the encoding direction currently displayed. The signal for each frame is the mean image intensity of the regions enclosed by the contours. The user is prompted to manually select a background noise sample, the standard deviation of which defines the noise level. Contours for each frame are a prerequisite.

Raw trajectories checkbox. This plots the unfitted trajectories commencing at each myocardial pixel centre on the first cardiac phase. The user can choose to sub-sample the density of the trajectories using the *Scale* pull-down menu. Tissue tracking is a prerequisite.

Fitted trajectories checkbox. This plots the fitted trajectories commencing at each myocardial pixel centre on the first cardiac phase. The user can choose to sub-sample the density of the trajectories using the *Scale* pull-down menu. Tissue tracking is a prerequisite.

Make avi button. Creates avi's of the selected checkboxes. For all avi's the user is prompted for the type of codec, the compression quality, the frame rate, and the number of frames to include.

Magnitude avi checkbox. Creates an avi of the magnitude-reconstructed images of the encoding direction currently displayed. This can be done directly after loading the images, and often provides a good indicator of the correct displacement field vector orientation.

k-space avi checkbox. Creates an avi of the k-space for encoding in the direction currently displayed. This is only available if raw k-space data has been loaded.

Eulerian displacement avi checkbox. Creates an avi of the cine DENSE displacement fields overlaid onto magnitude-reconstructed images.

Lagrangian displacement avi checkbox. Creates an avi of the trajectory positions relative to their starting points. Tissue tracking is a prerequisite.

Trajectory avi checkbox. Creates an avi of the evolving trajectory paths plotted on the magnitude-reconstructed image at the first cardiac phase. Tissue tracking is a prerequisite.

Trajectory&strain avi checkbox. Creates an avi portraying both Lagrangian displacement and *E*cc on the same images. Tissue tracking is a prerequisite.

The 2–D and 3–D radio buttons only affect the following image and movie features: *Displacement* plots; *Kinesia map* plots; and *Lagrangian displacement avi.* The rest of the features inherently display 2–D information.

Export checked data to Excel button. Exports the strain versus time relationships for all cardiac segments into a Microsoft Excel spreadsheet. Note that the strain-time curves should be plotted before running this.

Strain GUI button. This initialises the strain display GUI described in Hess (2006). This GUI provides an interactive tool for viewing all features of the cine DENSE strain, as well as the deformation gradient. It is also capable of creating movies of each strain component.

CineDENSEAnalysis for the LV and RV

A separate GUI was created for analysis of both the LV and RV. The majority of the features are the same as for the LV GUI. The additional features include the following, most of which are described in Chapter 8:

Buttons to draw both the RV epicardial and endocardial contours, as well as copy preceding and succeeding RV contours. The contours are drawn by clicking on points along the feature of interest, commencing within the LV epicardial border. The contour is deemed complete when it returns to lie within the LV epicardial border again. The contours are smoothed using splines.

Full 3–D unwrap button. This runs the fully-automated spatiotemporal phase unwrapping algorithm described in Chapter 3, but with the unwrapping confined to the contoured regions. This is slow and should only be run if a more robust phase unwrapping algorithm is required.

Track midwall button. This finds the LV and RV midwall contours at the first frame using a kd-tree algorithm, places pixel-spaced points along these contours, and then applies a 2–D/3–D cine DENSE tissue tracking algorithm to estimate the evolution of these points through time.

Ett (trajectories) button. Calculates tangential strain (Ett) at each frame directly from the midwall trajectories.

Ett (displacements) button. This calculates Ett from the displacement fields directly. LV and RV midwall contours for all frames are found using a kd-tree algorithm, and points are placed at a pixel spacing along these contours. Two dimensional linear interpolation is then used to estimate the DENSE displacement vectors at each of these points. Ett is then calculated for all midwall points at all frames.

Ett strain-time checkbox. Divides the LV into six tangentially-spaced segments, and the RV into two tangentially-spaced segments. *E*tt versus time relationships are then plotted for all segments. The measure of strain confidence described in Chapter 8 is calculated and represented on colour bars beneath each strain-time curve.

Trajectory&Ett checkbox. Creates a movie where both *E*tt and the midwall Lagrangian displacement are portrayed on the same images.

Export checked data to Excel button. This exports *E*tt versus time relationships for all LV and RV segments into a Microsoft Excel spreadsheet.

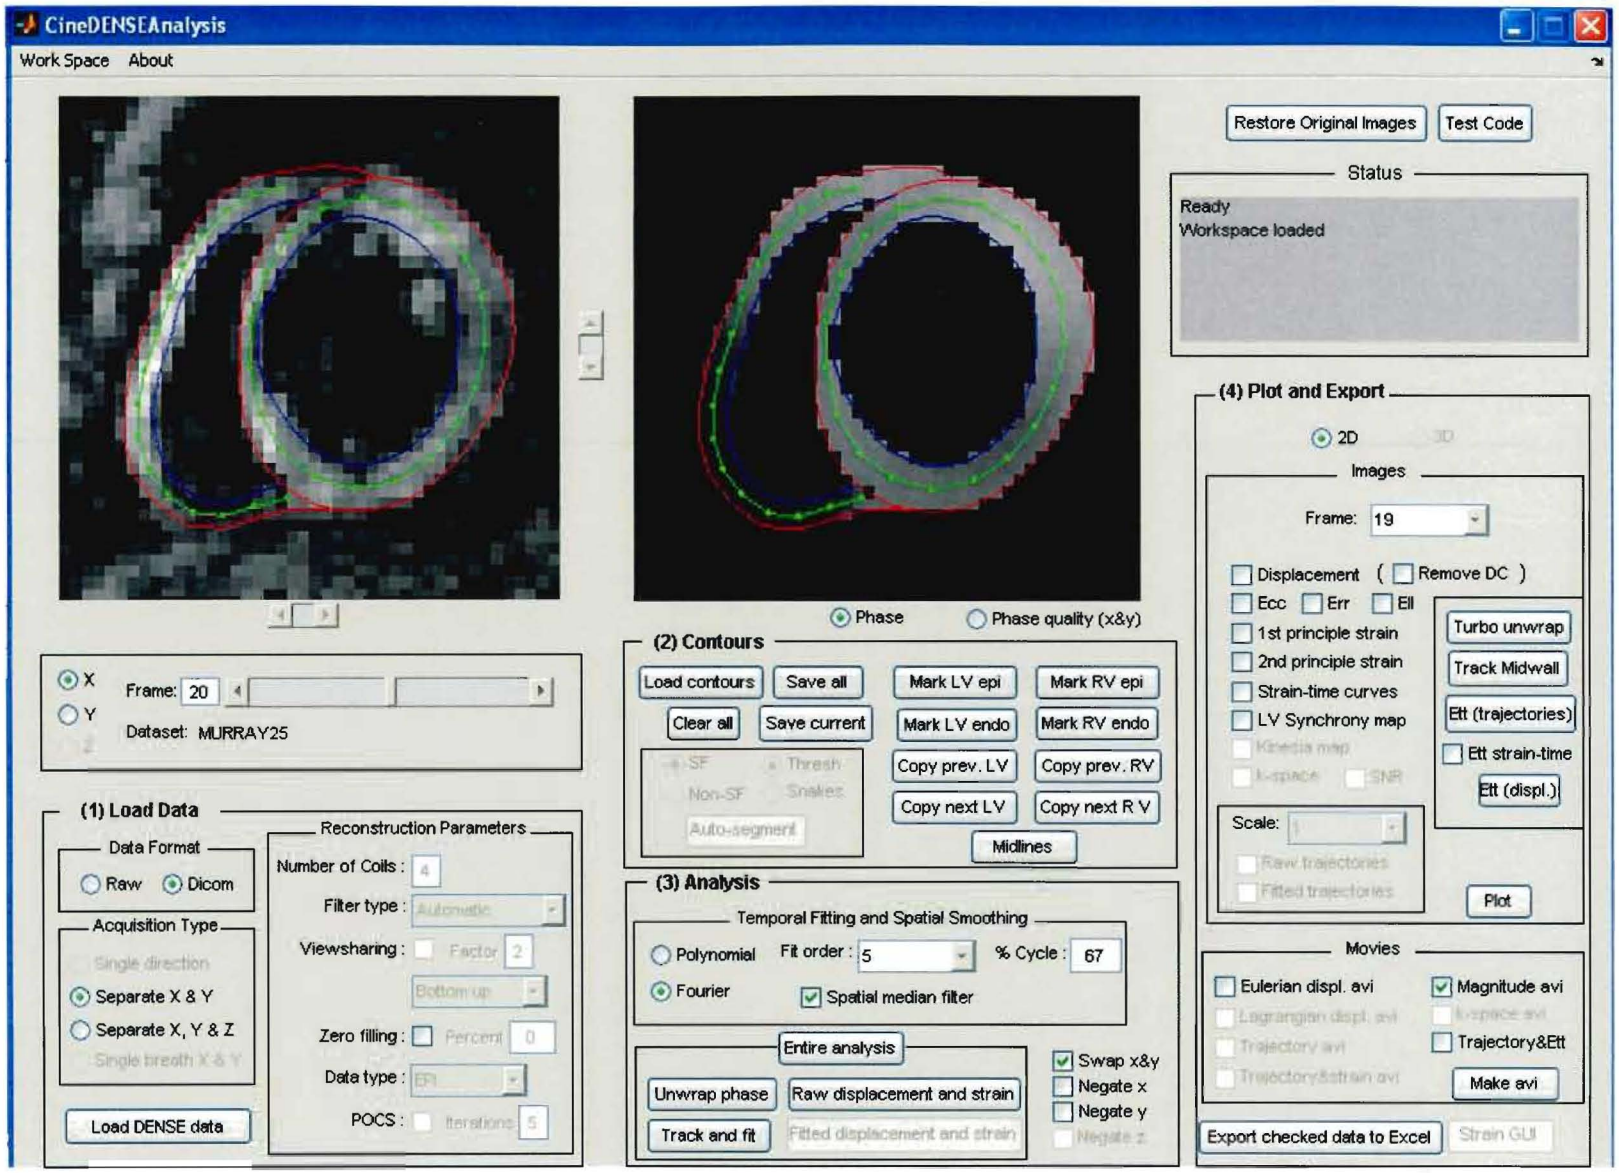

Figure A.2: CineDENSEAnalysis MATLAB graphical user interface for processing LV and RV cine DENSE data.

Appendix B

Spatial modulation of magnetisation (SPAMM)

The steps involved in a spatial modulation of magnetisation (SPAMM) sequence are described here based on the work by Axel and Dougherty (1989b) and Prince and McVeigh (1992). It is assumed that the tag pulse sequence is generated rapidly, so that the changes in longitudinal and transverse relaxation during the sequence are negligible. For convenience, the explicit time dependence is sometimes dropped in the notation. A SPAMM sequence consists of the following steps:

1. Assume an initial condition of uniform longitudinal magnetisation (Figure B.1(a)). Let $\mathbf{M}_0 = [M_{x0} M_{y0} M_{z0}]^T$ denote the net magnetisation at time t_0, where the z direction is aligned with the main magnetic field. Let $\mathbf{M}(t) = [M_x(t) M_y(t) M_z(t)]^T$ denote the net magnetisation at any time t.

2. The first RF pulse reduces this longitudinal magnetisation and introduces a transverse component (Figure B.1(b)). Consider the pulse to have a flip angle of θ degrees about the z-axis, and duration τ_θ. If the pulse is applied at $t = 0$, the resulting magnetisation can be described by

$$\mathbf{M}(\tau_{\theta_1}) = \begin{bmatrix} \cos\theta & 0 & \sin\theta \\ 0 & 1 & 0 \\ -\sin\theta & 0 & \cos\theta \end{bmatrix} \begin{bmatrix} M_{x0} \\ M_{y0} \\ M_{z0} \end{bmatrix} = \begin{bmatrix} M_{x0}\cos\theta + M_{z0}\sin\theta \\ M_{y0} \\ -M_{x0}\sin\theta + M_{z0}\cos\theta \end{bmatrix} \tag{B.1}$$

3. The gradient pulse produces a periodic spatial modulation of the phase of the transverse magnetisation (wraps it) along the direction of the gradient (Figure B.1(c)). Assume that a gradient of strength $G_x(t)$ and duration T_x is applied in the x direction immediately after the θ degree pulse. For gradient pulses of short duration, the solution to the Bloch equations (Appendix G) is given to good approximation (Prince and McVeigh, 1992) by

$$\mathbf{M}(\tau_{\theta_1} + T_x) = \begin{bmatrix} a\cos k_x x - M_{y0}\sin k_x x \\ a\sin k_x x + M_{y0}\cos k_x x \\ b \end{bmatrix} \tag{B.2}$$

where

$$a = M_{x0}\cos\theta + M_{z0}\sin\theta \tag{B.3}$$

$$b = -M_{x0}\sin\theta + M_{z0}\cos\theta \tag{B.4}$$

and

$$k_x = 2\pi\gamma \int_0^{T_x} G_x(\tau)d\tau \tag{B.5}$$

Here γ is the gyromagnetic ratio and τ is a dummy variable of integration.

4. The second RF pulse mixes the modulated transverse magnetisation with the longitudinal magnetisation, causing the total longitudinal magnetisation to be sinusoidally modulated along the gradient direction (Figure B.1(d)). The second RF pulse should have the same flip angle as the first, thus restoring the longitudinal magnetisation to its initial value at the peaks of the modulation. Mathematically, applying a θ degree pulse with duration τ_{θ_2} results in

$$\mathbf{M}(\tau_{\theta_1} + \tau_{\theta_2} + T_x) = \begin{bmatrix} a\cos\theta\cos k_x x - M_{y0}\cos\theta\sin k_x x + b\sin\theta \\ a\sin k_x + M_{y0}\cos k_x x \\ -a\sin\theta\cos k_x x + M_{y0}\sin\theta\sin k_x x + b\cos\theta \end{bmatrix} \tag{B.6}$$

This equation is derived by using the result in Equation B.2 in place of the vector $[M_{x0}M_{y0}M_{z0}]^T$ in Equation B.1.

5. The remaining transverse magnetisation can be spoiled by an optional second gradient pulse in the x direction. This is not, however, necessary since the magnetisation would be spoiled by the first gradient used in the imaging sequence that follows. If the spoiling gradient has a duration T_s, then

$$\mathbf{M}(\tau_{\theta_1} + \tau_{\theta_2} + T_x + T_s) = \begin{bmatrix} \approx 0 \\ \approx 0 \\ -a\sin\theta\cos k_x x + M_{y0}\sin\theta\sin k_x x + b\cos\theta \end{bmatrix} \tag{B.7}$$

Spoiling gradients are typically used after imaging pulse sequences, so one may safely assume that $M_{x0} = M_{y0} = 0$. Therefore, upon substitution for a and b, the z-magnetisation is

$$M_z(\tau_{\theta_1} + \tau_{\theta_2} + T_x + T_s) = M_{z0}(\cos^2\theta - \sin^2\theta\cos k_x x) \tag{B.8}$$

which is governed by a cosine function. The final modulated longitudinal magnetisation is shown in Figure B.1(e).

Image intensity is proportional to transverse magnetisation. The imaging sequence that follows the SPAMM sequence will cause a corresponding spatial modulation of the transverse magnetisation, thus rendering the tag lines visible.

The modulation *depth*, that is, the difference between maximum and minimum modulation intensities, is governed by the flip angle θ. The rate at which the spins dephase is proportional to their distance from the gradient centre. The line spacing is thus inversely proportional to the gradient amplitude and duration.

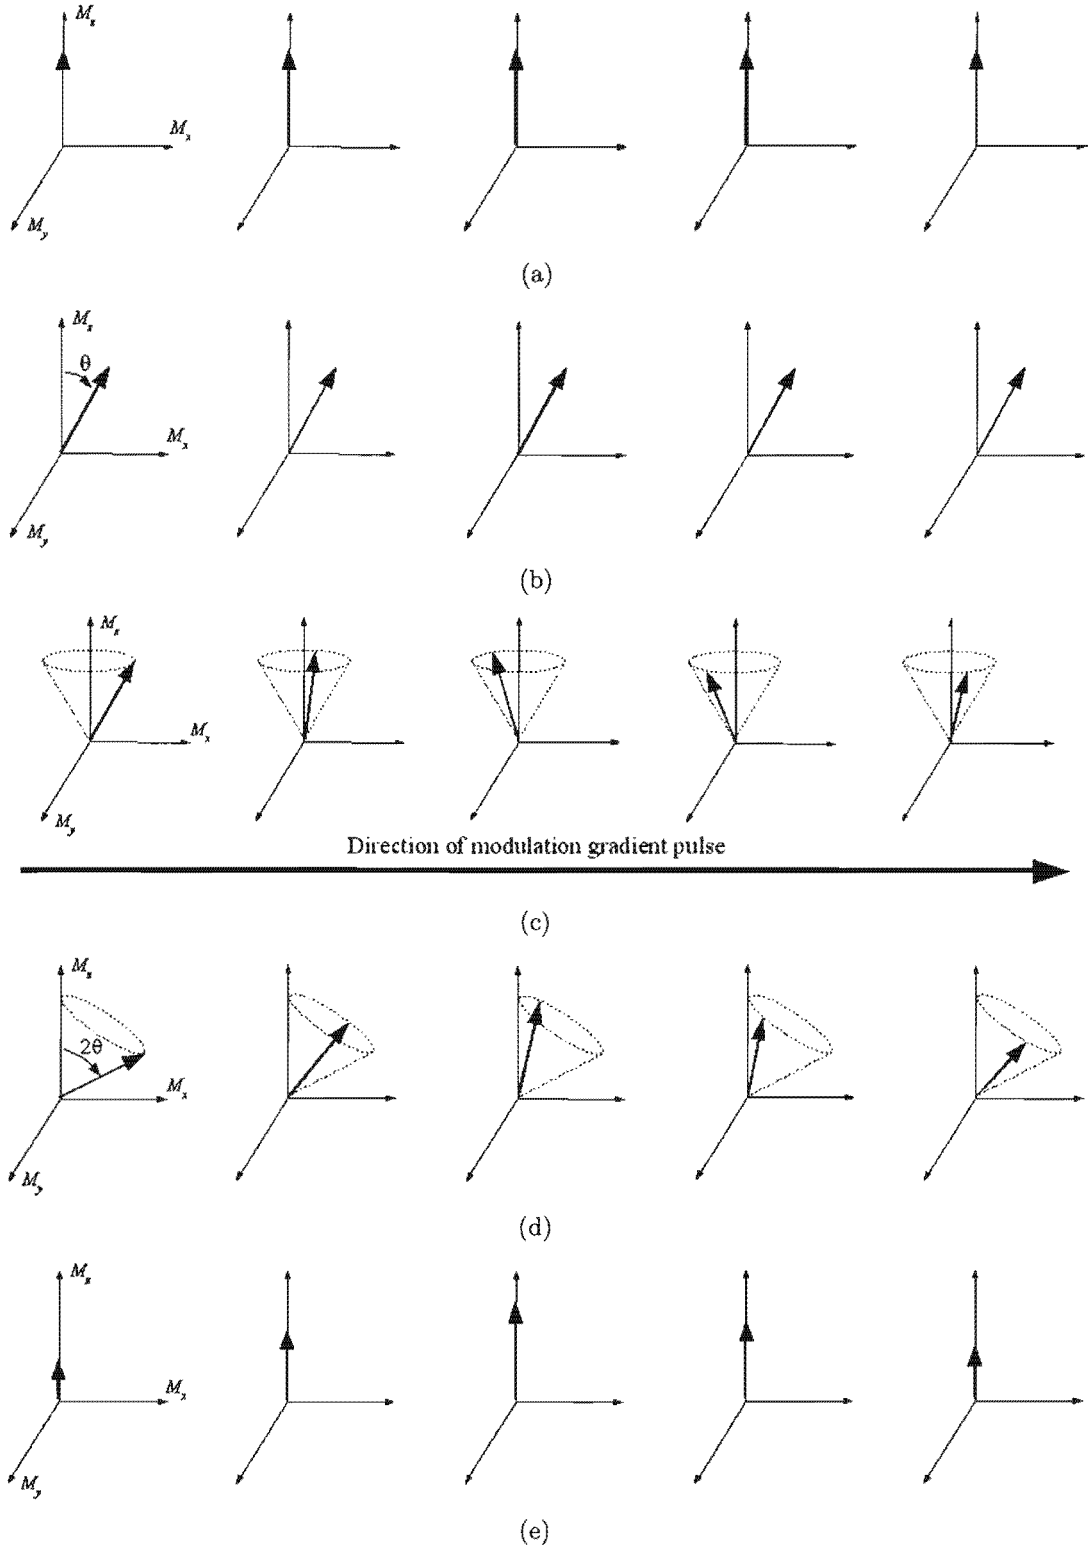

Figure B.1: State of magnetisation at different stages of a SPAMM pulse sequence. (a) Initial homogeneous magnetisation. (b) Magnetisation after the first RF pulse with flip angle θ. (c) Applied modulation gradient causes spins to dephase along the gradient. (d) Magnetisation after second RF pulse with flip angle θ. (e) Magnetisation after the remaining transverse magnetisation has been spoiled.

The basic 1–D SPAMM sequence just described can be extended to produce a grid pattern by repeating the series of events just described, but using a modulation gradient orthogonal to the first. This is introduced before the magnetisation is spoiled. Here the z–magnetisation component is given by

$$M_z(2\tau_{\theta_1}+2\tau_{\theta_2}+2T_s+T_x+T_y) = M_{z0}(\cos^2\theta-\sin^2\theta\cos k_x x)\cdot(\cos^2\theta-\sin^2\theta\cos k_y y) \quad \text{(B.9)}$$

The total duration of the SPAMM sequence proposed by Axel and Dougherty (1989b) is less than 3 ms, adding only a small delay to the initiation of imaging after R-wave detection. Axel and Dougherty (1989a) demonstrated in a further study that, with the substitution of more RF pulses with relative amplitudes distributed according to the binomial coefficients (eg. 1–3–3–1 or 1–4–6–4–1), sharper stripes are possible. The SPAMM sequence described above is thus classified as 1–1 SPAMM. Although longer sequences increase the duration to about 9 ms, they can still be used with conventional imaging sequences.

Appendix C

Myocardial tag tracking

Early attempts to monitor myocardial tissue deformation involved manually marking tag grid intersections on a frame-by-frame basis (Axel *et al.*, 1992). Even for very low tagging resolutions, a large number of tags must be detected and followed. Tag detection and tracking by manual inspection is a time consuming process and is prone to inaccuracies. A number of solutions have been proposed for automating tag detection, several of which are introduced here.

Tag tracking using active contour models

The most widely used method of tag line tracking is based on active contour models (or snakes), which were developed by Kass *et al.* (1988). A snake is an energy-minimising spline guided by a collection of forces. The position of the snake is represented parametrically by $v(s) = (x_1(s), x_2(s))$, where s denotes the arc length along the snake. The algorithm seeks to minimise a weighted combination of energy potentials

$$E_{snake} = \int_0^1 \{E_{int}(v(s)) + E_{image}(v(s)) + E_{con}(v(s))\} ds \tag{C.1}$$

Here E_{int} represents the internal energy of the spline due to bending, E_{image} gives rise to the image forces, and E_{con} gives rise to the external constraint forces. The internal spline forces serve to impose a piecewise smoothness constraint. The image forces push the snake toward salient features like lines, edges, and subjective contours. They attach

themselves into these features, localising them accurately. The external constraint forces could stem from, for example, image boundaries or user interaction.

Kumar and Goldgof (1994) were the first to adapt the snakes theory to myocardial tagging, and their method is briefly described here as an example. Kumar and Goldgof use a snake with a width of two pixels, corresponding to the width of the trough of a typical SPAMM sinusoid. The snake points are referred to as snaxels. The three specific energy functions, presented in the same format as in Equation C.1, are as follows: (1) A first-order continuity constraint that, when minimised, forces the snaxels to be uniformly spaced; (2) An image energy, defined as the product of the image intensities at two adjacent snaxels spanning the width of the snake. Minimising this term causes the snakes to seek pairs of adjacent pixels with minimum intensity. Since the SPAMM lines are typically the darkest objects in an image, they are tracked by the snakes; (3) A curvature energy, defined as the difference between two tangents of a curve, is used to provide a smooth contour for the snake.

The minimisation of the snake energy functions is an iterative process, and is done for each snaxel pair of each snake. The numerical methods to implement the minimisation process are presented in Kass *et al.* (1988). A brief description of the minimisation process, as applied by Kumar and Goldgof, is as follows:

1. Manually outline the region of interest as a square surrounding the heart.

2. Use template matching to find an initial approximate position for the snakes. This requires knowledge of the grid angle and spacing. To find the location of a grid line, a line template is placed on the image and slid over a distance equal to the spacing between the grid lines. The position that gives the least image energy corresponds to a grid line. Successive grid lines can be found by sliding the template further. The perpendicular grid lines are found in a similar manner.

3. For each snaxel pair, search a small neighbourhood and move the pair to a new location that will yield a smaller energy for the snake. The snaxel pairs are only allowed to move in a direction perpendicular to the snake. This is repeated for each snake and for a fixed number of iterations per snake.

4. The first few iterations are performed only with the image energy, resulting in the scenario in Figure C.1. Note that the snake has correctly identified the grid lines

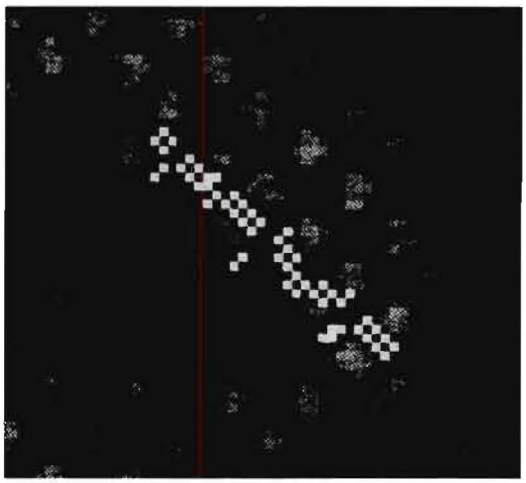

Figure C.1: Active contour model with image energy only. Image from Kumar and Goldgof (1994), © 1994 IEEE, reprinted with permission.

everywhere except at intersections. This is rectified by increasing the curvature and continuity coefficients and by neutralising small variations in intensity of the grid pixels. This entire process is repeated for each snake up to a limited number of iterations, at which point it is assumed that each point in the grid is located.

5. With the grid identified, the intersection points can be computed as the intersection points of the snakes.

6. Using the final contours of the snakes of one image as the starting point for the next image, the entire spatiotemporal motion of the grid can be tracked.

Other adaptations of the energy functions have subsequently been developed. Kraitchman *et al.* (1995) improved on previous techniques by allowing stray points within the snake to be intuitively guided into place to avoid incorrect local minima, and also introducing a more robust template matching method of identifying tag intersections in the first image of a series. Kraitchman *et al.* state that, prior to their improvements, an analysis of a hypertrophic cardiomyopathy patient would require more than 5 hours of manual user interaction time because of the increased number of tag intersections associated with the hypertrophy. This study could be evaluated in less than 25 minutes using their semi-automatic technique. In tests with a silicon gel phantom with controlled deformations, the average discrepancy between theoretically predicted and semi-automatically selected tag intersections was 0.3 ± 0.17 mm.

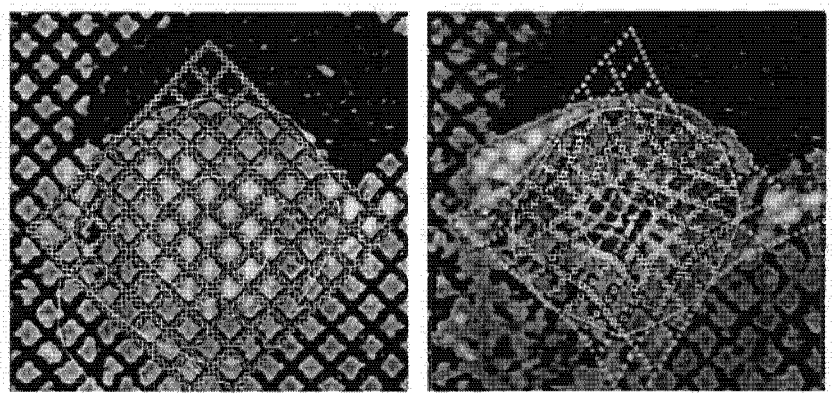

(a) (b)

Figure C.2: Example of short-axis tagged images with the superimposed active contour models during (a) end-diastole, and (b) end-systole. Image from Young *et al.* (1995), © 1995 IEEE, reprinted with permission.

Young *et al.* (1995) adapted the snakes algorithm to incorporate user interaction, whereby user-defined constraints can be set during the minimisation process. They also interpolated image forces between pixels to track stripes at a subpixel resolution. Figure C.2 shows the results of stripe tracking for two frames using the method by Young *et al.*. The snakes comprise a grid consisting of 10 stripes in each direction, connected at 100 points. Grid points outside the myocardium provide a weak form of continuity.

Tag tracking using optical flow

Optical flow is a method for analysing the motion of pixels between image pairs. It uses image brightness variations between consecutive frames to compute in-plane displacements or velocities of moving image features. A major advantage of optical flow methods is that they provide pixel level velocity information and are therefore not limited to the resolution of the tagging grid. Amartur and Vesselle (1993) applied optical flow methods to standard MR cine images of a single patient. Without anatomical landmarks, the estimated flow fields proved to be ambiguous. The application of optical flow theory to MR tagging is attractive because the high contrast stripes allow more accurate tracking of the material points from frame to frame.

A key assumption of the optical flow theory is that brightness is conserved between successive frames. This is not the case with tagged images, where tag fading takes place. Prince and McVeigh (1992) applied the method to tagged images, taking into account

the variable intensity. However, their method requires a priori knowledge of the $T1$, $T2$, and spin density of the myocardium. Dougherty *et al.* (1999) introduced an optical flow method that does not rely on prior knowledge of image content and overcomes the requirement of constant pixel intensity by pre-processing the images. They claim that their method is as fast as the snakes method and more accurate. It uses a fully automated tracking technique, which computes the flow field between sequential images and uses this to compute the displacement field. The tracking of the tag intersections takes less than 5 seconds per slice per phase of the cardiac cycle. The end systolic RMS error of 0.23 mm indicates accuracy within the 0.3 mm measurement error of the snakes method by Kraitchman *et al.* (1995).

Tag tracking using Fourier filtering

Groot-Koerkamp *et al.* (1999) describe a Fourier-based approach for automatically detecting the tag positions. This is done by using an appropriately matched bandpass filter, which detects both the grid crossings as well as the grid centres in one single filtering process. Because the tags are excited in a rectangular grid pattern, the Fourier spectra of the images have a grid of peaks. By filtering out a number of these spectral peaks, and thresholding the resulting images, the tag positions are rapidly identified.

Initially, the peaks can be found at known positions, defined by the spatial frequencies of the grid. As the grid deforms and rotates, the peaks spread out and move to slightly different frequencies. Groot-Koerkamp *et al.* applied their method to tagged images of the thigh with a computer simulated deformation. The average discrepancy between theoretically predicted and automatically detected tag positions was 0.04 ± 0.17 mm.

Recently, Osman *et al.* (1999) and Garot *et al.* (2000) introduced and evaluated a method for cardiac motion tracking using a Hilbert transform and the harmonic phase (HARP) of tagged images. This method is described in detail in Section 2.2.3.

Other techniques

A drawback of the majority of the above-mentioned models is a high demand in processing time. Allouche *et al.* (2001) proposed a fast method that uses a novel kinetic

deformation class. The method relies on very accurate image registration, but allows for automatic feature propagation and a sharp 2–D motion analysis of the walls. The processing time was of the order of a second. Chandrashekara *et al.* (2002) present a further fully automated 3–D technique based on nonrigid image registration using multi-level freeform deformations. The novel aspect of their technique is its integrated nature with tag localisation and deformation field reconstruction. Only a limited number of tests have been carried out using these two techniques.

Appendix D

Deformation gradient and finite strain tensors

Strain may be calculated relative to a Lagrangian or Eulerian description of motion. A reference map $\mathbf{p}(\mathbf{x}, t)$ relates the spatial position $\mathbf{x}$ of a point at time t to its corresponding reference position $\mathbf{p}$. This is known as a Eulerian description of motion, and the corresponding coordinates are known as spatial coordinates. The Lagrangian description involves a forward map $\mathbf{x}(\mathbf{p}, t)$, which relates the reference position $\mathbf{p}$ to its deformed position $\mathbf{x}$ at time t. These coordinates are known as material coordinates. Material coordinates are used in conjunction with the reference configuration only, and the spatial coordinates serve for all other configurations.

In the Eulerian description, attention is focused on what is happening to the individual particles during the motion, whereas in the Lagrangian description the emphasis is directed to the events taking place at specific points in space. For example, the displacement field in Lagrangian coordinates is $\mathbf{u}(\mathbf{p}, t) = \mathbf{x}(\mathbf{p}, t) - \mathbf{p}$. This describes the displacement that will occur to the particle that starts at $\mathbf{p}$. The displacement field in Eulerian coordinates is $\mathbf{u}(\mathbf{x}, t) = \mathbf{x} - \mathbf{p}(\mathbf{x}, t)$. This describes the displacement that the particle now at $\mathbf{x}$ has undergone.

Consider the Lagrangian description of motion $\mathbf{x} = \mathbf{x}(\mathbf{p}, t)$, where the $\mathbf{x}$ on the right hand side of the equation represents the function whose arguments are $\mathbf{p}$ and t, while the same symbol on the left hand side represents the value of the function, i.e. a point in space. This is common notation in continuum mechanics. For the sake of a deformation

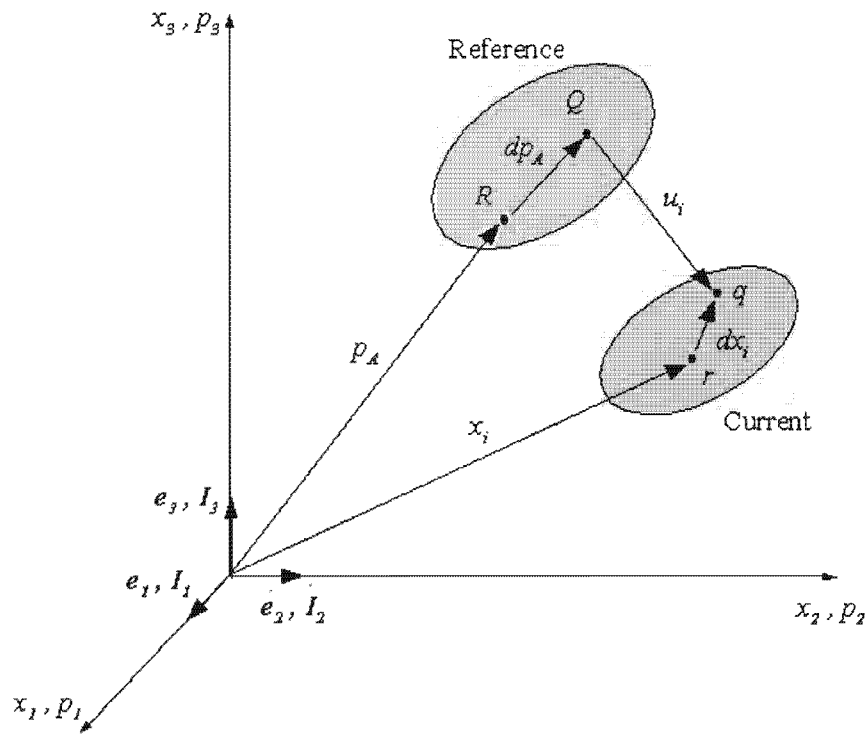

Figure D.1: Deformation between two sets of points in a continuum.

analysis, the mapping function need not be dependent on time, so $\mathbf{x} = \mathbf{x}(\mathbf{p})$. This is because only two stationary configurations are investigated, and no consideration is necessary for how the final deformed configuration is reached from the initial undeformed configuration.

The following description relates to the diagram in Figure D.1. Here, two points R and Q in the reference configuration, are deformed to positions r and q in the current configuration. The orthogonal unit vectors $[\hat{\mathbf{I}}_1, \hat{\mathbf{I}}_2, \hat{\mathbf{I}}_3]$ form a basis for the reference configuration $[p_1, p_2, p_3]$. Similarly, the vectors $[\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2, \hat{\mathbf{e}}_3]$ form a basis for the current configuration $[x_1, x_2, x_3]$.

The point Q is located with respect to R by the relative differential position vector

$$d\mathbf{p} = d\mathbf{p}_A \hat{\mathbf{I}}_A \tag{D.1}$$

After the particles move from R and Q to r and q, respectively, their relative position vector in the deformed configuration is

$$d\mathbf{x} = d\mathbf{x}_i \hat{\mathbf{e}}_i \tag{D.2}$$

Assuming that the mapping function $\mathbf{x} = \mathbf{x}(\mathbf{p})$ is continuous,

$$d\mathbf{x}_i = \frac{\partial \mathbf{x}}{\partial \mathbf{p}} d\mathbf{p}_A \tag{D.3}$$

A tensor is a mathematical entity which is used to represent the important physical quantities of continuum mechanics. Scalars are zeroth order tensors, vectors are first order tensors, stress and strain tensors are second order. The term $\frac{\partial \mathbf{x}}{\partial \mathbf{p}}$ characterises the local deformation at $\mathbf{p}$, and is known as the deformation gradient tensor F. It follows that $d\mathbf{x} = Fd\mathbf{p}$. For $\mathbf{x} = [x_1, x_2, x_3]$ and $\mathbf{p} = [p_1, p_2, p_3]$, the deformation gradient tensor is given by

$$F = \begin{bmatrix} \frac{\partial x_1}{\partial p_1} & \frac{\partial x_1}{\partial p_2} & \frac{\partial x_1}{\partial p_3} \\ \frac{\partial x_2}{\partial p_1} & \frac{\partial x_2}{\partial p_2} & \frac{\partial x_2}{\partial p_3} \\ \frac{\partial x_3}{\partial p_1} & \frac{\partial x_3}{\partial p_2} & \frac{\partial x_3}{\partial p_3} \end{bmatrix} \tag{D.4}$$

where the dependence of F on $\mathbf{p}$ and t is implicit. The smoothness conditions of the mapping function $\mathbf{x} = \mathbf{x}(\mathbf{p})$ signifies that F is invertible. Thus F^{-1} exists such that $d\mathbf{p} = F^{-1}d\mathbf{x}$.

There are a number of common measures of deformation based on the deformation gradient tensor. The first of these is Green's deformation tensor, which is a symmetric tensor based on the change during the deformation in the magnitude squared of the distance between the particles. It is given by

$$G = F^T \cdot F \tag{D.5}$$

From this, the Lagrangian finite strain tensor L can be immediately defined as

$$L = \frac{1}{2}(G - I) \tag{D.6}$$

where I is an identity matrix. The change during the deformation in the magnitude squared of the distance between the particles can also be represented in Eulerian form, that is, in spatial coordinates. The Cauchy deformation tensor is thus defined as

$$C = (F^{-1})^T \cdot (F^{-1}) \tag{D.7}$$

From this, the Eulerian finite strain tensor E can be given by

$$E = \frac{1}{2}(I - C) \tag{D.8}$$

Note that only a few of the concepts encompassing a strain analysis are presented here. Further detail can be found in any continuum mechanics text, for example Mase and Mase (1999).

Quantifying myocardial stress necessitates the knowledge of the spatial orientation of the muscle fibres, as well as components of Young's modulus and the shear modulus (Hu *et al.*, 2002). This information can not as yet be gauged *in vivo* using noninvasive imaging methods, although the use of diffusion imaging has yielded promising results in determining myocardial fiber orientation (Masood *et al.*, 2000). Studies on specimens have shown that the left ventricle has a well ordered distribution of fibre angles, varying from about 60 degrees from the circumferential direction near the endocardium, to about -60 degrees near the epicardium (Mirsky *et al.*, 1974).

Appendix E

SPAMM in the frequency domain

As described in Appendix B, 1–D 1–1 SPAMM generates a tag pattern by applying a θ degree RF pulse, followed by a gradient $\mathbf{G}$, followed by another θ degree RF pulse. The sequence is completed by crushing all lateral magnetisation.

Vector notation will be used here to provide a framework compatible with 2–D grid tagging. Let end-diastole be considered to be at time $t = 0$, and let every material point in the heart be marked by its position $\mathbf{p}$ at this reference time. If $I_0(\mathbf{p})$ is the intensity of the material points in the absence of the tag pulse sequence, then from Equation B.8, the application of 1–D 1–1 SPAMM will modulate this intensity to produce

$$I_1(\mathbf{p}) = I_0(\mathbf{p})T(\mathbf{p};\mathbf{k},\theta) \tag{E.1}$$

where

$$T(\mathbf{p};\mathbf{k},\theta) = \cos^2\theta - \sin^2\theta\cos(\mathbf{k}^T\mathbf{p}) \tag{E.2}$$

The variable $\mathbf{k}$ is the tag spatial frequency in radians/mm. Equation E.2 can also be written in the form

$$T(\mathbf{p};\mathbf{k},\theta) = \underbrace{C_1(\theta)\cos(\mathbf{k}_0^T\mathbf{p})}_{DC} + \underbrace{C_2(\theta)\cos(\mathbf{k}^T\mathbf{p})}_{fundamental\ frequency} \tag{E.3}$$

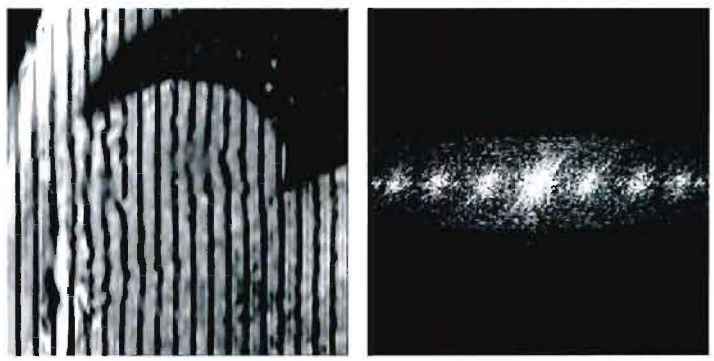

(a) (b)

Figure E.1: Example of (a) a 1–D tagged image, and (b) its corresponding 2–D Fourier transform.

where $C_1(\theta) = \cos^2\theta$, $C_2(\theta) = -\sin^2\theta$ and $\mathbf{k}_0 = 0$. Because a cosine has two symmetric spectral peaks in the Fourier domain, a 1–D 1–1 SPAMM pattern has three spectral peaks. More generally, a 1–D SPAMM pattern generated with K RF pulses has $2K - 1$ spectral peaks. Consider the 1–D (or planar) tagged image in Figure E.1(a) and the corresponding Fourier transform in Figure E.1(b). There are seven spectral peaks, implying that four RF pulses were delivered, i.e. the binomial coefficients 1–3–3–1.

Generating a 2–D 1–1 SPAMM pattern is done by applying two 1–D 1–1 SPAMM sequences in succession. For the two tag directions $\mathbf{k}_1$ and $\mathbf{k}_2$, the resulting modulated intensity is

$$I_2(\mathbf{p}) = I_0(\mathbf{p})T(\mathbf{p};\mathbf{k}_1,\theta)T(\mathbf{p};\mathbf{k}_2,\theta) \tag{E.4}$$

Substituting Equation E.2 into Equation E.4, and using trigonometric identities, results in the following weighted sum of cosines

$$\begin{aligned} I_2(\mathbf{p}) \;=\; & I_0(\mathbf{p})[2\cos^2\theta - \cos^2\theta\sin^2\theta\cos(\mathbf{k}_1^T\mathbf{p}) \\ & - \sin^2\theta\cos^2\theta\cos(\mathbf{k}_2^T\mathbf{p}) + \sin^2\theta\cos(\mathbf{k}_3^T\mathbf{p}) \\ & + \sin^2\theta\cos(\mathbf{k}_4^T\mathbf{p})] \end{aligned} \tag{E.5}$$

where $\mathbf{k}_3 = \mathbf{k}_1 + \mathbf{k}_2$ and $\mathbf{k}_4 = \mathbf{k}_1 - \mathbf{k}_2$. Writing $\mathbf{k}_0 = 0$, as in the 1–D case, Equation E.5 can also be written as

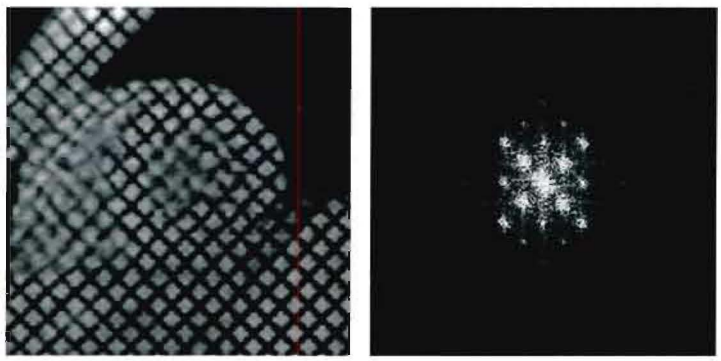

(a) (b)

Figure E.2: Example of (a) a 2–D tagged image, and (b) its corresponding 2–D Fourier transform.

$$I_2(\mathbf{p}) = I_0(\mathbf{p}) \sum_{n=0}^{4} D_n(\theta) \cos(\mathbf{k}_n^T \mathbf{p}) \tag{E.6}$$

Equation E.6 shows that a 2–D 1–1 SPAMM image is the summation of five images comprising a baseband pattern and four sinusoidally modulated patterns. The tagged image shown in Figure E.2 is obtained using 2–D 1–2–1 SPAMM. Because a 2–D SPAMM pattern is a product of two 1–D SPAMM patterns, it has $(2K-1)^2$ spectral peaks in k-space. Figure E.2 demonstrates this, with a single peak at DC, eight peaks in the primary sinusoid directions, and sixteen peaks in the direction of products of these sinusoids. Except for the case when $n = 0$, which corresponds to the spectral peak at the origin, each value of n corresponds to a pair of radially symmetric peaks in the direction $\mathbf{k}_n$. In general, tag fading will only affect the D_n functions.

The locations $\mathbf{k}_n$ of the spectral peaks in k-space are easily found from the directions of the gradient $\mathbf{G}$. Incorporating the reference map into Equation E.6 and modifying the equation for general SPAMM gives the following unified description of the tagged MR image during motion (Osman and Prince, 2000a)

$$I(\mathbf{p}(\mathbf{x},t)) = \sum_{n=1}^{N} I_0(\mathbf{p}(\mathbf{x},t)) C_n(t) e^{j\mathbf{k}_n^T \mathbf{p}(\mathbf{x},t)} \tag{E.7}$$

where $N = (2K-1)$ for 1–D SPAMM and $N = (2K-1)^2$ for 2–D SPAMM. The coefficients C_n are readily determined from D_n and are made functions of time to account for tag fading. The dependence of I on the tip angle θ has been ignored for notational

convenience. Equation E.7 shows that the SPAMM tagged image is the sum of N complex harmonic images, each corresponding to a distinct spectral peak identified by the frequency vector $\mathbf{k}_n$.

Appendix F

Calculating strain from HARP images

Strain can most simply be described as the normalised change in length in a particular direction. The most common strain measurement found in literature on cardiac motion is unit elongation in a direction defined by a vector $\mathbf{e}$ (Kerwin *et al.*, 1999). Using this notation, Lagrangian strain is given by

$$E(\mathbf{e}_p) = \frac{\|F\mathbf{e}_p\|}{\|\mathbf{e}_p\|} - 1 \tag{F.1}$$

and Eulerian strain is given by

$$E(\mathbf{e}_x) = \frac{\|\mathbf{e}_x\|}{\|F^{-1}\mathbf{e}_x\|} - 1 \tag{F.2}$$

Here $\mathbf{e}_p$ and $\mathbf{e}_x$ are vectors defining the direction of strain for Eulerian and Lagrangian descriptions, respectively, and F is the deformation gradient tensor. Appendix D provides background on frames of reference and strain tensors. The directions of strain in studies of the heart are usually orientated in a circumferential or radial direction relative to the centre of the left ventricle.

The HARP analysis lends itself well to determining Eulerian strain. During positive strain (stretch) the distance between two points increases, which in turn has the effect of decreasing the gradient of the phase between these points on an unwrapped HARP

image. Deformation at a point on the myocardium is thus inversely proportional to the gradient of the phase at that point.

Consider the 1–D case. Let ∇_x depict the phase gradient with respect to $\mathbf{x}$. This corresponds to the inverse of the deformation gradient tensor F, i.e. $F^{-1} = \nabla_x \mathbf{p}(\mathbf{x}, t)$. The gradient of the phase can thus be introduced in terms of the unwrapped phase image, $\phi^*(\mathbf{x}, t)$, as $\nabla_x \mathbf{p}(\mathbf{x}, t) = \mathbf{k}_1^{-1} \nabla_x \phi^*(\mathbf{x}, t)$, where $\mathbf{k}_1$ is the tag spatial frequency. Using this with Equation F.2 gives the equation for 1–D Eulerian strain

$$E(\mathbf{e}) = \|(\nabla_x \phi^*(\mathbf{x}, t))^{-1} \mathbf{k}_1 \mathbf{e}\| - 1 \tag{F.3}$$

where $\mathbf{e} \in \mathbb{R}^2$ is a unit vector defining the local direction of interest. The Euclidian norm here serves to give a scalar value at $(\mathbf{x}, t)$ and in the direction of $\mathbf{e}$. In the two dimensional case, the Eulerian strain is given by

$$E(\mathbf{e}) = \|(\nabla \phi^*(\mathbf{x}, t))^{-1} W \mathbf{e}\| - 1 \tag{F.4}$$

Here $\phi = [\phi_1 \ \phi_2]$, $\nabla \phi^*(\mathbf{x}, t)$ is the two dimensional representation of the deformation gradient, and $W = diag(\mathbf{k}_1, \mathbf{k}_2)$.

Measuring Lagrangian strain using HARP is only possible if adjacent points in the myocardium are tracked through time. This can be done in 2–D using a method described by Osman *et al.* (1999), where a Newton-Raphson iterative procedure was used to locate the closest set of identical harmonic phase coordinates in successive frames of a tagged series.

Appendix G

The Bloch equations

The Bloch equations describe the net magnetisation vector $\mathbf{M}$ of a sample of spins placed in a magnetic field $\mathbf{B}$. They were introduced in 1946 by Felix Bloch, and form the basis for nuclear magnetic resonance. Bloch assumed that a collection of spins represented as a net nuclear magnetisation $\mathbf{M}$ could be treated according to the laws of classical electromagnetism. In the presence of an external magnetic field $\mathbf{B}$, these spins experience a torque given by the vector cross product $\mathbf{M} \times \mathbf{H}$. The quantity $\mathbf{H}$ is known as the auxillary field, and it is given by $\mathbf{H} = \frac{1}{\mu_0}\mathbf{B} - \mathbf{M}$, where the constant μ_0 is the permeability of free space. The auxillary field $\mathbf{H}$ is a commonly-used measure in electromagnetics. In building an electromagnet, a certain current is passed through a coil. It is this current that is measured, and through a variation of Amperes law, describes $\mathbf{H}$. The magnetic field $\mathbf{B}$ depends on the specific materials that are used and even, if iron is present, on the history of the magnet.

Bloch noted that the torque of a system is equal to the time rate of change of its angular momentum, and, by the Larmor equation, that $\mathbf{H}$ is proportional to this angular momentum through the gyromagnatic ratio γ. From this he showed that the motion of $\mathbf{M}$ over time could be described by the vector equation

$$\frac{d\mathbf{M}}{dt} = \gamma(\mathbf{M} \times \mathbf{H}) \tag{G.1}$$

When $\mathbf{B}$ is static (denoted $\mathbf{B}_0$), this equation describes the motion of $\mathbf{M}$ as a simple precession around $\mathbf{B_0}$ with a frequency $\omega_0 = \mathbf{B}_0\gamma$. This relationship assumes that the

spins interact only with $\mathbf{B}_0$, and not with each other or their local environment. In reality this is not the case, and $\mathbf{M}$ will not precess indefinitely around $\mathbf{B}_0$, but rather progressively return to its initial alignment parallel to $\mathbf{B}_0$. Consider the magnetisation vector $\mathbf{M}$ to be comprised out of three components M_x, M_y and M_z, with initial positions M_{x0}, M_{y0} and M_{z0}, corresponding to their initial alignment with $\mathbf{B}_0$. Bloch described the loss of energy associated with this return to initial alignment by two first order systems with time constants $T1$ and $T2$, giving the relations

$$\frac{dM_z}{dt} = \frac{M_0 - M_z}{T1} \tag{G.2}$$

$$\frac{dM_y}{dt} = \frac{M_y}{T2} \tag{G.3}$$

$$\frac{dM_x}{dt} = \frac{M_x}{T2} \tag{G.4}$$

Including the components of the torques that the spins experience in the magnetic field (Equation G.1) results in the Bloch equations:

$$\frac{dM_z}{dt} = \gamma(\mathbf{M} \times \mathbf{H})_z + \frac{M_0 - M_z}{T1} \tag{G.5}$$

$$\frac{dM_y}{dt} = \gamma(\mathbf{M} \times \mathbf{H})_y + \frac{M_y}{T2} \tag{G.6}$$

$$\frac{dM_x}{dt} = \gamma(\mathbf{M} \times \mathbf{H})_x + \frac{M_x}{T2} \tag{G.7}$$

Note that $\mathbf{H}$ is parallel to the z-axis in MR imaging, which makes the cross product in the first equation equal to zero. The integration of the first equation thus yields the formula for the z-magnetisation as a function of time

$$M_z = M_{z0}(1 - e^{-t/T1}) \tag{G.8}$$

Equations G.6 and G.7 are coupled differential equations, which for the case of a 90° pulse, have the solution

$$M_x = M_{x0}\cos(\omega_0 t)e^{-t/T2} \tag{G.9}$$

$$M_y = M_{y0}\sin(\omega_0 t)e^{-t/T2} \tag{G.10}$$

Equations G.8, G.9 and G.10 predict that $\mathbf{M}$ will exhibit a spiralling precession around $\mathbf{B}_0$ at the Larmor frequency. Computation of the length of the magnetisation vector in the xy–plane yields the well-known $T2$ relaxation relation

$$M_{xy} = M_{xy0}e^{-t/T2} \tag{G.11}$$

Appendix H

Dyssynchrony tangential strain data for two normal volunteers, two patients with LBBB, and two patients with RBBB

Figures H.1 and H.2 show the Ett versus time curves for the normal volunteers, together with the corresponding end-systolic maps of Ett. The onset of Ett is relatively uniform in all segments for both volunteers.

Figure H.3(a) shows Ett versus time curves for a patient with a LBBB. A delay in onset to contraction is clear in the lateral LV free wall and posterior LV free wall. Figure H.3(b) shows the corresponding end-systolic map of Ett, where very little LV and RV contractility is evident, and the septum is actually bulging into the RV cavity. Figure H.3(c) shows a bullseye plot, where outer, middle and inner rings portray the percentage of Gd-enhancement at base, mid and apical short-axis LV slices, respectively.

Figure H.4(a) shows Ett versus time curves for the second patient with a LBBB. Here a delay to onset of contraction is also evident in the posterior LV wall, and the RV function is normal. Figure H.4(b) shows very little contractility in the LV, with the septum again bulging into the RV cavity. No infarcted tissue is evident in Figure H.4(c).

Figure H.5(a) shows Ett versus time for a patient with a RBBB. An infarct centred in the anterior septum is evident both in Figure H.5(b) and H.5(c). Delays in the onset of contraction are evident in both RV segments, but a slight delay is also evident in the lateral LV free wall. The lack of contractility in the infarcted anterior septum is also evident.

Figure H.6(a) shows the Ett versus time curves for a patient with a RBBB and a left anterior fascicular block (LAFB). An infarct centred in the posterior LV wall is depicted in Figure H.6(c), but this does not correlate with the Ett map in Figure H.6(b). The contractility of the LV free wall appears healthy in the strain-time curves, and a significant delay is present in the RV free wall.

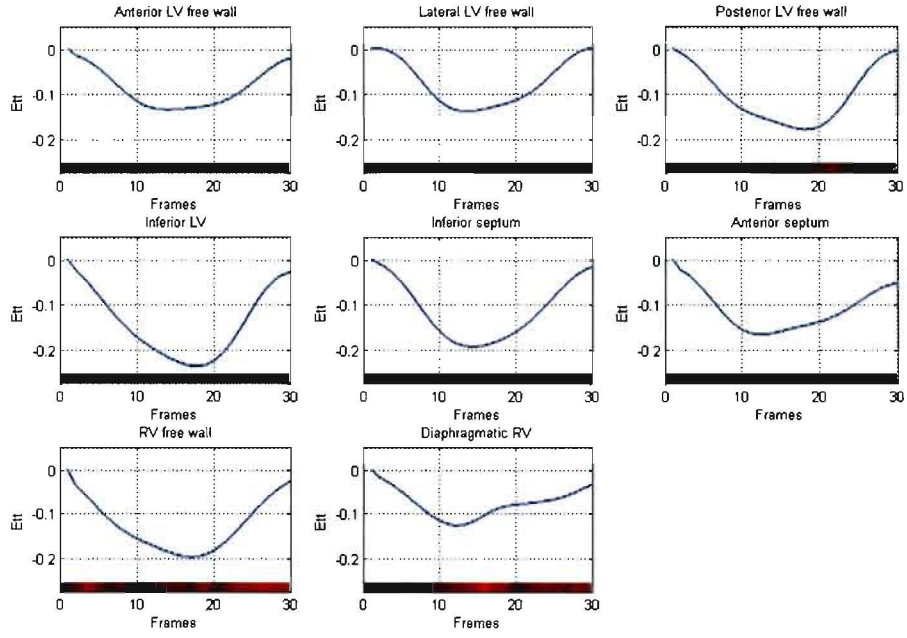

(a)

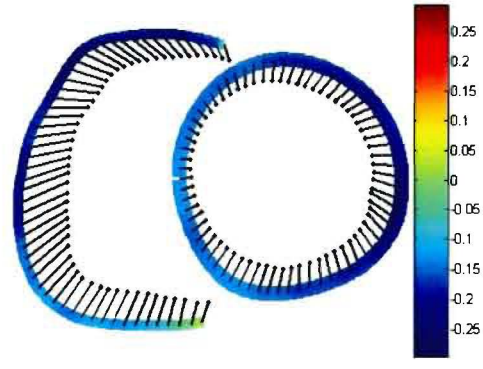

(b)

Figure H.1: Normal volunteer 1. (a) Ett versus time for the segments shown in Figure 8.4, and (b) end-systolic map of midwall Ett and 2–D trajectory position.

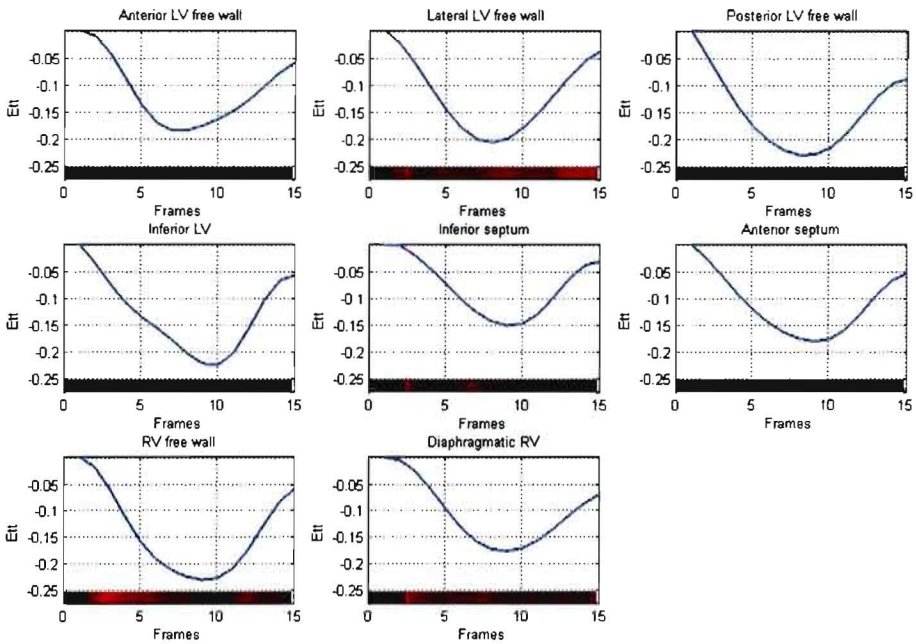

(a)

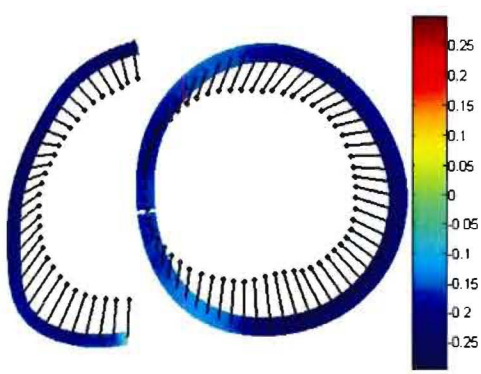

(b)

Figure H.2: Normal volunteer 2. (a) Midwall Ett versus time for the segments shown in Figure 8.4, (b) End-systolic map of midwall Ett and 2–D trajectory position.

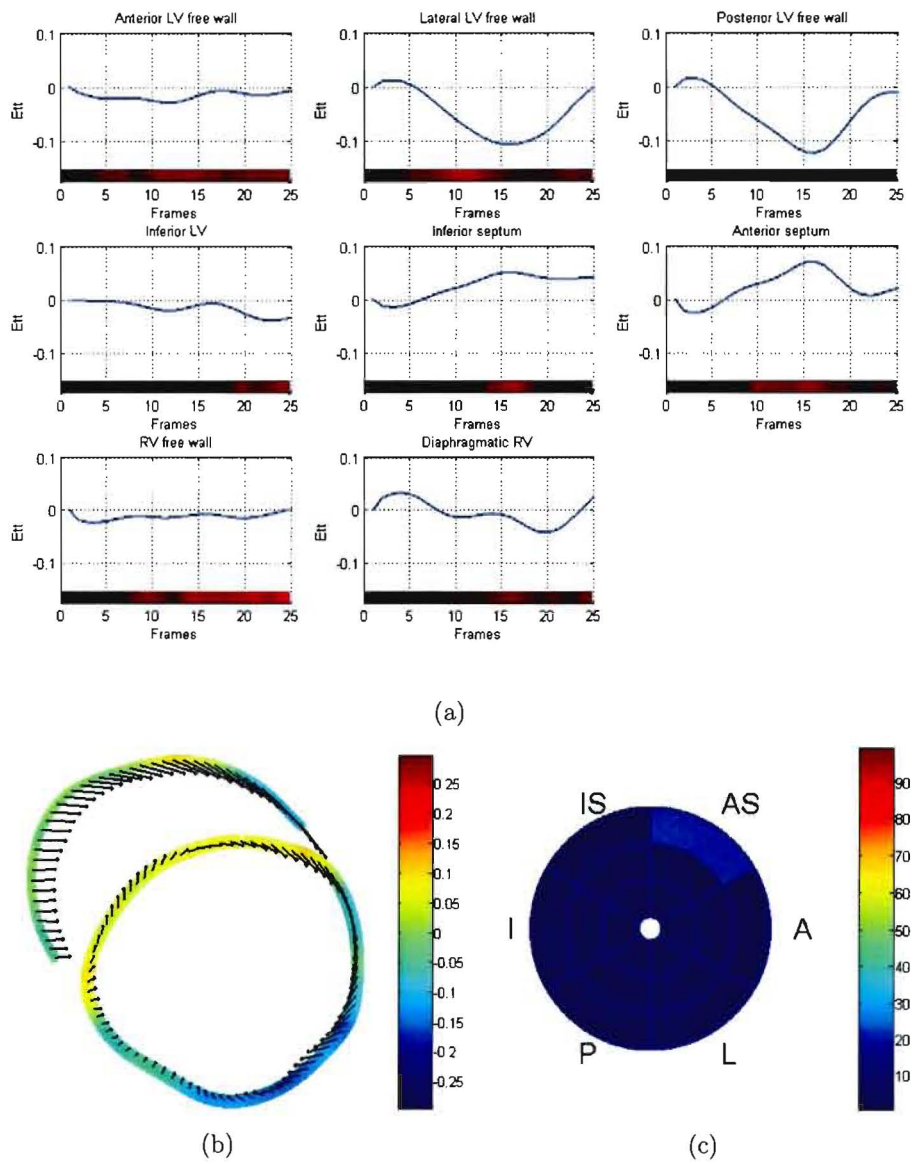

Figure H.3: Patient 1: LBBB, QRS=151 ms. (a) Ett versus time for the segments shown in Figure 8.4, (b) End-systolic map of midwall Ett and 2–D trajectory position, and (c) percent Gd delayed-enhancement of basal, mid, and apical short-axis slices (image reproduced with permission from Prof. F.H. Epstein). A - anterior LV; L - lateral LV; P - posterior LV; I - inferior LV; IS - inferior septum; AS - anterior septum.

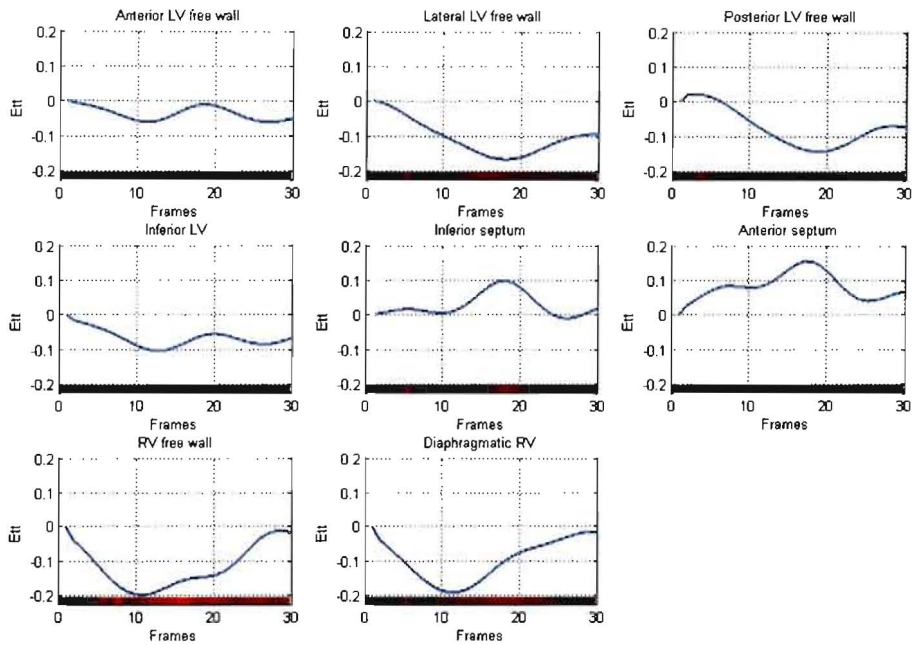

(a)

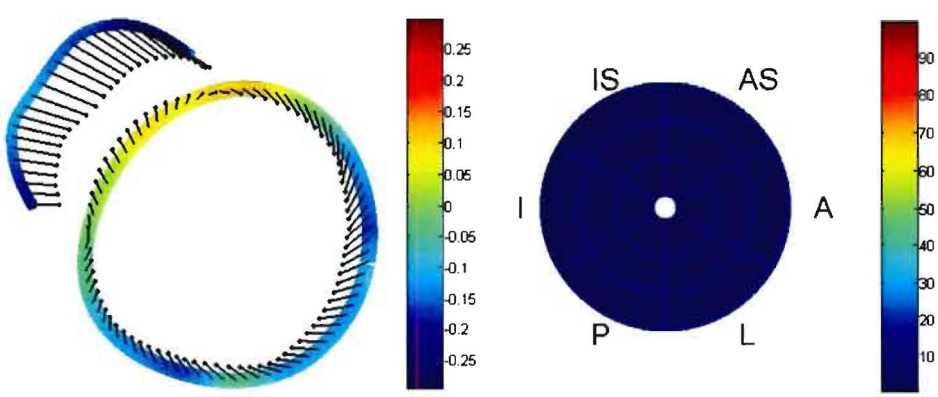

(b) (c)

Figure H.4: Patient 2: LBBB, QRS=133 ms (a) Ett versus time for the segments shown in Figure 8.4, (b) End-systolic map of midwall Ett and 2–D trajectory position, and (c) percent Gd delayed-enhancement of basal, mid, and apical short-axis LV slices (image reproduced with permission from Prof. F.H. Epstein). A - anterior LV; L - lateral LV; P - posterior LV; I - inferior LV; IS - inferior septum; AS - anterior septum.

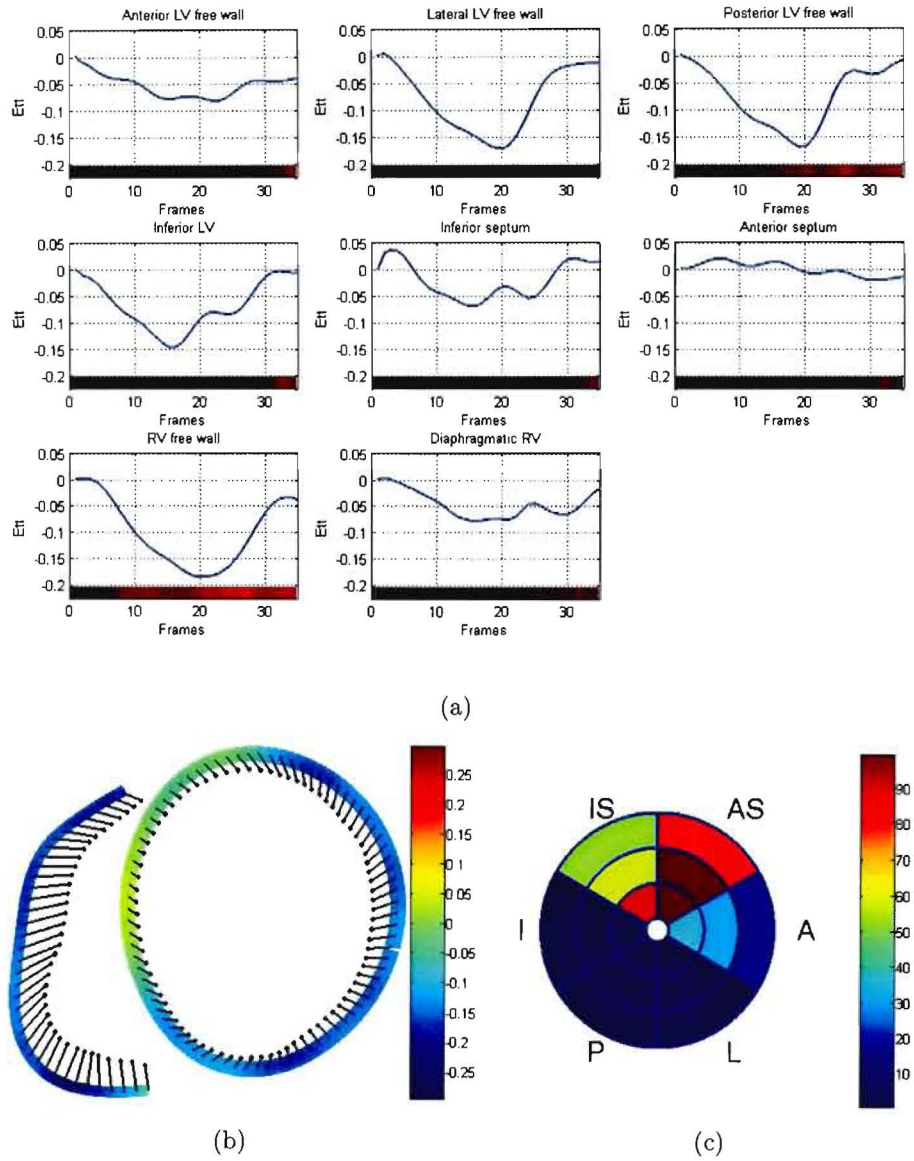

Figure H.5: Patient 3: RBBB, QRS=145 ms (a) Ett versus time for the segments shown in Figure 8.4, (b) End-systolic map of midwall Ett and 2–D trajectory position, and (c) percent Gd delayed-enhancement of basal, mid, and apical short-axis LV slices (image reproduced with permission from Prof. F.H. Epstein). A - anterior LV; L - lateral LV; P - posterior LV; I - inferior LV; IS - inferior septum; AS - anterior septum.

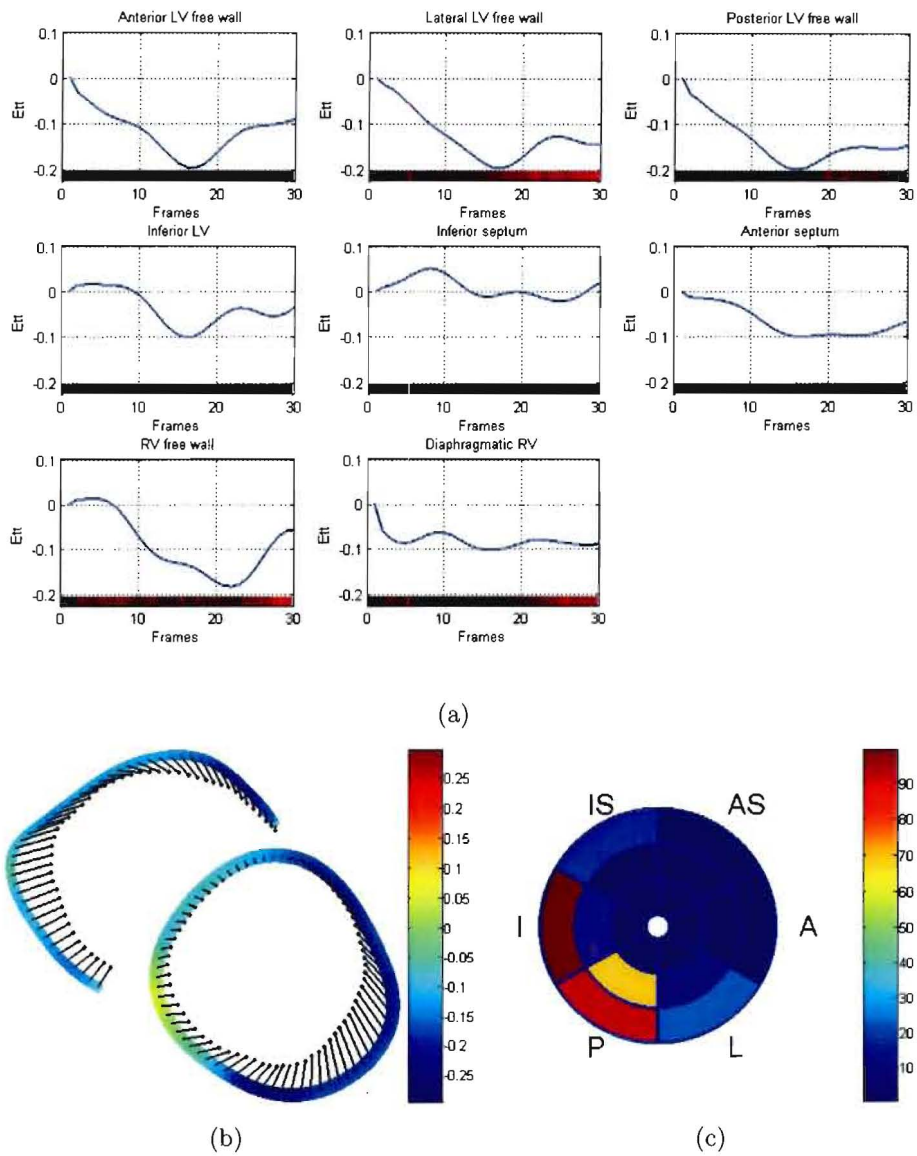

Figure H.6: Patient 4: RBBB and LAFB, QRS=147 ms (a) Ett versus time for the segments shown in Figure 8.4, (b) End-systolic map of midwall Ett and 2–D trajectory position, and (c) percent Gd delayed-enhancement of basal, mid, and apical short-axis LV slices (image reproduced with permission from Prof. F.H. Epstein; apex data was absent). A - anterior LV; L - lateral LV; P - posterior LV; I - inferior LV; IS - inferior septum; AS - anterior septum.

References

Abd-Elmoniem, K. Z., Osman, N. F., Prince, J. L., Stuber, M., 2005. zHARP: 3-D cardiac motion tracking from short-axis acquisitions. In: *Proceedings of the 13th ISMRM, Miami, USA.*

Abd-Elmoniem, K. Z., Stuber, M., Prince, J. L., 2006. True myocardial planar strain: Resolving through-plane rotation ambiguity in tagged MRI using ZHARP. *Journal of Cardiovascular Magnetic Resonance* 8 (1), 21–23.

AHA, 2003. American Heart Association webpage. Last accessed 20 September 2003. URL `http:/www.americanheart.org`

Aletras, A. H., Arai, A. E., 1999. Meta-DENSE complex acquisition for reduced intravoxel dephasing. *Journal of Magnetic Resonance* 196, 246–249.

Aletras, A. H., Balaban, R., Wen, H., 1999a. High-resolution strain analysis of the human heart with fast-DENSE. *Journal of Magnetic Resonance* 140, 41–57.

Aletras, A. H., Ding, S., Balaban, R. S., Wen, H., 1999b. DENSE: Displacement encoding with stimulated echoes in cardiac functional MRI. *Journal of Magnetic Resonance* 137, 247–252.

Allouche, C. F., Makram-Ebeid, S., Stuber, M., Ayache, N., Delingette, H., 2001. New methods and algorithms for the accurate, real-time motion analysis of the left ventricle with MRI-tagging. *International Congress Series* 1230 (1), 954–960.

Altar, E., McVeigh, E. R., 1994. Optimization of tag thickness for measuring position with magnetic resonance imaging. *IEEE Transactions on Medical Imaging* 13 (1), 152–160.

Amartur, S. C., Vesselle, H. J., 1993. A new approach to study cardiac motion: The optical flow of cine MR images. *Magnetic Resonance in Medicine* 29 (1), 59–67.

Axel, L., Bloomgarden, D., Chang, C. N., Kraitchman, D., Young, A. A., 1993. SPAMMVU: a program for the analysis of dynamic tagged MRI. *Society of Magnetic Resonance in Medicine, 12th Annual Scientific Meeting, New York, Book of Abstracts* 2, 724.

Axel, L., Dougherty, L., 1989a. Heart wall motion: Improved method of spatial modulation of magnetization for MR imaging. *Radiology* 172, 349–350.

Axel, L., Dougherty, L., 1989b. MR imaging of motion with spatial modulation of magnetization. *Radiology* 171, 841–845.

Axel, L., Gonçalves, R. C., Bloomgarden, D., 1992. Regional heart wall motion: Two-dimensional analysis and functional imaging with MR imaging. *Radiology* 183, 745–750.

Axel, L., Montillo, A., Kim, D., 2005. Tagged magnetic resonance imaging of the heart: a survery. *Medical Image Analysis* 9, 376–393.

Bennett, E., Spottiswoode, B., Lorenz, C., Wen, H., 2006. Optimal combination of phase cycling and gradient spoiling in DENSE displacement. In: *Proceedings of the 14^{th} ISMRM, Seattle, USA.*

BHF, 2003. British Heart Foundation statistics website. Last accessed 12 May 2003. URL `http://www.heartstats.org`

Bogaert, J., Rademakers, F. E., 2001. Regional nonuniformity of normal adult human left ventricle. *American Journal of Physiology, Heart and Circulatory Physiology* 280, H610–H620.

Bristow, M. R., Saxon, L. A., Boehmer, J., Krueger, S., Kass, D. A., Marco, T. D., Carson, P., DiCarlo, L., DeMets, D., White, B. G., DeVries, D. W., Feldman, A. M., 2004. Cardiac-resynchronization therapy with or without an implantable defibrillator in advanced chronic heart failure. *New England Journal of Medicine* 350 (21), 2140–2150.

Bryant, D. J., Payne, J. A., Firmin, D. N., Longmore, D. B., 1984. Measurement of flow with NMR imaging using a gradient pulse and phase difference technique. *Journal of computer assisted tomography* 8, 588–593.

Buckland, J. R., Huntley, J. M., Turner, S. R. E., 1995. Unwrapping noisy phase maps by use of a minimum-cost-matching algorithm. *Applied Optics* 34 (23), 5100–5108.

Callot, V., Bennett, E., Decking, U. K. M., Balaban, R. S., Wen, H., 2003. In vivo study of microcirculation in canine myocardium using the IVIM method. *Magnetic Resonance in Medicine* 50, 531–540.

Chandrashekara, R., Mohiaddin, R. H., Rueckert, D., 2002. Analysis of myocardial motion in tagged MR images using nonrigid image registration. *Proceedings of Medical Image Understanding and Analysis (MIUA), Portsmouth, United Kingdom* , 1–4.

Constable, R. T., Rath, K. M., Sinusas, A. J., Gore, J. C., 1994. Development and evaluation of tracking algorithms for cardiac wall motion analysis using phase velocity MR imaging. *Magnetic Resonance in Medicine* 32, 33–42.

Curry, C. W., Nelson, G. S., Wyman, B. T., Declerck, J., Talbot, M., Berger, R. D., McVeigh, E. R., Kass, D. A., 2002. Mechanical dyssynchrony in dilated cardiomyopathy with intraventricular conduction delay as depicted by 3–D tagged magnetic resonance imaging. *Circulation* 101 (1), 2.

Cusack, R., Papadakis, N., 2002. New robust 3-D phase unwrapping algorithms: application to magnetic field mapping and undistorted echoplanar images. *NeuroImage* 16, 754–764.

Delfino, J., Bhasin, M., Leon, A., Eisner, R., Oshinski, J., 2006. Longitudinal, radial, and circumferential myocardial velocities in patients with ventricular dyssynchrony. In: *Proceedings of the 14th ISMRM, Seattle, USA.*

Deng, X., Denney, Jr., T. S., 2004. Three-dimensional myocardial strain reconstruction from tagged MRI using a cylindrical B-spline model. *IEEE Transactions on Medical Imaging* 23 (7), 861–864.

D'hooge, J., Heimdal, A., Jamal, F., Sutherland, G., 2000. Regional strain and strain rate measurements by cardiac ultrasound: Principles, implementation and limitations. *European Journal of Echocardiography* 1, 154–170.

Dinsmore, R. E., Wismer, G. L., Levine, R. A., Okada, R. D., Brady, T. J., 1984. Magnetic resonance imaging of the heart: Positioning and gradient angle selection for optimal imaging planes. *American Journal of Roentgenology* 143, 1135–1142.

Dougherty, L., Asmuth, J. C., Blom, A. S., Axel, L., Kumar, R., 1999. Validation of an optical flow method for tag displacement estimation. *IEEE Transactions on Medical Imaging* 18 (4), 359–363.

Drangova, M., Zhu, Y., Bowman, B., Pelc, N. J., 1998. In vitro verification of myocardial motion tracking from phase-contrast velocity data. *Magnetic Resonance Imaging* 16 (8), 863–870.

Epstein, F. H., Gilson, W. D., 2004. Displacement-encoded cardiac MRI using cosine and sine modulation to eliminate (CANSEL) artifact-generating echoes. *Magnetic Resonance in Medicine* 52 (4), 774–781.

Fayad, Z. A., Ferrari, V. A., Kraitchman, D. A., Young, A. A., Palevsky, H. I., Bloomgarden, D. C., Axel, L., 1998. Right ventricular function using MR tagging: normals versus chronic pulmonary hypertension. *Magnetic Resonance in Medicine* 39, 116–123.

Fischer, S. E., McKinnon, G. C., Maier, S. E., Boesiger, P., 1993. Improved myocardial tagging contrast. *Magnetic Resonance in Medicine* 30, 191–200.

Fischer, S. E., McKinnon, G. C., Scheidegger, M. C., Prins, W., Meier, D., Boesiger, P., 1994. True myocardial motion tracking. *Magnetic Resonance in Medicine* 31, 401–413.

Fischer, S. E., Stuber, M., Scheidegger, M. B., Boesiger, P., 1995. Limitations of stimulated echo acquisition mode (STEAM) techniques in cardiac applications. *Magnetic Resonance in Medicine* 34, 80–91.

Foo, T. K. F., Bernstein, M. A., Aisen, A. M., Hernandez, R. J., Collick, B. D., Bernstein, T., 1995. Improved ejection fraction and flow velocity estimates with use of view sharing and uniform repetition time excitation with fast cardiac techniques. *Radiology* 195, 471–478.

Frangi, A. F., Niessen, W. J., Viergever, M. A., 2001. Three-dimensional modeling for functional analysis of cardiac images: A review. *IEEE Transactions on Medical Imaging* 20 (1), 2–25.

Garot, J., Bluemke, D. A., Osman, N. F., Rochitte, C. E., McVeigh, E. R., Zerhouni, E. A., Prince, J. L., Lima, J. A. C., 2000. Fast determination of regional myocardial strain fields from tagged cardiac images using harmonic phase MRI. *Circulation* 101, 981–988.

Ghiglia, D. C., Pritt, M. D., 1998. *Two-dimensional phase unwrapping: theory, algorithms and software*. Wiley-Interscience.

Gilson, W. D., 2004. *High-Resolution Myocardial Tissue Tracking in Post-Infarct Knockout Mice with Displacement-Encoded MRI*. Ph.D. thesis, University of Virginia.

Glover, G. H., Schneider, E., 1991. Three-point dixon technique for water/fat decomposition with B0 inhomogeneity correction. *Magnetic Resonance in Medicine* 18 (2), 371–383.

Gonzalez, R. C., Woods, R. E., 1992. *Digital Image Processing. Third Edition.*. Addison-Wesley, New York.

Groot-Koerkamp, M. A. M., Snoep, G., Muitjens, A. M. M., Kemerink, G., 1999. Improving contrast and tracking of tags in cardiac magnetic resonance images. *Magnetic Resonance in Medicine* 41, 973–982.

Guttman, M. A., Prince, J. L., McVeigh, E. R., 1994. Tag and contour detection in tagged MR images of the left ventricle. *IEEE Transactions on Medical Imaging* 13 (1), 74–88.

Guyton, A. C., 1981. *Textbook of Medical Physiology. Sixth Edition.*. W.B. Saunders Company, Philadelphia, London, Toronto.

Haacke, W. M., Brown, R. W., Thompson, M. R., Venkatesan, R., 1999. *Magnetic Resonance Imaging: Physical Principles and Sequence Design*. John Wiley and Sons, New York, Chichester, Weinheim, Brisbane, Singapore, Toronto.

Haber, I., Metaxas, D. N., Axel, L., 2000. Three-dimensional motion reconstruction and analyis of the right ventricle using tagged MRI. *Medical Image Analysis* 4, 335–355.

Haber, I., Metaxas, D. N., Geva, T., Axel, L., 2005. Three dimensional systolic kinematics of the right ventricle. *American Journal of Physiology - Heart and Circulatory Physiology* 289, H1826–H1833.

Heimdal, A., Stoylen, A., Torp, H., Skjaerpe, T., 1998. Real-time strain rate imaging of the left ventricle by ultrasound. *Journal of the American Society of Echocardiography* 11, 1013–1019.

Helm, R. H., Leclercq, C., Faris, O. P., Ozturk, C., McVeigh, E., Lardo, A. C., Kass, D. A., 2005. Cardiac dyssynchrony analysis using circumferential versus longitudinal

strain: Implications for assessing cardiac resynchronization. *Circulation* 111, 2760–2767.

Hess, A., 2006. *Calculating 3D intramyocardial strain tensors in a single slice of myocardium using MRI*. Master's thesis, University of Cape Town.

Hess, A. T., Zhong, X., Spottiswoode, B. S., Epstein, F. H., Meintjes, E. M., 2006. Combining 2–D DENSE and SENC to calculate the 3–D strain tensor. In: *Proceedings of the 14^{th} ISMRM, Seattle, USA.*.

HFSA, 2003. The Heart Foundation of South Africa. Last accessed 12 May 2003. URL `http://www.heartfoundation.co.za/`

Hu, Z., Metaxas, D., Axel, L., 2002. In-vivo strain and stress estimation of the left ventricle from MRI images. *International Society and Conference Series on Medical Image Computing and Computer-Assisted Intervention (MICCAI), LNCS Vol 2488* , 706–713.

Huntley, J. M., 2001. Three-dimensional noise-immune phase unwrapping algorithm. *Applied Optics* 40 (23), 3901–3908.

Ingels, N. B., Daughters, G. T., Stinson, E. B., Alderman, E. L., 1975. Measurements of midwall dynamics in intact man by radiography of surgically implanted markers. *Circulation* 52, 859–867.

Itoh, K., 1980. Analysis of the phase unwrapping problem. *Applied Optics* 21 (14), 2470.

Jenkinson, M., 2003. Fast, automated N-dimensional phase-unwrapping algorithm. *Magnetic Resonance in Medicine* 48, 193–197.

Jezzard, P., Balaban, R. S., 1995. Correction for geometric distortion in echo planar images from B0 field variations. *Magnetic Resonance in Medicine* 42, 65–73.

Kass, D. A., 2002. Ventricular dyssynchrony and mechanisms of resynchronization therapy. *European Heart Journal Supplements* 4 (Supplement D), D23–D30.

Kass, M., Witkin, A., Terzopoulos, D., 1988. Snakes: Active contour models. *International Journal of Computer Vision* 1, 321–331.

Kayser, H. W. M., van der Geest, R. J., van der Wall, E. E., Duchateau, C., de Roos, A., 2000. Right ventricular function in patients after acute myocardial infarction assessed with phase contrast MR velocity mapping encoded in three directions. *Magnetic Resonance in Medicine* 11 (5), 471–475.

Kerckhoffs, R. C. P., Faris, O. P., Bovendeerd, P. H. M., Prinzen, F. W., Smits, K., McVeigh, E. R., Arts, T., 2005. Electromechanics of paced left ventricle simulated by straightforward mathematical model: comparison with experiments. *American Journal of Physiology, Heart and Circulatory Physiology* 289, H1889–H1897.

Kerwin, W. S., Osman, N. F., Prince, J. L., 1999. *Handbook of Medical Image Processing.* Academic Press, New York. Editor: I. Bankman.

Kim, D., Bove, C. M., Kramer, C. M., Epstein, F. H., 2003. Importance of k-space trajectory in echo-planar myocardial tagging at rest and during dobutamine stress. *Magnetic Resonance in Medicine* 50, 813–820.

Kim, D., Gilson, W. D., Kramer, C. M., Epstein, F. H., 2004. Myocardial tissue tracking with two-dimensional cine displacement-encoded MR imaging: Development and initial evaluation. *Radiology* 230, 862–871.

Klein, S. S., Graham, T. P., Lorenz, C. H., 1998. Noninvasive delineation of normal right ventricular contractile motion with magnetic resonance imaging myocardial tagging. *Annals of Biomedical Engineering* 25 (5), 756–763.

Kraitchman, D. L., Young, A. A., Chang, C., Axel, L., 1995. Semi-automated tracking of myocardial motion in MR tagged images. *IEEE Transactions on Medical Imaging* 14 (3), 422–433.

Kuijer, J., Hofman, M. B., Zwanenburg, J. J., Marcus, J. T., Rossum, A. C., Heethaar, R. M., 2006. DENSE and HARP: two views on the same technique of phase-based strain imaging. In: *Proceedings of the 14^{th} ISMRM, Seattle, USA.*

Kuijer, J. P. A., Marcus, J. T., Götte, M. J. W., van Rossum, A. C., Heethar, R. M., 1999. Simultaneous MRI tagging and through-plane velocity quantification: A three-dimensional myocardial motion tracking algorithm. *Journal of Magnetic Resonance Imaging* 9, 409–419.

Kuijer, J. P. A., Marcus, J. T., Götte, M. J. W., van Rossum, A. C., Heethar, R. M., 2000. Three-dimensional myocardial strain analysis based on short- and long-axis magnetic resonance tagged images using a 1–D displacement field. *Magnetic Resonance Imaging* 18, 553–564.

Kumar, S., Goldgof, D., 1994. Automatic tracking of SPAMM grid and the estimation of deformation parameters from cardiac MR images. *IEEE Transactions on Medical Imaging* 13 (1), 122–132.

Lardo, A. C., Abraham, T. P., Kass, D. A., 2005. Magnetic resonance imaging assessment of ventricular dyssynchrony: current and emerging concepts. *Journal of the American College of Cardiology* 46 (12), 2223–2228.

Leclercq, C., Faris, O., Tunin, R., Johnson, J., Kato, R., Evans, F., Spinelli, J., Halperin, H., McVeigh, E., Kass, D. A., 2002. Systolic improvement and mechanical resynchronization does not require electrical synchrony in the dilated failing heart with left bundle-branch block. *Circulation* 106, 1760–1763.

Li, Y., Garson, C. D., Epstein, F. H., French, B. A., Hossack, J. A., 2006. High resolution 2–D quantification of myocardial motion abnormalities in mice using high resolution ultrasound with MRI validation. In: *Proceedings of ISBI*.

Lingamneni, A., Hardy, P. A., Powell, K. A., Pelc, N. J., White, R. D., 1995. Validation of cine phase-contrast MR imaging for motion analysis. *Journal of Magnetic Resonance Imaging* 5, 331–338.

Lorenz, C. H., Pastorek, J. S., Bundy, J. M., 2000. Delineation of normal human left ventricular twist throughout systole by tagged cine magnetic resonance imaging. *Journal of Cardiovascular Magnetic Resonance* 2 (2), 97–108.

Mase, G. T., Mase, G. E., 1999. *Continuum Mechanics for Engineers*. CRC Press, Boca Raton, London, New York, Washington D.C. Second edition.

Masood, S., 2003. *Analysis of myocardial contractility with magnetic resonance*. Ph.D. thesis, University of London.

Masood, S., Yang, G., Pennell, D. J., Firmin, D. N., 2000. Investigating intrinsic myocardial mechanics: The role of MR tagging, velocity phase magging, and diffusion imaging. *Journal of Magnetic Resonance Imaging* 12, 873–883.

McVeigh, E. R., 1996. MRI of myocardial function: Motion tracking techniques. *Journal of Magnetic Resonance Imaging* 14 (2), 137–150.

McVeigh, E. R., Prinzen, F. W., Wyman, B. T., Tsitlik, J. E., Halperin, H. R., Hunter, W. C., 1998. Imaging asynchronous mechanical activation of the paced heart with tagged MRI. *Magnetic Resonance in Medicine* 39, 507–513.

McVeigh, E. R., Zerhouni, E. A., 1991. Noninvasive measurement of transmural gradients in myocardial strain with MR imaging. *Radiology* 180, 677–683.

Mirsky, I., Ghista, D. N., Sandler, H., 1974. *Cardiac Mechanics: Physiological, Clinical, and Mathematical Considerations*. John Wiley and Sons Inc., New York, London, Sydney, Toronto.

Moaveni, S., 1999. *Finite element analysis; theory and application with ANSYS*. Prentice Hall, Upper Saddle River, New Jersey, USA.

Montillo, A., Metaxas, D., Axel, L., 2002. Automated segmentation of the left and right ventricles in 4–D cardiac SPAMM images. *International Society and Conference Series on Medical Image Computing and Computer-Assisted Intervention (MICCAI), LNCS Vol 2488* , 620–633.

Montillo, A., Metaxas, D., Axel, L., 2003. Automated deformable model-based segmentation of the left and right ventricles in tagged cardiac MRI. *International Society and Conference Series on Medical Image Computing and Computer-Assisted Intervention (MICCAI), LNCS Vol 2878* , 507–515.

Moore, C. C., Lugo-Olivieri, C. H., McVeigh, E. R., Zerhouni, E. A., 2000. Three-dimensional systolic patterns in the normal human left ventricle: characterization with tagged MR imaging. *Radiology* 214, 453–466.

Moore, C. C., Reeder, S. B., McVeigh, E. R., 1994. Tagged MR imaging in a deformable phantom: photographic validation. *Radiology* 190, 765–769.

Mosher, T. J., Smith, M. B., 1990. A DANTE tagging sequence for the evaluation of translational sample motion. *Magnetic Resonance in Medicine* 15, 334–339.

Nayler, G. L., Firmin, D. N., Longmore, D. B., 1986. Blood flow imaging by cine magnetic resonance. *Journal of Computer Assisted Tomography* 10, 715–722.

Netter, F. H., 1998. *Interactive Atlas of Human Anatomy*. ICON Learning Systems. Version 2.0.

O'Dell, W. G., Moore, C. C., Hunter, W. C., Zerhouni, E. A., McVeigh, E. R., 1995. Three-dimensional myocardial deformations: Calculation with displacement field fitting to tagged MR images. *Radiology* 195, 829–835.

Osman, N. F., Kerwin, W. S., McVeigh, E. R., Prince, J. L., 1999. Cardiac motion tracking using CINE harmonic phase (HARP) magnetic resonance imaging. *Magnetic Resonance in Medicine* 42, 1048–1060.

Osman, N. F., Prince, J. L., 1998. Direct calculation of 2–D components of myocardial strain using sinusoidal MR tagging. In: *Proceedings of SPIE Medical Imaging: Image Processing Conference, San Diego, California.*

Osman, N. F., Prince, J. L., 2000a. Imaging heart motion using harmonic phase MRI. *IEEE Transactions on Medical Imaging* 19 (3), 186–202.

Osman, N. F., Prince, J. L., 2000b. On the design of the bandpass filters in harmonic phase MRI. In: *Proceedings of the 2000 International Conference on Image Processing (ICIP 2000), Vancouver, Canada.*

Osman, N. F., Sampath, S., Atalar, E., Prince, J. L., 2001. Imaging longitudinal cardiac strain on short-axis images using strain-encoded MRI. *Magnetic Resonance in Medicine* 46, 324–334.

Pai, V. M., Wen, H., 2003. Developing a navigator-guided approach to non-breathold multiphase cardiac strain maps using DENSE. *Journal of Cardiovascular Magnetic Resonance* 5 (1), 239–242.

Pan, L., Lima, J. A. C., Osman, N. F., 2003. Fast tracking of cardiac motion using 3D-HARP. *Proceedings of Information Processing and Medical Imaging (IPMI), LNCS 2732* , 611–622.

Pan, L., Sampath, S., Prince, J. L., Stuber, M., Osman, N. F., 2005. Fast tracking of cardiac material points from SF-CSPAMM images using 3D SF-HARP. In: *Proceedings of the 13th ISMRM, Miami, USA.*

Park, J., Metaxas, D., Young, A. A., Axel, L., 1996. Deformable models with parameter functions for cardiac motion analysis from tagged MRI data. *IEEE Transactions on Medical Imaging* 15 (3), 278–289.

Parthasarathy, V., NessAiver, M., Prince, J. L., 2004. Dynamic range of harmonic phase magnetic resonance imaging (HARP-MRI). In: *IEEE International Symposium of biomedical imaging, Arlington.*.

Parthasarathy, V., Prince, J. L., 2004. Strain resolution from HARP-MRI. In: *Proceedings of the 12th ISMRM, Kyoto, Japan.*.

Patel, R. A., Zhong, X., Spottiswoode, B. S., Helm, P. A., Ferguson, J. D., Hummel, J. P., Kramer, C. M., Epstein, F. H., 2006. Cine DENSE MRI of left ventricular dyssynchrony: Development and initial clinical experience. In: *Proceedings of the 14th ISMRM, Seattle, USA.*.

Pelc, N. J., Drangova, M., Pelc, L. R., 1995. Tracking of cyclic motion with phase-contrast MR velocity data. *Journal of Magnetic Resonance Imaging* 5, 339–345.

Penicka, M., Bartunek, J., Bruyne, B. D., Vanderheyden, M., Goethals, M., Zutter, M. D., Brugada, P., Geelen, P., 2004. Improvement of left ventricular function after cardiac resynchronization therapy is predicted by tissue doppler imaging echocardiography. *Circulation* 109, 978–983.

Petersen, S. E., Jung, B. A., Wiesmann, F., Selvanayagam, J. B., Francis, J. M., Hennig, J., Neubauer, S., Robson, M. D., 2006. Myocardial tissue phase mapping with cine phase-contrast MR imaging: Regional wall motion analysis in healthy volunteers. *Radiology* 238 (3), 816–826.

Pettigrew, R., 1992. *Magnetic Resonance Imaging. Second edition.*. Mosby-Year book Inc, St. Louis, Baltimore, Boston, Chicago, London, Philadelphia, Sydney, Toronto. Editors: Stark D. D. and Bradley Jr W. G.

Prince, J. L., McVeigh, E. R., 1992. Motion estimation from tagged MR image sequences. *IEEE Transactions on Medical Imaging* 1 (2), 238–249.

Prinzen, F. W., Hunter, W. C., Wyman, B. T., McVeigh, E. R., 1999. Mapping of regional myocardial strain and work during ventricular pacing: experimental study using magnetic resonance imaging tagging. *Journal of the American College of Cardiology* 33 (6), 1735–1742.

Reichek, N., 1999. MRI myocardial tagging. *Journal of Magnetic Resonance Imaging* 10, 609–616.

Rutz, A. K., Ryf, S., Kozerke, S., Boesiger, P., 2006. Improved spatial resolution and reduction of artifacts in strain-maps with phase-unwrap harp. In: *Proceedings of the 14^{th} ISMRM, Seattle, USA.*

Salfity, M. F., Huntley, J. M., Graves, M. J., Marklund, O., Cusack, R., Beauregard, D. A., 2005. Extending the dynamic range of phase contrast magnetic resonance velocity imaging using advanced higher-dimensional phase unwrapping algorithms. *Journal of the Royal Society Interface* 5, In press.

Sampath, S., Derbyshire, J. A., Osman, N. F., Atalar, E., Prince, J. L., 2001. Real-time imaging of cardiac strain using ultra FastHARP sequence. In: *Proceedings of the 9^{th} ISMRM, Glasgow, Scotland.*

Sinele, K., Bennett, E., Wen, H., 2006. Automatic masking and phase unwrapping of DENSE myocardial tissue tracking images in human. *Journal of Cardiovascular Magnetic Resonance* 8 (1), 146–147.

Spragg, D. D., Akar, F. G., Helm, R. H., Tunin, R. S., Tomaselli, G. F., Kass, D. A., 2005. Abnormal conduction and repolarization in late-activated myocardium of dyssynchronously contracting hearts. *Cardiovascular Research* 67, 77–86.

Stuber, M., Nagel, E., Fischer, S. E., Spiegel, M. A., Scheidegger, M. B., Boesiger, P., 1998. Quantification of local heartwall motion by magnetic resonance myocardial tagging. *Computerized Medical Imaging and Graphics* 22, 217–228.

Stuber, M., Spiegel, M. A., Fischer, S. E., Scheidegger, M. B., Danias, P. G., Pedersen, E. M., Boesiger, P., 1999. Single breath-hold slice-followed CSPAMM myocardial tagging. *Magnetic Resonance Materials in Physics, Biology and Medicine* 9, 85–91.

Sutherland, G. R., Stewart, M. J., Groundstroem, K. W., Moran, C. M., Fleming, A., Guell-Peris, F., Riemersma, R. A., Fenn, L. N., Fox, K. A., McDicken, W., 1994. Color doppler myocardial imaging: a new technique for the assessment of myocardial function. *Journal of the American Society of Echocardiography* 7, 441–458.

Teceläo, S. R., Zwanenburg, J. J., Kuijer, J. P., Marcus, J. T., 2006. A quantitative comparison between 2-dimensional and 3-dimensional circumferential strain in the human heart. In: *Proceedings of the 14^{th} ISMRM, Seattle, USA.*

Tsao, J., Laurent, D., 2005. N-SPAMM for efficient displacement-encoded acquisition in myocardial tagging. In: *Proceedings of the 13^{th} ISMRM, Miami, USA.*

Tustison, N. J., Davila-Roman, V. G., Amini, A. A., 2003. Myocardial kinematics from tagged MRI based on a 4-D B-spline model. *IEEE Transactions on Biomedical Engineering* 50 (8), 1038–1041.

Udupa, J. K., LeBlanc, V. R., Schmidt, H., Imielinska, C., Saha, P. K., Grevera, G. J., Zhuge, Y., Currie, L. M., Molholt, P., Jin, Y., 2002. A methodology for evaluating image segmentation algorithms. *Proceedings of SPIE* 4684, 266–277.

Urayama, S., Matsuda, T., Sugimoto, N., Mizuta, S., Yamada, N., Uyama, C., 2000. Detailed motion analysis of the left ventricular myocardium using an MR tagging method with a dense grid. *Magnetic Resonance in Medicine* 44, 73–82.

van Dijk, P., 1984. Direct cardiac NMR imaging of heart wall and blood flow velocity. *Journal of computer assisted tomography* 8, 588–593.

Villarreal, F. L., Waldman, L. K., Lew, W. Y. W., 1988. A technique for measuring regional two-dimensional finite strains in canine left ventricle. *Circulation Research* 62, 711–721.

Wen, H., 2005. Anatomically guided shaped smoothing in myocardial strain data analysis. In: *22nd Annual Scientific Meeting of the European Society for Magnetic Resonance in Medicine and Biology (ESMRMB).*

Wen, H., Bennett, E., Epstein, N., Plehn, J., 2006a. Cardiac contractile dyssynchrony mapping using DENSE displacement imaging: preliminary tests in normal subjects. *Journal of Cardiovascular Magnetic Resonance* 8 (1), 96–97.

Wen, H., Rodriguez, I., Bennett, E., Vignaud, A., 2006b. Optimization of DENSE sequence for imaging regional strain distribution in the carotid artery wall and preliminary tests in humans. *Journal of Cardiovascular Magnetic Resonance* 8 (1), 153–154.

Wong, A. L. N., Liu, H., Shi, P., 2002. Segmentation of myocardium using velocity field constrained front propagation. In: *Proceedings of the 6th IEEE Workshop on Applications of Computer Vision (WACV)*, pp. 84–89, pp. 84–89.

Xu, Y., Haacke, E. M., 2001. Partial fourier imaging in multi-dimensions: A means to save a full factor of two in time. *Journal of Magnetic Resonance Imaging* 14, 628–635.

Yang, G. Z., Burger, P., Kilner, P. J., Karwatowski, S. P., Firmin, D. N., 1996. Dynamic range extension of cine velocity measurements using motion-registered spatiotemporal phase unwrapping. *Journal of Magnetic Resonance Imaging* 6, 495–502.

Young, A. A., 1991. *Theory of heart: biomechanics, biophysics, and nonlinear dynamics of cardiac function*. Springer-Verlag, New York. Editors: Glass L., Hunter P., McCulloch A.

Young, A. A., Axel, L., 1992. Three-dimensional motion and deformation of the heart wall: Estimation with spatial modulation of magnetization – a model-based approach. *Radiology* 185, 241–247.

Young, A. A., Fayad, Z. A., Axel, L., 1996. Right ventricular midwall surface motion and deformation using magnetic resonance tagging. *American Journal of Physiology* 271 (40), H2677–H2688.

Young, A. A., Hunter, P. J., Smaill, B. H., 1992. Estimation of epicardial strain using the motions of coronary bifurcations in biplane cineangiography. *IEEE Transactions on Biomedical Engineering* 39, 526–531.

Young, A. A., Kraitchman, D. L., Dougherty, L., Axel, L., 1995. Tracking and finite element analysis of stripe deformation in magnetic resonance tagging. *IEEE Transactions on Medical Imaging* 14 (3), 413–421.

Zerhouni, E. A., Parish, D. M., Rogers, W. J., Yang, A., Shapiro, E. P., 1988. Human HEART: Tagging with MR imaging–a method for noninvasive assessment of myocardial motion. *Radiology* 169, 59–63.

Zhong, X., Blemker, S. S., Spottiswoode, B. S., Helm, P. A., Hess, A. T., Epstein, F. H., 2006a. Application of Cine DENSE MRI to studying skeletal muscle mechanics during joint motion. In: *Proceedings of the 14th ISMRM, Seattle, USA*.

Zhong, X., Cowart, E. A., Spottiswoode, B. S., Epstein, F. H., 2005. Selective suppression of artifact-generating echoes in cine DENSE using through-plane gradients. In: *Proceedings of the 13th ISMRM, Miami, USA*.

Zhong, X., Spottiswoode, B. S., Epstein, F. H., 2006b. Improved efficiency in cine DENSE using low resolution phase reference images. In: *Proceedings of the 14th ISMRM, Seattle, USA*.

Zhu, Y., Drangova, M., Pelc, N. J., 1996. Fourier tracking of myocardial motion using cine-PC data. *Magnetic Resonance in Medicine* 35, 471–480.

Zhu, Y., Drangova, M., Pelc, N. J., 1997. Estimation of deformation gradient from cine-PC velocity data. *IEEE Transactions on Medical Imaging* 16, 840–851.

Zhu, Y., Pelc, N. J., 1999. Three dimensional motion tracking with volumetric phase contrast MR velocity imaging. *Journal of Magnetic Resonance Imaging* 9, 111–118.

Zhuge, Y., Udupa, J. K., Liu, J., Saha, P. K., Iwanaga, T., 2002. A scale-based method for correcting background intensity variation in acquired images. *Proceedings of SPIE, Vol 4684* , 1103–1111.

Zwanenburg, J. J. M., Götte, M. J. W., Kuijer, J. P. A., Heethaar, R. M., van Rossum, A. C., Marcus, J. T., 2004. Timing of cardiac contraction in humans mapped by high-temporal-resolution MRI tagging: early onset and late peak of shertening in lateral wall. *American Journal of Physiology, Heart and Circulatory Physiology* 286, H1872–H1880.

Zwanenburg, J. J. M., Götte, M. J. W., Marcus, J. T., Kuijer, J. P. A., Knaapen, P., Heethaar, R. M., van Rossum, A. C., 2005. Propagation of onset and peak time of myocardial shortening in ischemic versus nonischemic cardiomyopathy assessment by magnetic resonance imaging myocardial tagging. *Journal of the American College of Cardiology* 46 (12), 2215–2222.

Zwanenburg, J. J. M., Kuijer, J. P. A., Marcus, J. T., Heethaar, R. M., 2003. Steady-state free precession with myocardial tagging: CSPAMM in a single breathhold. *Magnetic Resonance in Medicine* 49, 722–730.