Spatio-temporal registration and modelling of the heart using cardiovascular MR imaging

Dimitrios Perperidis

A dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

of the

University of London

September 2005

Department of Computing

Imperial College London

Abstract

Cardiovascular diseases are the single most important cause of death in the developed world. Their early diagnosis and treatment is becoming crucial in order to reduce mortality and to improve the patient's quality of life. The detection of cardiac diseases requires advanced methods for the quantification and analysis of the cardiovascular system. Recent advances in image modalities allow the high resolution imaging of the heart. Among them, MR imaging plays an increasingly important role for the understanding of the heart and the detection of cardiac abnormalities. However, MR imaging remains a new technology limited to specialised imaging centers. Therefore, algorithms assisting with the interpretation of cardiac MR images are of high importance.

In this thesis, several novel approaches for the spatio-temporal alignment of cardiac MR image sequences have been developed. The registration algorithms have the ability to correct any spatial misalignment caused by differences in the acquisition of the hearts and by local shape differences. Furthermore, the registration techniques have the ability to correct temporal misalignment caused by differences in the length of the cardiac cycles and by differences in the dynamic motion patterns of the hearts.

The spatio-temporal registration methods are used for the development of novel statistical and probabilistic atlases describing the anatomy and function of a healthy heart. To our best knowledge, this is the first attempt to build atlases for cardiac MR image sequences describing the cardiac function as well anatomy. The probabilistic atlas provides statistical information in the form of tissue probability maps while the statistical atlas provides additional information regarding shape and function variability across the healthy population. During the construction of the statistical atlas the distribution of cardiac shapes is divided into two subspaces. One distribution subspace accounts for changes in cardiac shape caused by inter-subject variability. The second distribution subspace accounts for changes in cardiac shape caused by deformation during the cardiac cycle (i.e. intra-subject variability).

The atlases can be used as educational tools and for assisting the diagnosis of cardiac diseases. A possible use of the statistical atlases is demonstrated by using them to classify image sequences from normal subjects and subjects with hypertrophic myocardiopathy.

To my family and my fianceé

Acknowledgements

First I would like to thank my supervisor Prof. Daniel Rueckert. This thesis would not have been possible without his encouragement, support and enthusiasm during the last four years. It has been an honour to work with him and I look forward to continue collaborating with him in the future.

I would also like to thank Dr Raad Mohiaddin from the Royal Brompton and Harefield NHS Trust London, for his valuable input and for acquiring many of the image sequences used in this thesis.

Sincere thanks are also due to Prof. David Hawkes and Dr Elizabeth Berry for agreeing to be the examiners of this thesis. This thesis owes a great deal to Dr Anil Rao, Dr Raghavendra Chandrashekara, Dr Maria del Carmen Lorenzo-Valdés, Gerardo Ivar Sanchez-Ortiz and Dr Carlos Thomaz.

I would also like to express my gratitude to colleagues and friends of the VIP group for their support over the course of my research. Special thanks are due to Kanwal, George, Satheesh, Theodore, Maria, James, Fani and Paul for the numerous discussions in academic and non academic subjects. I also would like to thank all my friends in London and in Greece for their support and advice.

Finally, I would like to express my gratitude and love to my parents Stelios and Maika, my brother Antonis, my fianceé Constantina and all my family for their understanding and constant support. This thesis and the great happiness behind it is theirs as well as mine. You have given so much to me and I would like to dedicate this work to you.

Contents

List of Figures

List of Tables

Glossary of notation

$g(\mathbf{x})$	Grey level data at pixel $(\mathbf{x})$
$\bigcap$	Intersection
$\bigcup$	Union
∇	Gradient
$\phi_{i,j,k}$	Control point vector at the lattice position i, j, k in spatial domain
ϕ_t	Control point vector at the lattice position t in temporal domain
I'	Image used as a reference during registration
S'	Image sequence used as a reference during registration
$\mathbf{T}(\mathbf{x})$	Coordinate transformation of point vector $\mathbf{x}$
$\mathbf{T}(I)$	Image transformation of image I
$\mathbf{T}(S)$	Image transformation of image sequence S

Chapter 1

Introduction

Cardiovascular diseases (CVDs) are the single most important cause of death in the developed world [133]. According to a recent estimate of the *World Health Organization* 16.7 million deaths each year are caused by CVDs [133]. Their early diagnosis and treatment is becoming crucial in order to reduce mortality and to improve patients' quality of life. The detection of CVDs requires advanced methods for the quantification and analysis of the cardiovascular system. Novel cardiac imaging modalities provide researchers and clinicians with important tools for the diagnosis and treatment of CVDs. This chapter provides the motivation for this work, reviews the basic anatomy and function of the heart, describes the main modalities for the imaging of the cardiovascular system and finally analyses the contributions of this thesis.

1.1 Motivation

Recent advances in non-invasive imaging modalities allow for the high resolution imaging of the cardiovascular system. Among these modalities, magnetic resonance imaging (MR) is playing an increasingly important role. MR imaging allows not only the acquisition of high resolution 3D cardiac images which describe the cardiac anatomy but also the acquisition of 4D cardiac image sequences which describe the cardiac anatomy and function. Although the use of MR imaging in clinical practice is rapidly increasing, it

still remains a new technology primarily limited to specialised imaging centers. Therefore, applications assisting the interpretation of MR images are of high importance for increasing the clinical use of MR imaging.

The recent advances in MR imaging have led to an increased need for image registration and normalisation methods which are used in a large number of applications for calculating the cardiac motion [29], segmenting cardiac images [109] and building models which describe the cardiac anatomy [112]. However, most of these registration methods ignore any temporal information contained in the cardiac MR images and deal only with the cardiac anatomy. The development of techniques for the spatial and temporal normalisation and alignment of 4D cardiac MR images will enable comparison between the cardiac anatomy and function of number of subjects to be made. These spatial and temporal registration methods could also be used for the construction of computerised probabilistic and statistical models containing information regarding the variability of anatomy and function of a healthy heart. Furthermore, these models would assist the better interpretation of MR image sequences. They could also be used for classifying images and segmenting images. For example, Valdés *et al.* used a probabilistic atlas of the heart to enable the segmentation of cardiac MR images [108]. Moreover, the atlases will also enable statistical and computational comparisons between individuals and groups to be performed making them important clinical tools.

In order to build a cardiac model from MR images describing both the cardiac anatomy and function of a healthy heart the following research issues must be addressed:

- How to compare the cardiac anatomy and dynamics for an individual subject, between different subjects and between different group of subjects. Current approaches, enable the comparison of only the cardiac anatomy or the cardiac dynamics.

- How to build a model of the heart which captures statistical information regarding variability of both the cardiac anatomy and dynamics across a population of subjects. Current approaches enable the construction of cardiac models of only the cardiac anatomy or the cardiac dynamics. No approach presently exists which

combines information regarding the cardiac anatomy and dynamics into a single model.

1.2 Objectives of the thesis

The following are the main objectives of this thesis:

- To develop new registration methods which enable the mapping of cardiac MR image sequences into a common spatial and temporal coordinate system. Mapping only the cardiac anatomy, as most cardiac registration approaches do, is not enough since the heart is undergoing spatially and temporally a varying degree of motion during the cardiac cycle. The spatio-temporal mapping will allow the direct comparison between the cardiac anatomy and function of a number of image sequences to be made. In order to achieve this, novel methods for the spatial and temporal alignment of cardiac MR image sequences are required.

- To develop new methods for building probabilistic and statistical models of the cardiac anatomy and function. These models will contain statistical information regarding the variability of both the anatomy and function of the heart across a population of subjects. The spatio-temporal registration methods will be used in order to align a number of cardiac MR image sequences, used for the construction of the models, to the same spatial and temporal coordinate system enabling direct comparisons between their cardiac anatomy and function to be made.

1.3 The cardiovascular system

For a better understanding of the remainder of this dissertation, a brief review of the anatomy and function of the human cardiovascular system as well as its imaging modalities is given in the next two sections. For a more detailed analysis of the cardiovascular system see [18] and for MR imaging see [117].

The cardiovascular system is one of the most vital systems in the body. Its prime responsibility is the circulation of blood to cells throughout the body. The blood provides oxygen from the lungs to the cells and transfers carbon dioxide from the cells to the lungs. The circulation of the blood is achieved by the heart which forces the blood through the blood vessels. The circulation of the blood can be divided into two stages: the *pulmonary circulation* and the *systemic circulation*. During the *pulmonary circulation* the blood is carried out from the heart to the lungs. In the lungs oxygen is absorbed and carbon dioxide is removed from the blood. During the *systemic circulation* the blood is pumped by the heart to the cells in the body. It is obvious that if the heart stops pumping blood, the cells of the body will not be able to survive without the necessary oxygen and death will be the eventual consequence.
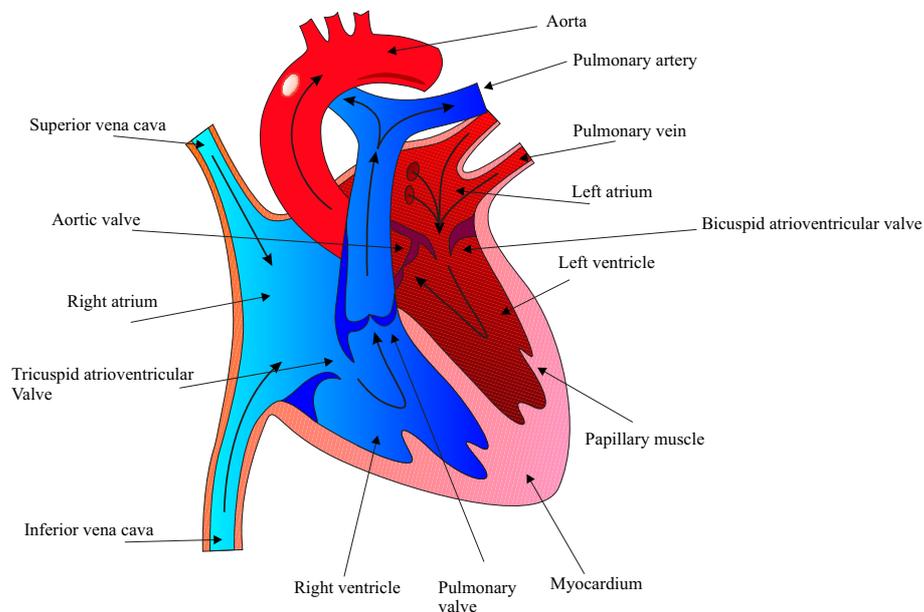
Figure 1.1: A schematic figure of the heart (adapted from [18].

The heart is composed of two pumps. A schematic diagram of the heart is shown in figure 1.1. The left part of the heart is responsible for the *systemic circulation* and the right part is responsible for the *pulmonary circulation*. Each side of the heart has two chambers (the *atrium* and the *ventricle*). The *atrium* receives blood from the veins and aids its flow into the *ventricle* which forces it into the arteries. Therefore, the blood coming from the body enters the right *atrium*, flows to the right *ventricle* (RV) and exits towards the

lungs. After absorbing the necessary oxygen, the blood returns to the heart through the left *atrium* and then it is propelled to the body by the left *ventricle* (LV). The walls of the *ventricles* are composed of a muscular tissue which contracts and thus pumps the blood out of the *ventricles*. The muscular tissue in the *ventricular* walls is the *myocardium*. The inner surface of the *myocardium* is the *endocardium*, while the outer surface is the *epicardium*. The *atria* can also contract like the *ventricles* but their main function is to act as reservoirs which are filled with blood flowing back through the veins to the heart. The heart is enclosed by the *pericardium* which separates it from the *epicardium* by the pericardial fluid.

The *ventricles* and the *atria* are connected by the *atrioventricular* (AV) valves (the *tricuspid* and the *mitral* valves). Attached to the free margins of these valves are the *tendinous cords* which are attached to projections of *ventricular* muscles known as *papillary muscles*. The exit from the right *ventricle* into the *pulmonary* artery is guarded by the *pulmonary semilunar valve* and the exit from the left *ventricle* into the aorta is guarded by the *aortic semilunar valve*.

The main role of the blood vessels is to carry blood through the entire body. The blood vessels are divided into *arteries* and *veins*. The *arteries* carry blood away from the heart while the *veins* carry blood towards the heart. The large arteries are named *elastic arteries* since their volume is dynamically changed due to the large volume of blood ejected from the heart. The elastic properties of these arteries have an important role in the cardiovascular haemodynamics. They transform the rhythmic pulsating high pressured blood flow coming from the heart into a more stable flow with significantly lower pressure.

During the circulatory process the blood flows from the systemic circulation into the right *atrium* via the superior and inferior *vena cava* and from the pulmonary circulation to the left *atrium* via the four *pulmonary veins*. When the pressure of each *atrium* is greater than the pressure of the corresponding *ventricle*, the AV valves open causing the blood to flow towards the *ventricles*. When the *ventricles* are about 80% full, the *atria* contract and propel blood into the *ventricles* to complete their filling. The high pressure increase in the

ventricles causes the AV valves to close preventing the blood flowing back to the *atria*. The contraction of the right *atrium* wall causes the blood to flow from the right *atrium* to the right *ventricle* via the *tricuspid* valve. The right *ventricle* contracts forcing the blood through the *semilunar* valve into the *pulmonary* artery leading to the lungs where the blood is saturated with oxygen. The high pressure increase in the right *ventricle* causes the *tricuspid* valve to close preventing the blood entering back to the right *atrium*. The oxygenated blood flows back into the left *atrium* through the *pulmonary* veins. The blood enters from the left *atrium* the left *ventricle* when the *mitral* valve opens. When the left *ventricle* contracts, the pressure of the oxygenated blood rises. When the pressure in the left *ventricle* exceeds the one in the *aorta*, the *aortic semilunar* valve opens ejecting blood with high pressure into the *aorta*.

1.3.1   Electrical properties of the heart

The cardiac muscle has the ability to contract without nervous input. The action potential for each heart beat is generated by a pacemaker node in the right *atrium* and is transmitted through the heart along specialised pathways (figure 1.2). The pacemaker is a small area in the wall of the right *atrium* known as *sinoatrial node* (SA). The action potential, generated in the SA node, is carried from cell to cell along the plasma membrane through the low electrical resistance of the intercalated disks. The conduction velocity throughout the *atrial* muscle is increased by three pathways: the *anterior* band, the *middle* band, and the *posterior* band. These bands merge near the next node, the *atrioventricular node* (AV node). The AV node is located in the *atrioventricular fibrous* ring on the right side of the *atrial septum*. From the AV node the action potential travels to the *bundle of His* and from there to the *Purkinje network*. From there the action potential is spread throughout the *ventricular* wall. Figure 1.2 provides a schematic description of the nodes and the conduction pathways of the heart.

The cardiac action potential has different characteristics in different regions of the heart. It has different characteristics in the *atrial* muscle, in the SA and AV nodes, in
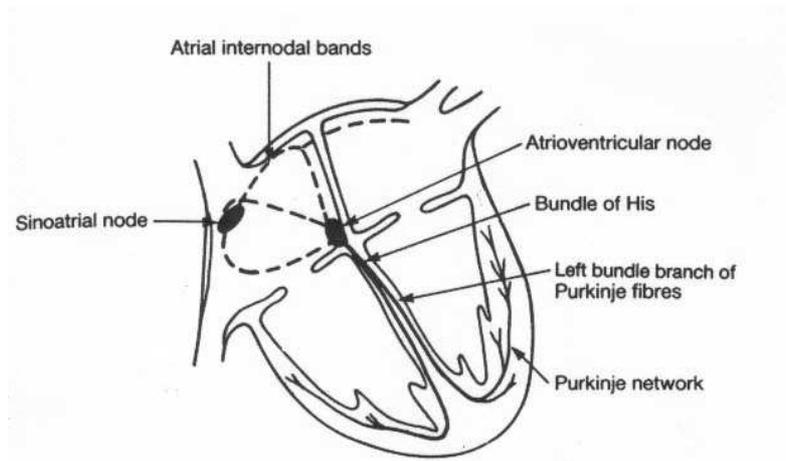
Figure 1.2: The conducting pathways of the heart(reproduced from [18]).
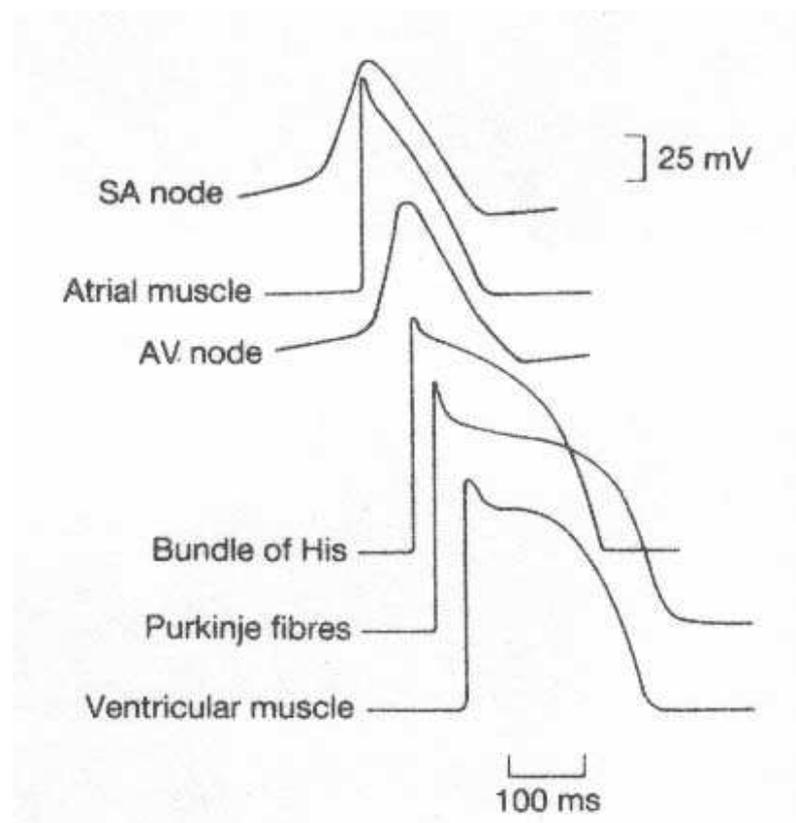
Figure 1.3: Shape, duration and sequence of cardiac output potential (reproduced from [18]).

Purkinje fibres and in the *ventricular* muscle (see figure 1.3). The duration of the action

potential (*action potential duration*) also varies in different regions of the heart (figure

1.3). The shortest duration occurs in the SA and AV nodes and also the *atrial* muscle
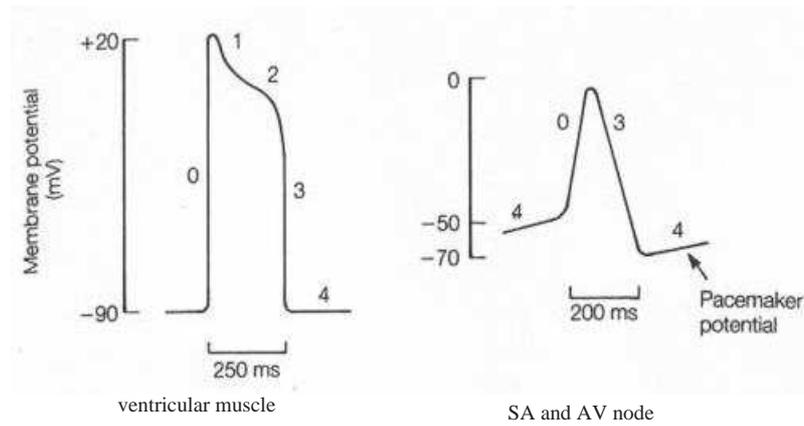
Figure 1.4: Cardiac potentials in the *ventricle* and SA or VA nodes (reproduced from [18]).

(200-250*msec*), while the longest duration occurs in the *Purkinje fibres* (300-400*msec*) [18]. In the *bundle of His* the duration of the action potential is within the range of 250-300*msec*. The action potential characteristics of the following areas are (figures 1.3 and 1.4) [18]:

- In the *ventricular* muscle the resting membrane potential is steady, approximately at -90*mV* (Phase 4 in figure 1.4). The action potential has a rapid depolarisation (Phase 0), reaches its peak at approximately +20*mV*, and afterwards it has a rapid but short decline (Phase 1). Then, the action potential has a prolonged shoulder (Phase 2). At the end, the membrane repolarises very quickly (Phase 3).

- The action potential of the *atrial* muscle has similar characteristics with the *ventricular* muscle. As figure 1.3 shows, it has less obvious plateau (phase 2) and longer repolarisation.

- In the *SA* and *AV* nodes the cells have less negative resting membrane potential than other heart cells. In between action potentials, their membrane potential (Phase 4 in figure 1.4) is also unstable and depolarises at approximately 20*mV* (from around -70*mV* to around -50*mV*). The cells in these nodes have slower depolarisation (phase 0 in figure 1.4) than the cells in the *ventricular* muscles. This slow depolarisation is called the *pacemaker potential*. The action potential reaches its peak at around

0*mV* and repolarises with speed similar to the *atrial* muscle (phase 3). This can be seen in figure 1.3.

1.3.1.1 The Electrocardiogram (ECG)

The synchronised deporalisation spreading through the heart causes currents in the extracellular fluid that establish field potentials over the entire body. These potential differences can be detected by placing electrodes on various places on the body's surface. The detected signal is called the *electrocardiogram* (ECG). Prior to recording, the detected signal needs to be amplified. Figure 1.5 provides a schematic figure of a standard ECG record at resting heart rate.
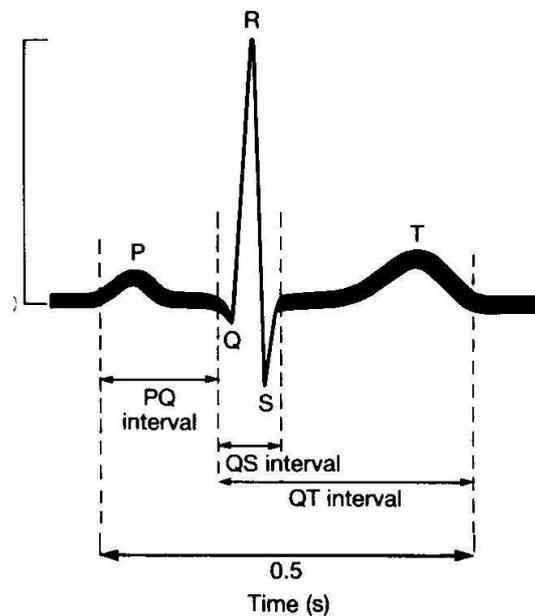
Figure 1.5: A standard ECG record (reproduced from [18]).

The pattern on the ECG signal depends in the position of the electrodes. However, certain features are always present (figure 1.5). The *P wave* is produced by the spread of electrical activity during the *atrial* depolarisation. The *QRS complex* is produced by *ventricular* depolarisation. The *T wave* is produced by *ventricular* repolarisation. When no repolarisation or depolarisation occurs, there is not potential difference in the ECG record (the *isometric line*). The *atrial* depolarisation occurs during the *QRS complex* and
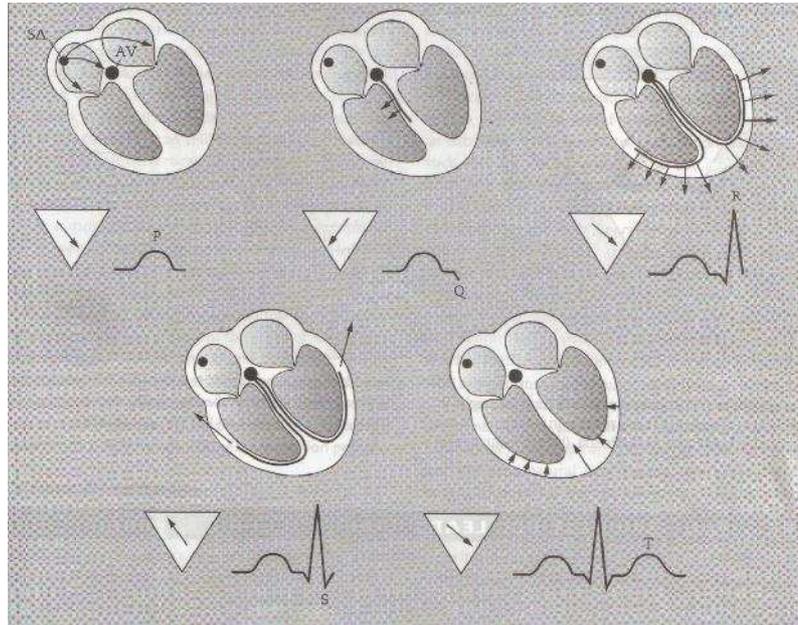
Figure 1.6: The depolarisation and repolarisation of the heart and the resulting ECG record (reproduced from [18]).

does not produce any visible wave in the ECG signal. The *T wave* is longer than the *QRS complex* but smaller in amplitude because the *ventricular* repolarisation is less well synchronised than the *ventricular* depolarisation. The *PQ* interval (or the *PR* interval) in figure 1.5 is the time required for excitation to spread through the *atria*, the AV node and the *bundle of His*. The time required for excitation to spread through the *ventricle* is denoted by the *QS* interval. The *QT* interval measures the duration of the *ventricular* action potential, while the *PS* interval measures the duration of the action potential. Figure 1.6 shows the depolarisation and repolarisation of the heart and the resulting ECG record [18].

1.3.2 The cardiac cycle

In normal conditions the human heart beats between 65 to 75 times per minute, which means that each heart beat lasts around 0.85*sec*. Each heart beat is considered as a cardiac cycle which is separated into a contraction phase (*systole*) and a relaxation phase (*diastole*) of the *atria* and *ventricles*. In order to analyse the cardiac cycle in more detail it latter is separated into seven phases ([91]). Figure 1.7 provides an example of the
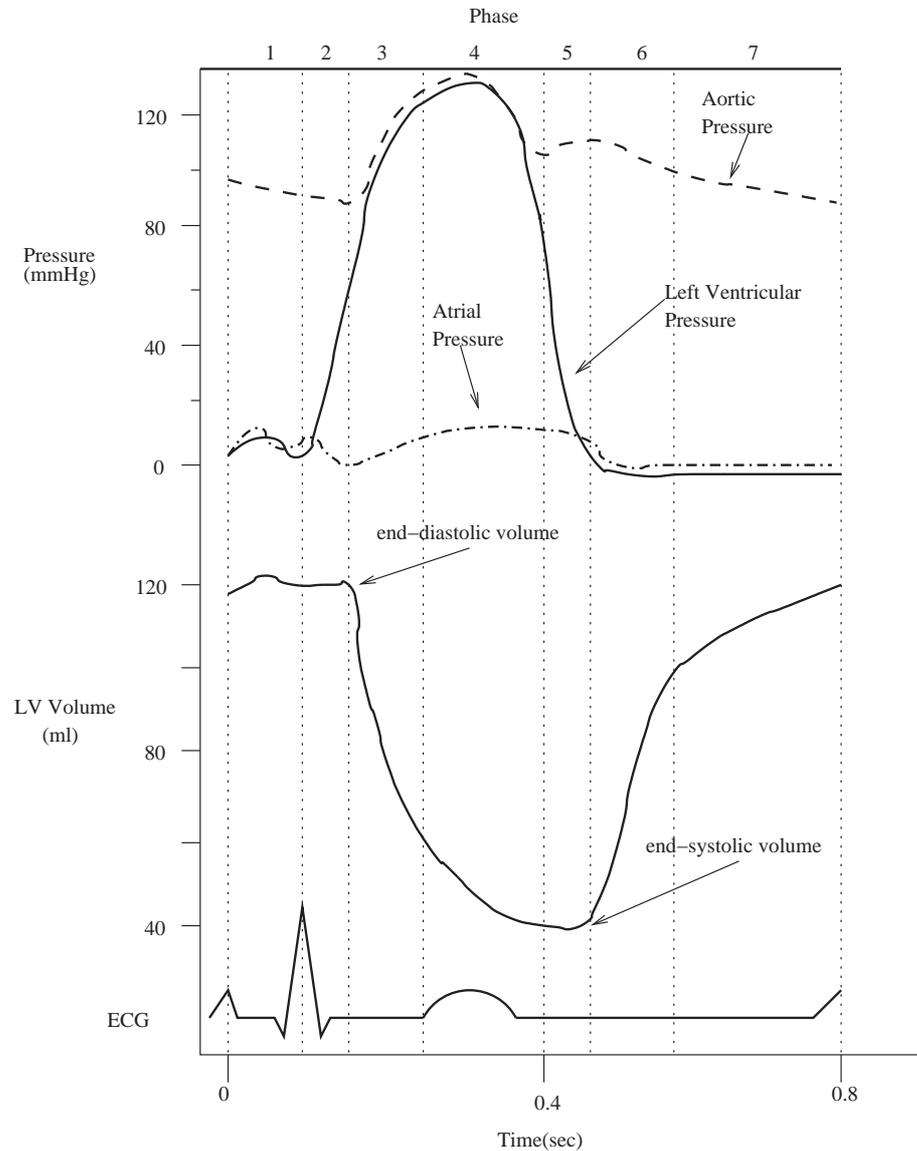
Figure 1.7: An example of the pressure and volume of the heart during the cardiac cycle (adapted from [91]).

volume of the left *ventricle* during each phase of the cardiac cycle. In addition, it shows the relationship between the ECG signal, the *ventricular* volume and the cardiac pressure during the cardiac cycle. The seven phases of the cardiac cycle are:

- The *atrial contraction*

- The *isovolumetric contraction*

- The *rapid ejection*

- The *reduced ejection*

- The *isovolumetric relaxation*

- The *rapid ventricular filling*

- The *reduced ventricular filling*

The *atrial contraction* is initiated by the electrical depolarisation of the *atria* (*P wave* of figure 1.5). As the *atria* contract, the pressure within the *atrial* chambers increases. The pressure gradient which is generated across the open AV valves causes a rapid flow of blood into the *ventricles* (figure 1.7).

The *isovolumetric contraction* is initiated by the *ventricular* depolarisation (*QRS complex* of figure 1.5). During the isovolumetric contraction all the valves are closed and the volume of the *ventricles* remains the same. However, there is a rapid increase in the *intraventricular* pressure due to depolarisation of the *ventricle*.

In the *rapid ejection* phase the *intraventricular* pressures exceed the pressures within the *aorta* and *pulmonary* artery. This causes the *aortic* and *pulmonic* valves to open and blood is ejected out of the *ventricles*.

After the *rapid ejection* phase, *ventricular* repolarisation occurs (*T wave* of figure 1.5) causing the *ventricular* tension to fall and the rate of *ventricular* emptying to be reduced. This is called the *reduced ejection* phase. In this phase the *ventricular* pressure falls below the blood pressure in the outflow tracts. However, the blood continues to flow due to its kinetic energy.

When the total energy of the blood within the *ventricles* is less than the energy of the blood in the outflow tracts, the *aortic* and *pulmonic* valves close (*isovolumetric relaxation* phase). During this phase, the *ventricular* pressures decrease. However, the *atrial* pressures continue to rise due to venous return (figure 1.7). The volumes of the *ventricles* remain constant since all the valves are closed. The volume of the blood which remains in the *ventricle* is called the *end-systolic* volume. In the left *ventricle* the *end-systolic* volume is about 50*ml*.

The *rapid ventricular filling* occurs when the *ventricular* pressures fall bellow the *atrial* pressures. Then, the AV valves open and the *ventricular* filling begins. The *ventricles* continue to relax, despite blood inflow, which causes *intraventricular* pressure to continue falling by a few additional *mmHg*. The opening of the AV valves and the rapid flow of blood cause a rapid fall in the *atrial* pressure (figure 1.7).

In *reduced ventricular filling* phase the *ventricular* pressure rises as the *ventricles* fill with blood. This reduces the pressure gradient across the AV valves so that the rate of the filling falls. The *aortic* pressure (and *pulmonary arterial* pressure) continue to fall.

1.3.3 Relationship between cardiac volume and pressure

The performance of the cardiac muscle is affected by the extent to which it is stretched during diastole, its current state of contractible energy and the *atrial* pressure against which the blood has to be ejected [18]. The performance of the heart can be studied by examining the relationship between the cardiac volume and the cardiac pressure. An example of a volume-pressure curve of the left *ventricle* is given in figure 1.8.
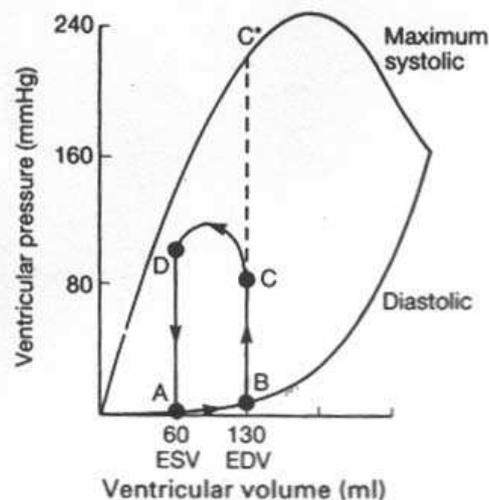
Figure 1.8: Left *ventricular* volume-pressure curve of a normal cardiac cycle (reproduced from [18]).

During the diastole the heart fills with blood, increasing its volume (from an end-systolic volume of 60*ml* to an end-diastolic volume of 130*ml*). At this phase there is an

increase in the pressure in the left *ventricle* from 5*mmHg* at point A (figure 1.8), where the AV valve opens (*rapid ventricular filling phase*), to around 10*mmHg* at point B where the AV valve closes (*isovolumetric contraction* phase). During the *isovolumetric contraction* there is a rapid increase in the pressure (point C on figure 1.8). If at this phase the *aorta* was clamped so that blood could not escape, the pressure would rise to point C^* on figure 1.8. This is the *maximum systolic ventricular* pressure. The pressure at C depends on the diastolic pressures of the *aorta*. At C the *aortic* valve opens (*rapid ejection* phase) and blood is ejected from the heart. During the ejection of the blood the ventricular pressure increases from about 80*mmHg* to about 120*mmHg* and then declines reaching point D (at this point the aortic valve closes) at about 100*mmHg*. The contraction phase from point C to D is referred to as *auxotonic* because it is occurring against the afterload of a varying aortic pressure. The aortic valve closes at D, where isovolumetric relaxation occurs and the pressure drops from D to A. The area enclosed under the volume-pressure loop (ABCDA) is a measure of the *external work* done by the heart.

A number of measurements are important in clinical practice. The *stroke volume* is the difference of the *ventricular* volumes at the end-diastole (EDV) and end-systole (ESV). The *stroke volume* multiplied by the heart rate is the *cardiac output*. The ratio of the *stroke volume* to the *end diastolic* volume is the *ejection fraction* [4]. The volume of the left *ventricle*, the *ejection fraction* and the wall thickness are the most important indices of cardiac performance [132]. Table 1.1 provides representative values for these cardiac measurements at rest and after exercising [18].

	Cardiac output $(\mathrm{L}min^{-1})$	Heart rate (beats min^{-1})	Stroke volume (*ml*)	EDV (*ml*)	ESV (*ml*)
	Non-athlete				
Rest	5	70	70	130	60
Max. exercise	21	190	110	130	20
	Trained athlete				
Rest	5	40	120	200	80
Max. exercise	34	190	180	200	20

Table 1.1: Representative values of cardiac measurements at rest and after exercising. Values for trained athletes and non-athletes are provided (adapted from [18]).

1.3.4 Cardiovascular diseases

According to the American Heart Association, 54% of deaths in United States in 2002 resulting from *cardiovascular diseases* (CVDs) are due to the *coronary heart disease* [2]. Furthermore, *stroke* results to 18% of deaths caused by CVDs, while *congestive heart failure* and high blood pressure cause the 6% and 5% of deaths resulting from CVDs. Finally, diseases of the arteries account for 4% of the deaths caused by CVDs. Figure 1.9 provides the percentage breakdown of deaths caused by CVDs [2].
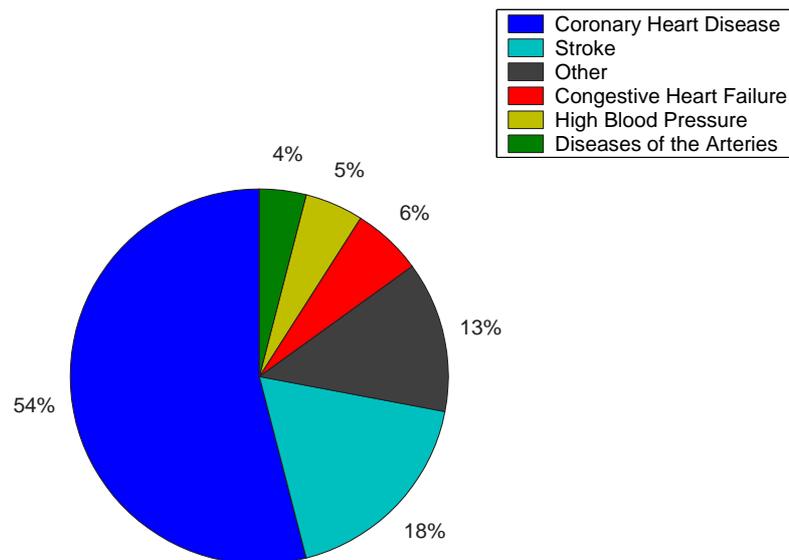
Figure 1.9: Analysis of deaths caused by cardiovascular diseases (adapted from [2]).

Coronary heart diseases are caused by the occlusion of coronary arteries. Occlusion in the arteries can be caused from the build up of fatty tissue. If a coronary artery is occluded, part of the cardiac muscle receives a restricted supply of oxygen. This causes the loss of cardiac function which can even lead to death. The *stroke* does not affects the heart but the arteries leading to and from the brain. A *stroke* is caused when a vessel carrying blood to the brain is blocked by a clot. The consequences of the *stroke* are very serious and result in the death of a part of the brain. *Congestive heart failure* is the result of the weakening of the heart by a chronic disease. Diseases of the arteries are most often caused by a condition known as *atherosclerosis*. *Atherosclerosis* is caused by the buildup of a plaque in the arteries. The plaque consists of fatty substances, cholesterol, cellular

waster products and other substances. *Atherosclerosis* reduces the blood flow through the artery and can cause it to lose its elasticity.

1.4   Imaging of the cardiovascular System

The main imaging modalities of the cardiovascular system are: *X-Ray*, *Ultrasound (US)*, *Positron Emission Tomography (PET)*, *Single Photon Emission Computed Tomography (SPECT)*, *Computed Tomography (CT)* and *Magnetic Resonance Imaging (MR)*. A more detailed description of the imaging methods can be found by Suetens [188] and by Webb [203]. Each imaging modality provides a different type of information and is used for specific purposes [122].

1.4.1   X-Ray

In *X-Ray* imaging of the heart a radio-opaque medium is injected into the blood in the heart. The contrast medium is injected directly into the *atrial* muscle, *ventricles* or vessels of the heart by means of the cardiac catheter. After introducing the medium a recording of the chambers or the vessels of the heart is made. The recording (*angiocardiogram*) is made by using a rapid-sequence digital subtraction technique. In this technique, one image of the heart is acquired before the contrast medium is added and one after the injection of the medium. Then, the first image is subtracted from the contrast enhanced image allowing the anatomy of the blood vessels and the blood supply of an organ to be imaged with high spatial (0.1*mm*) and temporal (20-50*msec*) resolution. An example of coronary angiography is shown in figure 1.10.

Angiocardiography is commonly used in diagnosis of cardiac diseases and in planning the surgical treatment of the heart. It allows clinicians to assess the status of the coronary arteries in *ischemic* disease, the function of the cardiac valves, the congenital heart lesions and the abnormalities of the great vessels. Furthermore, it allows clinicians to measure the volume of the cardiac chambers.

Figure 1.10: An example of coronary angiography.

1.4.2 Ultrasound

Ultrasound (US) employs pulsed, high frequency sound waves. When an ultrasound wave meets an interface of differing echogenicity, the wave is reflected, refracted and absorbed. The reflected sound waves are sensed by the transducer (which also emits the sound waves) and then, processed to produce the image. The ultrasound beam can be aimed at specific directions and obeys the laws of optics with regards to refraction, reflection and transmission. The intensity of the ultrasound beam decreases as it travels away from the ultrasound source due to divergence, absorption, scatter and reflection of the wave energy at tissue interfaces. Stronger reflections are returned when the ultrasound beam is perpendicular to the imaged structure. Fig 1.11 shows an example of a cardiac ultrasound image.

Ultrasound produces high resolution images and is capable of resolving fine anatomical structures. The thickness, size and location of various soft-tissue structures in relation to the origin of the ultrasound beam can be calculated at any point in time.

Ultrasound is frequently used for the imaging of the heart (*echocardiography*). There are three types of *echocardiography*: *M-mode* US, *B-mode* (brightness) US and *Doppler* mode US. The *M-mode* provides one dimensional view (depth) into the heart. The *B-*

Figure 1.11: An example of cardiac ultrasound.

mode US uses an array of transducers which allows a plane of tissue to be imaged, thus making the anatomical relationships between different structures easier to appreciate than with the *M-mode* imaging. *Doppler* imaging allows evaluation of blood flow patterns (direction and velocity) by detecting frequency changes occurring when ultrasound waves reflect off individual blood cells moving towards or away from the transducer.

The recent advances in the development of *Ultrasound* imaging enable the acquisition of 3D images which overcome some limitations of the conventional 2D *Ultrasound* imaging. 3D *Ultrasound* imaging [57] has the potential to provide real time 3D visualisations of the heart.

Echocardiography is more frequently used to evaluate cardiac chamber size, wall thickness, wall motion, valve configuration and motion and the proximal great vessels. Ultrasound has a high temporal resolution and is widely used because it is more portable and less expensive than other modalities. However, the signal to noise ratio of the US images is low. The fact that it is operator dependent also leads to significant artifacts especially in 3D imaging.

1.4.3 Nuclear imaging

In nuclear imaging a tracer containing radioactive isotopes is administrated to the patient prior to the imaging procedure. The isotopes reach the organs as part of its metabolic process and emit gamma ray photos which are used to reconstruct maps of the concentration of the tracer over time. Two kinds of imaging can be performed depending on the tracer module: *Positron Emission Tomography* and *Single Photon Emission Computed Tomography*.

1.4.3.1 Positron Emission Tomography

Positron Emission Tomography (PET) has improved the understanding of the biochemical basis of normal and abnormal functions within the body. Positrons ($\beta+$) are positively charged electrons. They are emitted from the nucleus of some positively charged radioisotopes (excessive number of protons). In positron emission the nucleus is stabilised by converting a proton into a neutron (i.e. by removing a positive charge). During the conversion the element formed from positron decay is not radioactive. When a positron comes in contact with an electron, the mass of the two particles is turned into two 511-KeV gamma rays which are emitted with 180 degree angle to each other. These rays escape from the human body and can be recorded by a number of detectors. The detection of the two rays is called a coincidence line. The coincidence line provides a unique detection scheme for forming tomographic images with PET.

Positron Emission Tomography (PET) provides valuable information regarding the physiology of the heart. It provides a unique tool for biochemistry and physiology measure of the cardiac muscles. It can be used to localise and describe coronary artery diseases and also to identify injured but zoetic myocardium.

1.4.3.2 Single Photon Emission Computed Tomography

In *Single Photon Emission Computed Tomography* (SPECT), radioactive tracers are injected to the subject. The radioactive tracers emit radiation within the organs. The SPECT

camera is a large scintillation crystal connected to multiple photo-multiplier tubes which can detect the radiation emitted from the body. The technology of SPECT arises from positioning the camera head at a large number of angles around the body accumulating as many as 180 views of a specific organ.

There are techniques for the direct imaging of the myocardium or the imaging of the blood pool. Nuclear-tagged compounds are used for selection of the correct radioactive material for a specific type of tissue. For example for myocardial imaging, compounds such as Thallium 201 are used. SPECT imaging is similar to PET but only one gamma-ray is emitted per nuclear disintegration while in PET imaging two gamma-rays are emitted.

SPECT is used to assess the location and extent of *ischemia* in the heart resulting from coronary heart disease. This modality provides a 3D density map of blood in the myocardium. In order to identify *ischemic*, *infarcted* and *normal* tissues, a stress-rest study is performed. In the stress-rest study images are acquired while the patient is at rest and while the patient heart is under stress (by exercising). If the density distributions in the two perfusion maps are normal in the stress and rest images the state of the myocardium is *normal*. The state of the myocardium is *ischemic* if the density distribution in the stress perfusion map is low but normal in the rest perfusion map. The state of the myocardium is *infarcted* if the density distribution in both the rest and stress images are low.

1.4.4   Computed Tomography

Computed Tomography (CT) is a technique based on x-rays. The resulting images are formed by x-ray absorption of tissue. The technique uses a narrowly collimated x-ray beam to irradiate a slice of the body. The amount of radiation transmitted along each projection line is collected by a number of photo-multiplier tubes. The image is formed by rapidly acquiring a large number of views by rotating the tube and the detectors around the body. In cardiac imaging the advantage of CT is its ability to image the cardiac arteries. CT provides anatomical images with a high spatial resolution and with a good contrast between bone structures and soft-tissue structures. However, the contrast between

different types of tissue is very poor and can be only enhanced using contrast materials. Figure 1.12 shows an example of a CT image of the heart.
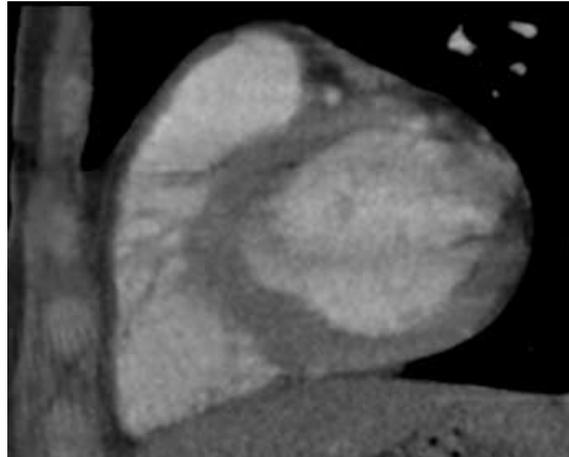
Figure 1.12: An example of cardiac CT.

The motion artifacts due to respiratory movement affect the quality of CT images. These artifacts occur when one or more slices acquired in different breathholds are displaced due to the breathing motion of the patient. One major disadvantage of CT modalities is the use of x-ray radiation. Even though a large number of measures have been taken to reduce the amount of radiation during a CT scan, it is generally not considered acceptable to expose a patient to such an amount of radiation unless when it is absolutely necessary. The main applications of CT in cardiac imaging are the evaluation of cardiac masses and the evaluation of aortic and pericardial diseases.

1.4.5   Magnetic resonance imaging

Magnetic Resonance Imaging (MR) plays an increasingly important role for the imaging of the cardiovascular system. The physical principles of MR are more complex than any of the previous tomographic modalities. MR uses strong magnetic fields to produce maps of atomic nuclei. Hydrogen is the most commonly atomic nucleus in the human body. In the human body, hydrogen atoms are contained in water and fat molecules. The spin of the atomic nuclei could be considered as a magnetic vector, causing the proton to behave like a magnet. The image acquisition involves an initial sequence of exciting pulses and

the recording of the emitted signal. The amplitude of the signal is used to generate maps showing the anatomy of the heart.

The advantage of MR modality is the high resolution imaging of the cardiac anatomy in tomographic planes of any desired position and orientation. MR imaging allows not only the acquisition of 3D images describing the anatomy of the heart but also the acquisition of 3D image sequences (3D + time) describing the cardiac anatomy and function. The high tissue contrast enables the assessment and measurement of different cardiac structures. A detailed discussion regarding the principles of MR imaging is outside the scope of this thesis. For more detailed analysis of the MR imaging see [24, 195, 117, 97, 118] .

The quality of cardiac MR images can be affected by a large number of factors:

- The constant motion in all three dimensions that the heart undergoes during the cardiac cycle.

- The motion artifacts due to respiratory movement. These can occur when one or more slices acquired in different breathholds are displaced due to the breathing motion of the patient.

- The partial volume effects. This problem occurs when a voxel contains two or more types of tissue and as a result the edges of the images are blurred. This problem is usually caused due to the anisotropic resolution of the images. The in-plane resolution can be high ($1mm$) but the through-plane resolution is usually low ($10mm$).

Conventional MR imaging techniques require a sequence of excitations to reconstruct the image. These excitations need to be triggered away at the same part of the cardiac cycle. In order to achieve this the electrocardiogram is used resulting in a technique known as *cardiac gating*. There are two main techniques for synchronising the measurement of MR data with the cardiac cycle. *Prospective gating* techniques wait for the trigger signal to start acquisition of data. The data are acquired using excitations at a fixed time points in the cardiac cycle. The data collection is paused at the end of the cardiac cycle and and the sequence waits for the next trigger signal. As mentioned in section 1.3.1.1 the *R wave* of

the ECG signal corresponds to the beginning of the contraction phase of the heart. After a small time interval, t_{delay}, the first frame of the image sequence is acquired (figure 1.13). After the acquisition of the first frame, frames are acquired every $t_{offset}msec$. In *retrospective gating* data collection is done continuously throughout the entire cardiac cycle. Then, the recorded trigger signal is used to retrospectively assign the data to the different positions in the cardiac cycle.
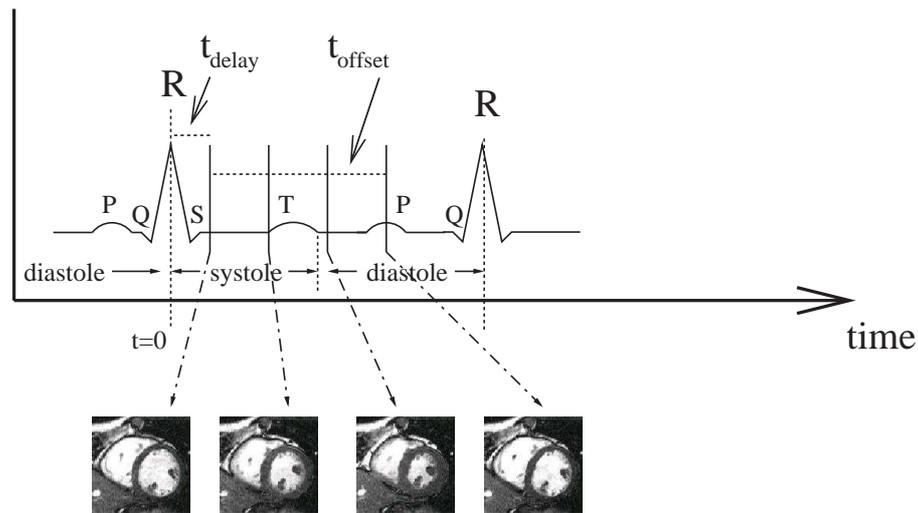
Figure 1.13: An example of an ECG-gated acquisition of a cardiac MR image sequence: The parameter t_{delay} describes the delay after the *R wave* after which the MR acquisition starts while t_{offset} describes the temporal resolution of the image sequence.

Conventional ECG-gated *spin echo* and *gradient echo* techniques were the first to be applied in the imaging of the heart. However, these techniques are slow and cause degradation in the image resolution. Rapid acquisition and improvement in the image contrast are of growing importance. *Dual Echo*, *Fast Spin Echo* (FSE) and *Turbo Spin Echo* (TSE) are variants of the first *spin echo* techniques. They are normally used for the anatomic delineation of mediastinum and great vessels. *Fast Low-Angle Shot* (FLASH), *Turbo-FLASH*, *Turbo Field Echo* (TFE) and *Echo Planar Imaging* (EPI) are also variants of the first *gradient echo* techniques. They are commonly used for coronary artery imaging, *ventricular* function assessment, myocardial perfusion assessment, valvular motion and for valvular regurgitation assessment. The differences between these techniques and the conventional ECG-gated techniques is that they use lower flip angles.

In order to enable cardiac motion studies to be performed, magnetisation tags within the heart walls are introduced in the images (tagged images) [7, 4, 5]. Tagging is achieved by reducing the magnetisation in limited regions within the cardiac wall. These regions appear different if the image data are acquired before the magnetisation of the perturbed areas. In addition, if the cardiac wall moves between the tagging and imaging times, the magnetisation tag will follow the tissue [4]. The most commonly used magnetisation sequences are: *Spatial Modulation of Magnetisation* (SPAMM) [7, 6, 207], *Complementary Spatial Modulation of Magnetisation* (CSPAMM) [58] and *Delays Alternating with Nutations for Tailored Excitation* (DANTE) [129]. MR imaging can also be used for analysis of the blood flow. In *phase-contrast* MR imaging velocity data is additionally acquired in the form of three different images, V_x, V_y, V_z, corresponding to the Cartesian coordinates of the the velocity vector V [142].

Ultrafast methods such as *Simultaneous Acquisition of Spatial Harmonics* (SMASH) [184] and *sensitivity encoding* (SENSE) [164] have been investigated to reduce the effects of the cardiac motion to the image quality. These methods use combination of signals from an RF coil array to acquire multiple lines of k-space.

MR imaging is also used to access blood perfusion in the myocardium (*perfusion* MR imaging). Blood perfusion in tissues can by studied by using contrast agents. Perfusion MR techniques use both exogenous tracers and endogenous contrast. Chen *et al.* [31] developed *T1 Fast Acquisition Relaxation Mapping* (T1-FARM) to obtain single-slice T1 maps of the heart using exogenous tracers. In techniques using endogenous contrast only indirect measures of the blood flow can be obtained. The endogenous contrast techniques are categorised as *spin-labeling* [205], *magnetisation transfer contrast* [161, 9] and *blood-oxygen dependent contrast* [106].

A recent development, the steady state free precession imaging with balanced gradient (TrueFISP), has been shown to be less susceptible to artifacts caused by slow flow [155, 156]. TrueFISP also provides significantly enhanced blood-myocardial contrast in comparison with the conventional *gradient echo* mechanism.

1.4.5.1 Imaging planes

MR images are acquired in standard imaging planes (coronal, sagittal and transverse). However, these planes are not suitable for the imaging of the LV of the heart because its location can lead to varying obliquity which may cause significant artifacts. When imaging the heart, planes that are oriented parallel or perpendicular to the long-axis of the left *ventricle* are used instead of the standard planes. These planes are called horizontal and vertical long-axis and short-axis views. The short-axis and the vertical and horizontal long-axis planes are shown in figure 1.14. Examples of such planes can been seen in figure 1.15. In this figure, the left image is a short-axis view of the heart, the image in the middle is the horizontal long-axis and the image on the right is the vertical long-axis. The space correspondence between the short-axis and the long-axis is also displayed in the figure.
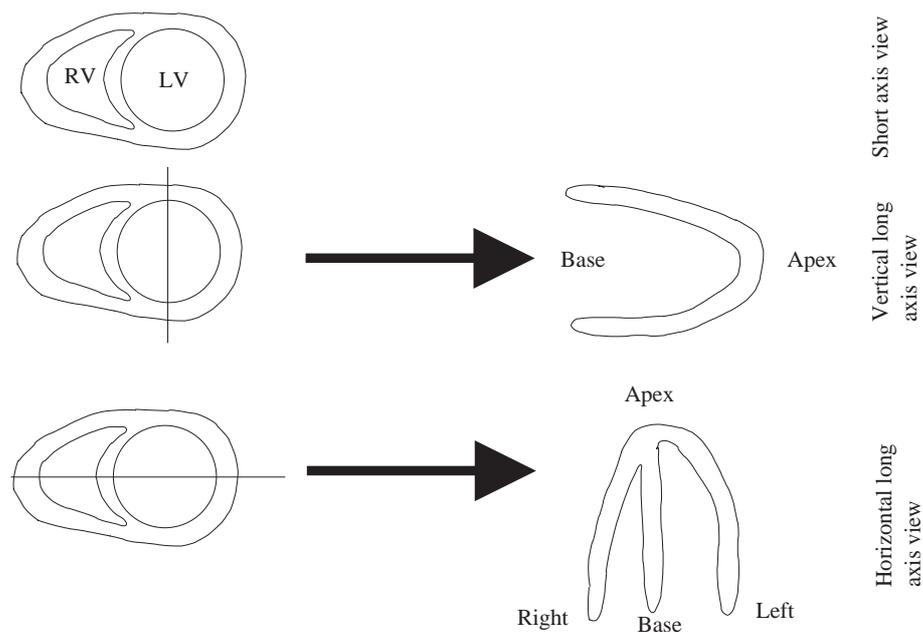
Figure 1.14: The cardiac imaging planes [27].

1.4.6 Comparison of imaging modalities

This chapter has provided a description of the cardiovascular system and its imaging modalities. Each image modality has its own properties and provides different types of

Short axis   Long axis horizontal   Long axis vertical
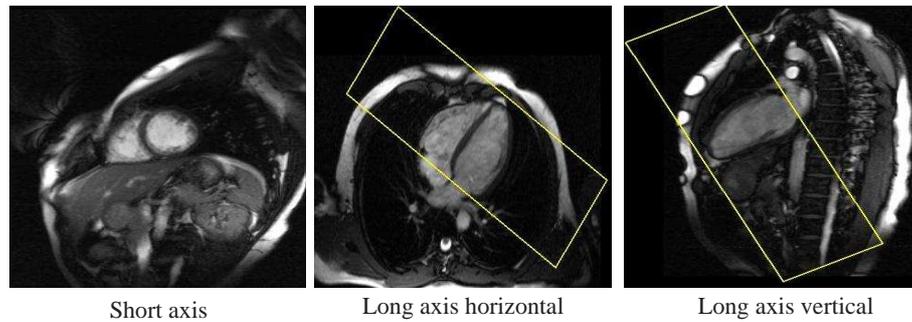
Figure 1.15: Two different MR imaging views of the heart. On the left is the short-axis view, on the middle is the horizontal long-axis view and on the right is the vertical long-axis view. The space correspondence between the short-axis and the long-axis views is also displayed by the rectangular boxes.

information. O 'Dell *et al.* describe a number of criteria for the objective comparison between image modalities [132]. These criteria include the signal quality (indicated by signal to noise ratio (SNR)), the degree of difficulty in distinguishing the myocardium from its neighbouring tissue (indicated by the contrast to noise ratio (CNR)), the temporal and spatial resolution, the susceptibility to image blurring and artifact, the acquisition and analysis time, the relative cost and ease of use. Table 1.2 provides a comparison between the properties of the above imaging modalities.

	Imaging properties					
	X-Ray	**US**	**PET**	**SPECT**	**CT**	**MR**
Invasive	Yes	No	Yes	Yes	No	No
2D/3D	2D	2D/3D	3D	3D	3D	3D
Resolution	$0.1mm$	$\leqslant 1.5mm$	$\leqslant 5mm$	$\leqslant 15mm$	$\leq 1mm$	$\simeq 1mm$ (in x,y) and 5-10mm (in z)
CNR	Low	Low			High	High

Table 1.2: Comparison between different cardiac image modalities.

1.5 Contributions

In recent years, with the development of new sufficient computational methods for the analysis of cardiac MR images, MR imaging has emerged as an important modality for

the imaging and understanding of the cardiovascular system. MR imaging is becoming the modality of choice due to its flexibility and versatility. It is safe, non-invasive and can acquire both 3D and 4D images with high spatial and temporal resolution. Even though the use of MR imaging in clinical practice is rapidly increasing, cardiac MR imaging a new technology primarily limited to specialised imaging centers.

The focus of the research presented in this thesis is the development of spatio-temporal registration methods and methods for modelling the cardiac anatomy and function. The work presented in this thesis makes the following main contributions:

- A new method for the simultaneous spatial and temporal alignment of cardiac MR image sequences to the same coordinate system is presented. Initially, the presented registration algorithm had the ability to correct spatial misalignment of affine nature between the image sequences and also temporal misalignment which could be the result of differences in the length of the cardiac cycles of the subjects and in the temporal acquisition parameters. Then the registration method is extended by the introduction of a *deformable* spatial transformation model which not only corrects global spatial shape differences but also local differences in the cardiac anatomy.

- The registration method for the simultaneous spatial and temporal alignment of cardiac MR image sequences is extended by the introduction of a *deformable* temporal transformation part. Therefore, this spatio-temporal registration method has the ability to correct spatial misalignment between the images caused by global and local shape differences. Furthermore, it has the ability to correct temporal misalignment caused by differences in the length of the cardiac cycles and in the dynamic properties of the hearts.

- Another method for the spatio-temporal alignment of cardiac MR image sequences is presented. As with the earlier method, this method also corrects spatial misalignment between the images caused by global and local shape differences and temporal misalignment caused by differences in the length of the cardiac cycles and in the dynamic properties of the hearts. The major difference compared to the

previous method is that this method optimises the temporal and spatial components separately. The temporal mapping of the image sequences is based on a normalised cross-correlation measure, while the spatial mapping is based on image information from only the first frames of the sequences. This registration method is significantly faster than the previous methods. Moreover, this registration method can be used to enable only the temporal alignment of cardiac MR image sequences, since it does not require the cardiac image sequences to be aligned in the spatial domain in order to find their temporal correspondence.

- A 4D probabilistic atlas describing the cardiac anatomy and function (only in terms of how the cardiac anatomy changes over the cardiac cycle) is developed. The probabilistic atlas contains information regarding the anatomy and function of a healthy heart in the form of tissue probability maps. Separate probabilistic atlases for the *left ventricle*, the *right ventricle* and the *myocardium* have been developed. Modelling the cardiac anatomy function addresses the limitations of current probabilistic atlases of the heart which are limited only to the cardiac anatomy and not the cardiac function.

- A statistical atlas containing information regarding the anatomy of a healthy heart and how the anatomy of the heart changes during the cardiac cycle has been also constructed. In order to build the statistical atlas the distribution space of the cardiac shape is subdivided into two subspaces. The first distribution space describes the changes in the cardiac shape caused by different subjects. The second distribution space describes the changes in the cardiac shape caused by the cardiac cycle. Two separate *principal component analysis* (PCA) have been performed in order to calculate the most significant modes of variation of each subspace. As in the probabilistic atlas, separate statistical atlases for each of the three anatomical structures have been constructed.

- A possible use of the statistical atlases for the classification of cardiac data is demonstrated. The statistical atlases of the myocardium are used for the classifica-

tion of image sequences from normal subjects as well as subjects with hypertrophic cardiomyopathy.

1.6   Overview of the thesis

The dissertation is divided in two main parts and the final conclusions. The first part of the dissertation includes this chapter and two chapters describing the areas of registration of cardiac images and modelling of anatomy. Chapter 2 describes the underlying theories behind image registration and presents an overview of different registration techniques and their application in cardiac imaging. Chapter 3 describes the underlying theories for anatomy modelling and presents an overview of different modelling methods.

The second part presents the contributions of this dissertation. Chapter 4 presents a method for the spatio-temporal alignment of cardiac MR image sequences. The method corrects spatial misalignment of affine nature between the image sequences and also temporal misalignment which may be the result of differences in the length of the cardiac cycles of the subjects. Later in the chapter, a *deformable* spatial transformation model is introduced enabling the correction of local differences in the shape of the hearts as well as global differences. The methods described in chapter 4 have been published in [152], [153]:

- D. Perperidis, A. Rao, R. Mohiaddin, and D. Rueckert. Non-rigid spatio-temporal alignment of 4D cardiac MR images. *In Second International Workshop on Biomedical Image Registration (WBIR '03)*, Lecture Notes in Computer Science, LNCS 2717, pages 191-200, 2003.

- D. Perperidis, A. Rao, M. Lorenzo-Valdés, R. Mohiaddin, and D. Rueckert. Spatio-temporal alignment of 4D cardiac MR images. *In Functional Imaging and Modeling of the Heart (FIMH '03)*, Lecture Notes in Computer Science, LNCS 2674, pages 205-214, 2003.

Chapter 5 presents two registration methods for the *deformable* spatio-temporal alignment of cardiac MR image sequences. These methods correct spatial misalignment between the image sequences caused by global and local shape differences. In addition, they have the ability to correct temporal misalignment caused by differences in the length of the cardiac cycles and in the dynamic properties of the hearts. The methods described in chapter 5 have been published in [151], [150], [147], [148]:

- D. Perperidis, R. Mohiaddin, and D. Rueckert. Spatio-temporal free-form registration of cardiac MR image sequences. *In Medical Image Analysis*, 9(5):441-456, 2005.

- D. Perperidis, R. Mohiaddin, and D. Rueckert. Fast spatio-temporal registration of cardiac MR image sequences. *In Functional Imaging and Modeling of the Heart (FIMH'05)*, Lecture Notes in Computer Science, LNCS 3504, pages 414-424, 2005.

- D. Perperidis, R. Mohiaddin, and D. Rueckert. Spatio-temporal free-form registration of cardiac MR image sequences. *In Seventh Int. Conf. on Medical Image Computing and Computer- Assisted Intervention (MICCAI '04)*, Lecture Notes in Computer Science, LNCS 3217, pages 911-919, 2004.

- D. Perperidis, R. Mohiaddin, and D. Rueckert. Spatio-temporal free-form registration of cardiac MR image sequences. *In Proc. of Medical Image Understanding and Analysis '04*, pages 157- 160, 2004.

Chapter 6 presents a method for building a probabilistic atlas of the cardiac anatomy and function as well as a method for building a 4D statistical atlas of the cardiac anatomy. The methods described in chapter 6 have been published in [149], [145], [146]:

- D. Perperidis, R. Mohiaddin, and D. Rueckert. Construction of a 4D statistical atlas of the cardiac anatomy and its use in classification. *In Eight Int. Conf. on Medical Image Computing and Computer- Assisted Intervention (MICCAI '05)*, Lecture Notes in Computer Science, 2005.

- D. Perperidis, M. Lorenzo-Valdés, R. Chandrashekara, A. Rao, R. Mohiaddin, G. I. Sanchez- Ortiz, and D. Rueckert. Building a 4D atlas of the cardiac anatomy and motion using MR imaging. *In IEEE International Symposium on Biomedical Imaging*, pages 412-415, 2004.

- D. Perperidis, M. Lorenzo-Valdés, R. Chandrashekara, A. Rao, R. Mohiaddin, G.I. Sanchez- Ortiz, and D. Rueckert. Building a 4D atlas of the cardiac anatomy and motion using MR imaging. *In Proc. of Medical Image Understanding and Analysis '04*, pages 9-12, 2004.

Chapter 7 discusses the presented algorithms, their limitations and the conclusions which can be drawn from this work and proposes potential directions for future work in this area.

Chapter 2

Image registration

Image registration is an active research area in computer vision, image processing and medical image processing. Reviews on image registration techniques can be found in [21, 208], on medical image registration in [77, 104] and on cardiac image registration in [116]. This chapter provides a description of medical image registration techniques and in particular cardiac image registration techniques.

2.1 Image registration

The goal of image registration is to calculate a mapping, $\mathbf{T}$, which relates each point of one image I to a corresponding anatomical point in the reference image I'. Figure 2.1 provides an example of the image registration. Images usually have two dimensions $\mathbf{x} = (x, y)$, three dimensions $\mathbf{x} = (x, y, z)$ or four dimensions $\mathbf{x} = (\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{t})$. If the images can be defined as arrays (two, three or four dimensional) of a given size denoted by I' and I, then the mapping between them can be expressed as ([21]):

$$I'(\mathbf{x}) = g(I(\mathbf{T}(\mathbf{x}))) \tag{2.1}$$

where $\mathbf{x}$ is a vector of the point's location, $I'(\mathbf{x})$ and $I(\mathbf{x})$ is the intensity of image I' and I in position $\mathbf{x}$, g is the intensity transformation and $\mathbf{T}$ is the coordinate transformation.

The transformation $\mathbf{T}$ can have several forms. In this thesis the following types of
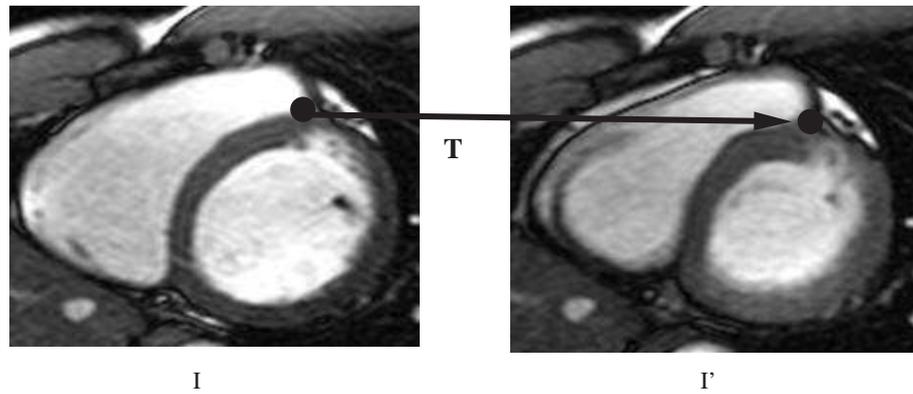
Figure 2.1: A schematic figure of the image registration.

transformation are considered: rigid, affine and non-rigid transformations. A rigid transformation allows only rigid operations to be performed (rotation and translation) i.e. operations that maintain distance between points. An affine transformation (rotation, translation, scaling and shearing) maps parallel lines to parallel lines. A non-rigid transformation allows the object additionally to deform. Finding the best transformation, $\mathbf{T}$, can be broken into the following tasks [21]:

- Selection of the feature space

- Selection of the similarity measure

- Selection of the type of the transformation

- Selection of the optimisation strategy

Maintz *et al.* described a classification approach for registration methods [114]. The main classification criteria are the following:

1. Dimensionality of the images: 1D/1D, 2D/2D, 2D/3D, 3D/3D, 4D/4D, etc.

2. Feature space

    (a) Extrinsic

    (b) Intrinsic

    (c) Non-Image Based

3. Nature of the transformation

   (a) Rigid

   (b) Affine

   (c) Non-rigid

4. Domain of the transformation

5. Interaction

6. Optimisation procedure

7. Modalities involved

   (a) Monomodal

   (b) Multimodal

8. Subject

   (a) Intra-subject

   (b) Inter-subject

   (c) Atlas

Monomodal registration methods use images acquired by the same imaging modality (for example MR, SPECT, CT) while *multimodal* registration methods combine images acquired by two different types of modalities (for example MR/SPECT). *Intra-subject* registration approaches register images from the same subject (for example before and after treatment), while *inter-subject* registration approaches register images from two different subjects. *Intersubject* registration methods are frequently used in the construction of models of anatomical structures. *Atlas based* registration approaches register an atlas to a specific subject or a medical atlas to another atlas. *Atlas based* registration methods are frequently used for the segmentation of medical images.

2.2   Feature space of the registration

One of the first steps when registering two images is to decide upon the feature space to use during the approach. This may be the image itself (i.e. intensity values), other features of the images (edges, contours, surfaces, etc.) and foreign objects inserted to the image space. The registration methods can be classified according to their feature space as:

- *Extrinsic* registration methods. These methods calculate the optimal transformation using information provided by foreign objects which are introduced into the image space. For example medical imaging methods use artificial fiducials attached to the patients. The fiducials need to be well visible in order to be very accurately detectable. The most commonly used fiducials in brain imaging is a *stereotactic frame* which is screwed rigidly to the patient's outer skull. The stereotactic frame is considered as the "gold-standard" for registration accuracy. Other artificial markers include objects screwed to the bones. The advantage of *extrinsic* registration methods is that is easy and fast to calculate the image transformation mapping. However, the quality of the transformation mapping will always depend on the correct placing of the markers. Furthermore, in cardiac imaging it is impossible to use *extrinsic* registration methods since invasive markers cannot be placed on the heart.

- *Intrinsic* registration methods. These methods calculate the optimal transformation using information contained within the images. In these methods the registration procedure can be based on identified salient features (landmarks and edges), on segmented structures, on extracted surfaces or directly on measures computed from the intensities of the images. Calculating salient features and obtaining segmented structures can be done manually or by using automated approaches. Manual land-marking and segmenting images is a very time consuming task prone to errors. A number of approaches have been developed which use image registration methods to automatically segment cardiac images [109] and to automatically identify landmarks in cardiac images [63].

- *Non-image based* registration methods. These methods calculate the image mapping without using image information. For example in cases where the coordinate systems of two images are calibrated with respect to each other, the mapping can be found without using image information.

2.3 Transformations

As mentioned above, the transformation model can have a large number of forms depending on the specific application. The transformation model imposes mathematical constraints on the type of geometric distortions that can be imposed during the registration procedure. The number of parameters needed to describe a transformation are often called *degrees of freedom*. The number of *degrees of freedom* depends on the type of transformation type and also on the dimensionality of the images. For example if the transformation utilises only translation and the images are 3D, then the transformation will have three *degrees of freedom* (translation along the x-axis, y-axis and z-axis). The domain of the transformation is *global* if the transformation applies to the entire image and *local* if the transformation applies to a part of the image.

2.3.1 Linear transformations

Using homogeneous coordinates a linear transformation can be expressed as matrix multiplications. In 3D, a linear transformation has the following form [206]:

$$\mathbf{T}_{linear}(x', y', z') = \begin{pmatrix} x' \\ y' \\ z' \\ 1 \end{pmatrix} = \begin{pmatrix} a_{01} & a_{02} & a_{03} & a_x \\ a_{11} & a_{12} & a_{13} & a_y \\ a_{21} & a_{22} & a_{23} & a_z \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix} \tag{2.2}$$

The linear transformation is defined by the parameters a_{ij}.

2.3.1.1 Rigid transformation

The rigid body model is the most constrained transformation model in the area of medical image registration. It ensures that distances and internal angles within an image structure will not change during the registration. As the name suggests, the rigid body model allows only operations which maintain the shape and size of the anatomical structures, i.e. translation and/or rotation. The rigid body transformation model in 3D has 6 *degrees of freedom* and has the following form:

$$
\mathbf{T}_{rigid}(x', y', z') = \begin{pmatrix} x' \\ y' \\ z' \\ 1 \end{pmatrix} = \begin{pmatrix} r_{01} & r_{02} & r_{03} & t_x \\ r_{11} & r_{12} & r_{13} & t_y \\ r_{21} & r_{22} & r_{23} & t_z \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix}
\tag{2.3}
$$

The parameters r form a 3×3 matrix which rotates the image around in the x-, y- and z-axis while the parameters t translate the image along the x-, y- and z-axis. A transformation utilising only translation can be expressed as:

$$
\mathbf{T}_{translation}(x', y', z') = \begin{pmatrix} x' \\ y' \\ z' \\ 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & t_x \\ 0 & 1 & 0 & t_y \\ 0 & 0 & 1 & t_z \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix}
\tag{2.4}
$$

Furthermore using the same matrix notation, a rotation around the x-axis can be expressed as:

$$
\mathbf{T}_{rotation}(x', y', z') = \begin{pmatrix} x' \\ y' \\ z' \\ 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & cos\theta & sin\theta & 0 \\ 0 & -sin\theta & cos\theta & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix}
\tag{2.5}
$$

where the parameter θ is the rotation angle. Similarly, rotations around the y- and z-axis
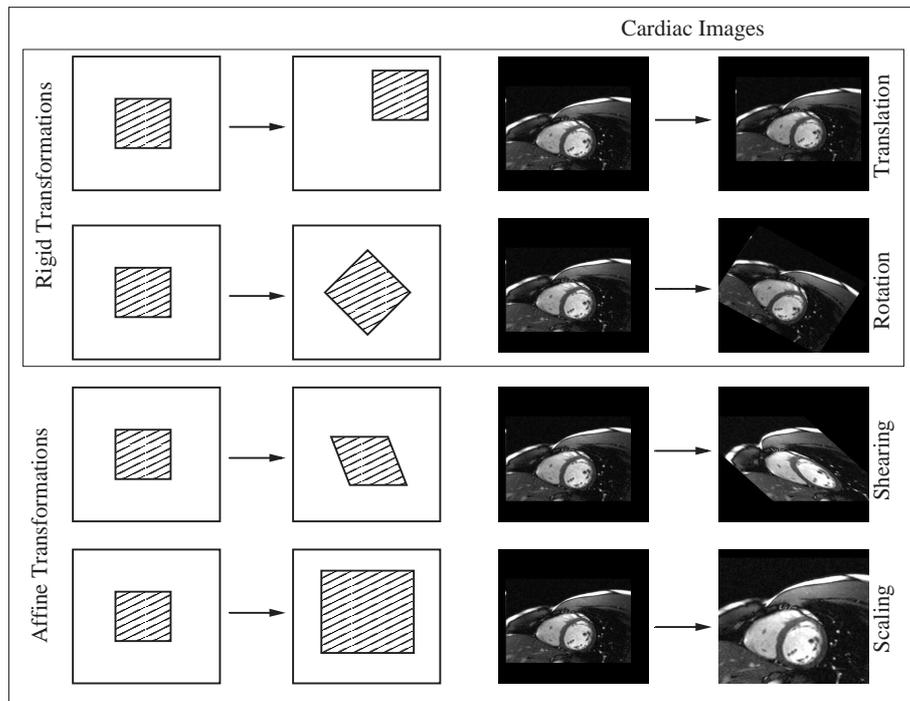
Figure 2.2: Examples of 2D linear transformations.

can be expressed in matrix form.

2.3.1.2 Affine transformations

A more general class of transformations is the affine transformation. The affine transformation model allows rotation, translation, scaling and shearing. In 3D, it has the following form:

$$
\mathbf{T}_{affine}(x', y', z') = \begin{pmatrix} x' \\ y' \\ z' \\ 1 \end{pmatrix} = \begin{pmatrix} \theta_{11} & \theta_{12} & \theta_{13} & \theta_{14} \\ \theta_{21} & \theta_{22} & \theta_{23} & \theta_{24} \\ \theta_{31} & \theta_{32} & \theta_{33} & \theta_{34} \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix} \tag{2.6}
$$

where the coefficients θ parameterise the twelve *degrees of freedom* of the transformation. The affine transformation model ensures that lines which are parallel before the transformation will remain parallel after the transformation. Figure 2.2 provides examples of 2D linear transformations.
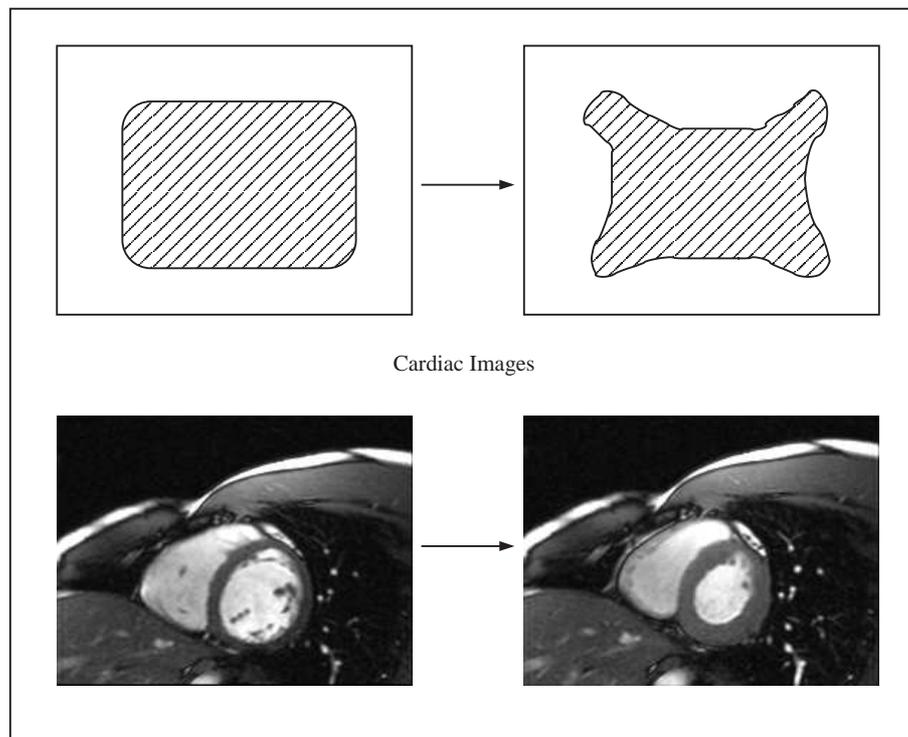
Figure 2.3: Example of a 2D non-linear transformation.

2.3.2 Non-linear transformations

The above transformations (rigid and affine) preserve straightness of lines but they cannot model the change in the shape of an object which has undergone local deformation. For example several applications require a transformation which can accommodate tissue deformation or the substantial anatomical variability across individual subjects (e.g. the variability in the shape of a healthy heart). Linear transformations preserve straightness of lines while non-linear transformations deform the image structures. Hence, non-rigid transformations are suitable for this kind of applications because they allow the objects to deform. Figure 2.3 provides an example of a 2D non-linear transformation.

Unlike the rigid and affine transformations, non-rigid transformations are still subject of ongoing research activity. An overview on hierarchical non-linear registration approaches for medical images can be found in [104] while an overview of common registration approaches is given in [74].

2.3.2.1 Spline based transformations

The term splines originated in engineering. In order to model ships and planes, the engineers used to employ long flexible strips of wood or metal, the splines. The strips were bent to the desired shape by applying a set of weights at particular places. This is the main idea behind the use of splines to model an image transformation. In spline based transformations, the applied weights correspond to the displacement of a particular point.

A detailed description of the usage of splines in signal and image processing can be found in [196]. Many registration techniques which use splines based transformations assume that a set of corresponding landmarks, the *control points* of the spline based transformation, can be identified in both images. The transformation either interpolates or approximates the displacement required to map each control point of the reference image to the corresponding control point in the other image. Furthermore, a spline based transformation provides a smooth varying displacement field between these control points. The interpolation condition can be written as:

$$\mathbf{T}(\phi_i) = \phi_i' \qquad i = 1, ..., n \tag{2.7}$$

where ϕ_i is the location of the control point in the reference image and ϕ_i' is the location of the corresponding control point in the other image. There is a large number of ways to determine the control points. For example, anatomical landmarks which can be identified in both images can be used to define the spline based transformation. An alternative approach is to identify the control points to have equal spacing along the two images forming a rectangular mesh [45]. In such case, the control points are referred as *pseudo-landmarks* and they are only used as a parameterisation of the transformation.

Thin-plate splines are a family of splines based on radial-basis functions. They have been formulated by Duchon [52] and Meinguet [124] for the surface interpolation of scattered data and they have also been used in a large number of image registration approaches

[68, 15, 16]. They can be defined as a combination of n radial functions $\theta(s)$:

$$t(x, y, z) = \alpha_1 + \alpha_2 x + \alpha_3 y + \alpha_4 z + \sum_{j=1}^{n} b_j \theta(|\phi_j - (x, y, z)|) \qquad (2.8)$$

The transformation is defined as three separate splines, $\mathbf{T} = (t_1, t_2, t_3)^T$. The coefficients α characterise the affine part of the transformation while the coefficients b characterise the non-affine part. There are $3n$ interpolation equations (equation 2.7). In order to determine $3(n+4)$ coefficients uniquely, 12 additional equations are required. These equations must guarantee that the sum of the coefficients b is equal to zero and that their cross-product is equal to zero as well. This can be expressed with the following matrix form:

$$\begin{pmatrix} \boldsymbol{\Theta} & \boldsymbol{\Phi} \\ \boldsymbol{\Phi}^{\mathbf{T}} & 0 \end{pmatrix} \begin{pmatrix} \mathbf{b} \\ \alpha \end{pmatrix} = \begin{pmatrix} \boldsymbol{\Phi}' \\ 0 \end{pmatrix} \qquad (2.9)$$

In the above equation, b is a $n \times 3$ matrix of non-affine coefficients, α is a 4×3 vector of affine coefficients and $\boldsymbol{\Theta}$ is the kernel matrix with $\Theta_{ij} = (|\phi_i - \phi_j|)$. The solution for α and b is a thin-plate spline transformation which interpolates the displacement at the control points. The radial basis function of a thin-plate spline is defined as:

$$\theta(s) = \begin{cases} |s|^2 \log(|s|), & \text{in 2D} \\ |s|, & \text{in 3D} \end{cases} \qquad (2.10)$$

2.3.2.2   Free-form deformations

Radial basis functions have infinite support. Therefore, each control point has a global influence on the entire transformation. This is undesirable in cases where local deformations need to be modelled. Furthermore, the computational complexity of a thin-plate spline is high and not efficient.

Free Form Deformations (FFDs) have been introduced by Sederberg and Parry [179] and are used to model local deformations. This approach requires a rectangular mesh of control points with uniform spacing to be placed on the image. An FFD deforms an

object by manipulating the underlying mesh of control points. To define an FFD the spatial domain of the image volume is denoted as $\Omega_I = \{(x, y, z) \mid 0 \leq x < X, 0 \leq y < Y, 0 \leq z < Z\}$. Let Φ denote a $n_x \times n_y \times n_z$ mesh of control points $\phi_{i,j,k}$ with uniform spacing δ. Then, the displacement field **u** can be evaluated by the trivariate tensor product of Bernstein polynomials:

$$\mathbf{u}(x, y, z) = \sum_{i=0}^{l} \sum_{j=0}^{m} \sum_{k=0}^{n} B_l^i B_m^j B_n^k (1-x)^{l-i} x^i (1-y)^{m-j} y^j (1-z)^{n-k} z^k \phi_{i,j,k} \quad (2.11)$$

where $\phi_{i,j,k}$ is a vector containing the Cartesian coordinates of the control points. The polynomials are defined by $B_j^i = \binom{i}{l} x^i (1-x)^{i-l}$ where $\binom{i}{l}$ is a binomial coefficient (figure 2.4):

$$
\begin{aligned}
B_{0,0}(x) &= 1 \\
B_{0,1}(x) &= (1-x) \\
B_{1,1}(x) &= x \\
B_{0,2}(x) &= (1-x)^2 \\
B_{1,2}(x) &= 2(1-x)x \\
B_{2,2}(x) &= x^2 \\
B_{0,3}(x) &= (1-x)^3 \\
B_{1,3}(x) &= 3(1-x)^2 x \\
B_{2,3}(x) &= 3(1-x)^2 \\
B_{3,3}(x) &= x^3
\end{aligned}
$$

An FFD could also be formulated as the 3D tensor product of B-Splines instead of the non-tensor product of Bernstein polynomials [179]. The most commonly used B-Spline function is the cubic B-Spline [99, 100]:

$$\mathbf{T}_{local}(x, y, z) = \sum_{l=0}^{3} \sum_{m=0}^{3} \sum_{n=0}^{3} B_l(u) B_m(v) B_n(w) \phi_{i+l,j+m,k+n} \quad (2.12)$$
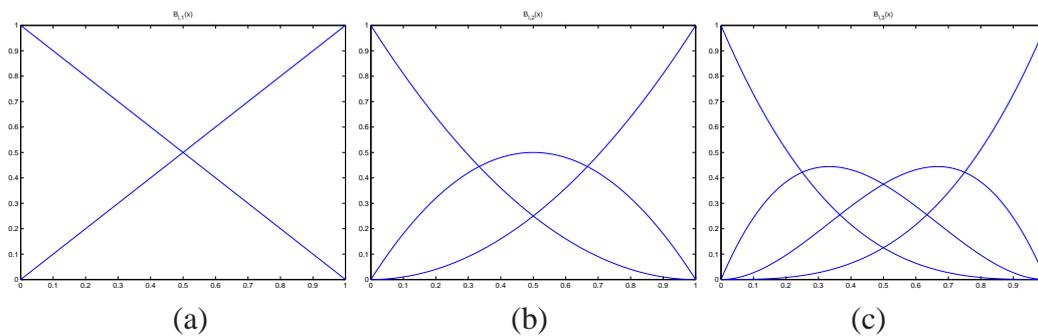
Figure 2.4: Graphical representation of some of the first Bernstein polynomials
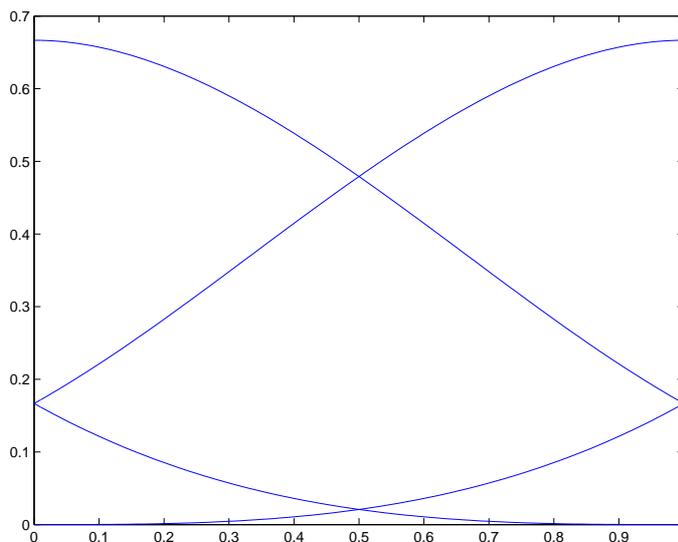
Figure 2.5: Graphical representation of B-Splines.

where $i = \lfloor \frac{x}{n_x} \rfloor - 1, j = \lfloor \frac{y}{n_y} \rfloor - 1, k = \lfloor \frac{z}{n_z} \rfloor - 1, u = \frac{x}{n_x} - \lfloor \frac{x}{n_x} \rfloor, v = \frac{y}{n_y} - \lfloor \frac{y}{n_y} \rfloor, w = \frac{z}{n_z} - \lfloor \frac{z}{n_z} \rfloor$ and where B_l represents the l-th basis function of the B-Spline (figure 2.5):

$$B_0(u) = (1 - u)^3/6$$

$$B_1(u) = (3u^3 - 6u^2 + 4)/6$$

$$B_2(u) = (-3u^3 + 3u^2 + 3u + 1)/6$$

$$B_3(u) = u^3/6$$

B-Splines are locally controlled which makes them computationally efficient even for a large number of control points. In particular, the basis functions of cubic B-Splines have

a limited support, i.e. changing a control point $\phi_{i,j,k}$ affects the transformation only in the local neighborhood of that control point.

Rueckert *et al.* introduced a transformation model which consists of a global part and a local part [174]:

$$\mathbf{T}(x', y', z') = \mathbf{T}_{global}(x, y, z) + \mathbf{T}_{local}(x, y, z) \tag{2.13}$$

The global transformation, $\mathbf{T}_{global}$, describes the overall differences between the two subjects (i.e. differences due to translation, rotation, shearing and scaling) [174]. An affine transformation (equation 2.6) was used as a global part. The local differences in the shape of the two subjects are modelled by $\mathbf{T}_{local}$ which is an FFD based on cubic B-Splines (equation 2.12).

An example of the FFD can be seen in figure 2.6. In figure 2.6, (a) is the identity FFD, (b) is a deformed FFD and (c) is the same deformed FFD with smaller control spacing δ.

Denton *et al.* [48] compared and evaluated a number of registration methods for breast MR images based on rigid, affine and FFD transformations. The study contained 54 MR breast scans from which 27 MR scans were reported as normal and 27 with an abnormality. The images were registered using the above registration methods and the results were assessed by two experienced radiologists. Their results showed significant improvement on the registration of the data with the FFD model compared to other two models. They also showed that there was no significant improvement when using the affine model instead of the rigid model.

2.3.3 Physical models of deformation

A large amount of research interest has been focused on non-rigid registration techniques based on FFDs. In addition to B-Splines based registration techniques, there is significant research interest on non-rigid registrations methods which use elastic and fluid transformation models.

(a)   (b)   (c)
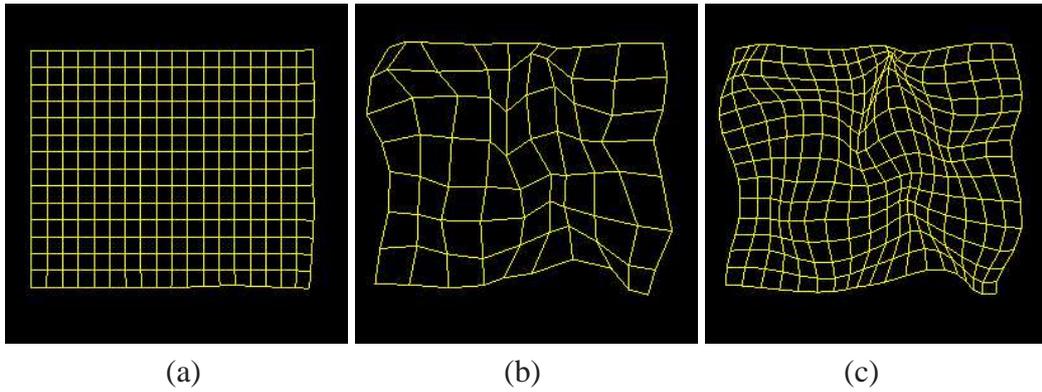
Figure 2.6: Examples of FFDs. (a) shows the initial FFD (identity), (b) shows how the FFD deforms after the registration of two subjects and (c) shows the same FFD with smaller control point spacing.

2.3.3.1   Elastic transformation

The use of elastic transformation models for image registration was first proposed by Broit [20] and has been extended by Bajcsy and Kovacic [8, 74]. They have been initially used for matching a brain atlas with a CT image of a human subject. The underlying idea behind the elastic transformation models is to represent the deformation of an image into the reference image by a process which is similar to stretching an elastic material like rubber. This process is driven by two separate forces, the *internal* and the *external* forces. The *internal* force models any forces which deform the elastic body from its equilibrium state. The *external* force models any forces acting on the elastic body. The deformation of the elastic body stops only when the two acting forces form an equilibrium solution. The behavior of the elastic body can be described by the Navier linear elastic partial differential equation:

$$\mu\nabla^2\mathbf{u}(x,y,z) + (\lambda + \mu)\nabla(\nabla \cdot \mathbf{u}(x,y,z)) + \mathbf{f}(x,y,z) = 0 \qquad (2.14)$$

where $\mathbf{u}$ describes the displacement field, $\mathbf{f}$ is the external force, ∇ is the gradient operator while ∇^2 denotes the Laplace operator. The behavior of the elastic body is defined by the constants λ and μ (Lame's elasticity constants). These constants are often interpreted in terms of Young's modulus E_1, which relates the strain and the stress of the object,

and the Poisson's ratio E_2, which is the ratio between lateral shrinking and longitudinal stretching:

$$E_1 = \frac{3\lambda + 2\mu}{\lambda + \mu} \qquad E_2 = \frac{\lambda}{2(\mu + \lambda)} \qquad (2.15)$$

The equation 2.14 can be solved by finite differences and successive over relaxation (SOR) which yields a discrete displacement for each voxel. Davatzikos [43] proposed an extension of this elastic registration in which certain anatomical structures are allowed to deform more freely than others.

2.3.3.2 Viscous fluid transformation

One disadvantage of the elastic transformations is that they cannot model large localised deformations since the deformation energy increases proportionally with the stress of the deformation. Viscous fluid transformation models can handle a large amount of deformation while preventing any folding in the resulting deformation fields. These transformation models are popular in intersubject registration tasks where large shape variability exists. However, they have a large number of *degrees of freedom* which makes it more difficult to find the optimal transformation and the scope for misregistration becomes larger.

Fluid deformations are described in a *Eulerian* reference frame, i.e. with respect to their final position, while elastic deformations are described with respect to their initial position, the *Lagrangian* reference frame. In the *Eulerian* reference frame, the deformation of fluid is characterised by the Navier-Stokes differential equation:

$$\mu \nabla^2 \mathbf{v}(x, y, z) + (\lambda + \mu)\nabla(\nabla \cdot \mathbf{v}(x, y, z)) + \mathbf{f}(x, y, z) = 0 \qquad (2.16)$$

Equation 2.16 is similar to equation 2.14 except that the differentiation is carried out on the velocity field $\mathbf{v}$ rather than on the displacement field $\mathbf{u}$ and it is solved for each time step. The relationship between the displacement field and the Eulerian velocity is given by:

$$\mathbf{v}(x, y, z, t) = \frac{\partial \mathbf{u}(x, y, z, t)}{\partial t} + \mathbf{v}(x, y, z, t) \cdot \nabla \mathbf{u}(x, y, z, t) \qquad (2.17)$$

Christensen *et al.* [32] proposed to solve equation 2.16 by using a successive over relaxation method. However, this approach is slow and requires significant processing time. A faster approach has been suggested by Bro-Nielsen *et al.* [19]. In this approach the equation 2.16 is solved by deriving a convolution filter from the eigenfunctions of the linear elasticity operator which is similar to a regularisation by a convolution with a Gaussian filter [19].

2.4 Similarity measures

The similarity measure provides a way to assess how similar two images are. This section describes the basic similarity measures which have been widely used in image registration.

2.4.1 Point based methods

Point registration methods rely on the identification of corresponding sets of points in the two images $\mathbf{x}_i : i \in 0, 1..N$ and $\mathbf{y}_i : i \in 0, 1..N$. The sets of landmarks of $\mathbf{x}$ and $\mathbf{y}$ could be *extrinsic* (foreign objects which have been introduced to the image space) or *intrinsic* (anatomical features). The identification of the point sets can be manual or automatic. Alignment is achieved by minimising the distance between the corresponding landmarks of $\mathbf{x}$ and the landmark of set $\mathbf{y}$:

$$D = \sum_i w_i^2 \|\mathbf{x}_i - \mathbf{T}(\mathbf{y}_i)\|^2 \tag{2.18}$$

where $\mathbf{T}(\mathbf{y}_i)$ is the transformed landmark $\mathbf{y}_i$ and w_i^2 are weight coefficients that measure the degree of confidence in which the points features have been located. For rigid transformations a least square fitting method can be used [3]. Furthermore, methods exist for spline based transformations and for affine transformations (with isotropic and non-isotropic scaling) [59].

Landmark based registration of the heart is a very difficult task because there are

few spatial anatomical landmarks which could be accurately identified. Moreover, these landmarks can be less visible in certain modalities and as well as under a number of pathological conditions.

2.4.2   Surface based methods

Surface based registration methods require the identification and extraction of surface features in the images. In these methods registration is achieved by finding the transformation, $\mathbf{T}$, which aligns corresponding points of surface features present in the images. Surfaces tend to be more distinct than landmarks and can be accurately identified by a number of segmentation methods. The *marching cubes* algorithm [107] is often used to extract contours from images and to generate a triangulation of the surface. Schroeder *et al.* presented an approach for the decimation of the triangle meshes [178]. This approach is useful in cases where the generated surfaces contain a large number of vertices. Pelizzari and colleagues [143, 105] proposed a method for surface based registration using the *head and hat* algorithm. In this method two surfaces are identified in the two images. The high resolution surface is represented as a stack of disks (*head*). The second surface is represented by a list of unconnected 3D points (*hat*). The optimal rigid transformation is calculated as the one that minimises the square distance between each point of the *hat* and the closest point of the head surface towards the direction of the *head's* centroid.

The *iterative closest point algorithm* (ICP) is another surface based registration method which finds the optimal transformation between two surfaces [12]. In ICP the surfaces being registered are represented as point sets: one surface point set is kept fixed during the registration while the other surface point set is moving. The optimal transformation is found using an iterative approach. In each iteration, the algorithm establishes point correspondence by identifying for each point of the moving surface the closest point in the fixed surface. Then, the two surfaces are registered (using a point based registration method) and the transformation is applied to the moving surface point set.

2.4.3 Intensity based methods

2.4.3.1 Cross-correlation

Cross-correlation is one of the first intensity based similarity methods used on image registration methods [170]. The *cross-correlation* measure of two images is defined by [162]:

$$R = \sum_{x \in \Omega} I'(\mathbf{x}) \cdot I(\mathbf{T}(\mathbf{x})) \tag{2.19}$$

where Ω is the spatial domain of overlap of the images I' and I and $\mathbf{x}$ is a column vector of a particular image position (e.g. in 3D case $\mathbf{x} = (x, y, z)^T$). As this measure stands it is affected by changes in the image contrast and brightness. These are linear intensity transformations of the type:

$$I'(\mathbf{x}) = \alpha \cdot I(\mathbf{x}) + \beta \tag{2.20}$$

In order to make the cross-correlation insensitive to contrast changes it can be normalised resulting in the *normalised cross-correlation*:

$$R = \frac{\sum_{\mathbf{x} \in \Omega}(I'(\mathbf{x}) - \bar{I}') \cdot (I(\mathbf{T}(\mathbf{x})) - \bar{I})}{\sqrt{\sum_{\mathbf{x} \in \Omega}(I'(\mathbf{x}) - \bar{I}')^2} \cdot \sqrt{\sum_{\mathbf{x} \in \Omega}(I(\mathbf{T}(\mathbf{x})) - \bar{I})^2}} \tag{2.21}$$

where $\bar{I}'$ is the average intensity of the reference image I' and $\bar{I}$ is the average intensity of the source image I. Equation 2.21 will have a maximum value when the two images are the same.

2.4.3.2 Sum of squared differences

The idea behind the use of the *sum of squared differences* (SSD) or the *Euclidean distance* as a similarity measure is simple. If the two images are the same then the sum of the squared differences between their pixels' intensity will be equal to zero. In this case the optimal transformation, $\mathbf{T}$, is found by minimising:

$$E^2 = \sum_{\mathbf{x} \in \Omega}(I'(\mathbf{x}) - I(\mathbf{T}(\mathbf{x})))^2 \tag{2.22}$$

The above equation is expanded to:

$$E^2 = \sum_{\mathbf{x}\in\Omega}(I'(\mathbf{x}))^2 + \sum_{\mathbf{x}\in\Omega}(I(\mathbf{T}(\mathbf{x}))^2 - 2(\sum_{\mathbf{x}\in\Omega}(I'(\mathbf{x}) \cdot I(\mathbf{T}(\mathbf{x})))) \qquad (2.23)$$

The first term of equation 2.23 is independent of $\mathbf{T}$ and can be eliminated since the goal is the minimisation of E with respect to $\mathbf{T}$. This means that in order to minimise E the third term of the equation 2.23 has to be maximised. However, the last term of the equation is the cross-correlation of the images (equation 2.19). Therefore, minimising the sum of the squared differences is similar to maximising the cross-correlation of the images. The SSD is the optimum similarity measure when the images differ only by Gaussian noise [199, 59].

Another similarity measure is the *sum of absolute differences* (SAD). In this case the absolute intensity difference is calculated instead of the squared intensity difference. Hoh *et al.* compared the two methods by simulating various defects and misalignments using cardiac PET images [80]. No significant differences in the resulting errors, obtained by the two similarity measures, in translation and rotation were found [80].

Similar to equation 2.19, equation 2.22 is also affected by contrast and brightness changes. It can be normalised in the same ways as the cross-correlation, resulting in:

$$E^2 = \frac{\sum_{\mathbf{x}\in\Omega}(I'(\mathbf{x}) - I(\mathbf{T}(\mathbf{x}))^2)}{\sqrt{\sum_{\mathbf{x}\in\Omega}(I'(\mathbf{x}))^2} \cdot \sqrt{\sum_{\mathbf{x}\in\Omega}(I(\mathbf{T}(\mathbf{x})))^2}} \qquad (2.24)$$

2.4.3.3   Entropy and mutual information

Entropy based similarity measures do not depend on a specific functional relationship between the intensity distributions in the images, making them particularly suitable for *multimodal* registration methods. The *entropy* similarity measure is closely related to predictability. A predictable random variable has low entropy while a totally unpredictable random variable has high entropy [198]. The average information provided by the intensity values of an image can be measured by the *marginal entropy*. The *marginal entropy*
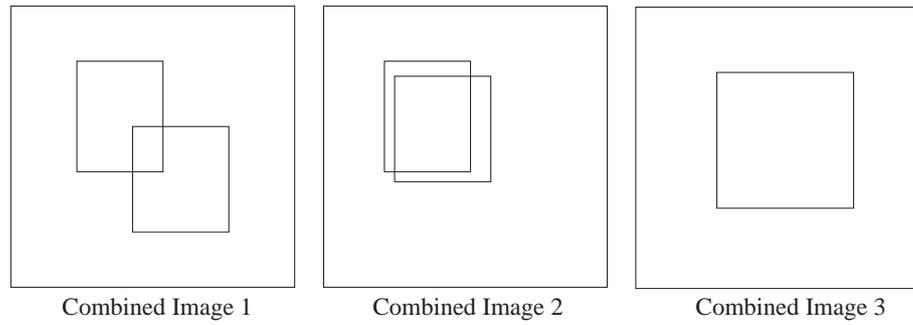
Figure 2.7: A schematic figure of the entropy of the combined image.

of a pair of images I and I' is given by [186]:

$$H(I) = -\sum_{i \in I} p\{i\} log(p\{i\}) \tag{2.25}$$

$$H(I') = -\sum_{i' \in I'} p\{i'\} log(p\{i'\}) \tag{2.26}$$

where $p\{i\}$ and $p\{i'\}$ are the marginal probabilities of I and I' (i.e. the probability of a voxel of image I to have a value i).

When the transformation, mapping corresponding features, combines two images, then the information content of this combined view is given by the *joint entropy* of the images [186]:

$$H(I', I) = -\sum_{i' \in I'} \sum_{i \in I} p\{i', i\} log(p\{i', i\}) \tag{2.27}$$

where $p\{i', i\}$ is the joint distribution probability of I' and I .

The combined image can be considered as an image in which pairs of intensity values occur together, where each pair of values corresponds to two values occurring in the two images. The concept of the image correlation using the *entropy* as a similarity measure is simple. The information which the combined image provides will be minimum only when the two images are registered correctly. Figure 2.7 shows that when the images are correctly registered (combined image 3), there are no duplicated regions and hence the amount of information provided by the combined image is minimum.
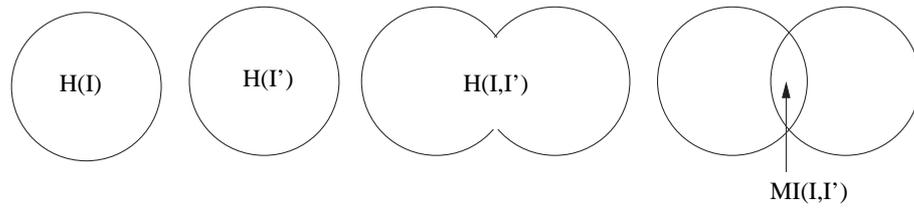
Figure 2.8: Venn diagrams illustrating the relationship between *entropy* and *mutual information* (adapted from [186]).

The *entropy* of the combined image will be low under two conditions: when the images is predictable from the model and when the image is by itself predictable [199]. There are a number of problems using the *entropy* as a similarity measure. When the images have a limited field of view, then any measure of information content of the combined image will be a function of the extent to which the fields of view overlap.

A commonly used similarity measure in medical image registration is the *Mutual Information* (MI) [199, 198, 113, 186]. *Mutual information* was initially used in information theory as a measure of information between the transmitter and the receiver at either end of a communication channel. It relates the changes in the value of the *joint entropy*, $H(I, I')$, back to the *marginal entropies*, $H(I')$ and $H(I)$, of the two images :

$$MI(I'; I) = H(I') + H(I) - H(I', I) \tag{2.28}$$

If the *entropies* $H(I')$, $H(I)$ and $H(I', I)$ are substituted by equations 2.25, 2.26 and 2.27, the following equation can be derived [113]:

$$MI(I'; I) = \sum_{i' \in I} \sum_{i \in I} p\{i', i\} log \frac{p\{i', i\}}{p\{i'\}p\{i\}} \tag{2.29}$$

The relationship between the *mutual information* and the joint entropy is provided by the Venn diagrams of figure 2.8 [186]. There are two ways to calculate the joint probability distribution values $p\{i, i'\}$. One approach is to form a continuous mathematical estimate of the distribution by, for example, fitting a function to the data. The most commonly used continuous approach is using a Parzen Window [51]. Alternatively, the joint

distribution probability values can be estimated by using a joint discrete histogram of the images' intensity values. The second approach is computationally less expensive since no functional fit is required [79].

The *mutual information* has the following properties [113]:

- **Non negativity:** $MI(I'; I) \geq 0$

- **Independence:** $MI(I'; I) = 0 \Rightarrow p\{i', i\} = p\{i'\} \cdot p\{i\}$

- **Symmetry:** $MI(I'; I) = MI(I; I')$

- **Self Information:** $MI(I; I) = H(I)$

- **Boundedness:**

$$
\begin{aligned}
MI(I'; I) &\leq min(H(I'), H(I)) \\
&\leq \frac{(H(I') + H(I))}{2} \\
&\leq max(H(I'), H(I)) \\
&\leq H(I', I) \\
&\leq H(I') + H(I)
\end{aligned}
\tag{2.30}
$$

- **Data Processing:** $MI(I'; I) \leq MI(I'; \mathbf{T}(I))$

The representation of the *mutual information* in equation 2.29 relates the size of the overlap of a pair of regions ($p\{i', i\}$) to their total size ($p\{i'\}$ and $p\{i\}$). A change in the proportion of the image overlap may result in higher *mutual information* values even if the image registration is not improved. In order to take into account changes in the proportion of the image overlap, the amount of *mutual information* with respect to the information provided by the individual images is required. An approach for the normalisation of the measure is to evaluate the ratio of the *joint* and *marginal entropies* [186]:

$$
MI(I'; I) = \frac{H(I') + H(I)}{H(I', I)}
\tag{2.31}
$$

Other approaches for achieving the same goals have also been proposed. They include the *entropy correlation coefficient* [186]:

$$Ce(I'; I) = \frac{2MI(I'; I)}{H(I', I)} \tag{2.32}$$

Another approach for measuring image similarity would be instead of using the *joint entropy* for image correlation, to use the *entropy* of the difference image $Idif$ [144, 22]. In this approach the *entropy* operates on a single difference image which is created by subtracting the reference image from the other image by using a suitable scale factor s. Then, the *entropy* of the difference image is used instead of the *joint entropy*:

$$H(Idif) = -\sum p\{i\} log(p\{i\}) \tag{2.33}$$

2.4.3.4   Comparison of intensity based similarity measures

Penney *et al.* [144] performed a comparison of a number of similarity measures used in 2D/3D medical image registration. The measures under consideration included *normalised cross correlation*, *entropy* of the difference image, *mutual information* and gradient correlation. The similarity measures were evaluated by performing rigid registration (6 *degrees of freedom*) between a CT scan and a fluroscopy image of a spine phantom. In their research, the *mutual information* had the worst performance among the considered similarity measures while the *entropy* of the difference image was one of the best considered measures. A similar evaluation of the same similarity measures was later performed by Russakoff *et al.* [175]. The results of their research were substantially different from those reported by Penney *et al.* [144]. Russakoff *et al.* found that the *mutual information* is one of the most reliable similarity measures. One possible explanation for the difference in the performance of *mutual information* is that Russakoff *et al.* used higher resolution images. As a result a larger number of samples are used for the calculation of *mutual information* leading to more accurate estimates of the probability density functions and image entropies. Therefore, *mutual information* may fail if the 2D histogram is relatively

sparsely populated while the *difference entropy* may succeed because the corresponding 1D histogram will be much better populated. It should be noted that while Penney *et al.* used fluroscopy images of a spine phantom, Russakoff *et al.* used only clinical data which may have an effect on the results. Holden *et al.* also performed an evaluation of eight different similarity measures used for rigid body registration of serial MR brain scans [81]. In this evaluation it has been shown that the similarity measures based on *entropy* performed more consistently than the other similarity measures. In particular, *entropy* based similarity measures proved to be the least sensitive methods in the presence of extradural tissue.

Despite its good performance and wide use, *mutual information* has also been shown to lack robustness for certain registration problems [158]. Problems can arise when the images are of a low resolution [158], when the overlapping part of the images is small or as a result of the interpolation method used in the registration approach [159]. To overcome this problem, Pluim *et al.* [158] suggested to include spatial information to the measure by combining *mutual information* with a term based on the gradient of the two images. The gradient term will seek to align locations which have a gradient with a high magnitude and also similar orientation. Rueckert *et al.* proposed the use of higher order *mutual information* [171], which incorporates spatial information by forming four dimensional intensity histograms. A survey of medical image registration methods based on *mutual information* can be found in [160].

2.5 Optimisation approach

The optimisation approach attempts to find the optimal transformation, $\mathbf{T}$, that maximises the similarity of the two images. This is achieved by finding the optimal parameters of the transformation, $\mathbf{T}$, which control the correspondences between the images. Finding the correct parameters of a transformation is not an easy task, especially when the transformation has a large number of *degrees of freedom*. Usually the optimisation approach is an iterative process which in each step tries to improve the correspondence of the images.

Figure 2.9 shows a schematic representation of the optimisation approach.
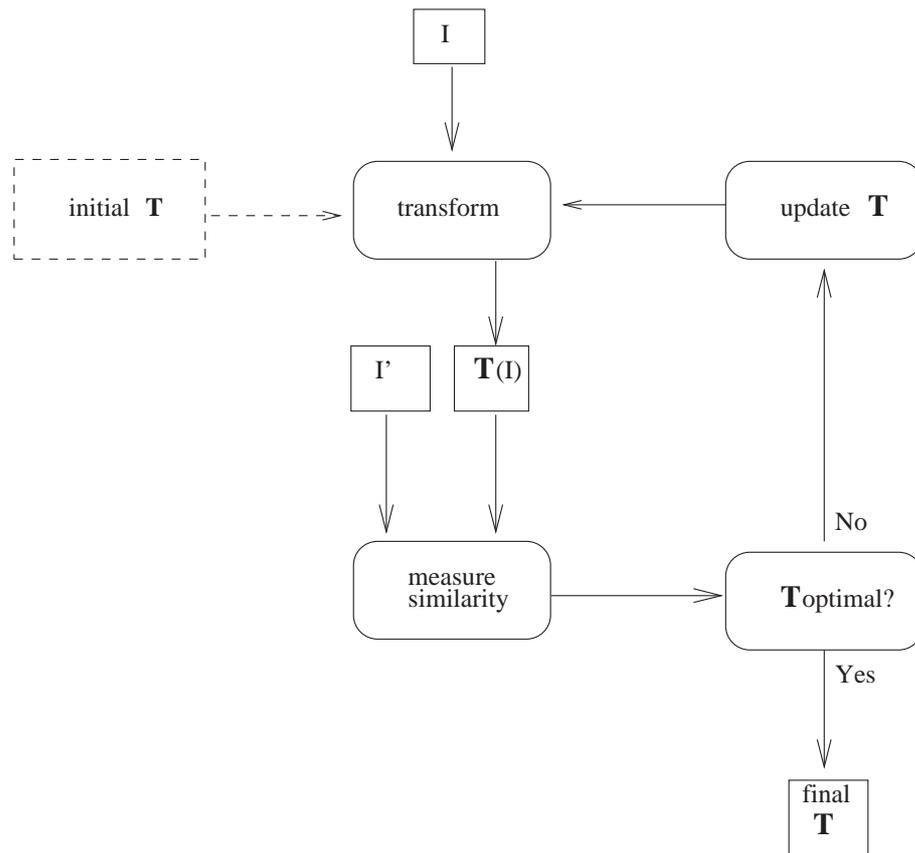
Figure 2.9: A schematic representation of the optimisation approach (adapted from [81]).

The selection of the appropriate search space strategy is very crucial in order for the optimal transformation to be found. For example, the search space may have a large number of local minima which could trap the optimisation approach or the starting point may be far away from the correct solution. The shape of the search space is affected by the type of the similarity measure, the properties of the images and the type of the transformation. There is extensive research on techniques used for searching such spaces. For more details on optimisation methods and search strategies see [163]. Hierarchical strategies are often used in order to increase the likelihood of finding a global optimum match [104]. The multi-scale space techniques are a widely used type of hierarchical strategies. These methods use stacks of images which contain increasingly simplified versions of the initial images [104]. The optimisation approach starts using images in the top of the stacks (the most simplified version of the initial images) to calculate a rough

estimate of the transformation. Then the estimate of the transformation is improved by using each successive level of the image stacks. In cardiac image registration, hierarchical strategies are often used to find the optimal transformation [134, 14, 115].

2.6 Applications of cardiac image registration

This section provides an analysis on several cardiac image registration approaches for a number of applications. It is worth noticing that cardiac image registration is an active research area but no reference to a 4D cardiac MR image registration approach can been found in the literature. A 4D cardiac MR image registration approach is considered to be the one which provides a 4D mapping between two 4D cardiac MR images where the fourth dimension is the time.

2.6.1 Image registration for analysis of the cardiac motion

Chandrashekara *et al.* [28, 29] used an extension of the FFD model (equation 2.13) to analyse the motion of the myocardium using tagged MR cardiac images. The extension of the FFDs is introduced by Schnabel *et al.* [177] where a number of single level FFDs are combined to a multi-level one (MFFD):

$$\mathbf{T}(\mathbf{x}, t) = \sum_{h=1}^{T} T_{local}^{h}(\mathbf{x}) \tag{2.34}$$

The estimation of the myocardial motion requires a sequence of registration steps. Each image, $V_1, V_2, ..., V_t$, of the sequence is registered to the end-diastolic image V_0. In order to recover the long-axis motion both short-axis and long-axis images were registered at the same time. The registration involves the optimisation of a cost function based on the *normalised mutual information* between the registered short-axis and long-axis images. Because the similarity measure is evaluated on both short-axis and long-axis image sets, a complete 3D motion field of the myocardium can be recovered. After registering V_1 to V_0, a multi-level FFD (MFFD) is obtained consisting of a single level representing the

motion of the myocardium at time $t = 1$. The registration of V_2 to V_1 provides the next level of the MFFD representing the motion of the myocardium at time $t = 2$. The process continues until all volumes in the sequence are registered.

The method was validated using a cardiac motion simulator to produce cardiac images with known motion and also by reconstructing the deformation field within the myocardium on images acquired from 11 healthy volunteers. The technique was tested on 9 image sequences produced by the motion simulator. The relative RMS error in the reconstruction of the deformation fields reaches a maximum between 4.2-6.5% at the last frame [29]. In order to assess how well the registration algorithm performed in tracking the motion of the myocardium for the volunteer data, the tag-intersection points in three different SA slices and one LA slice for all time points between end-diastole and end-systole were manually identified by an expert. Then, the RMS error in the in-plane displacements estimated by the registration method and by the expert was measured. The results showed that for all the image sequences the RMS error was smaller than a voxel for most of the cardiac cycle. In seven image sequences the RMS error was approximately 2*mm* at end-systole while in the other four it was approximately 3*mm*.

Carbayo *et al.* proposed a similar method for calculating the cardiac displacement field in 2D ultrasound image sequences [98]. The basic idea of the method is similar to the one introduced by Chandrashekara *et al.* The cardiac displacement field is calculated by registering all the images of the sequence to the first frame using a *deformable* registration model. However what makes this method different is the use of a spatio-temporal semi-local deformation model. In this work the displacement field **g** is represented using a time-space separable linear model $d_{j,l}$:

$$\mathbf{g}(t, \mathbf{x}) = \mathbf{x} + \sum_{l \in \mathbb{Z}} \sum_{j \in \mathbb{Z}^N} d_{j,l} \phi_j(\mathbf{x}) \psi_l(t) \tag{2.35}$$

where $\phi_j(\mathbf{x})$ influences the spatial deformation and $\psi_l(t)$ influences the temporal coherence of the deformation. B-Splines are used to model both the spatial and the temporal

components:

$$\mathbf{g}(\mathbf{x}) = \mathbf{x} + \sum_{l \in \mathbb{Z}} \sum_{j \in \mathbb{Z}^N} d_{j,l} \beta_{\mathbf{x}}(\mathbf{x}/h - j) \beta_t(t/s - l) \qquad (2.36)$$

where the basis functions β_x are placed on a rectangular grid in the spatial domain and β_t at regularly spaced time points. Moreover, the scale parameters h and s govern the total number of parameters $d_{j,l}$, the smoothness of the mapping and the knot spacing between the control points of the B-Splines. The optimal displacement field is found by optimising a global pixel-based criterion assessing the quality of the registration over the entire image sequence at once. The method was tested using 4 simulated image sequence models, 4 image sequences from healthy volunteers and 4 image sequences from *ischemic* patients. In the simulated images, the mean square error over the entire sequence for 85 selected points within the myocardium was found to be between 0.3*mm* and 0.83*mm*. No quantitative analysis on the method's performance for the real sequences was reported.

Rao *et al.* [167, 166] used the FFD model in order to enable the direct comparison between the motion extracted from cardiac MR image sequences within or across patients. This is achieved by using two registrations. Initially, the cardiac motion field is calculated by using the method introduced by Chandrashekara *et al.* [28, 29]. Then, the end-diastolic untagged images of each patient (the untagged images are acquired shortly after the tagged images) are registered together using a non-rigid registration method based on FFDs. This resulting transformation provides a way to map the motion fields of each subject to the same coordinate system. This method provides a good approach for the comparison of the cardiac motion patterns within subjects. Its use has potential advantages including the comparison of changes in cardiac motion in patients and the assessment of pharmacological or surgical intervention. The potential of the method was demonstrated by visually assessing a small number of images. Figure 2.10 demonstrates the method.

Petijean *et al.* have also developed a similar approach for the direct comparison between motion extracted from cardiac MR image sequences within or across patients [154]. The key difference between the two approaches is that Rao *et al.* used a parametric non-
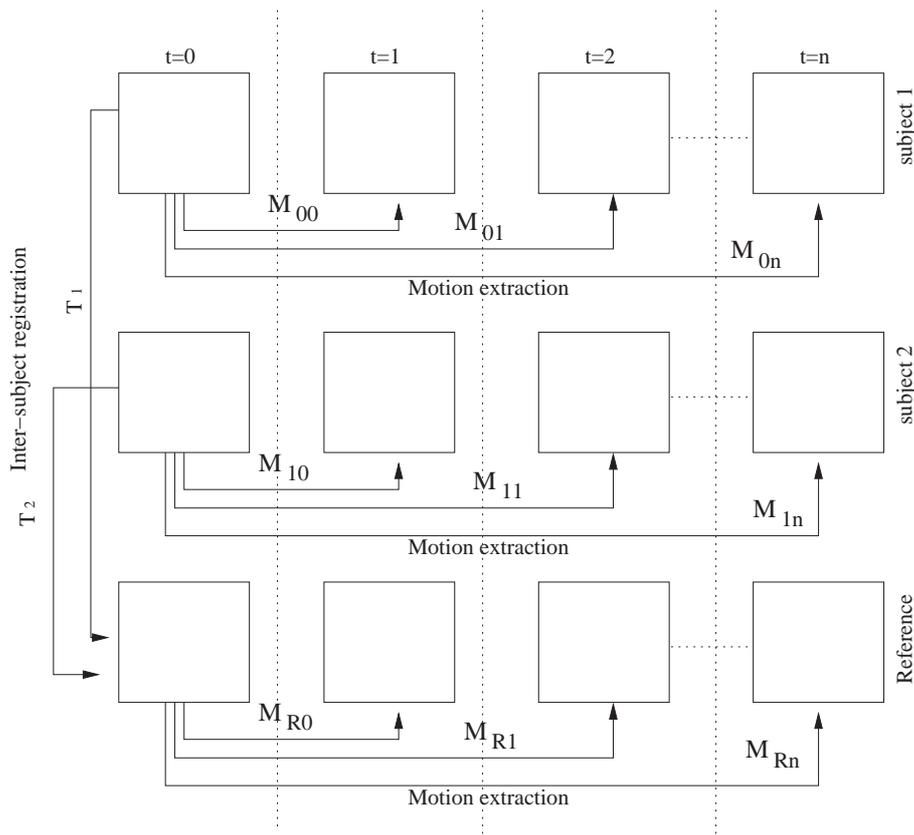
Figure 2.10: Mapping cardiac motion fields across subjects (adapted from [166]).

rigid registration method over a free-form deformation space, while Petijean *et al.* used f-information based non-rigid registration over a non-parametric transformation [154].

2.6.2 Image registration for cardiac motion correction

Turkington *et al.* developed a registration model to align images from dynamic cardiac $N-13$ ammonia positron emission tomography scans [194] used for measuring regional myocardial blood flow. This type of study requires to acquire images for the first 2 minutes after the injection of $^{12}NH_3$. The registration technique was used to correct artifacts in the images introduced from breathing, motion of the heart and the overall motion of the patient. The method corrected only translation differences across the images. It was assumed that the orientation of the heart remained the same during the scanning. This assumption is appropriate to a certain degree, but a small rotational motion could occur during the acquisition of the images [194]. The registration was based on a number of

templates. These templates are derived from a frame acquired around 10 minutes after the injection of the contrast media. Three different templates were constructed: a myocardial template, a blood pool template (which was used to match frames where the left ventricle's blood pool is dominant) and finally a modified blood pool template (which reduced the affects of the right ventricle activity). Each image frame was assigned to a template by visual inspection of the cardiac activity and registration was performed. The method was tested on phantom data and resulted in translations of less than one voxel which is consistent with the fact that the phantom was scanned at the same position. The technique was also tested in a small number of human data (3 subjects). Even if quantitative analysis in the performance of the method has not been provided, it has been shown that the technique corrects for cardiac motion to a large degree. The fitting of the blood pool template in the direction of the z-axis was not as good as in the direction of the x- and y-axis. The authors explained that this may be the result of the fact that the blood pool extends beyond the base of the myocardium [194].

In myocardial perfusion studies, a contrast media is injected before the acquisition of the MR images. Due to the dynamic nature of the data and because the acquisition has to be gated, it usually takes more than 3 minutes to obtain reliable information about the perfusion and distribution of the contrast media. During this time it is impossible for the subjects to hold their breaths. Therefore, motion artifacts are introduced due to breathing during the acquisition procedure. Bidaut *et al.* developed an automated registration approach for the correction of these artifacts during acquisition of dynamic MR images [14]. The approach uses sequences of short-axis views of the heart. A rigid two-dimensional transformation model with 3 *degrees of freedom* (2 for translation and 1 for rotation) is used for the correction of motion artifacts. In order to constrain the registration only to the region of interest, a mask was applied to each frame encompassing the most likely position of the heart. The mask was calculated by taking the maximum value for each pixel over the entire sequence. A single early slice from the sequence was selected as a reference image. Then, all slices were registered to the reference image by minimising the pixel based squared differences between each frame and the reference frame. After

the registration of the entire sequence, a new reference image was defined as the average of all registered slices. The registration was then repeated in order to handle larger displacements.

The method was evaluated by calculating the relative motion of anatomical landmarks on individual slices before and after registration. For the evaluation of the method images from eight ambulatory patients with stable coronary artery disease were used [14]. The results show that the method improved significantly the overlap of the images even when compared with manually aligned images. In addition the method was evaluated by using a compartment model for estimating two myocardial perfusion parameters: the blood to myocardium transfer constant, K_1, and the Gd-DTPA distribution volume, V_d. After applying the method to the data, the variability of both parameters was reduced compared to the variability obtained from the uncorrected images.

Ablitt *et al.* presented a technique for predictive cardiac motion modelling and correction [1]. The method uses a registration approach based on FFDs (equation 2.13) to recover the cardiac deformation due to respiration. Then, it uses partial least square regression to extract intrinsic relationships between the 3D cardiac deformation due to respiration and multiple one-dimensional measurable real-time intensity tracers at chest. This model is used to predict cardiac motion due to respiration. The method was tested using cardiac MR images from 10 normal subjects. The accuracy of the motion prediction method was assessed by performing cross-validation experiments (i.e. the data were divided into two parts, one for training the model and one for testing the model). The errors of the motion model were normalised with the residual errors achieved by 3D free-form registration (the results of the free-form registrations were considered as the gold standard for this study). These measurements demonstrated that with the model the maximum/minimum error is consistently very small.

Klein *et al.* introduced an affine registration model for correction of respiratory motion on respiratory-gated PET data [93, 92]. In this approach the use of an affine (12 dof) global transformation model for the registration of different end-diastolic respiratory gates in a PET sequence is investigated. In order to correct the respiratory motion, all images

of the sequence are registered to the image acquired during end-expiration. The cost function used during the registration procedure consisted of a least square difference term and a temporal term. In order to ensure smoothly varying motion between adjacent image frames, a priori knowledge model is used which assumes that the motion of the heart from one frame to the next is likely to follow a smooth progression. The temporal part of the cost function penalizes any departure from the prediction model. The method was tested on images obtained from 10 subjects as well as on noisy phantom data [93] and was compared to two similar methods. One of the methods used only a rigid transformation model while the other method used an affine model without the temporal constraints. The results indicated the use of an affine transformation and the temporal smoothing constraint provides better performance than the other two methods.

McLeish *et al.* performed a study of the motion and deformation of the heart due to respiration [123]. In their study, 3D cardiac MR images were used. The study included images from 8 healthy volunteers and 10 patients. The healthy volunteers were scanned into various time points between the end-expiratory and end-inspiratory positions. The patients were scanned only at the end-expiratory and end-inspiratory positions. The images at maximum exhalation were selected as references and all the other images were registered to them using rigid registration followed by non-rigid registration. The results indicated that during inhalation the average movement of the heart is in craniocaudal direction (CC) (by $12.4 \pm 5.9mm$), in the anterior-posterior direction (AP) (by $4.3 \pm 3.7mm$ towards the anterior direction) and in the left-right (RL) direction (by 2.0 ± 2.1 to the right) [123]. Furthermore, the left ventricle deforms more at the apex than at its base. The left ventricle has larger deformations than the right atrium and the right coronary artery. Moreover, when going from an inhale to an exhale position, the outer right atrial wall and the right coronary artery move outwards, while the left ventricle moves upwards.

2.6.3   Image registration for cardiac segmentation

Image registration has also been used for the segmentation of cardiac MR images [109, 130]. In cardiac imaging, segmentation is used in a large number of applications including the calculation of cardiac volume and blood ejection fraction, the analysis of contraction and wall motion and the visualisation of the cardiac anatomy. In order to be clinically valuable, the segmentation procedure has to be automatic. Lorenzo-Valdés *et al.* developed a fully automated approach for the segmentation of the myocardium and the ventricles of a cardiac MR image sequence [109]. Their approach uses a non-rigid registration algorithm based on free form deformations (equation 2.12).

The key idea of the proposed algorithm is to reduce the segmentation of the entire sequence problem to one of manually segmenting the end-diastolic frame of the sequence and then propagate the segmentation to the rest of the sequence's frames. In order to propagate the segmentation, each frame of the image sequence is registered to the end-diastolic frame using a multilevel non-rigid registration method based on B-Splines. Furthermore, an approach for the automatic segmentation of the end-diastolic frame is also presented. In order to automatically segment the end-diastolic frame, two atlases are used.

A population specific atlas of the end-diastolic frame was constructed by registering the manual segmentations of 14 diastolic cardiac MR images. The population specific labelled 3D atlas contained the left ventricle, the right ventricle and the myocardium. If the mapping between the atlas and a specific end-diastolic image is known, the segmentation can be propagated to the end-diastolic image. The robust and accurate segmentation of this frame is crucial since it is propagated to the rest of the sequence's images. A subject specific atlas of the heart corresponding to the end-diastolic time frame is constructed. The subject specific atlas is used instead of the end-diastolic image during the registration with the population specific atlas. The subject specific atlas is constructed by registering all time frames to the end-diastolic time frame, transforming each frame with the resulting transformation and calculating the average image. This enables a better alignment with the population specific since the subject specific atlas reduces the effect of image noise

and poor contrast to the registration procedure.

The automatic segmentation approach was validated against the manual segmentations of nine 3D MR image sequences. All the segmentations contained the left ventricle (LV), the right ventricle (RV) and the myocardium (MYO). The volumes of the ventricles and the myocardium were calculated and regression analysis was used to compare the two methods. Two types of experiments were performed assessing how well the segmentation can be propagated to the subsequent frames as well as the entire automated approach. In the first experiment the transformations between different time frames were applied to the manual segmentations of the end-diastolic frame instead of the subject-specific atlas. The resulting volumes yield a good correlation between the manual and the automated segmentation (LV=0.99, MYO=0.98, RV=0.96). In the second experiment the image sequences were segmented by using the entire automated approach. In this case the results showed a good correlation between the volumes of the corresponding structures (LV=0.94, MYO=0.83, RV=0.96).

Noble *et al.* also used non-rigid registration for the segmentation of the endo-cardial and epi-cardial surfaces of cardiac MR images [130]. Their method required the end-diastolic endocardial and epicardial surfaces to be manually segmented. In this method only three slices corresponding approximately to basal, mid and apical positions of the heart are used instead of the entire image. The first step of the segmentation approach is to calculate the centre of area of the ventricular blood pool in the segmented images for all slices. Then all the images were re-sampled in a polar fashion around each slice's centre. Two segmentation approaches have been developed [130]. The first registers all the time frames to the manually segmented image (all to one), while the second registers each adjacent frame (piecewise). To evaluate the method, images from 10 patients undergoing cardiac MR for the investigation of coronary artery disease were used. The performance of the two methods was compared to manual segmentations of the images and also to a commercially available software package. The correlation coefficient between the epi- and endo-cardial volumes produced by the manual segmentation and the above two segmentation methods were compared. The results showed that the piecewise segmentation

method correlates with the manual segmentation better than the other method.

2.6.4   Image registration for alignment of cardiac stress and rest SPECT images

Declerck *et al.* proposed a method to enable the better visual or quantitative interpretation of myocardial perfusion studies using SPECT imaging [47]. In a stress-rest study, two perfusion maps of the cardiac muscle in the left ventricle are acquired. One image is acquired before the injection of the tracer (at rest) and the other after the injection of the tracer during maximal exercise.

The approach proposed by Declerck *et al.* involves two major steps [47]. First, the stress and rest pair of images are aligned using a point-based registration method. During the registration of the rest and stress pair of images an affine transformation which defines correspondence between a point in the stress image and a point in the rest image is used. Then a non-rigid registration method based on B-Splines (equation 2.12) is used to map the stress images to a template image. The template image is a single selected normal image characterised by good contrast and low intensities on all non-cardiac features. Finally, the obtained transformation maps are used to resample both images to the geometry of the template image.

2.6.5   Spatial and temporal registration of cardiac SPECT and MR images

In cardiology more than one type of image modality can be acquired for a single patient. Images from different modalities provide different kinds of information to clinicians. The goal of multimodal image registration is to merge the information provided from these images. Faber *et al.* developed an approach for the spatial and temporal registration of single photon emission computed tomography (SPECT) and magnetic resonance images (MR) [55]. They have used high resolution MR images containing the left ventricle (LV) of the heart and lower resolution SPECT images containing information regarding my-
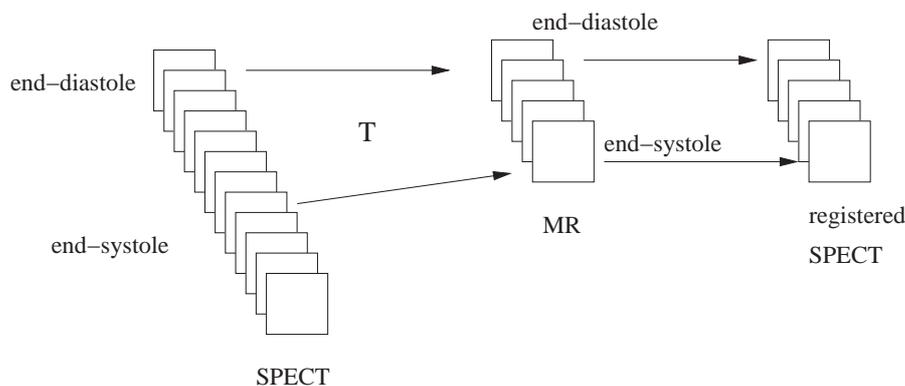
Figure 2.11: Method for the spatial and temporal registration of SPECT and MR images (adapted from [55]).

ocardial perfusion. When the images are registered, the exact anatomic location of the perfusion can be identified and its effects on the cardiac motion and wall thickening can be studied.

The first step of the algorithm involves the application of a surface detection algorithm to both the MR images and the SPECT images. The surface detection algorithm fits the images' intensity gradients to a model of the LV. The output of the detection algorithm is a set of 288 points on the endocardial surface of the LV. The end-systolic (ES) and end-diastolic (ED) frames are determined from the segmented surfaces. A rigid registration method is used to determine the best transformation $\mathbf{T}$ which maps both the end-diastolic and end-systolic SPECT surfaces to the end-diastolic and end-systolic MR surfaces. The registration method aims to minimise the distance between corresponding surfaces. After the registration of the ED and ES surfaces, linear interpolation is used to create SPECT frames corresponding to the MR frames [55]. A graphical representation of the registration is presented in figure 2.11.

The algorithm was evaluated using images from 3 healthy subjects and one subject with coronary artery disease. The subjects were studied in both MR and SPECT perfusion imaging while three of them underwent stress perfusion imaging as well (7 studies in total). The mean distance between the surfaces before the registration was 4.00mm while after the registration it was 2.7mm. No evaluation of how much the surface detection algorithm affects the performance of the registration method is provided. In addition,

the temporal mapping between the SPECT and MR images is provided only for the end-systolic and end-diastolic frames. This does not mean that the same temporal relationship will be valid for the frames between them since the dynamic properties of the hearts may change from one scanning session to the other.

2.7 Conclusions

This chapter has provided an overview of medical image registration. Image registration is an essential tool for the analysis of cardiac images. The following issues affect the registration procedure and need to be addressed in order to accomplish the registration of two images:

- **The nature of the data to be registered.** The nature of the data are described in sections 2.1 and 2.2. Depending on the imaging modalities used to acquired the images, the data can be points, surfaces or intensities. In case of point or surface based registration, features need to be extracted during the registration procedure. In the case where image intensities are used, a preprocessing step may be used to enhance image features and to improve image quality. The nature of the data is a very important factor on deciding on the similarity measure as well as on the optimisation approach.

- **The similarity measure to use during the registration.** A large number of similarity measures have been described in section 2.4. The selection of the similarity measure depends on the nature of the data to be registered. Euclidean distance measures can be used when registering surfaces and points. A large number of similarity measures which could be used for intensity based registration approaches have been presented in section 2.4.

- **Type of transformation.** Details regarding the different types of transformations are provided in section 2.3. Depending on the application, the type of transformation can be linear or non-linear. Linear transformations maintain the shape of the

structures and are usually used to combine information from a number of image modalities. Non-linear transformations provide a better mapping between anatomical structures by deforming the images. They are commonly used in applications in the analysis of cardiac images (for example tracking of cardiac motion), due to the variability in the shape of the heart of different subjects.

- **The optimisation procedure.** The selection of the correct optimisation procedure is crucial in order to find the optimal parameters of the transformation with the least number of calculations. Usually the optimisation approach is an iterative process which in each step tries to improve the correspondence of the images. Details on the optimisation approach are provided in section 2.5.

Since the heart is undergoing spatially and temporally a varying degree of motion during the cardiac cycle, 4D cardiac image registration methods are required when registering cardiac MR image sequences. Contrary to most of the methods reviewed in this chapter, the work presented in this thesis aims to align a number of cardiac MR image sequences not only to the same spatial but also to the same temporal coordinate system. Thus, the transformation model used during the spatio-temporal registration addresses differences in the spatial domain as well as differences in the temporal domain of the image sequences. Furthermore, the registration methods presented in this thesis are automatic (except initialisation) and require no segmentation in order to determine temporal features in the cardiac cycles. Table 2.1 summarises the cardiac image registration methods considered in this section.

Method	Dimensionality	Modality	Registration	Area
Chandrashekara *et al.* **[28, 29]**	3D	MR	Non-rigid	Analysis of cardiac motion
Carbayo etl al. **[98]**	2D+time	Ultrasound	Non-rigid	Analysis of cardiac motion
Rao *et al.* **[167, 165]**	3D	MR	Non-rigid	Comparison of cardiac motion across subjects
Petijean *et al.* **[154]**	3D	MR	Non-rigid	Comparison of cardiac motion across subjects
Turkington *et al.* **[194]**	3D	PET	Only translation	Cardiac motion correction
Ablitt *et al.* **[1]**	3D	MR	Non-rigid	Cardiac motion correction
Bidaut *et al.* **[14]**	2D	MR	Rigid	Cardiac motion correction
Klein *et al.* **[93, 92]**	3D	PET	Affine	Cardiac motion correction
McLeish *et al.* **[123]**	3D	MR	Non-rigid	Motion and deformation of the heart due to respiration
Lorenzo-Valdés *et al.* **[109]**	3D	MR	Non-rigid	Image segmentation
Noble *et al.* **[130]**	3D	MR	Non-rigid	Image segmentation
Declerck *et al.* **[47]**	3D	SPECT	Non-rigid	Comparison of stress and rest images
Faber *et al.* **[55]**	3D+time	SPECT/MR	Rigid	Spatial and temporal registration of two modalities

Table 2.1: Overview of cardiac image registration applications.

Chapter 3

Modelling of anatomy

Applications assisting the automatic interpretation and understanding of MR images are of high importance for increasing the clinical use of MR imaging. Computational anatomy enables the construction of models describing the anatomy and function of anatomical structures. Moreover, these anatomical models can represent information regarding anatomical and functional variability in the population. Models of anatomical structures enable clinicians not only to interpret medical images but also to better understand anatomical structures. A model provides a better visualisation of the anatomical structure since medical images suffer from artifacts (e.g. noise, poor contrast to signal ratio, etc). Furthermore, different visualisation methods will allow clinicians to visualise different properties of the modelled anatomy and function. In addition, collecting information from a large number of subjects to a single model will enable the better understanding of the anatomical structure and its variation within the population. Finally, modelling anatomical structures with certain pathologies will enable the better understanding of the pathology, i.e. how the anatomical structure is affected, how the pathology progresses, etc. This chapter reviews on methods for modelling anatomical structures.

3.1 Anatomical atlases

Computational anatomy is an active research area [70]. Anatomical atlases are very similar to geographical atlases as both contain information regarding the anatomy of a physical reality. However, geographical atlases describe a constant physical reality (for example a continent) which can be described by a large number of abstract representations (e.g. population, rainfall, temperature). On the other hand, anatomical atlases do not describe a single constant reality since the shape, size and function of a particular anatomical structure differs across the population. In order for anatomical atlases to be meaningful and representative of the population, they must deal with the fact that an anatomical structure might have a large number of physical realities. Traditional anatomical atlases contain anatomical and functional information from a single subject and focus primarily on the human brain [176, 190]. These atlases do not contain any information regarding the anatomical and functional variability across the entire population and are suitable for anatomical structures without large inter-subject variability. However, they are less suitable for anatomical structures with large variability across the population.

3.2 Probabilistic atlases

Population based atlases provide a solution to this problem by incorporating information from a lager number of subjects. Population based atlases can be used to guide knowledge based image analysis algorithms and also to support pathology detection in individual subjects or groups [192]. In order for an atlas to be representative of the population, the variability must be captured in an appropriate framework. Probabilistic atlases retain information regarding variability in the form of tissue probability maps (i.e. each voxel of the atlas is assigned with a value that describes its probability to belong to a certain structure). They have a number of advantages over conventional atlases, most importantly their ability to retain information regarding anatomical and functional variance. Furthermore, the probabilistic framework enables calculations between the morphometry

of different subjects or atlases to be performed efficiently. Moreover, it also enables statistical and computational comparisons between individuals and groups making the atlas an important clinical research tool. Building a probabilistic atlas should be thought as of an ever-evolving process. It should be relatively easy to add new data to the atlas. The more data is added to the atlas the more representative of the population the atlas will be.

During the last few years a number of approaches have been developed for the creation of probabilistic atlases describing the anatomy and function of anatomical structures as well as the variability across the entire population. The majority of these atlases focus on the anatomy and function of the human brain [121], [119], [120], [168], [50], [131], [35]. Probabilistic atlases of the human brain have been successfully used to investigate the structural and functional differences in the human brain as parts of the *International Consortium for Brain Mapping* [121], [119], [120]. The aim of the *International Consortium for Brain Mapping* (ICBM) is to develop a voxel based probabilistic atlas of the human brain. The atlas will contain information from a large number of subjects with a wide ethnic and racial distribution and various imaging modalities. Furthermore, it will describe the brain anatomy and function in a 3D spatial domain as well as a temporal one modelling the age of subjects. The images acquired during this work have been separated into target brain and reference brain sets. The target brain set is the dataset, derived from a small number of individuals from whom the richest collection of data exists. The target brains have been segmented into several anatomical structures. They are used in several applications including automatic segmentation of brain subjects (by registering them with unlabelled data). Contrary to the target brain set, the reference brain set is derived from a large number of subjects (the aim of this work is to include more than 7000 subjects). This dataset provides information regarding the variance of the population and could be used for calculation of population and sub-population statistics.

Figure 3.1 provides illustrations of probabilistic atlases developed as part of the *International Consortium of Brain Mapping* [60]. These atlases are constructed using T1-weighted MR scans from 452 subjects. The reference space of the atlas is the average position, scale and shear from all the individual subjects. Figures 3.1 (a), (b), (c) show

the probabilistic atlases of the gray Matter (GM), the white matter (WM) and the CSF.

(a)   (b)   (c)
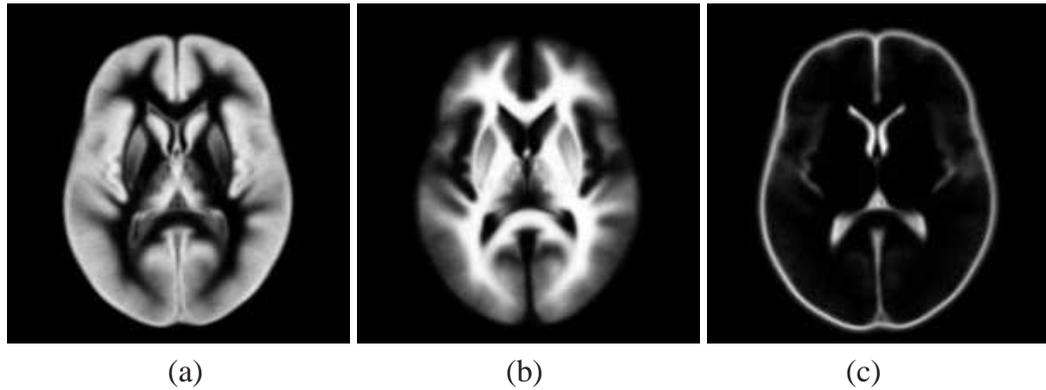
Figure 3.1: Examples of probabilistic atlases of the human brain (a) gray matter atlas, (b) white matter atlas, (c) CSF Atlas. These atlases are developed as part of the *International Consortium for Brain Mapping* [60].

The following steps are used for the construction of the probabilistic brain atlases (analysis pipeline) [119, 120]:

- Screening of the data to find incomplete studies or studies with artifacts.

- Intensity normalisation of the data (in all three dimensions).

- Registration of the data across studies within the same subject.

- Tissue classification (i.e. GM, WM, CSF).

- Removal of extracerebral structures.

- Spatial normalisation of each subject to a target subject.

- Extraction of surface features.

- Visualisation of the atlas.

The intensity of the data is normalised in order to assist the tissue classification procedure.Spatial normalisation of each subject to a target brain enables the automatic segmentation of the subjects. Different image segmentation methods could be used to segment the brain images to particular anatomical structures. The development of novel image

registration methods is very important for the construction of probabilistic atlases. Registration is not as simple as equating the origin of similar coordinate systems. The registration methods must accommodate for diverse types of image data each with different spatial resolution and coverage [192]. Furthermore, image registration is also an essential for the further use of the atlases [120], [192]:

- Mapping all the images to a common coordinate system enables the construction of the atlas. The type (rigid, affine, non-rigid) of the registration method depends on the application of the atlas. Registration methods which maintain the shape (i.e. methods which use linear transformations) of the anatomical features are used during the construction of most atlases.

- *Deformable* registration methods enable the registration of the atlas to an individual subject. Adapting the shape of an atlas to an individual subject allows the construction of individualised atlases. Information from an atlas can be transferred to the individual subject while maintaining the intricate patterns of structural variation in the subject's anatomy. Hence, non-rigid registration of an atlas to a subject can be used to produce valuable information regarding abnormalities [191]. Furthermore, non-rigid registration methods can be used to determine morphometric variability which exists in the data [36]. Moreover, registration of an atlas to an individual image enables the segmentation of the image [109].

- Registration between individual subjects can be used to transfer physiological data from different individuals to a single anatomical template enabling their comparison without the confounding effects of their anatomical shapes. For example, rigid registration [66], [103], [81] and non-rigid registration [65], [169] can be used for the comparison of humans brains in a normalised reference space.

- Image registration can be used to compare atlases from different populations and help identify significant anatomical and functional differences between different population groups.

The visualisation of the atlas can be separated into two forms [119, 120]. The first concerning the visualisable aspect of the atlas (3D renderings, etc) and the second one relating the atlas to a database containing clinical information regarding the region of interest.

The *Montreal Neurological Institute* [35] has developed a probabilistic atlas of the human brain. In order to build this atlas, MR images from 305 subjects were mapped into the same stereotactic space then, intensity normalised and averaged on a voxel-by-voxel basis. More recently, Mazziotta *et al.* have built a four-dimensional probabilistic atlas of the human brain which includes both macroscopic and microscopic information on the function and structure of the human brain [119]. At the current stage, the brain atlas is constructed from more than 1000 subjects, aged 18 to 90 years old. However, the aim of this work is to include more than 7000 subjects of different ages and countries.

Dinov *et al.* have used a probabilistic atlas of the human brain in order to develop a subvolume thresholding method for the analysis of positron emission tomography (PET) and single photon emission CT data of the brain [50]. The atlas is also used to determine the statistical significance of the effects of motor simulation on brain perfusion. Nowinski *et al.* have also used a brain atlas to assist a method for localisation analysis of functional images [131]. This technique has a number of limitations mostly due to the nature of the atlas. The brain atlas they used, an enhanced and extended electronic Talairach-Tournoux brain atlas [189, 190], is one of the first electronic brain atlases. However, it has significant limitations due to the fact that it is based on a single subject. Rasser *et al.* have developed a nonlinear registration technique to project the Brodmann areas of the brain onto 3D co-registered functional MR datasets [168]. Similar to the Nowisnksi's *et al.* approach, this method uses a single subject based MR atlas.

Park *et al.* produced a probabilistic atlas of the abdomen using 32 noncontrast abdominal computed tomography scans acquired from patients [140]. The probabilistic atlas consisted of four organs (liver, kidneys and spinal cord). In order to construct the atlas all the images were manually segmented. One subject was selected as reference and all the other images were registered to the reference subject using a *deformable* registra-

tion method based on thin plate splines [140]. Care was taken to ensure that the reference subject was a normal representative of the population. The information provided by the atlas was incorporated into a segmentation framework in order to aid the automatic segmentation of abdominal images.

In brain imaging there is very complex structural variability between normal individuals and particularly between different population groups [192]. Therefore, a single brain atlas may fail to serve as a faithful representation of the population. In this case, customised atlases for specific population groups could be more faithful representations of the group. Hill *et al.* have recently presented a novel approach for building dynamic atlases [78]. In this approach a dynamic atlas can be customised to meet particular criteria, for example: age, sex, etc. The method enables the specification a number of criteria for the subjects used to build the atlas. The advantages of dynamic atlases is that they can be tailored to meet particular needs of the research question of interest.

In cardiac image analysis, probabilistic cardiac atlases have been developed by Lötjönen *et al.* [111]. In this work a probabilistic atlas of the cardiac anatomy using MR images has been constructed by combining information from standard short- and long-axis images. The atlas was built from 25 healthy subjects. The images were segmented into atria, ventricles and epicardium by fitting a 3D surface model to both short- and long-axis images simultaneously . The main steps in the construction of the atlas were: the affine registration of the segmented subjects to the reference subject, the blurring of the registered images with a Gaussian kernel and the averaging of the blurred images [111]. Lorenzo-Valdés *et al.* have also constructed probabilistic atlases of the cardiac anatomy and function from 14 MR image sequences of healthy volunteers [108]. Lorenzo-Valdés *et al.* developed separate probabilistic atlases of the left and right ventricles as well as for the myocardium. The main differences between these two approaches is that Lötjönen *et al.* address the issue of the poor image resolution in the direction orthogonal to the slice by using both short- and long-axis images, while Lorenzo-Valdés *et al.* used only short-axis images. Moreover, Lötjönen *et al.* model only the cardiac anatomy at end-diastole while Lorenzo-Valdés *et al.* model both the cardiac anatomy and function. During the con-

struction of the probabilistic atlas, Lorenzo-Valdés *et al.* map the image sequences only to the same spatial coordinate system and not to a spatio-temporal coordinate one. This introduces artifacts into the atlas since image frames from different temporal positions in the cardiac cycles would be averaged.

3.2.1 Reference space of the atlas

The selection of the reference space of the atlas is very crucial. In the case where one subject of the population is randomly selected to be the reference subject, care must be taken to ensure that it is a normal representative of the entire population. The Talairach coordinate system [189, 190] was derived from the brain of a 60 year old woman and was one of the first coordinate systems used in brain atlases. The Talairach coordinate system was initially developed to help interpret brain stem and ventricular studies acquired using pneumoencephalography [189]. It has become an international standard for reporting functional activation sites in PET studies.

In order to avoid bias towards a specific subject, iterative registration approaches can be used to calculate the reference space of the atlas. Such a reference space can be constructed by [192]: (1) using automated linear registration to align the data to a randomly selected image; (2) intensity averaging the aligned data; and (3) recursively re-registering the data to the resulting average image. The resulting average image is then adjusted to have the mean affine shape for the group [192]. Lötjönen *et al.* used a similar method to calculate the reference space of the atlas [111], while Valdés *et al.* have randomly selected one of subjects to be the reference space [108]. Guimond *et al.* presented a convergence study for calculating average brain models [71]. Their approach consists of the following steps:

1. Evaluation of the global intensity and shape differences by calculating an affine transformation $\mathbf{T}_i^{global}$ between each subject I_i and the reference subject I'.

2. Use of non-rigid registration to map each image, I_i, to the reference, I', using the obtained affine transformation, $\mathbf{T}_i^{global}$, as an initial estimate. This step provides the

resulting matched images $\mathbf{T}_i^{local}(I_i)$ and the resulting non-rigid displacement field $\mathbf{T}_i^{local}$.

3. Calculation of the mean intensity image by averaging all $\mathbf{T}_i^{local}(I_i)$.

4. Production of the mean deformation field by calculating the vectorwise average of the deformation fields $\mathbf{T}_i^{local}$.

5. Application of the mean deformation to the average intensity image to produce the mean average intensity and shape model.

These steps are repeated using an iterative approach in order for the method to converge to a reference image close to the centroid of the population.
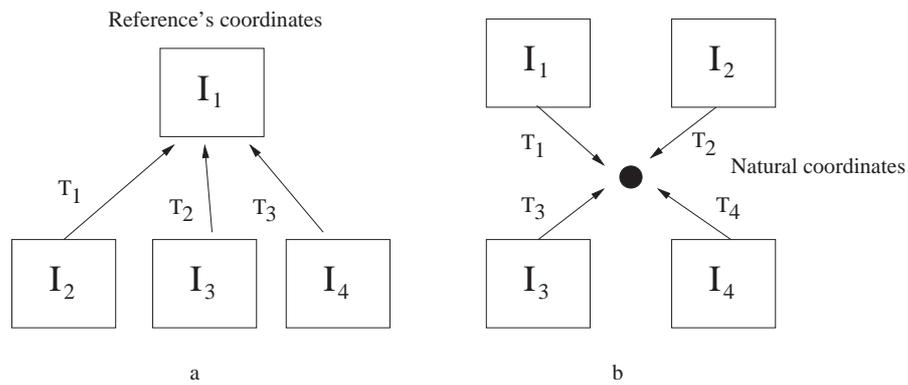
Figure 3.2: Examples of an atlas coordinate systems. (a) The atlas is defined in the coordinate system of I_1. (b) The atlas is defined in the natural coordinate system.

Rueckert *et al.* constructed an atlas of the brain using pairwise registration between each subject and the reference subject [173]. The mean deformation is then applied to the atlas to obtain a model in its *natural coordinates* (the coordinates of the mean shape) [173]. Bhatia *et al.* created an atlas directly in the *natural coordinates* by using a method for groupwise non-rigid registration of brain MR images [13]. The natural coordinate system is calculated implicitly by constraining the sum of all deformations from the reference space to each subject to be equal to zero [13]. Figure 3.2 shows how an atlas is

constructed using an individual subject as a coordinate system (a) and the natural coordinate system (b).

3.2.2   Encoding variability

The methods for creating probabilistic atlases fall into three main categories [192]. Each category differs in its conceptual foundations. The three categories are:

- **Intensity based approaches**. In these approaches the average representation of anatomy is calculated by averaging the intensity of corresponding voxels.

- **Segmented based approaches**. In these approaches the data are segmented into anatomical structures. The probability map for each segmented structure is constructed by determining the proportion of subjects assigned to a given anatomic label at each voxel position.

- **Deformation based approaches**. In these approaches probabilistic information is locally encoded from the deformation maps. The deformation maps are produced by using non-rigid registration and enable determination of the magnitude and directional biases of anatomic variation.

The main difference in the three categories for the construction of probabilistic atlases is how the statistical distribution is modelled and analysed. Random vector fields are analysed on deformation based approaches, while random scalar fields are used to model intensity statistics in the intensity based approaches and binary labels in space in the segmentation based approaches [192]. One problem with the intensity and segmentation based methods is that averaging after linear registration introduces blurring in the boundary definition of structures with spatial variability in the population. This could destroy information regarding small structures (e.g. the gyral feature of the cortex in the brain) [192]. A solution to this problem is to use a non-linear registration method which will enable a better boundary definition of anatomical structures.

3.3   Statistical modelling

Probabilistic atlases contain information about the degree of variability at every voxel of the atlas. However, they cannot provide information regarding the type of variability. Statistical atlases enable the calculation of additional information regarding the type variability which exists in anatomical structures. The construction of statistical models of shape usually requires the identification of a set of landmarks on every image of the population. The landmarks are placed on strong image features like boundaries and points with high curvature. Statistical analysis is performed on the landmark locations enabling the calculation of the average shape and also its significant modes of variation. Figure 3.3 provides an example of statistical shape modelling. In order to perform statistical analysis on the location of the landmarks, correspondence between the landmarks in each image has to be established.
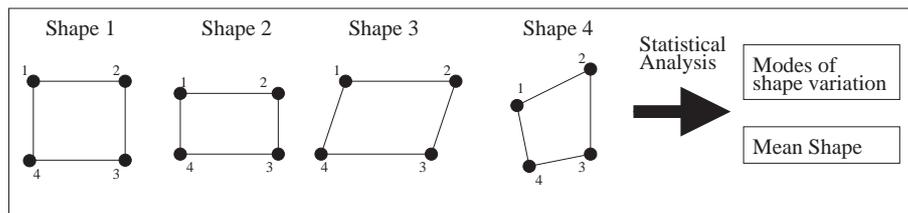
Figure 3.3: Statistical modelling of shapes.

3.3.1   Statistical shape models

There have been a number of attempts to build statistical shape models of the cardiac anatomy [40, 38, 64, 112] and statistical models of the appearance of the heart [17, 128, 126, 197]. Active Shape Models (ASM) and Active Appearance Models (AAM) are examples of statistical approaches for modelling cardiac anatomy and appearance. Active shape models have been introduced by Cootes *et al.* [40] for modelling the shape of anatomical structures by gathering statistical information from a large set of images. The construction of an active shape model requires all sets of landmarks to be aligned with

the corresponding reference landmark set. The alignment aims to correct scaling, translation and rotation differences and is performed using an iterative approach based on the Procrustes method [69]. The alignment method is the following [40]:

- **Align** (utilising rotation, translation and scaling) each shape with the first shape of the population.

- **Repeat:**

  - Calculate the mean shape from the aligned shapes.

  - Normalise the orientation, scale and origin of the mean to suitable defaults.

  - Realign every shape of the population with the current mean.

- **Until:** the process converges.

After the shape alignment, a correspondence between each point of each set can be established. The resulting alignment is the model space. Let $\{\mathbf{x_i}; i = 0...N\}$ denote N shapes. Each shape consists of m 3D landmarks, $\{\mathbf{p_j} = (p_{1j}, p_{2j}, p_{3j}); j = 1....m\}$. Each vector $\mathbf{x_i}$ consists of the landmarks $(p_{11}, p_{21}, p_{31}, p_{12}, p_{22}, p_{32}, ..., p_{1m}, p_{2m}, p_{3m})$. After the alignment into a common coordinate system, each shape $\mathbf{x_i}$ can be represented by a single point in a $3m$-dimensional space.

A *Point Distribution Model* (PDM) which models the variation in the coordinates of the aligned shapes' landmarks can be generated. The variation is assumed to be ellipsoidal and its centre and major axis are calculated. Calculating the centre provides the mean shape, while finding the major axes of the ellipsoidal gives a way of moving around the shape distribution space. The mean shape is calculated as:

$$\bar{\mathbf{x}} = \frac{1}{N} \sum_{i=1}^{N} \mathbf{x_i} \tag{3.1}$$

The aim of the statistical analysis is to approximate the distribution of the landmarks with a linear model of the form:

$$\mathbf{x} = \bar{\mathbf{x}} + \Phi_s \mathbf{b_s} \tag{3.2}$$

where $\bar{\mathbf{x}}$ is the average landmark vector, $\mathbf{b_s}$ is the shape parameter vector of the model, and Φ_s is a matrix of eigenvectors. The matrix Φ_s is obtained by performing a *Principal Component Analysis* (PCA) [87] to the $3m \times 3m$ covariance matrix $\mathbf{S}$:

$$\mathbf{S} = \frac{1}{N} \sum_{i=1}^{N} (\mathbf{x_i} - \bar{\mathbf{x}})(\mathbf{x_i} - \bar{\mathbf{x}})^T \tag{3.3}$$

During the *principal component analysis*, the principal components of $\mathbf{S}$ are calculated as its eigenvectors $\phi_{\mathbf{i}}$ and the corresponding eigenvalues λ_i are sorted (such that $\lambda_i < \lambda_{i+1}$).

New shape examples can be generated by varying the parameters $\mathbf{b_s}$ of equation 3.2. Assuming that the distribution of the data follows a multidimensional Gaussian distribution, the variance of the ith parameter of $\mathbf{b_s}$ across the training set is given by λ_i. If limits in the variation of b_{si} are applied such that $b_{si} \leq \pm 3\sqrt{\lambda_i}$, then it is ensured than the generated shape is similar to those contained in the training class.

A similar technique could be used to model the appearance as well as the shape. Active Appearance Models (AAM) have also been introduced by Cootes *et al.* [37], [39]. In order to build a statistical model of the appearance, each image of the population is warped so that its control points (landmarks) match the mean shape. The images are warped to match the mean shape using a triangulation algorithm [37, 39]. Intensity information from the shape normalised images is sampled over the region covered by the mean shape. The samples can be further normalised to reduce the effect of variation in global lighting. Then PCA can be applied to obtain a linear model of appearance similar to the one defined by equation 3.2:

$$\mathbf{g} = \bar{\mathbf{g}} + \Phi_g \mathbf{b_g} \tag{3.4}$$

where $\bar{\mathbf{g}}$ is the mean normalised grey-level vector, $\Phi_{\mathbf{g}}$ is the eigenvector of appearance variation and $\mathbf{b_g}$ is the model's parameters [37, 39]. Therefore, the shape and the appearance of a subject can be described by the vectors $\mathbf{b_s}$ and $\mathbf{b_g}$. Since there might be correlations between the shape and the appearance of a subject, the vectors are concatenated by applying an additional PCA to both shape and intensity. The following model is

obtained [37, 39]:

$$\mathbf{b} = \begin{pmatrix} \mathbf{W_s b_s} \\ \mathbf{b_g} \end{pmatrix} = \begin{pmatrix} \mathbf{Q_s} \\ \mathbf{Q_g} \end{pmatrix} \mathbf{c} = \mathbf{Qc} \tag{3.5}$$

where $\mathbf{W_s}$ is a diagonal matrix of weights for each shape parameter, allowing for the difference in units between the shape and intensity models. $\mathbf{Q}$ is a set of orthogonal modes and $\mathbf{c}$ is a vector of parameters controlling both the shape and the appearance of the model.

An alternative approach for modelling shape and appearance is to use *Independent Component Analysis* [85], [86] (ICA) instead of PCA. In PCA the objective is to find modes of shape variation which explain the maximal amount of variance in the population. In ICA on the other hand, the objective is to find modes of shape variation which are statistically independent. ICA enables the modelling of input data which do not have Gaussian distribution and can also describe localised variation. ICA is widely used for separation of mixed signals.

Assume that $\mathbf{x}$ is a shape vector which come from a mixture of signals of the form:

$$\mathbf{x} = \mathbf{A} \cdot \mathbf{S} \tag{3.6}$$

where $\mathbf{A}$ is a matrix containing the mixing parameters and $\mathbf{S}$ the source shape. Then, the goal of ICA is to calculate the original shape from the mixed signal. Therefore, ICA calculates the de-mixing matrix $\mathbf{U}$:

$$\hat{\mathbf{S}} = \mathbf{U} \cdot \mathbf{x} \tag{3.7}$$

The matrix $\mathbf{U}$ is found by calculating a cost function. The *Joint Approximated Diagonalization of Eigenmatrices* algorithm (which is based on the joint diagonalisation of the cumulant matrices) can be used to calculate the matrix $\mathbf{U}$ [23]. Üzümcü *et al.* [197] used ICA to build an Active Appearance Model of the heart using 2D MR images. The model was used for cardiac MR segmentation. One major disadvantage of ICA compared

to PCA is that the resulting vectors of shape deviation are not ordered and a method for ordering the independent components is required [197].

Mitchell *et al.* have developed a multistage hybrid active appearance model of cardiac MR images. They use the model for the segmentation of the left and right ventricles of cardiac MR images. Hemarneh *et al.* have developed 2D spatio-temporal active shape models (ASMM) [75], while Sonka *et al.* presented an active appearance motion model [102, 127, 185] (AAMM) which captures the cardiac dynamics as well as the image appearance of the heart. In their framework, shape correspondence was defined in a similar way to the conventional ASM approaches, while temporal correspondence was defined by normalising the cardiac cycle of the images. The temporal normalisation is achieved by selecting a fixed number of frames covering the entire cardiac cycle and using nearest neighbour interpolation to generate image information between these frames. The contour points from the phase normalised images were then concatenated to form a vector. Standard PCA analysis and AAM intensity normalisation schemes were applied to capture the model statistics. Bosch *et al.* have also developed AAMs which capture the spatial and temporal information of echocardiographic sequences [17]. In their framework, correction for non-Gaussian intensity distribution of the appearance is used prior to the construction of the model.

Shape modelling needs a large number of landmarks to be identified in all the images used for the construction of the model. This is a very difficult task which is prone to errors. An exemption to this is the work by Frangi *et al.* [63], [64], [61]. In this approach, a set of *pseudo-landmarks* are used instead of real anatomical landmarks. The *pseudo-landmarks* are generated using the *marching cubes* algorithm [107] to generate a dense triangulation (*pseudo-landmarks*) of the boundary surfaces of each anatomical structure. The automatic landmarking of each image is achieved by using a non-rigid registration algorithm based on B-Splines to propagate the *pseudo-landmarks* from the landmarked atlas to each image of the population. The resulting model included the left and right ventricles.

In order to evaluate how well the propagation of landmarks is performed, Frangi *et al.* used images from 14 healthy volunteers [64]. Three experienced observers were asked to identify seven landmarks on each subject and also on the atlas. The observers identified the seven landmarks twice in two different sessions. The algorithm was able to automatically place the landmarks with an average accuracy of about 2.2*mm* and a precision of about 1.5*mm*. In the same experiments they found that the precision of manual landmarking was about 0.8*mm* [64].

Lötjönen *et al.* used the method developed by Frangi *et al.* [64] to develop statistical shape models of atria, ventricles and epicardium [112]. Their statistical models were constructed using short- (SA) and long-axis (LA) MR images from 25 healthy volunteers. The use of long-axis images in the construction of the atlas provides more information in the direction orthogonal to the short-axis slice. A typical short-axis cardiac MR image has out-of-plane resolution several times larger than in-plane resolution. Therefore, the use of both long- and short-axis cardiac images enables the more accurate localisation of the ventricles in the basal and apical levels. The mapping between the short- and long-axis images was calculated using information, contained in the image files, regarding the coordinate system of the images. An algorithm for the correction of movements in the images was applied prior to the construction of the atlas. The movement correction algorithm optimises the normalised mutual information (NMI) between the SA and LA volumes. The algorithm assumes that the displacement of each slice is independent from the displacements of the other slices. One slice is randomly selected from the SA and LA stack and it is moved to the direction which optimises the NMI [112]. The atria, ventricles and epi-cardium were manually segmented from each image by a clinician. In this approach, instead of registering the grey-level images, the segmented data were utilised. Intensity volumes were generated from the surface data. One volume was selected as the reference and other volumes were aligned to the reference using translation, rotation and isotropic scaling. The variability in the shape was modelled using several analysis methods: PCA, ICA, LPDs (landmark probability distribution).

Horkaew *et al.* developed an approach for the construction of dynamic statistical models for complex topological shapes. Their method uses a harmonic embedding method to establish optimal global correspondence between a set of dynamic surfaces [82], [84]. The proposed method eliminates the need for shape partitioning and introduction of artificial structures to the parameterisation domain. The strength of the method was demonstrated by constructing a statistical model of the left side of the heart that includes the left ventricle, left atrium, aortic outflow track and the pulmonary veins.

3.3.2 Statistical deformation models

Statistical deformation models (SDMs) are very similar to statistical models of shape. One of the key ideas here is to carry out statistical analysis directly on the deformation fields which describe a dense correspondence between the anatomies of two images. There are two main advantages of performing statistical analysis on the deformation fields rather than the shape [173]. Firstly, the resulting statistical model is not limited to a single anatomical structure. It can instead describe the intra- and inter-structure variability across a population. Secondly, the deformation fields can be obtained by non-rigid registration, eliminating the need for segmentation of the images. Rueckert *et al.* used non-rigid registration based on B-Splines to construct a 3D statistical deformation model of the brain using MR images from 25 different subjects [173]. The registration algorithm uses a transformation, $\mathbf{T}$, consisting of a global and a local part (this transformation model is described in detail in section 2.3.2.1):

$$\mathbf{T}(x', y', z') = \mathbf{T}_{global}(x, y, z) + \mathbf{T}_{local}(x, y, z) \tag{3.8}$$

The global transformation describes the overall differences of the images and is represented by an affine transformation. The local transformation describes the local shape differences and is represented by a *deformable* model based on B-Splines:

$$\mathbf{T}_{local}(x, y, z) = \sum_{l=0}^{3} \sum_{m=0}^{3} \sum_{n=0}^{3} B_l(u) B_m(v) B_n(w) \phi_{i+l, j+m, k+n} \tag{3.9}$$

where ϕ denotes a $m_x \times m_y \times m_z$ lattice of control points. The resulting transformation $\mathbf{T}$ maps each point of the reference subject I' to the corresponding point in the anatomy of subject I_i. The goal of SDMs is to perform statistical analysis of the deformations required to map I' to each I_i. However, the effects of the global transformation component, $\mathbf{T}_{global}$, which are the result of differences in position, orientation and overall size of each subject's anatomy, have to be removed prior to statistical analysis. In order to remove any dependency of the local transformation to the global transformation, the following displacements are calculated:

$$
\begin{aligned}
\mathbf{d}(\mathbf{x}, \mathbf{y}, \mathbf{z}) &= \mathbf{J}^{-1} \circ \mathbf{T}_{\mathbf{local}}(\mathbf{x}, \mathbf{y}, \mathbf{z}) \\
&= \sum_{l=0}^{3} \sum_{m=0}^{3} \sum_{n=0}^{3} B_l(u) B_m(v) B_n(w) \mathbf{J}^{-1} \phi_{i+l,j+m,k+n}
\end{aligned}
\tag{3.10}
$$

where $\mathbf{J}$ is the Jacobian matrix of the global transformation. The Jacobian matrix of the affine transformation, with coefficients a_{ij} is:

$$
\mathbf{J} = \begin{pmatrix} a_{01} & a_{02} & a_{03} \\ a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \end{pmatrix}
\tag{3.11}
$$

Suppose that there are n free-form deformations described as control points $\mathbf{C_1}, ..., \mathbf{C_n}$. Each vector of control points $\mathbf{C_i}$ corresponds to a concatenation of $m_x \times m_y \times m_z$ 3D control points, $\mathbf{C_i} = (c_1, ..., c_m)$. These points produce a free-form deformation mapping the anatomy of the reference subject I' to subject I_i. The goal of SDMs is to approximate the distribution of $\mathbf{C}$ using a parameterised linear model [173], similar to the one defined by equation 3.2:

$$
\mathbf{C} = \bar{\mathbf{C}} + \mathbf{\Phi_d b_d}
\tag{3.12}
$$

where $\bar{\mathbf{C}} = \frac{1}{n} \sum_{i=1}^{n} \mathbf{C_i}$ is average control point vector, $\mathbf{b_d}$ is the model's parameter vector and $\mathbf{\Phi_d}$ is the matrix of eigenvectors. As in the statistical shape modelling, the columns

of $\mathbf{\Phi_d}$ are formed by performing PCA analysis on the covariance matrix:

$$\mathbf{S_d} = \frac{1}{n} \sum_{i=1}^{N} (\mathbf{C_i} - \bar{\mathbf{C}})(\mathbf{C_i} - \bar{\mathbf{C}})^T \tag{3.13}$$

Joshi [88] and Gee *et al.* [67] have also perform statistical deformation modelling of the brain. These approaches are closely related to the approach developed Rueckert *et al.* [173] but differ in one important aspect. Rueckert *et al.* performs statistical analysis on the control points of the free-form deformation rather than directly on the deformation fields. The control points parameterise the deformation fields providing a very compact representation. Cseransy *et al.* have used SDMs for the analysis of the hippocampal shape [42] and Wang *et al.* for the analysis of hippocampal asymmetry [202]. Moreover, Davatzikos *et al.* employed a non-rigid registration algorithm to calculate 2D deformation fields of the corpus callosum. Then, statistical analysis of these deformation fields is used to quantify changes between two population groups (in this case male and female subjects) [44]. In a similar way, Bookstein *et al.* studied the shape variability of corpus callosum in patients with schizophrenia and normal control subjects by analysing the deformation maps based on thin-plate splines [15], [49].

Lapp *et al.* combined the concepts of SDMs to AAMs to develop a 4D Active Appearance model of the heart. In this method all corresponding frames of the image sequences are registered to the same coordinate system using non-rigid registration. The resulting deformation fields are assembled to a pseudo 4D deformation fields. Then PCA analysis is performed on the obtained deformation fields [94, 95]. The appearance of the heart is modelled similarly to the AAMs. The images are warped to same coordinate system using the obtained deformation fields and then PCA analysis is performed on their appearance vectors [94, 95]. The temporal misalignment between the image sequences has not been corrected during the construction of the models.

3.4 Atlases of the cardiac function

Rao *et al.* [167, 165, 166] have developed an atlas of the cardiac motion by using MR image sequences of the heart from nine volunteers. The method for extracting the cardiac motion and mapping the motion fields across subjects is described in detail in section 2.6.1. As described in section 2.6.1, the cardiac motion fields are calculated using the method described by Chandrashekara *et al.* [29]. Then a non-rigid registration algorithm is used to calculate the mapping between end-diastolic frames of each untagged image sequence. Petijean *et al.* [154] developed a similar atlas of the motion of the heart. The key difference between the two approaches is that while Rao *et al.*'s approach produces a parametric atlas, this approach produces a non parametric motion atlas which has the ability to preserve the statistical diversity of the motion content of MR data during the contraction phase of the heart [154]. Moreover, Petijean *et al.* performed PCA analysis on the motion data calculating not only the average deformation field but also the significant modes of motion variation. Both approaches provided a temporal correction only for the length of the cardiac cycles [154]. This may not be enough since each heart may have different dynamic properties (one heart may have a longer contraction and a shorter relaxation phase). In another publication Chandrashekara *et al.* presented a new technique for tracking the movement of the myocardium by the use of a statistical model derived from the cardiac motion fields of several healthy volunteers [30]. The approach for the construction of the statistical model consists of three steps. Firstly, the motion fields are extracted for all subjects between end-diastole and end-systole using the previously mentioned approach [29]. Secondly, the extracted motion fields are mapped to the same coordinate system using Rao *et al.*'s approach [167]. Finally, (PCA) is performed on the motion fields in order to calculate the significant modes of variation in the motion fields.

3.5   Other modelling approaches

In recent years a large number of approaches have also been developed for the geometric and biomechanical modelling of the heart. A comprehensive review of these approaches can be found in Frangi *et al.* [62]. A major difference between geometric approaches and probabilistic and statistical approaches for modelling anatomical structures is that geometric models do not contain any information regarding the shape and function variability within the population. This section provides a brief description of such methods. Detailed analysis of geometric modelling methods is outside the scope of this thesis. Frangi *et al.* categorise the geometric modelling methods depending on the way in which the models are geometrically represented [62]:

- Surface modelling

- Volumetric modelling

- Deformable modelling

In early studies of 2D echocardiography images, the left-ventricle is modelled as a simple ellipsoidal [200], [53]. More advanced approaches which use 3D acquisition to reconstruct the LV using global and hierarchical parameterised models have also been developed. Approaches which use global parameterised models are based on simple geometric models which enable a rough shape approximation of the anatomical structure [62]. Cauvin *et al.* used a combination of an ellipsoid and a cylinder (a *truncated bullet*) to model the left ventricle [26], while Metaxas and Terzopoulos proposed *superquadrics* to model simple objects with a limited number of parameters [125]. Park *et al.* extended the *superquadrics* model by introducing parameter functions like radial and longitudinal contraction, twisting and long-axis deformation [141]. Bardinet *et al.* used a combination of a *superquadric* and free-form deformations to model the LV [10, 11]. In their approach, a *superquadric* model is used to fit the inner and outer surfaces of the left ventricle and then a free-form deformation model is used to refine the crude approximation of the *superquadric* model. The method is applied in order to estimate the LV wall motion

[10]. This is achieved by deforming the entire model (*superquadric* and FFD) to match the anatomy of the first frame and modifying only the FFD in the subsequent frames.

Hierarchical parameterised models contain hierarchical representation of the model, where each level describes the model in more detail. Gustavsson *et al.*, for instance, used a truncated ellipsoid to obtain a coarse position of the left ventricle from contours drawn into the short-axis and three apical views of ultrasound images [72]. The model is refined using cubic B-Splines curves which approximate the manually segmented contours in multiple views.

There are also surface models which incorporate temporal and spatial variation of the LV shape. Tu *et al.* introduced a 4D model-based LV boundary detector for 3D CT image sequences [193]. The method applies a spatio-temporal gradient operator in spherical coordinates which is only sensitive to moving edges. An iterative approach refines the boundaries of the model by discarding edge points which are far away from the global model. Spherical harmonics are used to parameterise the model as the refinement approach proceeds. Faber *et al.* [56] uses a discrete 4D model to segment the left ventricle from SPECT and MR images using a relaxation labelling scheme [90]. In this approach, the endocardial and epicardial surfaces are modelled as a discrete template which is defined in a mixed spherical/cylindrical system co-axial with the long-axis of the left ventricle. Each point in the template represents a radius connected to the long-axis. During the segmentation procedure, information from neighbouring points in both the spatial and temporal domains is used. More recently, Declerck *et al.* introduced a 4D continuous planispheric transformation which enables the tracking of LV motion [46]. The planispheric transformations have the advantage that they are continuous in both spatial and temporal domains. Sermesant *et al.* have developed a *deformable* biomechanical model of the heart by combining information from various imaging modalities [180, 181]. The construction of the model involved three main stages: generation of the geometrical mesh, non-rigid registration of the mesh in images of various modalities and finally retrieving information from each volumetric image to the volumetric mesh.

In recent years the trend is moving away from surface based modelling of the heart to more accurate volumetric based modelling of the cardiac motion [54]. A review on imaging three-dimensional cardiac function and on analysis of cardiac deformation can be found by O'Dell and McGulloch [132]. Shi *et al.* extended an earlier surface modelling approach [182] to produce a biomechanical model by combining surface and motion information from magnitude and phase-contrast MR images respectively [183]. Haber *et al.* developed a model of biventricular geometry using finite elements [73]. Papademetris *et al.* developed a cardiac biomechanical modelling approach using a deformation model inspired by continuum mechanics [136, 138, 135, 137]. In this approach a dense triangulation field is calculated using point correspondence which is obtained by using an extension of the work developed by Shi *et al.* [183]. A linear elasticity model is used for the estimation of a dense motion field. The model accounts for the muscle fiber directions in the left ventricle. The motion field is used to calculate the deformation of the left ventricle's wall in terms of strain in cardiac specific directions [137].

3.6 Conclusions

There are three main methods for modelling the shape and function of the heart. *Probabilistic atlases* contain information regarding the shape and function of structures in the form of tissue probability maps (i.e. the probability of a voxel to belong to a certain anatomical structure). A disadvantage of probabilistic atlases is that they can not provide information regarding the type of variability. *Statistical atlases* enable the calculation of additional statistical information regarding the type of variability. The main categories of statistical models which have been reviewed in this chapter are: the statistical shape and appearance models and statistical deformation models. In statistical shape models, analysis on the shape of the heart is performed. Statistical models of appearance go one step further by analysing not only the shape of the heart but also its appearance. Statistical deformation models are very similar to statistical models of shape. One of the key ideas here is to perform statistical analysis directly on the deformation fields which describe a dense

correspondence between the anatomies of two images. Furthermore, there are approaches for the construction of statistical atlases of the cardiac motion. Finally, in recent years a large number of approaches have also been developed including *geometric modelling* and *biomechanical modelling* of the heart. A major difference between these models is that they do not contain information regarding the shape and function variability within the population.

Contrary to most of the methods reviewed in this chapter, the work presented in this thesis aims to build 4D probabilistic and statistical atlases of the cardiac anatomy and function. In this chapter a number of approaches for building atlases of the cardiac anatomy and function have been presented. However, none of the above approaches captures information regarding anatomical and functional variability into a single model. The probabilistic and statistical atlases presented in this thesis contain information regarding anatomical variability and functional variability. These models describe the shape of a healthy heart and how the shape of the heart changes over the cardiac cycle.

Chapter 4

Spatio-temporal alignment of cardiac MR image sequences

This chapter presents a novel method for the spatio-temporal registration of 3D cardiac MR image sequences. The algorithm uses a 4D transformation model which is separated into decoupled spatial and temporal components. Firstly, a registration algorithm is presented which has the ability to correct spatial misalignment of affine nature between the image sequences. It also has the ability to correct temporal misalignment which may be the result of differences in the length of the cardiac cycles of the subjects and in the temporal acquisition parameters [152]. Secondly, a *deformable* spatial transformation model is introduced which enables the better spatial registration of the image sequences. With the introduction of the spatial *deformable* transformation model, the algorithm corrects not only global spatial shape differences but also local differences in the shape of the hearts [153].

4.1 Why spatio-temporal registration is needed

Since the heart is undergoing spatially and temporally a varying degree of motion during the cardiac cycle, 4D cardiac image registration algorithms are required when registering two cardiac MR image sequences. Spatial alignment of corresponding frames of the im-

age sequences (e.g. the second frame of one image sequence with the second frame of the other) is not sufficient since these frames may not correspond to the same position in the cardiac cycle of the hearts. This is due to differences in the acquisition parameters (trigger offset from R-wave and different intervals in the acquisition of consecutive frames), differences in the length of cardiac cycles (e.g. one cardiac cycle may be longer than the other) and differences in the dynamic properties of the hearts (e.g. one heart may have a longer contraction phase and a shorter relaxation phase). Figure 4.1 shows an example of how differences in the trigger offset (t_{delay}), in the frequency of the acquisition of consecutive frames and in the length of the cardiac cycles affect the temporal alignment of two image sequences. An affine temporal transformation can correct for differences in the acquisition parameters and in the length of the cardiac cycles. However, the correction for the differences in the dynamic properties of the hearts requires more complex transformations. The spatio-temporal alignment enables comparison between corresponding anatomical positions and corresponding positions in the cardiac cycle of the hearts. This can be seen from the volume curves of the left ventricles in figure 4.2. In 4.2(a) there is no temporal registration between the image sequences. This may lead to the comparison of frames in different positions in the cardiac cycle (the position were one heart is at peak contraction while the other heart is still during the contraction phase). What is needed is to find the temporal relationship between the two image sequences in order to compare corresponding positions in the cardiac cycles (figure 4.2 (b)).

The method developed by Rao *et al.* [167, 165, 166] for the direct comparison of motion fields between different subjects suffers from the problem mentioned above. The deformation fields calculated for each subject contain the deformation between adjacent frames. When non-rigid registration is used to map the deformation fields to the same reference system, the mapping is performed only in the spatial domain. However, the size of motion fields from one frame to the next will depend on the offset (t_{offset}) in the acquisition of consecutive frames. The larger this acquisition offset is, the larger the cardiac motion between adjacent frames would be. Therefore, comparing a number of motion fields without knowing the temporal relation between the cardiac cycles introduces

errors in the observations.
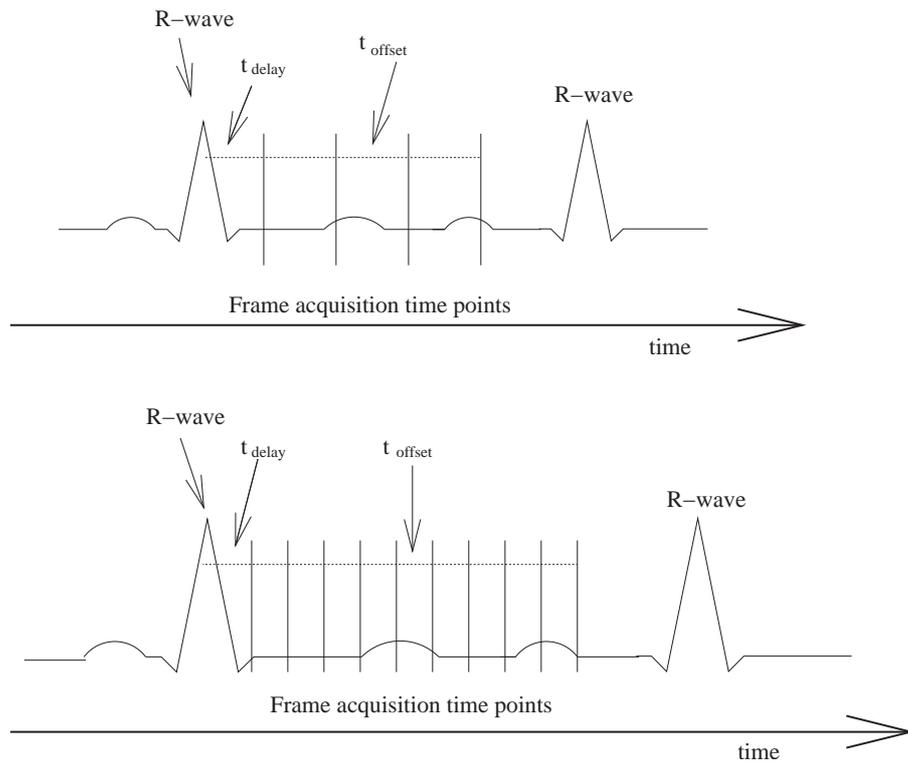
Figure 4.1: An example of two MR image acquisitions with different initial delay in the acquisition of the first frame and offset in the acquisition of consecutive frames. The image sequences have different numbers of frames. It is clear that comparison between corresponding frames is not sufficient since these frames correspond to different positions in the cardiac cycles of the heart. Temporal registration is needed to establish correspondence between these frames.

4.2 Contributions

This chapter makes the following contributions:

- The development of a new method for the spatial and temporal alignment of a number of cardiac MR image sequences to the same coordinate system. The registration method uses a 4D transformation mapping decoupled into separate temporal and spatial components. This spatio-temporal transformation mapping enables the direct comparison of the anatomy and function of cardiac MR image sequences to be
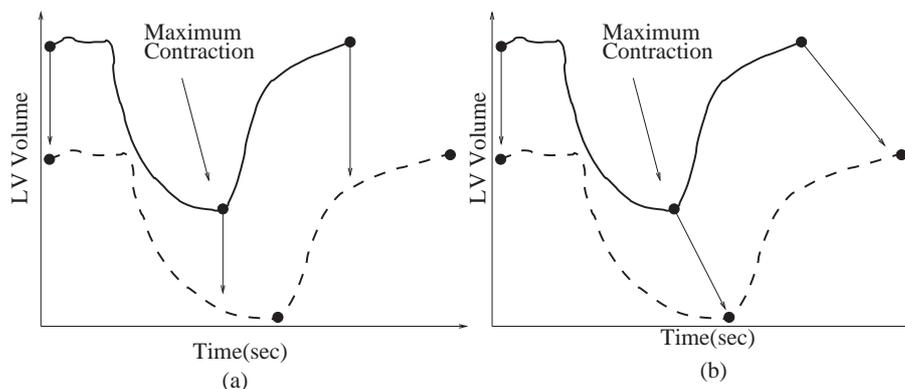
Figure 4.2: An example of why spatio-temporal registration is needed. (a) the identity temporal relation has been assumed, (b) the temporal relation provides a mapping between corresponding positions in the cardiac cycle.

made. The registration method aims to minimise any patient specific temporal and spatial variability which is caused as a result of spatial differences in the position, size and orientation of the cardiac anatomies and temporal differences in the length of the cardiac cycle of the hearts.

- The extension of the spatio-temporal registration method with the introduction of a *deformable* spatial transformation. This extension enables the correction of any specific temporal and spatial variability which is caused as a result of spatial local shape differences between the cardiac anatomies in addition to the global spatial and temporal differences.

4.3 Spatio-temporal registration of cardiac MR image sequences

The proposed registration approach uses information provided only by the intensity of the image sequences. There is no need to identify feature positions in the cardiac cycles (e.g. the positions of the peak contraction) of the hearts and also no need to segment the images.

A 3D cardiac image sequence can be represented as an ordered sequence of n 3D images $S_k(x, y, z)$ with a fixed field of view Ω_{S_k} and an acquisition time $t_k, t_k < t_{k+1}$, in the temporal direction. The resulting image sequence can be viewed as 4D image $S(x, y, z, t)$ defined on the spatio-temporal domain $\Omega_{S_k} \times [t_1, t_n]$. The goal of 4D image registration is to relate each point in one image sequence to its corresponding point in the reference image sequence. In this case the transformation $\mathbf{T} : (x, y, z, t) \rightarrow (x', y', z', t')$ where $(x', y', z', t') = (x + u, y + v, z + w, t + \tau)$ maps any point of one image sequence $S(x, y, z, t)$ into its corresponding point in the reference image sequence $S'(x', y', z', t')$. In general, the 4D mapping $\mathbf{T}$ can have the following forms:

1. $\mathbf{T}(x, y, z, t) = (x'(x, y, z, t), y'(x, y, z, t), z'(x, y, z, t), t'(x, y, z, t))$

2. $\mathbf{T}(x, y, z, t) = (x'(x, y, z, t'), y'(x, y, z, t'), z'(x, y, z, t'), t'(t))$

3. $\mathbf{T}(x, y, z, t) = (x'(x, y, z), y'(x, y, z), z'(x, y, z), t'(t))$

The first type of 4D mapping is under-constrained and it would be very difficult to find the optimal transformation mapping without introducing additional constraints or employing heuristic algorithms to reduce the search space. Furthermore, this mapping is undesirable due to the fact that two points from the same temporal frame can correspond to two points in two different time frames in the source image sequence. This is not intuitive or desirable. The second type of 4D mapping allows a different spatial mapping for each frame of the image sequences. Such a 4D mapping is preferable when the differences in the shape of the hearts are not constant within the cardiac cycle. However, finding the optimal transformation of such type would have a very high computational complexity. The last form of 4D mapping separates the temporal alignment of the image sequences from the spatial alignment. Hence, this form of 4D mapping decouples the misalignment caused by spatial differences from misalignment caused by temporal differences. Furthermore, its computational complexity is substantially lower than that of the previous type of transformation mapping since a common spatial transformation for each frame of the sequences is required.
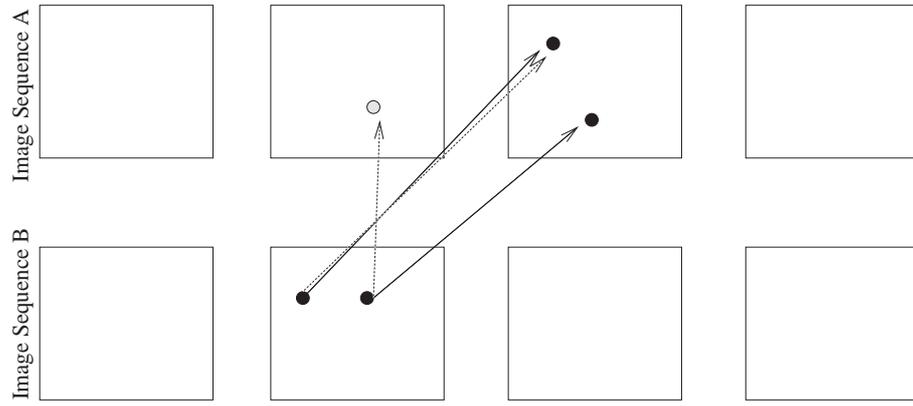
Figure 4.3: The mapping has to ensure that different regions in a 3D image will not be warped differently in the temporal direction (dashed line). Each voxel in the temporal frame t of the image sequence B will map to another voxel in frame t' of image sequence A.

The 4D mapping used in this thesis is of the form:

$$\mathbf{T}(x, y, z, t) = (x'(x, y, z), y'(x, y, z), z'(x, y, z), t'(t)) \tag{4.1}$$

It can be of a subvoxel displacement in the spatial domain and of a sub-frame displacement in the temporal domain.

Since the temporal and the spatial domains are different, the 4D mapping can be resolved into decoupled spatial and temporal components $\mathbf{T}_{spatial}$ and $\mathbf{T}_{temporal}$ respectively, where

$$\mathbf{T}_{spatial}(x, y, z) = (x'(x, y, z), y'(x, y, z), z'(x, y, z)) \tag{4.2}$$

and

$$\mathbf{T}_{temporal}(t) = t'(t) \tag{4.3}$$

One consequence of this decoupling is that each temporal frame t in image sequence S will map to another temporal frame t' in image sequence S', ensuring causality and preventing different regions in a 3D image $S_k(x, y, z)$ to be warped differently in the temporal direction by $\mathbf{T}_{temporal}$ (figure 4.3). Another advantage of the decoupled model is the ability to interchange easily the type of the temporal transformation $\mathbf{T}_{temporal}$ and spatial transformation $\mathbf{T}_{spatial}$.
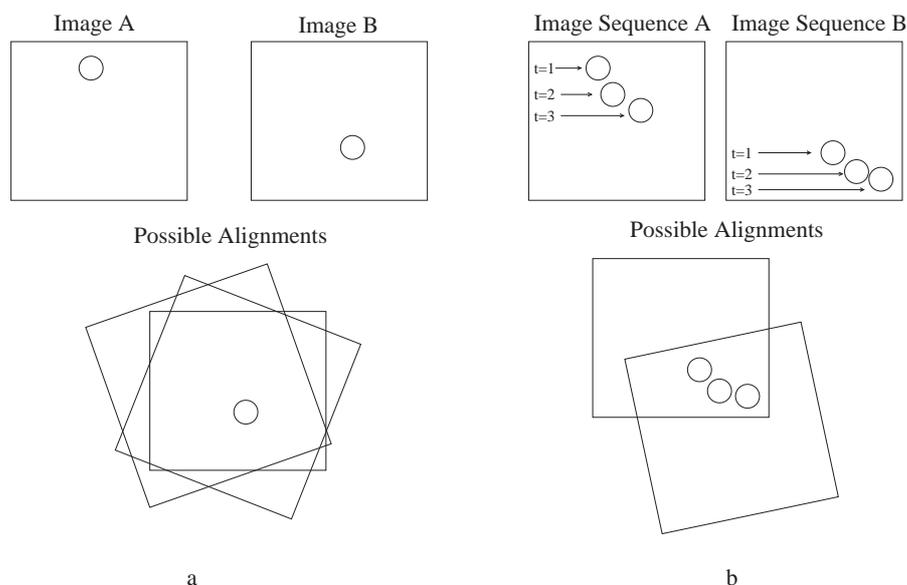
Figure 4.4: The spatio-temporal mapping can help to solve spatial ambiguities (adapted from [25]).

A similar spatio-temporal registration approach has also been used by Caspi *et al.* for the registration of different video sequences [25]. In this approach a decoupled spatio-temporal transformation model is used to establish spatial and temporal correspondence between two 2D video sequences (2D+time). The algorithm uses an affine temporal model correcting for different offsets between the sequences and different frame rates. The spatial model is a 2D projective transformation correcting for different internal and external camera parameters. Two optimisation methods have been developed. The first is based on the trajectories of segmented objects while the second one on intensity information only.

The use of spatio-temporal registration has a number of potential advantages. The registration could be based not only on information provided by each frame but also on subframe information. Furthermore, introducing a temporal transformation can help to solve a number of spatial ambiguities. An example of spatial ambiguities is shown in figure 4.4. In figure 4.4 (a) there are two images of two spherical objects. There is a number of possible alignments between them. In figure 4.4 (b) the same spherical objects are displayed over time. It can be seen how much easier it is to find a mapping between the objects if temporal information is available. However, introducing the temporal trans-

formation can also result in spatio-temporal ambiguities. These occur when the temporal alignment corrects for spatial differences or the spatial alignment corrects for temporal differences [25].

As mentioned above, in this thesis the spatio-temporal alignment of image sequences is enabled using only image information. Alternative ways for the temporal alignment of image sequences include the use of the ECG signal and the ventricular pressure. The ECG signal records the electrical activations of the heart which can be related to the phases of the cardiac cycle (figure 1.7). However, during the MR image acquisition the ECG signal is distorted due to the presence of a strong magnetic field. Recently, Kachelriess *et al.* presented a technique for the extraction of information about the cardiac motion directly from the measured raw data of spiral CT of the heart [89]. The measure is called the *kymogram* and is a local measure of the heart motion as a function of time t or as a function of projection angle α. The *kymogram* can be used to detect the beginning and the end of the contraction but it does not provide enough information regarding the other phases of the cardiac cycle. The ventricular pressure could also be used to find temporal correspondence between two cardiac cycles. However, the pressure information tends not to be recorded and is usually not available in clinical practice.

4.4 Affine spatio-temporal registration

4.4.1 Spatial alignment of 4D image sequences

The aim of the spatial part of the transformation is to relate each spatial point of an image to a point of the reference image, i.e. $\mathbf{T}_{spatial} : (x, y, z) \rightarrow (x', y', z')$ maps any point (x, y, z) of a particular time frame t in one image sequence into its corresponding point (x', y', z') of another particular time frame t' of the reference image sequence. The simplest choice of $\mathbf{T}_{spatial}$ is a rigid transformation which has six *degrees of freedom* corresponding to translation and rotation. In this approach an affine transformation is selected which is a more general class of transformations with 12 *degrees of freedom*

utilising scaling and shearing in addition to translation and rotation (3 *degrees of freedom* for rotation, 3 for translation, 3 for scaling and 3 for shearing). Such a 3D affine transformation can be written as:

$$\mathbf{T}_{spatial}^{global}(x, y, z) = \begin{pmatrix} \theta_{11} & \theta_{12} & \theta_{13} \\ \theta_{21} & \theta_{22} & \theta_{23} \\ \theta_{31} & \theta_{32} & \theta_{33} \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} + \begin{pmatrix} \theta_{14} \\ \theta_{24} \\ \theta_{34} \end{pmatrix} \tag{4.4}$$

where the coefficients θ parameterise the twelve *degrees of freedom* of the transformation.

4.4.2 Temporal alignment of 4D image sequences

The temporal alignment of two image sequences aims to find the transformation function $\mathbf{T}_{temporal}$ which establishes a correspondence between time t in one image sequence and the corresponding time t' in the reference image sequence, thereby establishing a correspondence between corresponding time points in two cardiac cycles. In addition to differences in the length of the cardiac cycle, the temporal alignment of two image sequences is further complicated by the fact that the acquisition of cardiac MR image sequences typically depends on two parameters (figure 4.1): the first parameter describes the delay, t_{delay}, after which the MR acquisition starts while the second parameter, t_{offset}, describes the temporal resolution of the image sequence.

In the current approach the temporal transformation, $\mathbf{T}_{temporal} : (t) \rightarrow (t')$, is represented as an affine transformation of the following form:

$$\mathbf{T}_{temporal}(t') = \alpha t + \beta \tag{4.5}$$

Here α accounts for scaling differences between the two image sequences (different length of cardiac cycles or different temporal resolution) while β accounts for translation differences between the two image sequences. Translation differences may be introduced by the acquisition parameter t_{delay} and by missing frames.

4.4.3 Similarity and optimisation

The optimal transformation $\mathbf{T}$ is found by maximising a voxel similarity measure, normalised mutual information (NMI) [187], as a measure of spatio-temporal alignment of the image sequences. Normalised mutual information has previously been used successfully in mono- and multi-modal image registration approaches [33, 113, 157, 204]. The use of a voxel-based similarity measure eliminates any need for any feature detection and segmentation of structures as the epi- or endo-cardial surfaces. The normalised mutual information between two image sequences is defined as:

$$NMI(S', S) = \frac{H(S') + H(S)}{H(S', S)} \tag{4.6}$$

where $H(\cdot)$ denotes the marginal entropy of each image sequence and $H(\cdot, \cdot)$ is the joint entropy of two image sequences. The normalised mutual information can be calculated directly from the joint intensity histogram of the two sequences over their spatio-temporal domain of overlap $\Omega_{S'} \times [t_{S'_1}, t_{S'_n}] \bigcap \mathbf{T}(\Omega_S \times [t_{S_1}, t_{S_n}])$. Figure 4.5 provides an example of the joint intensity histogram over the spatio-temporal domain of overlap of two image sequences. The temporal transformation, $T_{temporal}$, has only 2 *degrees of freedom* while the spatial one has 12. Therefore, the spatio-temporal registration has 14 *degrees of freedom* in total. A simple iterative uphill method can be used to optimise the transformation. The optimisation is carried out to calculate the optimal transformation :

$$\arg \max_{\alpha, \beta, \Theta} NMI(S', \mathbf{T}(S)) \tag{4.7}$$

where S' is the reference image sequence and $\mathbf{T}(S)$ is the transformed image sequence S. The method is the same as the one described by Studholme [186].

Listing 1 describes the algorithm. The idea behind this optimisation approach is to use a large step size to obtain a rough estimate of the optimal transformation and then, reduce the step size to obtain a better estimate. However, care must be taken when selecting the step size. A small step size may not be enough to reach the optimal solution and

the optimisation procedure may become stuck into a local minima, while a large step may move the transformation far away from the optimal solution (figure 4.6) . In figure 4.6 is shown that a small step size is not enough to overcome the local minima, while a large step size moves the transformation far away from the correct solution, trapping the optimisation procedure in another local minima.

Listing 1 The uphill descent optimisation approach.

1: Given a current estimate of the transformation, $\mathbf{T}_n = \mathbf{T}_0$.
2: **repeat**
3:     **repeat**
4:         **evaluate** the similarity measure for a set of 29 transformations $\mathcal{T}(\mathbf{T}_n)$. This set consists of the transformations resulting by increasing and decreasing each *degree of freedom* of the current estimate by a certain step, s, and the current transformation $(14 dofs \times 2 + 1)$
5:         **select** the best estimate of the transformation with respect to the similarity NMI:

$$\mathbf{T}_{n+1} = \max_{\mathbf{T} \in \mathcal{T}(\mathbf{T}_n)} \{(NMI(S', \mathbf{T}(S))\} \tag{4.8}$$

6:         **if** $\mathbf{T}_{n+1}$ results to a greater value in the similarity measure than $\mathbf{T}_n$, **then**
7:           $\mathbf{T}_{n+1}$ becomes the current estimate: $\mathbf{T}_n = \mathbf{T}_{n+1}$
8:         **end if**
9:     **until** maximum number of iterations
10:     **subdivide** step size s
11: **until** maximum number of subdivisions

During the optimisation new voxel values are generated in the temporal domain using linear interpolation and trilinear interpolation in the spatial domain. Linear interpolation has high computational efficiency. The use of a more advanced interpolation method, like sinc interpolation and B-Spline interpolation, will make the registration process slower without substantially improving its accuracy. A review on interpolation methods in medical image processing can be found in [101].
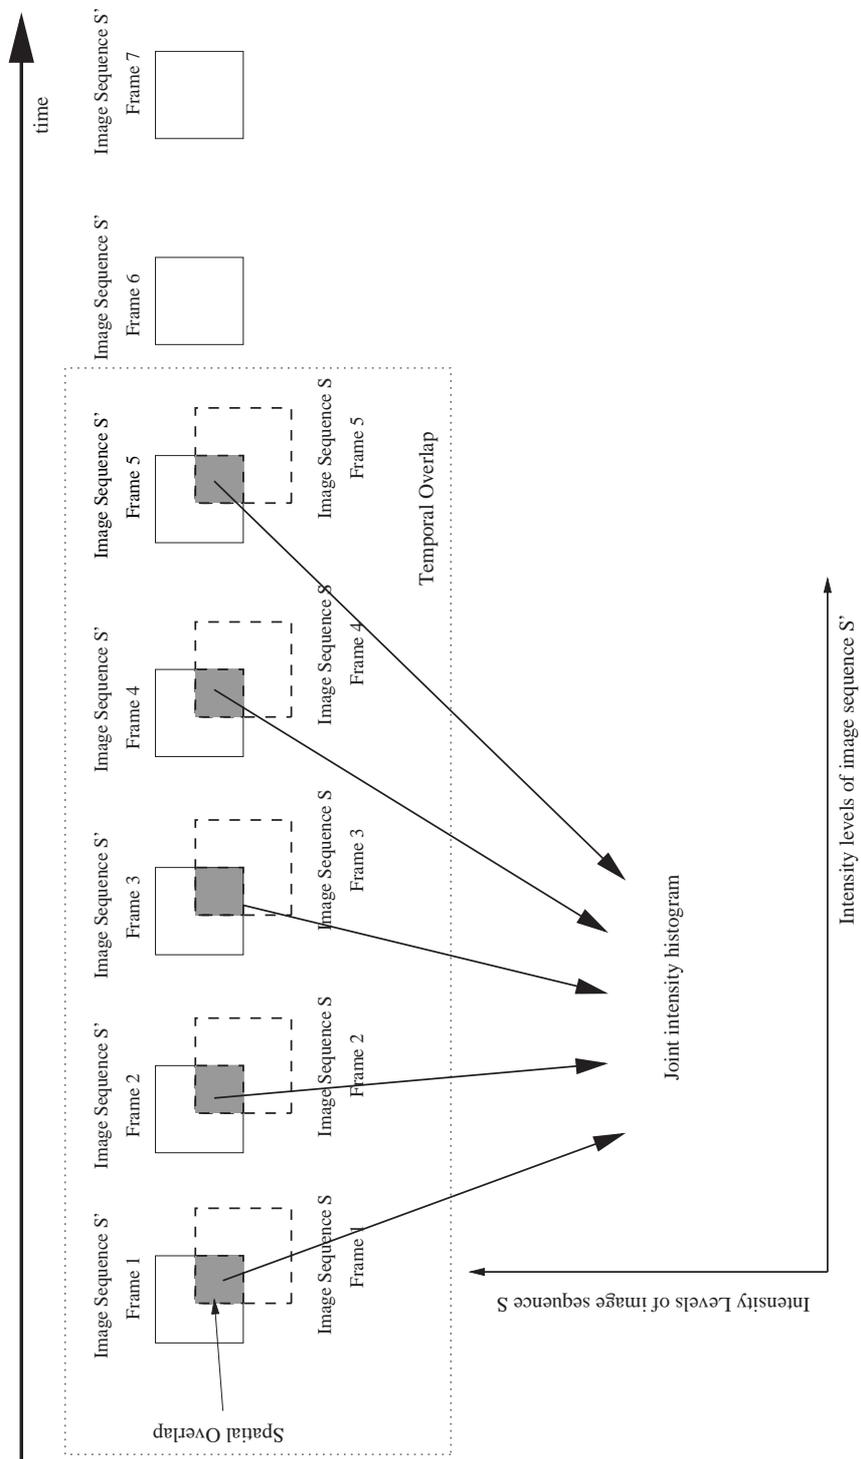
Figure 4.5: The joint intensity histogram over the spatio-temporal domain of overlap of two image sequences. The joint intensity histogram is formed by using the intensity levels of the spatial (grey areas in the frames) and temporal (dotted box) overlap of the two image sequences

Parameter	Setting
No. of bins	64
No. of iterations	20
No. of steps	4
Length of spatial steps	5mm
Length of temporal steps	5msec

Table 4.1: Registration parameters used for the affine spatio-temporal registration.

Table 4.1 contains the parameters used for the registrations:

- **Number of bins:** the number of partitions of the histogram for estimating the similarity measure. Practical evaluation has shown that the value of 64 bins is a good choice for cardiac MR images. A larger number of bins will make the registration process slower without improving accuracy. Furthermore, a smaller number of bins will adversely affect the results since intensity information from a number of different types of tissue could correspond to the same bin in the histogram.

- **Number of steps:** the number of times the optimisation process has to be repeated. In each iteration the length of step is halved. If the length of the final step is less than the dimensions of the sequence's voxels then, subvoxel accuracy can be achieved.

- **Number of iterations:** the maximum number of times that the optimal transformation estimate, $\mathbf{T}_{n+1}$, is calculated for a certain step size.

- **Length of steps:** the initial length of the step. The length of the step is defined in *mm* in the spatial domain and *msec* in the temporal domain.

4.4.4 Evaluation of the method

To evaluate the affine spatio-temporal registration algorithm, cardiac MR image sequences from seven volunteers have been acquired. All image sequences used in the experiments were acquired on a Siemens Sonata 1.5T scanner using a TrueFisp pulse sequence. For the reference subject 32 different time frames were acquired (the length
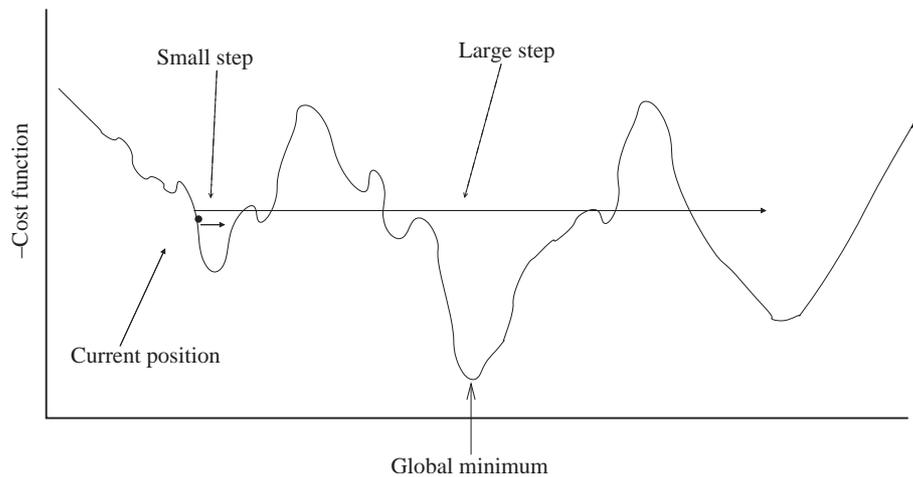
Figure 4.6: An example of cost function and the optimisation step. A small step is not enough to pass the local minima. A large step will trap the optimisation in another local minima.

of the cardiac cycle was 950*msec*). The frames of the reference sequence had a resolution of $256{\times}192{\times}46$ with a pixel size of $0.97mm{\times}0.97mm$ and a slice thickness of $3mm$. The other six 4D cardiac MR images were registered to the reference subject. These sequences had a pixel size between $1.36mm{\times}1.36mm$ and $1.48mm{\times}1.48mm$ and a slice thickness of $10mm$. For these subjects 15-20 different time frames were acquired (cardiac cycles' length between $600msec$ to $800msec$). In one image sequence the temporal acquisition parameters were estimated because the real parameters were unknown. An initial estimate of transformation was provided due to the large differences in the length of the cardiac cycle of each image and also due to the large variety in the position and orientation of the cardiac anatomies. The temporal part of the transformation was initially calculated in order to match the temporal ends of the 4D cardiac MR images. The initial estimate of the spatial part was calculated from the rigid registration of three manually selected cardiac anatomical positions in the first frame of the image sequences. These anatomical positions are the apex of the left ventricle, the anteroseptal at the base of the heart and the inferoseptal at the base of the heart.

Figures 4.7 and 4.8 show examples of the affine spatio-temporal registration of two random subjects . In the first row of figures 4.7 and 4.8 are the short-axis (a) and the long-axis (b) views of particular slices of the reference sequence and a temporal view of a

short-axis line of the reference sequence (c). The second row of figures 4.7 and 4.8 shows the corresponding views before the registration (without mapping the temporal ends of the image sequences). The third row of figures 4.7 and 4.8 shows the corresponding views before the registration (after mapping the temporal ends of the image sequences). Finally, the bottom row of the figures shows the corresponding views after affine spatio-temporal registration. The isolines of intensity of the reference subject are overlayed on every image. The figures show a large improvement in the alignment of the images after the registration. In figures 4.7 and 4.8 (j)-(l) the spatio-temporal overlap of the two image sequences has been substantially improved after their spatio-temporal affine registration. Figure 4.9 shows another example of registration of a subject to the reference image sequence (the figure shows the registration over the cardiac cycle). Figure 4.9 (a) shows different frames of the two image sequences (the top part of each frame is the reference image sequence) before the registration, while 4.9 (b) shows the corresponding frames from the same two image sequences after affine registration. Before the registration the two image sequences follow different motion patterns, while after affine spatio-temporal registration they follow similar motion patterns.

In order to evaluate how well the temporal alignment has been performed, the volume of the left ventricle in each frame was measured before the registration and after the registration. Figure 4.10 (a) contains the same volume curves of the left ventricles before the registration, while figure 4.10 (b) contains the volumes of the left ventricle over time after spatio-temporal registration. Before the registration, the contraction and relaxation properties of each heart are completely different (figure 4.10 (a)). For example, the position of the peak contraction has large variability. In contrary, after affine 4D registration the hearts have similar contraction and relaxation properties (figure 4.10 (b)).

The quality of the registrations in the spatial domain has also been assessed quantitatively by calculating the volume differences of the left ventricle between each subject and the reference subject (before and after registration). The volume difference of the left

ventricle between two images is defined by [34]:

$$\Delta_{volume}(I', I) = \frac{V_{I'} - V_I}{V_{I'}} \times 100\%$$ (4.9)

where Δ_{volume} is the absolute volume difference between the two image, $V_{I'}$ is the volume of the left ventricle of the reference image and V_I is the volume of the left ventricle of the other image. In order to calculate the volume difference for the entire image sequences, $\Delta_{volume}(S', S)$, the mean $\Delta_{volume}(S'_i, S_i)$ for each pair of corresponding frames is calculated. The volumes of the left ventricle are calculated by the use of the manually segmented images (the segmentation was performed by an expert). The quality of the registration was also evaluated by calculating the volume overlap between the ventricles and the myocardium of each subject and the reference subject. The volume overlap for an object O is defined as:

$$\Delta(S', S) = \frac{2 \times |S' \bigcap S|}{|S'| + |S|} \times 100\%$$ (4.10)

Here S' denotes the voxels in the reference (target) image part of object O and S denotes the voxels in the other image part of object O. In order to calculate the volume overlap for the entire image sequences, the mean $\Delta(S'_i, S_i)$ for each pair of corresponding frames is calculated.

Volume overlap	Before registration	After registration
Left ventricle	54.43%	79.94%
Right ventricle	54.32%	74.53%
Myocardium	47.39%	67.93%

Table 4.2: The mean volume overlap before and after the 4D affine registration.

Table 4.2 shows the mean volume overlap before and after the registration for the left ventricle, the right ventricle and the myocardium (the mean volume overlap is calculated over their temporal domain of overlap). The mean volume difference of the left ventricle (calculated over their temporal domain of overlap) between the image sequences was 29.55% before the registration and has been reduced to 13.90%. The mean error in mapping the positions of the left ventricles' peak contraction was 80*msec* before the reg-

istration (taking into account that the temporal resolution of the reference image sequence was 30*msec* this error corresponds to 2.6 frames) while after the registration it is reduced to 40*msec* (which corresponds to 1.3 frames). This can be also observed in figure 4.10. The figure clearly shows that the temporal features of the volume curves are substantially better aligned after the spatio-temporal registration even though no segmentation of the left ventricle was used during the registration.

It is not expected that the images would be perfectly aligned either in the temporal domain or in the spatial domains. This is due to the nature of the temporal and spatial components of the 4D transformation. For example mapping the cardiac anatomy of two different subjects most likely requires a non-rigid transformation. These spatial differences can be modelled to a certain extent by the current $\mathbf{T}_{spatial}$. Furthermore, the temporal transformation (equation 4.5) can only be applied to the entire cardiac cycle and not parts of it. Therefore, it is not possible to address the temporal misalignment caused by differences in the dynamic properties of the contraction and relaxation phases of the cardiac cycle (e.g. one heart may have a longer contraction phase than relaxation phase). An example of a temporal misalignment that cannot be recovered by the current $\mathbf{T}_{temporal}$ can been seen in the volume curves (figure 4.10). In this figure, the volume of the left ventricle of the reference image sequence (figure 4.10 (a), (b)) appears stable after a certain time period which means this heart has a long isovolumetric relaxation phase. However, this is not true for all the other images (figure 4.10 (a)).
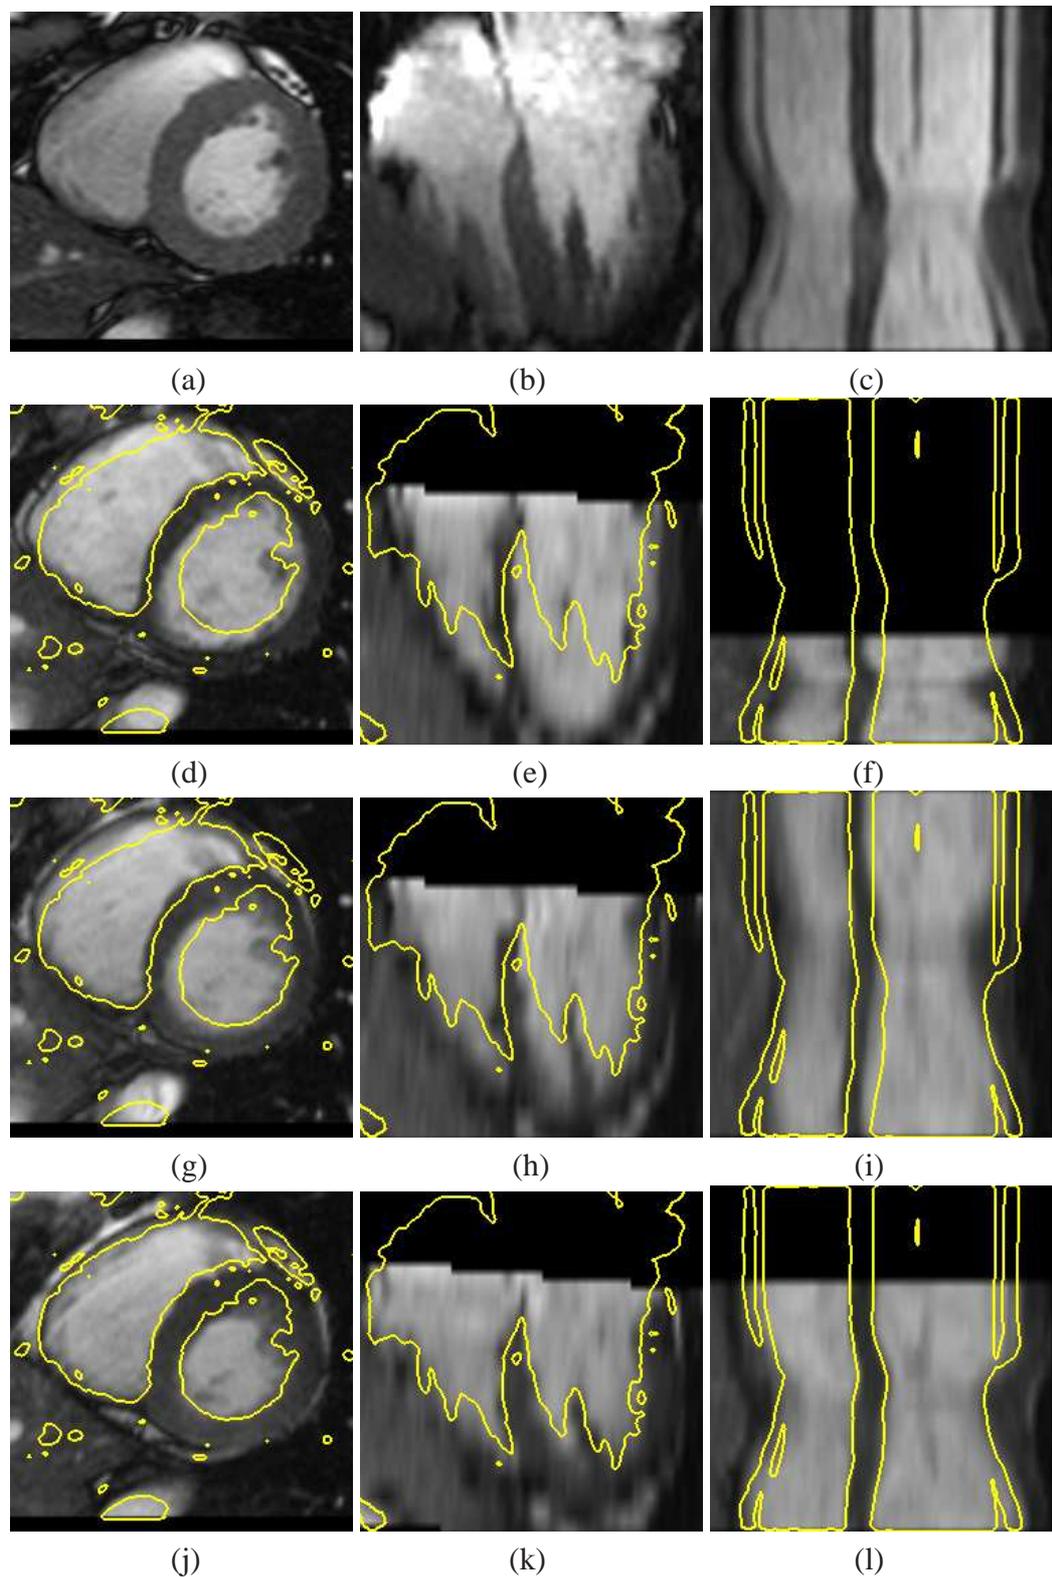
Figure 4.7: Results of the 4D affine cardiac MR registration algorithm: (a) the short-axis view of the reference subject, (b) the long-axis view of the reference subject, (c) a temporal view of a short-axis line of the reference subject, (d)-(f) the corresponding views before the registration (without mapping the temporal ends of the image sequences), (g)-(i) shows the corresponding views before the registration (after mapping the temporal ends of the image sequences), (j)-(l) show the corresponding images after the 4D registration.
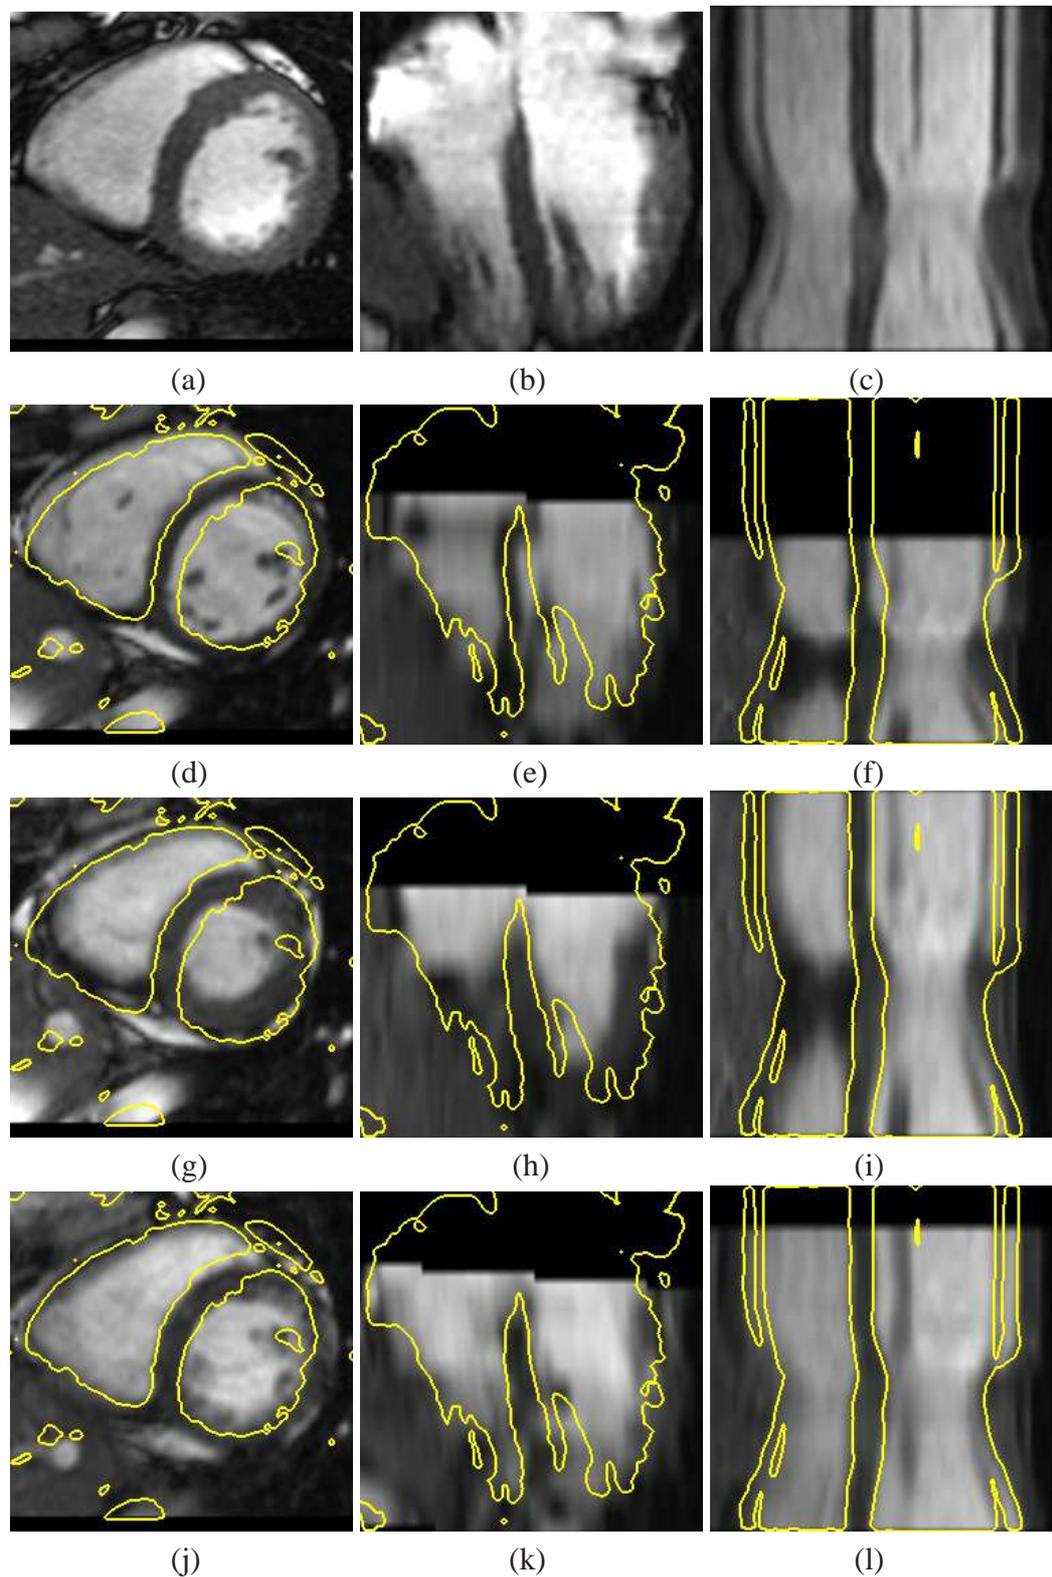
Figure 4.8: Results of the 4D affine cardiac MR registration algorithm: (a) the short-axis view of the reference subject, (b) the long-axis view of the reference subject, (c) a temporal view of a short-axis line of the reference subject, (d)-(f) the corresponding views before the registration (without mapping the temporal ends of the image sequences), (g)-(i) the corresponding views before the registration (after mapping the temporal ends of the image sequences), (j)-(l) the corresponding images after the 4D registration.
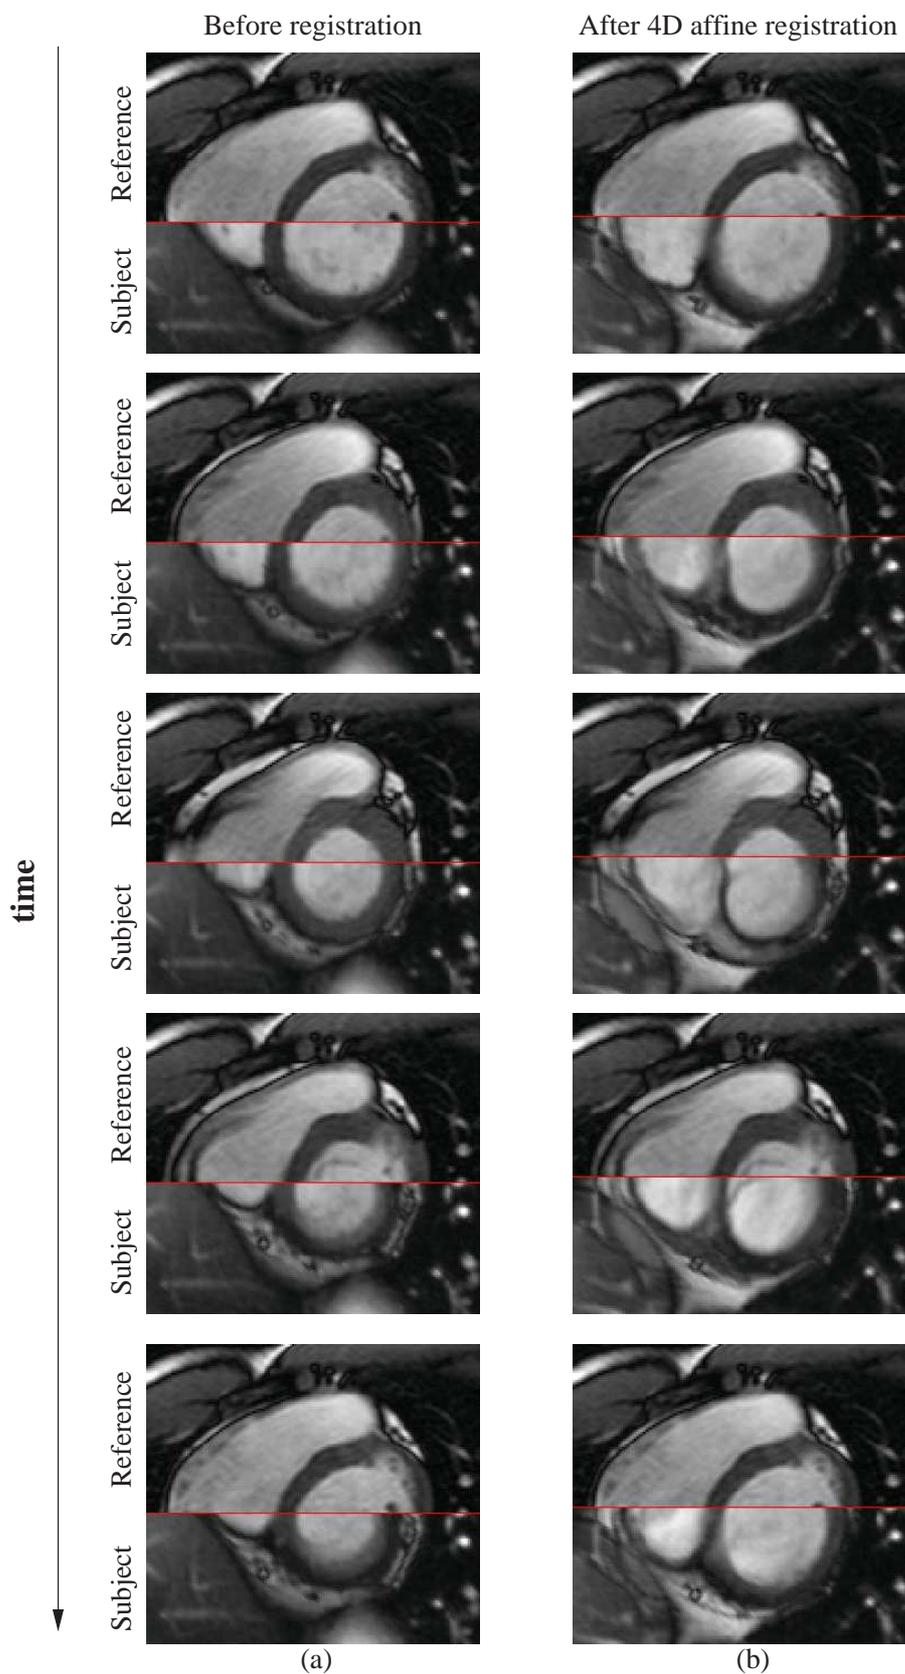
Figure 4.9: An example of affine spatio-temporal registration of two image sequences: (a) frames of the reference image sequence (above half) with a particular subject before the registration; (b) the corresponding frames after the registration.
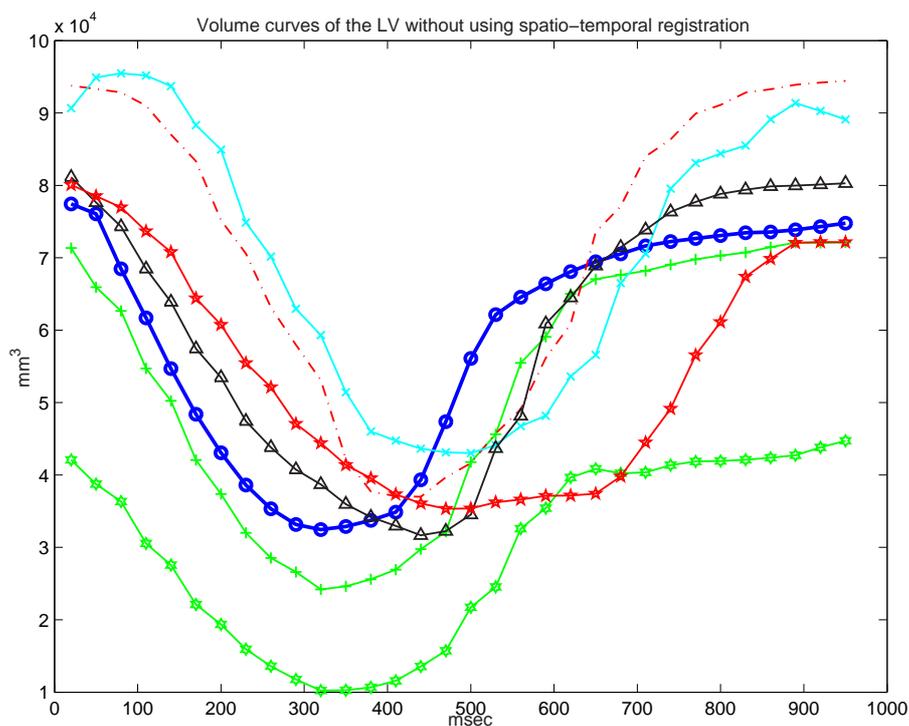
(a)
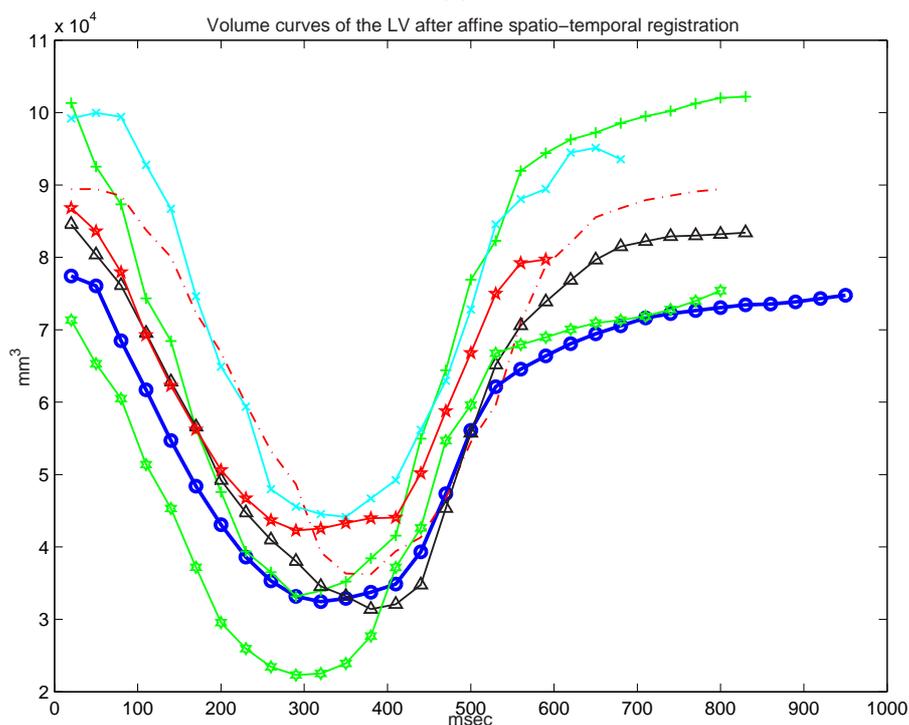
(b)

Figure 4.10: The volume of the left ventricle of the reference image and the other subjects over time: (a) before the 4D affine registration (mapping the temporal ends of the sequences), (b) after the 4D affine registration . The thick line is the volume curve of the reference subject.

4.5   Introduction of a non-rigid spatial translation

As mentioned above, the affine spatial alignment allows only translation, scaling, rotation and shearing operations to be performed. These operations are not enough to correct the misregistration caused by local shape differences of the cardiac anatomy. In this section the affine spatio-temporal cardiac MR image registration method [152] is extended by the introduction of a non-rigid spatial part, $\mathbf{T}_{spatial}$, based on B-Splines [153]. The temporal part of the spatio-temporal registration approach remains the same. i.e. an affine temporal transformation (equation 4.5) is used correcting scaling and translation differences between the image sequences.

4.5.1   Non-rigid spatial transformation

The spatial part of the 4D transformation $\mathbf{T}_{spatial}$ contains a global and a local part:

$$\mathbf{T}_{spatial}(x, y, z) = \mathbf{T}_{spatial}^{global}(x, y, z) + \mathbf{T}_{spatial}^{local}(x, y, z) \qquad (4.11)$$

The global part, $\mathbf{T}_{spatial}^{global}$, will correct differences in the size, orientation and alignment of the hearts while the local part, $\mathbf{T}_{spatial}^{local}$, will address the differences in the shape of the cardiac anatomy. An affine transformation was selected as the global part (equation 4.4).

A free-form deformation (FFD) model based on B-Splines is used in order to describe the differences in the local shape of the hearts. To define a spline based FFD, the spatial domain of the image volume is denoted as $\Omega_S = \{(x, y, z) \mid 0 \leq x < X, 0 \leq y < Y, 0 \leq z < Z\}$. Let Φ denote a $n_x \times n_y \times n_z$ mesh of control points $\phi_{i,j,k}$ with uniform spacing δ. Then, the FFD can be written as the 3D tensor product of the familiar 1D cubic B-Splines [174]:

$$\mathbf{T}_{spatial}^{local}(x, y, z) = \sum_{l=0}^{3} \sum_{m=0}^{3} \sum_{n=0}^{3} B_l(u) B_m(v) B_n(w) \phi_{i+l,j+m,k+n} \qquad (4.12)$$

where $i = \lfloor \frac{x}{n_x} \rfloor - 1, j = \lfloor \frac{y}{n_y} \rfloor - 1, k = \lfloor \frac{z}{n_z} \rfloor - 1, u = \frac{x}{n_x} - \lfloor \frac{x}{n_x} \rfloor, v = \frac{y}{n_y} - \lfloor \frac{y}{n_y} \rfloor, w =$

$\frac{z}{n_z} - \lfloor \frac{z}{n_z} \rfloor$ and where B_l represents the l-th basis function of the B-Spline ([99, 100]):

$$
\begin{aligned}
B_0(u) &= (1-u)^3/6 \\
B_1(u) &= (3u^3 - 6u^2 + 4)/6 \\
B_2(u) &= (-3u^3 + 3u^2 + 3u + 1)/6 \\
B_3(u) &= u^3/6
\end{aligned}
$$

One advantage of B-Splines is that they are locally controlled which makes them computationally efficient even for a large number of control points. In particular, the basis functions of cubic B-Splines have a limited support, i.e. changing a control point affects the transformation only in the local neighbourhood of that control point. This spatial transformation model based on free-form deformations has been introduced by Rueckert *et al.* [174]. It has also been used for a number of applications [29, 174, 172, 63, 61, 28, 167, 165, 109].

4.5.2 Optimisation approach

There is no need to optimise $\mathbf{T}_{temporal}$ and $\mathbf{T}_{spatial}^{global}$ since they have been previously optimised. The only part of the transformation which needs to be optimised is the $\mathbf{T}_{spatial}^{local}$. In order to find the optimal $\mathbf{T}_{spatial}^{local}$ information from the entire image sequences is used. As before, normalised mutual information is used as a similarity measure. The normalised mutual information was calculated directly from the joint intensity histogram of the two image sequences over their spatio-temporal domain of overlap $\Omega_{S'} \times [t_{S'_1}, t_{S'_n}] \bigcap \mathbf{T}(\Omega_S \times [t_{S_1}, t_{S_n}])$. The optimisation approach calculates the optimal transformation:

$$
\arg \max_{\alpha,\beta,\Theta,\Phi} NMI(S', \mathbf{T}(S)) \tag{4.13}
$$

where S' is the reference image sequence and $\mathbf{T}(S)$ is the transformed image sequence S. As mentioned above, α, β, Θ have been already optimised using a simple iterative uphill method.

The free-form deformation has a large number of *degrees of freedom* (in this case: number of control points $\times$ 3). Thus, a simple iterative uphill method will not be sufficient in order to find the optimal parameters Φ. An iterative gradient descent method is used to find the optimal $\mathbf{T}_{spatial}^{local}$. The method is similar as the one used by Rueckert *et al.* for the optimisation of the FFDs [174].

In order to find the optimal parameters, Φ, a cost function containing two competing terms is minimised. The first term represents the cost associated with the image similarity measure NMI. The second goal of the cost function ensures the smoothness of the transformation:

$$\mathcal{C}(\Theta, \Phi) = -NMI(S', \mathbf{T}(S)) + \lambda \mathcal{C}_{smooth}(\mathbf{T}_{spatial}) \qquad (4.14)$$

where $\mathcal{C}_{smooth}$ is a penalty term regularising the transformation [174] and λ is the weighting parameter which defines the tradeoff between maximising the alignment of the images and the smoothness of the transformation. The smoothness constraint is described by Wahba [201].

Rueckert *et al.* suggest that a value of $\lambda = 0.01$ provides a good compromise between the two terms in the cost function [174]. The regularisation term in the cost function becomes important when the spacing of the control points is small. This is due to the fact that the FFD's ability to model localised deformation increases as the the spacing of the control points of the B-Spline function decreases [174]. In the experiments reported in this chapter, $\lambda = 0$. Therefore, the optimisation approach relies on the intrinsic smoothness of the B-Spline deformation fields. This is sufficient for this type of application since the spacing of the control points used during the registration is not very small. The optimisation procedure steps in the direction of the gradient vector $\nabla \mathcal{C} = \frac{\partial \mathcal{C}(\Theta, \Phi^l)}{\partial \Phi^l}$ with a certain step size μ [174]. Listing 2 describes the algorithm.

As mentioned above, $\mathbf{T}_{spatial}^{local}$ has a large number of *degrees of freedom*. Therefore, using very small spacing between the control points and also very high resolution images will make the registration approach very slow. Hence, in this method the space of the

---

Listing 2 The gradient descent optimisation approach

1: **initialise** the control points Φ
2: **repeat**
3:     **repeat**
4:        **repeat**
5:           **calculate** the gradient vector of the cost function in equation 4.14 with respect to the non-rigid transformation parameters Φ:

$$\nabla\mathcal{C} = \frac{\partial\mathcal{C}(\Theta, \Phi^l)}{\partial\Phi^l}$$

6:           **while** $||\nabla\mathcal{C}|| > \epsilon$ **do**
7:              **recalculate** the control points $\Phi = \Phi + \mu\frac{\nabla\mathcal{C}}{||\nabla\mathcal{C}||}$
8:              **recalculate** the gradient vector $\nabla\mathcal{C}$
9:           **end while**
10:          **increase** the control point resolution by calculating new control points Φ^{l+1} from Φ^l
11:          **increase** the image resolution
12:       **until** finest level of resolution is reached
13:    **until** maximum number of iterations
14:    **subdivide** steps size μ
15: **until** maximum number of subdivisions

FFD's control point was set to 20*mm* (resulting in 1176 *degrees of freedom*). This control spacing allows to use $\lambda = 0$ without affecting the results. Table 4.3 contains the parameters used for these registrations.

Parameter	Setting
No. of bins	64
No. of iterations	20
No. of steps	4
δ **for finite difference approximation of gradient**	5*mm*
λ	0

Table 4.3: Registration parameters used for the spatio-temporal registration with a non-rigid spatial part.

4.5.3   Evaluation of the method

The method has been evaluated using the same seven image sequences as the ones used in the spatio-temporal affine registration method.

Figures 4.11 and 4.12 show examples of the registration. In the first row of figures
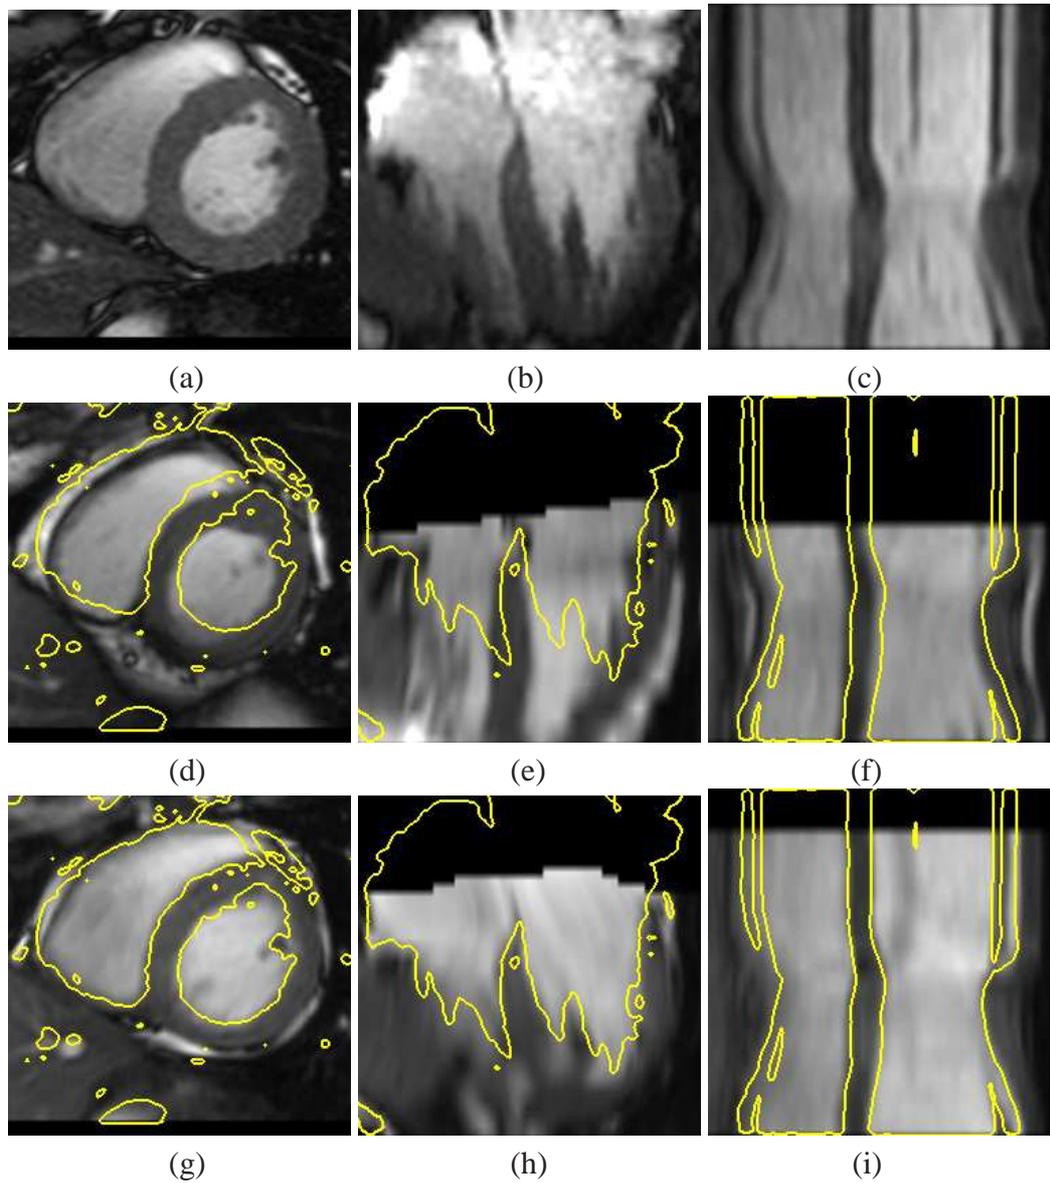
Figure 4.11: Results of the non-rigid spatial 4D MR registration algorithm: (a) the short-axis view of the reference subject, (b) the long-axis view of the reference subject, (c) a temporal view of a short-axis line of the reference subject, (d-f) the images after the affine registration, (g)-(i) the corresponding images after the non-rigid registration.

4.11 and 4.12 are the short-axis (a) and the long-axis (b) views of particular slices of the reference sequence and a temporal view of a vertical short axis line (c) of the reference image sequence. The middle and the bottom rows contain the corresponding views of the corresponding slices after the optimisation of $\mathbf{T}_{temporal}$ and $\mathbf{T}_{spatial}^{global}$ (i.e. after spatio-temporal affine registration) and after the optimisation of $\mathbf{T}_{spatial}^{local}$. On the images of the middle and bottom rows isolines of intensity of the reference subject are also overlayed.
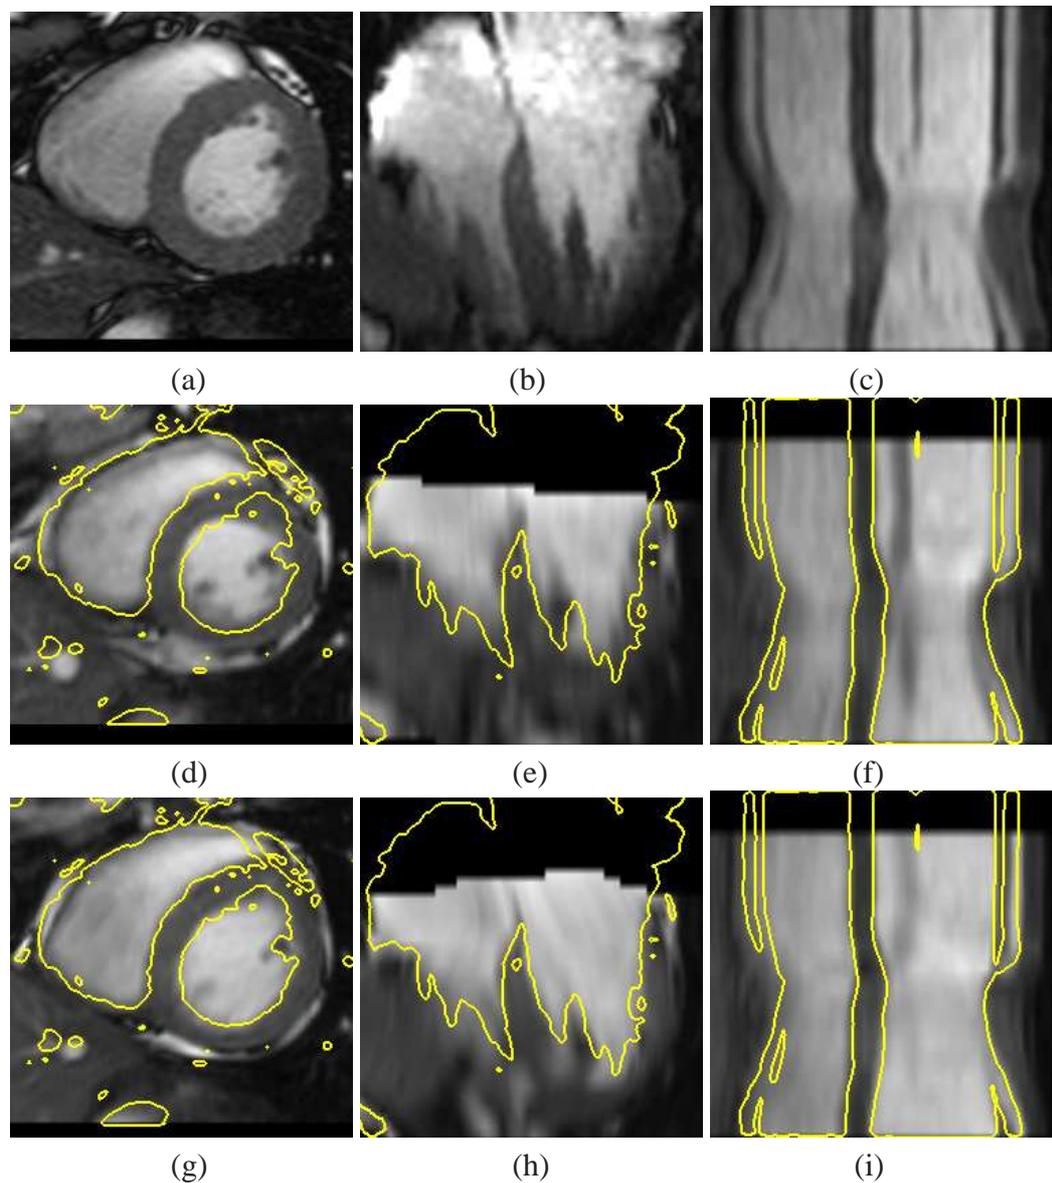
Figure 4.12: Results of the non-rigid spatial 4D MR registration algorithm: (a) the short-axis view of the reference subject, (b) shows the long-axis view of the reference subject, (c) a temporal view of a short-axis line of the reference subject, (d-f) the corresponding images after affine registration, (g)-(i) the corresponding images after the non-rigid registration.

It can be seen from the figures that the introduction of a non-rigid spatial part improves substantially the spatial alignment of the image sequences.

The volume overlap (equation 4.10) was also used to measure the quality of the registrations. Table 4.4 shows the mean volume overlap for each anatomical region after non-rigid registration and after affine registration. These measures also indicate that the

introduction of the non-rigid spatial part results in a substantial improvement in the overlap of the image sequences.

Anatomical region	Non-rigid 4D registration	Affine 4D registration
Left ventricle	86.68%	79.94%
Right ventricle	77.16%	74.53%
Myocardium	71.15%	67.93%

Table 4.4: The mean volume overlap after the affine 4D registration and after non-rigid 4D registration.

The introduction of the non-rigid spatial part, $\mathbf{T}_{spatial}$, provides a significant improvement in the spatio temporal cardiac MR image registration. However, it is still not expected that the image sequences would be perfectly aligned due to the nature of the temporal transformation. The temporal transformation cannot address misalignment caused by differences in the dynamic properties of the contraction and relaxation phases of the cardiac cycle (for example one heart may have a longer contraction phase than relaxation phase). The optimisation approach may also try to use the spatial transformation to correct for temporal differences. The registration of the image sequences will be significantly improved by the introduction of a temporal transformation which can address temporal misalignment due to different cardiac motion patterns.

4.6   Conclusions

In this chapter a spatio-temporal registration method for the alignment of cardiac MR image sequences has been presented. The method has a number of advantages over the registration methods presented in chapter 2. It enables comparison between the anatomy and function of a number of cardiac image sequences to be made. The method (except initialisation) is automatic. Furthermore, contrary to the method introduced by Faber *et al.* [55], no segmentation is required in order to determine the end-diastolic and end-systolic frames.

A 4D transformation model which consists of decoupled spatial and temporal components is used by the registration method. Initially, an affine spatial transformation model and an affine temporal transformation model was used. This spatial model corrects global shape differences (translation, rotation, scale and shearing) while the temporal transformation model corrects global differences in the length of the cardiac cycles and also in the temporal acquisition parameters. Later in the chapter, a *deformable* spatial transformation model based on B-Splines which enables the better spatial registration of the image sequences has been introduced. Both registration methods were tested using MR cardiac image sequences from normal subjects. The results demonstrate that the spatio-temporal registration method not only enables the spatial mapping between two cardiac images but also the temporal mapping between their cardiac cycles. However, the current temporal mapping cannot address temporal differences due to different temporal dynamics. The introduction of a temporal transformation in the next chapter addresses temporal misalignment due to different temporal dynamics will provide a better spatio-temporal alignment between the image sequences.

Chapter 5

Spatio-temporal free-form registration of cardiac MR image sequences

In this chapter two registration algorithms for the spatio-temporal alignment of cardiac MR image sequences are presented. Both algorithms have the ability to correct spatial misalignment between the image sequences caused by global and local shape differences. In addition, they have the ability to correct temporal misalignment caused by differences in the length of the cardiac cycles and in the dynamic properties of the hearts. The algorithms use a 4D *deformable* transformation model which is separated into spatial and temporal components. The first registration algorithm optimises the spatial and temporal transformation models simultaneously, while the second registration algorithm optimises the temporal transformation component before optimising the spatial component. The combined optimisation of the transformation components provides better accuracy than the method which optimises the transformation components separately. However, the method which optimises the transformation components separately has substantially lower computational complexity. Furthermore, two experiments are performed aiming to test how the use of an initial estimation of the temporal transformation may affect the performance (in terms of accuracy and computational complexity) of the spatio-temporal registration which optimises the transformation components simultaneously.

5.1 Why *deformable* spatio-temporal registration is needed

In the previous chapter, a method for the spatial and temporal alignment of cardiac MR image sequences was presented. The method has the ability to correct spatial misalignment caused by both global and local differences in the shape of the cardiac anatomy. It addition, it can correct temporal misalignment caused by differences in the length of the cardiac cycles and in the temporal acquisition parameters. However, this method cannot correct any temporal misalignment caused by different contraction and relaxation patterns.

In chapter 1 it has been presented that the cardiac cycle is separated into 7 phases (the *atrial systole*, *isovolumetric contraction*, *rapid ejection*, *reduced ejection*, *isovolumetric relaxation*, *rapid ventricular filling* and *reduced ventricular filling*) [91]. Figure 5.1 shows an example of how the pressure and volume of the heart changes during each phase.

The length of each of these phases varies from heart to heart. Therefore, using only an affine temporal transformation to provide a mapping between the temporal characteristics of two image sequences will not be enough in most cases. Figure 5.2 provides an example of such a case. The figure shows the volume curves of the LV of two subjects over their cardiac cycle. The length of each phase of the cardiac cycle varies in each subject resulting in different contraction and relaxation patterns. There is no temporal transformation utilising only scaling and translation which will provide correct temporal mapping between these two image sequences. What is required is a *deformable* temporal transformation to deform the cardiac motion of one subject in order to map the cardiac motion of the other subject.

The registration methods presented in this chapter extend the 4D cardiac MR image registration method presented in chapter 4 [152, 153, 145], by introducing a *deformable* temporal transformation model. Hence, both temporal and spatial components use more sophisticated *deformable* transformation models which allow the better spatio-temporal registration of cardiac MR image sequences. This transformation model corrects spatial
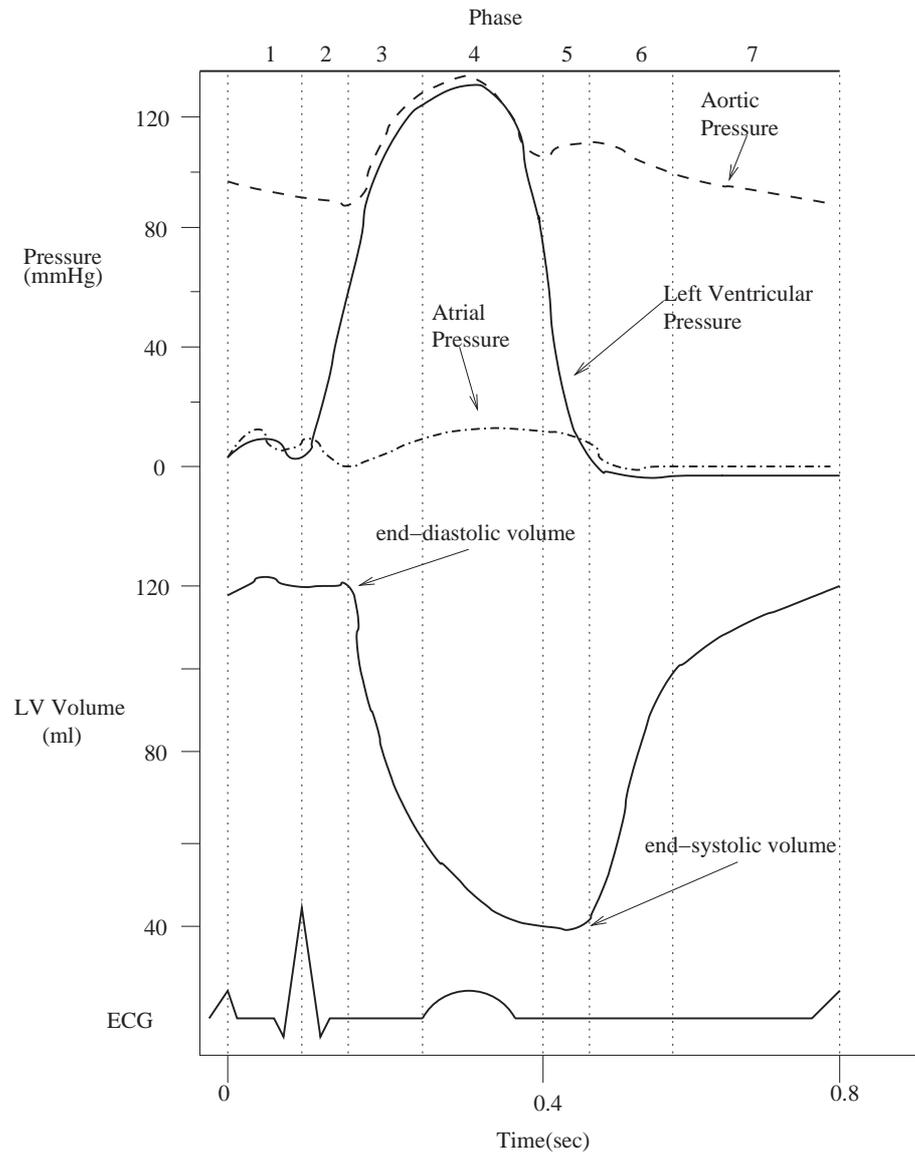
Figure 5.1: An example of the pressure and the volume of the heart during the cardiac cycle (adapted from [91]).

misalignment between the image sequences caused by global and local shape differences. Furthermore, it also corrects temporal misalignment caused by differences in the length of the cardiac cycles and in the dynamic properties of the hearts. This allows a direct comparison between both the cardiac anatomy and function of different subjects to be made. Two different approaches for the optimisation of the non-rigid spatio-temporal registration of cardiac MR image sequences are presented. The first approach finds the optimal spatio-temporal mapping by optimising of the spatial and temporal components simultaneously. The second approach optimises each component separately (first the temporal
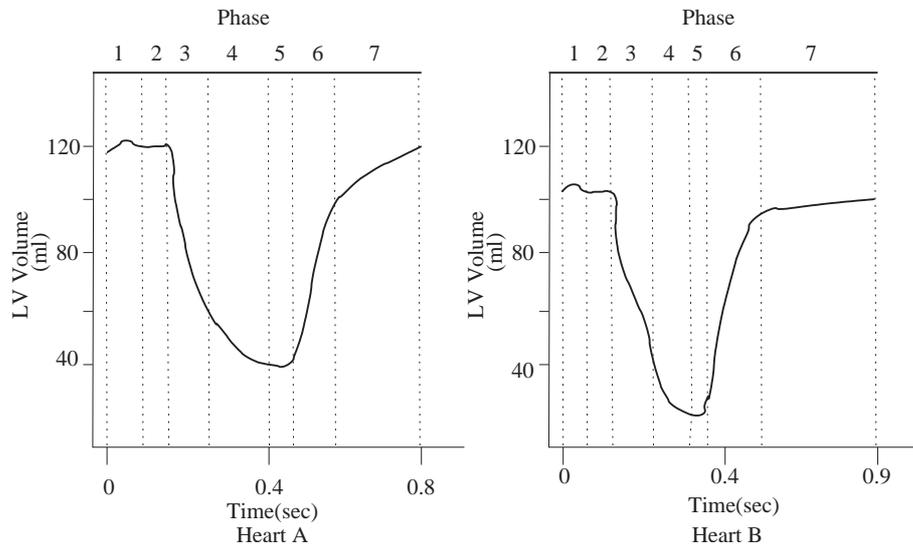
Figure 5.2: An example where a *deformable* spatio-temporal registration is needed.

component and then the spatial component).

5.2 Contributions

The contributions of this chapter are:

- The development of a method for the simultaneous spatial and temporal *deformable* registration of MR image sequences. This registration method has the ability to correct spatial misalignment between the image sequences caused by global and local shape differences. Furthermore, it has the ability to correct temporal misalignment caused by differences in the length of the cardiac cycles and in the dynamic properties of the hearts.

- The development of another method for the *deformable* spatio-temporal alignment of cardiac MR image sequences. The major difference compared to the previous method is that this method optimises the temporal and spatial components separately. In this method, the temporal registration of the image sequences is based on a normalised cross-correlation measure, while the spatial mapping is based on image information from only the first frames of the sequences. This registration

method is substantially faster than the previous method. Moreover, it can be used to enable only the temporal alignment of cardiac MR image sequences, since it does not require the cardiac image sequences to be aligned in the spatial domain in order to find their temporal correspondence. Furthermore, two experiments are performed aiming to test how the use of this method to calculate the temporal transformation may affect the performance (in terms of accuracy and computational complexity) of the first spatio-temporal registration method.

5.3   Spatio-temporal registration

As before, a 4D cardiac image sequence can be represented as an ordered sequence of n 3D images $S_k(x, y, z)$ with a fixed field of view Ω_{S_k} and an acquisition time t_k with $t_k < t_{k+1}$, in the temporal direction. The resulting image sequence can be viewed as 4D image $S(x, y, z, t)$, defined on the spatio-temporal domain $\Omega_{S_k} \times [t_1, t_n]$. The goal of 4D image registration described in this section is to relate each point of one image sequence to its corresponding point of the reference image sequence. In this case the transformation $\mathbf{T} : (x, y, z, t) \rightarrow (x', y', z', t')$ maps any point of one image sequence $S(x, y, z, t)$ onto its corresponding point in the reference image sequence $S'(x', y', z', t')$. The mapping used in this section is the same as in chapter 4 and it has the following form:

$$\mathbf{T}(x, y, z, t) = (x'(x, y, z), y'(x, y, z), z'(x, y, z), t'(t)) \tag{5.1}$$

This mapping can be of a subvoxel displacement in the spatial domain and of a sub-frame displacement in the temporal domain. The 4D mapping can be resolved into decoupled spatial and temporal components $\mathbf{T}_{spatial}$ and $\mathbf{T}_{temporal}$ respectively where:

$$\mathbf{T}_{spatial}(x, y, z) = (x'(x, y, z), y'(x, y, z), z'(x, y, z))$$

and

$$\mathbf{T}_{temporal}(t) = t'(t)$$

5.3.1 Spatial alignment

The aim of the spatial part of the transformation is to relate each spatial point in an image to a point of the reference image, i.e. $\mathbf{T}_{spatial} : (x, y, z) \rightarrow (x', y', z')$ maps any point (x, y, z) of a particular time frame t in one image sequence into its corresponding point (x', y', z') of another particular time frame t' in the reference image sequence. As in section 4.5.1, the transformation $\mathbf{T}_{spatial}$ consists of a global transformation and a local transformation:

$$\mathbf{T}_{spatial}(x, y, z) = \mathbf{T}_{spatial}^{global}(x, y, z) + \mathbf{T}_{spatial}^{local}(x, y, z) \tag{5.2}$$

The global transformation addresses differences in the size, orientation and translation of the hearts while the local part addresses differences in the shape of the cardiac anatomies. $\mathbf{T}_{spatial}^{global}$ is an affine transformation with 12 *degrees of freedom* utilising scaling, shearing, translation and rotation:

$$\mathbf{T}_{spatial}^{global}(x, y, z) = \begin{pmatrix} \theta_{11} & \theta_{12} & \theta_{13} \\ \theta_{21} & \theta_{22} & \theta_{23} \\ \theta_{31} & \theta_{32} & \theta_{33} \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} + \begin{pmatrix} \theta_{14} \\ \theta_{24} \\ \theta_{34} \end{pmatrix} \tag{5.3}$$

A free-form deformation (FFD) model based on B-Splines is used to describe the differences in the local shape of the cardiac anatomies. To define a spline based FFD the spatial domain of the image volume is denoted as $\Omega_S = \{(x, y, z) \mid 0 \leq x < X, 0 \leq y < Y, 0 \leq z < Z\}$. Let Φ denote a $n_x \times n_y \times n_z$ mesh of control points $\phi_{i,j,k}$ with uniform spacing δ. Then the FFD can be written as the 3D tensor product of the familiar 1D cubic B-Splines [174]:

$$\mathbf{T}_{spatial}^{local}(x, y, z) = \sum_{l=0}^{3} \sum_{m=0}^{3} \sum_{n=0}^{3} B_l(u) B_m(v) B_n(w) \phi_{i+l,j+m,k+n} \tag{5.4}$$

where $i = \lfloor \frac{x}{n_x} \rfloor - 1, j = \lfloor \frac{y}{n_y} \rfloor - 1, k = \lfloor \frac{z}{n_z} \rfloor - 1, u = \frac{x}{n_x} - \lfloor \frac{x}{n_x} \rfloor, v = \frac{y}{n_y} - \lfloor \frac{y}{n_y} \rfloor, w = \frac{z}{n_z} - \lfloor \frac{z}{n_z} \rfloor$ and where B_l represents the l-th basis function of the B-Spline [99, 100](see

section 4.5.1) .

5.3.2 Temporal alignment

In contrast to chapter 4.4.2, the temporal part of the transformation consists of a temporal global part, $\mathbf{T}_{temporal}^{global}$, and a temporal local part, $\mathbf{T}_{temporal}^{local}$:

$$\mathbf{T}_{temporal}(t) = \mathbf{T}_{temporal}^{global}(t) + \mathbf{T}_{temporal}^{local}(t) \tag{5.5}$$

As before, the global part of the transformation, $\mathbf{T}_{temporal}^{global}$, addresses global differences in the size of the cardiac cycles and differences in the acquisition parameters. The local part of the temporal transformation, $\mathbf{T}_{temporal}^{local}$, addresses local differences between the motion patterns of the hearts. An affine transformation is used as a global part:

$$\mathbf{T}_{temporal}^{global}(t') = \alpha t + \beta \tag{5.6}$$

Here α accounts for scaling differences between the two image sequences (different length of cardiac cycles or different frequency in the acquisition of each frame) while β accounts for translation differences between the two image sequences. Translation differences may be introduced by different acquisition parameters (for example the trigger offset) and by missing frames.

The local temporal transformation, $\mathbf{T}_{temporal}^{local}$, is modelled by a free-form deformation using a 1D B-Spline and corrects for temporal misalignment caused by different cardiac dynamic properties (differences in the length of each of the cardiac phases (figure 5.1), e.g. one heart may have a longer contraction phase and a shorter relaxation phase, different motion patterns, etc.). To define a spline based temporal free-form deformation, the temporal domain of the image sequence is denoted as $\Omega_t = \{(t) \mid 0 \leq x < T\}$. Let Φ_t denote a set of n_t control points ϕ_t with a temporal spacing δ_t. Then the temporal free-form deformation can be defined as a 1D cubic B-Spline:

$$\mathbf{T}^{local}_{temporal}(t) = \sum_{l=0}^{3} B_l(u)\phi_{t_{i+l}} \tag{5.7}$$

where $i = \lfloor \frac{t}{n_t} \rfloor - 1$, $u = \frac{t}{n_t} - \lfloor \frac{t}{n_t} \rfloor$ and B_l represents the l-th basis function of the B-Spline.

$\mathbf{T}^{local}_{temporal}$ deforms the temporal characteristics of each image sequence in order to follow the same motion pattern as the reference image sequence. The combined 4D transformation model (equation 5.1) is the spatio-temporal free-form deformation (STFFD) based on a 4D B-Spline model.

The temporal transformation $\mathbf{T}_{temporal}(t)$ (equation 5.5) needs to be non-decreasing, i.e. whenever $t_1 \leq t_2$ then $\mathbf{T}_{temporal}(t_1) \leq \mathbf{T}_{temporal}(t_2)$, in order to preserve causality of events between the two cardiac cycles. In order to ensure that $\mathbf{T}_{temporal}(t)$ is non-decreasing it is ensured that the derivative of $\mathbf{T}_{temporal}(t)$ at the positions of control points, ϕ_t, is always positive or zero. In practice this restriction will ensure that equation 5.5 is non-decreasing.

5.3.3 Optimisation of the transformation components

Two registration algorithms for finding the optimal transformation $\mathbf{T}$ have been developed. The first optimisation algorithm optimises the spatial and temporal transformation components simultaneously, while the second registration algorithm optimises the temporal transformation component before optimising the spatial component.

5.3.3.1 Combined optimisation of the spatial and temporal components

In this registration method the optimal transformation $\mathbf{T}$ is found by optimising the temporal $\mathbf{T}_{temporal}$ and the spatial $\mathbf{T}_{spatial}$ transformation components at the same time using image information only.

The optimal transformation $\mathbf{T}$ is found by maximising a voxel based similarity measure. The use of a voxel-based similarity measure eliminates the need for any feature detection and segmentation of structures such as the epi- or endo-cardial surfaces dur-

ing the registration procedure. As before, normalised mutual information (NMI) [187] is selected as a measure of the spatio-temporal alignment. This similarity measure has been previously used successfully for mono- and multi-modality image registration. The normalised mutual information of two image sequences can be written as:

$$NMI(S', S) = \frac{H(S') + H(S)}{H(S', S)} \tag{5.8}$$

where $H(\cdot)$ denotes the marginal entropy of an image sequence and $H(\cdot, \cdot)$ is the joint entropy of two image sequences. The normalised mutual information of the two image sequences can be calculated directly from the joint intensity histogram of the two sequences over the spatio-temporal domain of overlap $\Omega_{S'} \times [t_{S'_1}, t_{S'_n}] \cap \mathbf{T}(\Omega_S \times [t_{S_1}, t_{S_n}])$. During the optimisation new voxel values are generated in the temporal domain using linear interpolation and trilinear interpolation in the spatial domain.

In the first part of the optimisation procedure, NMI is optimised as a function of $\mathbf{T}^{global}_{spatial}$ and $\mathbf{T}^{global}_{temporal}$ using an iterative uphill descent algorithm (described in section 4.4.3). In the second part, NMI is optimised as a function of $\mathbf{T}^{local}_{spatial}$ and $\mathbf{T}^{local}_{temporal}$. In order to find the optimal parameters Φ_t and Φ a cost function similar to equation 4.14 is minimised. As in equation 4.14, the cost function contains two competing goals.

$$\mathcal{C}(\Theta, \Phi, \alpha, \beta, \Phi_t) = -NMI(S', \mathbf{T}(S)) + \lambda \mathcal{C}_{smooth}(\mathbf{T}_{spatial}) \tag{5.9}$$

The terms λ and $\mathcal{C}_{smooth}$ have been described in section 4.4.3. The optimisation procedure steps in the direction of the gradient vectors $\nabla \mathcal{C} = \frac{\partial \mathcal{C}(\Theta, \Phi^l)}{\partial \Phi^l}$ and $\nabla \mathcal{C}_t = \frac{\partial \mathcal{C}(\alpha, \beta, \Phi^l_t)}{\partial \Phi^l_t}$ with certain step sizes μ and μ_t. Listing 3 describes the optimisation method.

As previously mentioned, $\mathbf{T}^{local}_{spatial}$ has a large number of *degrees of freedom*. Therefore by using very small spacing between the control points and image sequences with high resolution image sequences the registration approach will be very slow. Hence, in this method the space of the FFD's control points was set to 15*mm* and 10*mm* (resulting to 2400 and 7560 *degrees of freedom* respectively). This control spacing allows to use

$\lambda = 0$ without affecting the results. Table 5.1 contains the parameters used for these registrations.

Listing 3 The spatio-temporal gradient descent optimisation approach.

1: **initialise** the control points Φ and Φ_t
2: **repeat**
3:     **repeat**
4:         **repeat**
5:             **calculate** the gradient vectors of the cost function in equation 5.9 with respect to the non-rigid transformation parameters Φ and Φ_t:

$$\nabla \mathcal{C} = \frac{\partial \mathcal{C}(\Theta, \Phi^l)}{\partial \Phi^l}$$

$$\nabla \mathcal{C}_t = \frac{\partial \mathcal{C}(\alpha, \beta, \Phi_t^l)}{\partial \Phi_t^l}$$

6:             **while** $||\nabla \mathcal{C}|| > \epsilon$ or $||\nabla \mathcal{C}_t|| > \epsilon_t$ **do**
7:                 **recalculate** the control points $\Phi = \Phi + \mu \frac{\nabla \mathcal{C}}{||\nabla \mathcal{C}||}$
8:                 **recalculate** the control points $\Phi_t = \Phi_t + \mu_t \frac{\nabla \mathcal{C}_t}{||\nabla \mathcal{C}_t||}$
9:                 **recalculate** the gradient vectors $\nabla \mathcal{C}$ and $\nabla \mathcal{C}_t$
10:             **end while**
11:             **increase** the control point resolution by calculating new control points Φ^{l+1} from Φ^l and Φ_t^{l+1} from Φ_t^l
12:             **increase** the image resolution
13:         **until** finest level of resolution is reached
14:     **until** maximum number of iterations
15:     **subdivide** steps size μ and μ_t
16: **until** maximum number of subdivisions

Parameter	Setting
No. of bins	64
No. of iterations	20
No. of steps	4
δ for finite difference approximation of gradient	$10mm$
δ_t for finite difference approximation of gradient	$50msec$
λ	0

Table 5.1: Registration parameters used for the spatio-temporal free-form registration of image sequences (combined optimisation of the transformation parts).

5.3.3.2 Separate optimisation of the spatial and temporal components

The computational complexity of the previous spatio-temporal *deformable* registration method is very high. However, it can be reduced by optimising each transformation component (the temporal and the spatial one) separately. This optimisation approach optimises the temporal component $\mathbf{T}_{temporal}$ of the transformation $\mathbf{T}$ first and then the spatial component $\mathbf{T}_{spatial}$. The global temporal component $\mathbf{T}_{temporal}^{glocal}$ is calculated to align the temporal ends of the image sequences while the local temporal component $\mathbf{T}_{temporal}^{local}$ is a temporal free-form deformation (equation 5.7) which aligns limited temporal feature positions of the cardiac cycles. $\mathbf{T}_{temporal}^{local}$ aligns the feature position of the peak contraction of the left ventricle, the end-diastolic feature position and the beginning and end of the cardiac cycles (as in figure 5.3). In order to detect these temporal feature positions in each image sequence, the normalised cross-correlation coefficient of the first and each of the frames of the sequence is calculated:

$$CC = \frac{\sum_x \sum_y \sum_z (S_0(x,y,z) - \bar{S}_0) \cdot (S_i(x,y,z) - \bar{S}_i)}{\sqrt{\sum_x \sum_y \sum_z (S_0(x,y,z) - \bar{S}_0)^2} \cdot \sqrt{\sum_x \sum_y \sum_z (S_i(x,y,z) - \bar{S}_i)^2}} \quad (5.10)$$

where S_0 is the first frame, $\bar{S}_0$ the mean intensity of the first frame, S_i is the ith frame of the sequence and $\bar{S}_i$ the mean intensity of that frame.

The idea behind this approach is that during the contraction phase of the cardiac cycle each consecutive image will appear less similar to the first image and during the relaxation phase of the cardiac cycle each consecutive image will appear more similar to the first image. The end-systolic image should have the highest degree of dissimilarity with the first image since the heart has a different shape and size due to the contraction. Similarly, assuming periodicity of the cardiac cycle, the end-diastolic image should have a high degree of similarity with the first image since the heart at these positions of the cardiac cycle has similar size and shape. There is a wide choice of similarity measures which could be used for this application. In practice it is found that the normalised cross-correlation coefficient works well for this application. Figure 5.4 provides an example of

the displacements of the control points of $\mathbf{T}^{local}_{temporal}$ required to align the 4 feature positions of the cardiac cycles. The position of the control points, ϕ_t, of $\mathbf{T}^{local}_{temporal}$ is calculated by using linear interpolation between the displacements required to align the 4 identified feature positions. Thus $\mathbf{T}^{local}_{temporal}$ also enables the alignment of temporal features between these 4 determined feature positions.
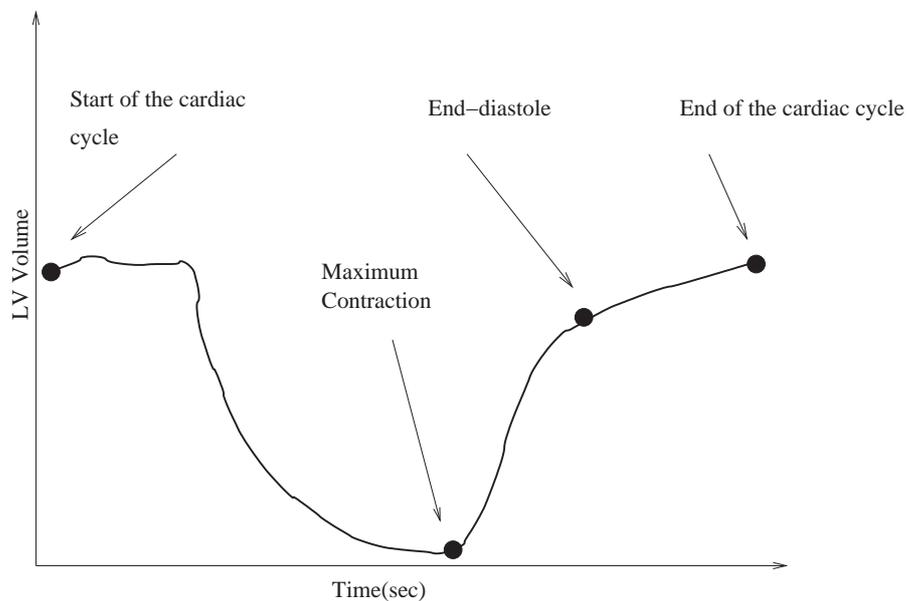
Figure 5.3: The temporal positions in the cardiac cycle used for the temporal alignment between two image sequences.

A similar technique based on cross-correlation has been used by Larson *et al.* [96] for the recovery of temporal information from cardiac cine MRI. The purpose of this work was to develop a new *self-gated* acquisition technique by extracting the motion synchronisation signal directly from the cardiac MR images. Three different strategies using radial k-space sampling are proposed for deriving temporal information from the MR images [96]. Among these techniques are the peak magnitude, the kymogram [89] and the 2D correlation. A comparison of the results showed that the image quality obtained by these techniques is similar to the image quality obtained by conventional ECG gating techniques.

Figure 5.5 (a) shows the plot of the calculated normalised cross-correlation curve for a particular image sequence, figure 5.5 (b) shows the second derivative of the cross-
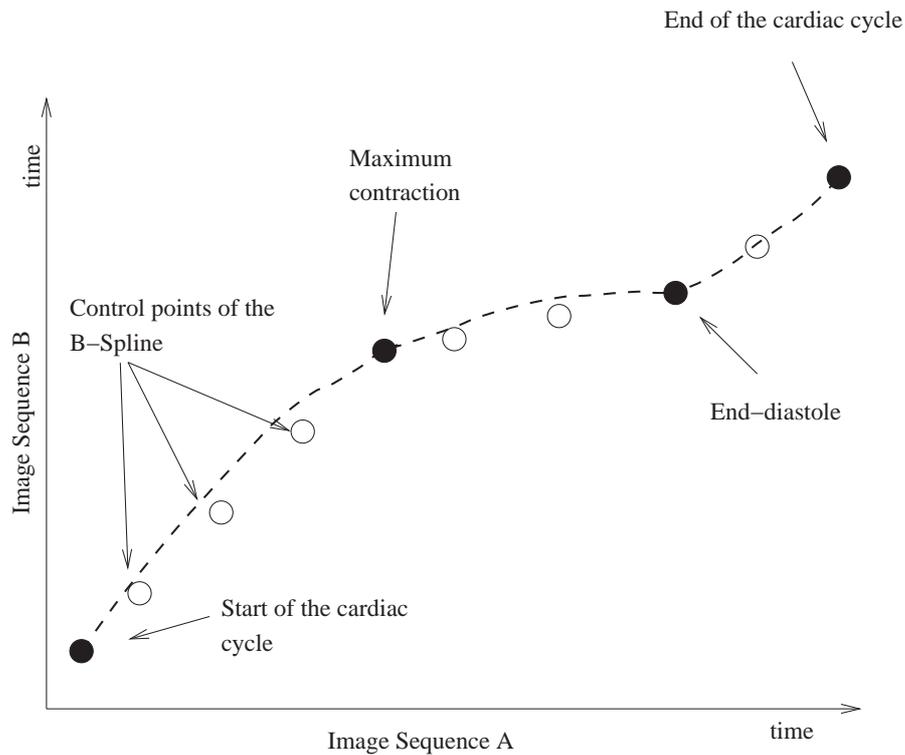
Figure 5.4: An example of the temporal mapping between two image sequences using $\mathbf{T}^{local}_{temporal}$ which aligns the 4 positions of the cardiac cycles.

correlation curve while figure 5.5 (c) shows the volume of the left ventricle of the same image sequence over time. It is clear from figure 5.5 that the cross-correlation and the volume curves are very similar. The feature position of peak contraction is found by the minimum cross-correlation value. In order to find the end-diastolic position, the minimum value of the second derivative after the location of peak contraction is used (figure 5.5 (b)). The second derivative is calculated using finite differences between neighbouring time frames.

The optimal spatial transformation $\mathbf{T}_{spatial}$ is calculated using the non-rigid 3D registration (equation 5.2) of the first frames of the image sequences. In this case, $\mathbf{T}^{global}_{spatial}$ is an affine transformation correcting translation, rotation, shearing and scaling differences between the first frames. $\mathbf{T}^{local}_{spatial}$ is a free-form deformation (equation 5.4) deforming the sequence's first frame to map the reference sequence's first frame. Both $\mathbf{T}^{global}_{spatial}$ and $\mathbf{T}^{local}_{spatial}$ are optimised using normalised mutual information (NMI). However, in this case NMI is based on the intensity histogram of the spatial domain of overlap of the two first
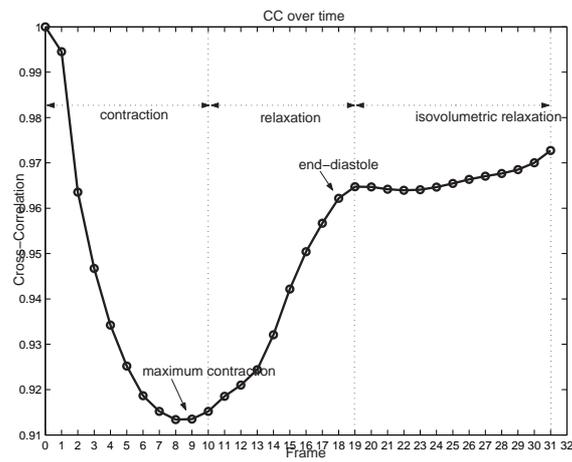
(a)
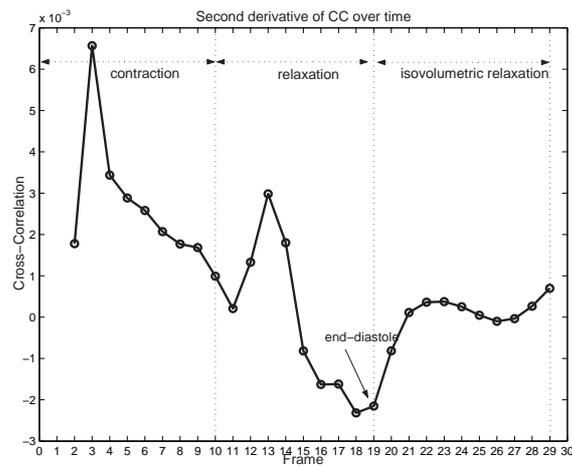
(b)
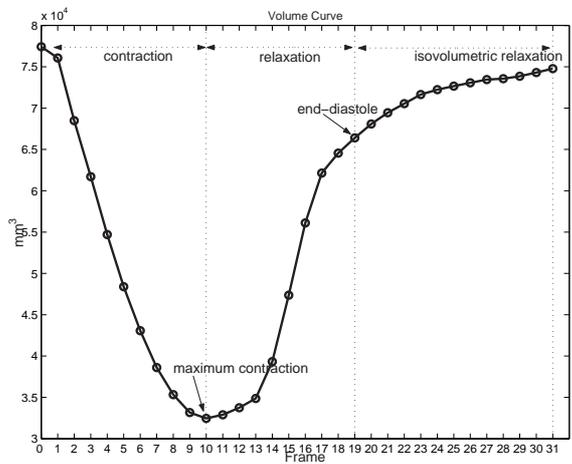
(c)

Figure 5.5: The cross-correlation between the first frame and each consecutively frame (a), the second derivative of the cross-correlation (b) and the volume of the left ventricle of the same subject over time (c).

frames and not on the spatio-temporal domain of overlap. The normalised mutual information of the first frames of two image sequences, S' and S, can be written as:

$$I(S_1', S_1) = \frac{H(S_1') + H(S_1)}{H(S_1', S_1)} \tag{5.11}$$

$\mathbf{T}_{spatial}^{global}$ is optimised using a simple iterative uphill method. The optimisation is carried out to calculate the optimal transformation:

$$\arg \max_{\Theta} NMI(S_1', \mathbf{T}_{spatial}(S_1)) \tag{5.12}$$

$\mathbf{T}_{spatial}^{local}$ is optimised using an iterative gradient descent method which is the same as the one described in section 4.4.3. The following cost function, is minimised:

$$\mathcal{C}(\Theta, \Phi) = -NMI(S_1', \mathbf{T}(S_1)) + \lambda \mathcal{C}_{smooth}(\mathbf{T}_{spatial}) \tag{5.13}$$

Equation 5.13 is the same as equation 4.14 with the exception of the domain where the normalised mutual information is calculated. Table 5.2 contains the parameters used in these registrations.

Parameter	Setting
No. of bins	64
No. of iterations	20
No. of steps	4
δ **for finite difference approximation of gradient**	$10mm$
λ	0

Table 5.2: Registration parameters used for the spatial free-form registration of the first frames.

5.4 Results

To evaluate the spatio-temporal *deformable* registration algorithm, cardiac MR image sequences from fifteen healthy volunteers have been acquired. All image sequences used for these experiments were acquired on a Siemens Sonata 1.5T scanner using TrueFisp pulse sequence. For the reference subject 32 different time frames were acquired (cardiac cycle of length 950$msec$). Each 3D image of the sequence had a resolution of $256 \times 192 \times 46$ with a pixel size of 0.97$mm \times$0.97mm and a slice thickness of 3mm. Fourteen 4D cardiac MR images were registered to the reference subject. These image sequences had a pixel size between 1.36$mm \times$1.36mm and 1.48$mm \times$1.48mm and a slice thickness of 10mm. For these subjects 15-20 different time frames were acquired (cardiac cycles' length between 600$msec$ to 800$msec$). In one image sequence the temporal acquisition parameters were estimated because the real parameters were unknown. Figure 5.6 contains a short-axis and a long-axis views of the reference image and of another subject. In both registration methods the global temporal transformation was calculated in order to compensate the differences in length of the cardiac cycles of the subjects (by matching the first and the last time frames of the image sequences). This is a fair assumption since all image sequences almost contained the entire cardiac cycles. Furthermore, an initial estimate of the global spatial transformation was also provided due to the large variety in the position and orientation of the hearts. The initial estimate was calculated from rigid registration of six manually selected cardiac anatomical positions in the first frames of the image sequences. These anatomical positions are: the apex of the left ventricle, the apex of the right ventricle, two landmarks on the myocardium at the base of the heart, the anteroseptal at the base of the heart and the inferoseptal at the base of the heart (figure 5.7).

The registrations were qualitatively evaluated by visual inspection. Furthermore, the quality of the registration in the spatial domain was measured by calculating the volume overlap for the left and right ventricles as well as for the myocardium (as in section 4.4.4).

(a)   (b)   (c)   (d)
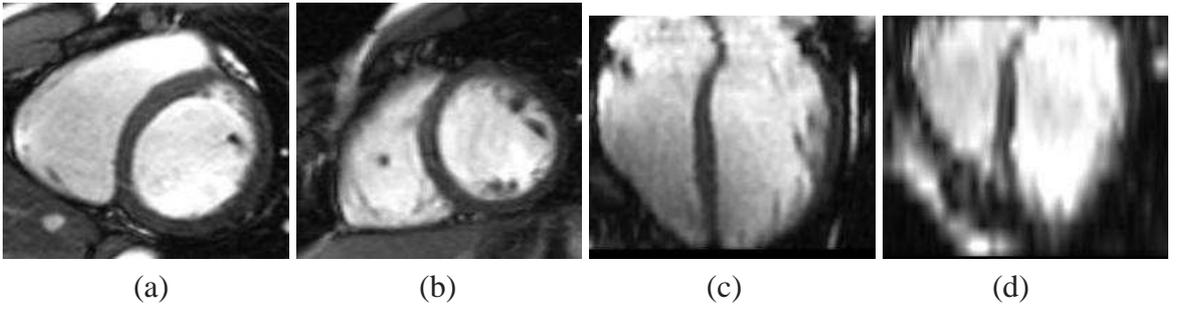
Figure 5.6: The short-axis (a) and (c) and the long-axis (c) and (d) views of the reference and of an another subject.
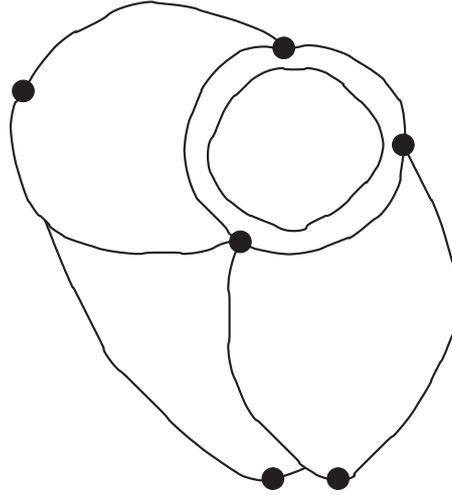
Figure 5.7: The six manually selected landmarks which are used to calculate the initial estimate of the spatial transformation.

The volume overlap for an object O is defined as:

$$\Delta(S', S) = \frac{2 \times |S' \bigcap S|}{|S'| + |S|} \times 100\% \qquad (5.14)$$

Here S' denotes the voxels in the reference (target) image part of object O and S denotes the voxels in the other image part of object O. The mean surface distance of the above anatomical regions was calculated after the affine and the *deformable* 4D registration. In order to calculate the overlap of the anatomical structures, the surface distance and the volume curves, segmented images were used. The segmented images were obtained using the EM algorithm developed by Valdés *et al.* [110]. In order to calculate the mean surface distance between an anatomical structure of two image sequences (the reference

and another one), the boundary surfaces (of the anatomical structures) in all of frames of the segmented image sequences are generated. Then, the Euclidean distance transform of the anatomical structure in all of the frames of the reference sequence is also calculated. Finally, the position of the boundary surface is compared to the corresponding position in the calculated distance map (for all frames of the image sequences).

5.4.1 Separate optimisation of the transformation components

5.4.1.1 Qualitative evaluation

Figure 5.8 (a) shows the volume curves of the left ventricles after affine spatio-temporal registration, while figure 5.8 (b) shows the corresponding volume curves after separate optimisation of the transformation components $\mathbf{T}_{spatial}^{temporal}$ and $\mathbf{T}_{spatial}^{local}$. It is clear from this figure that with the introduction of the *deformable* transformation components the hearts are substantially better aligned in the temporal domain.

Figures 5.9 - 5.14 provide examples of the *deformable* spatio-temporal registration. The images in the top rows are the short-axis (a), the long-axis (b) and the temporal (c) views (the temporal view is a short-axis line over time) of a frame in the middle of the image sequence after optimisation of the global transformation components (i.e. affine spatio-temporal registration). The lines in the images represent iso-contours of the reference image sequence. The images in the bottom rows (d-f) of the figures are the same images after spatio-temporal *deformable* registration (separate optimisation of the transformation components). It is clear from the figures that with the introduction of the *deformable* temporal and spatial transformations there is a substantial improvement in the alignment of the image sequences both in the spatial and in the temporal domain. The dark areas in the long-axis views are caused by the a smaller field of view in the current subject.
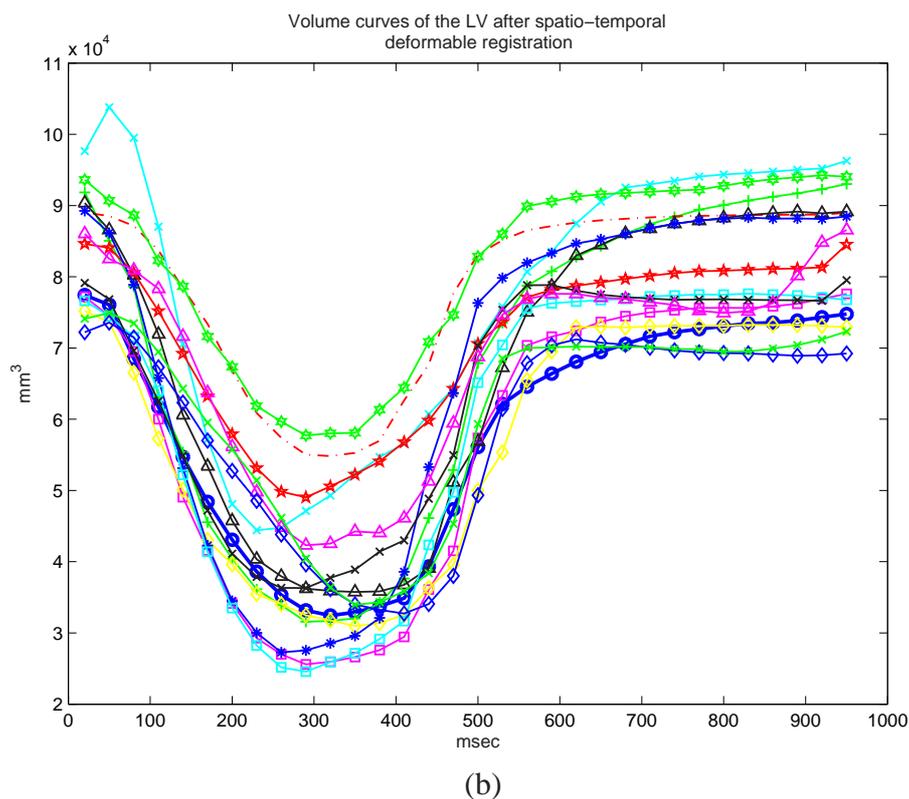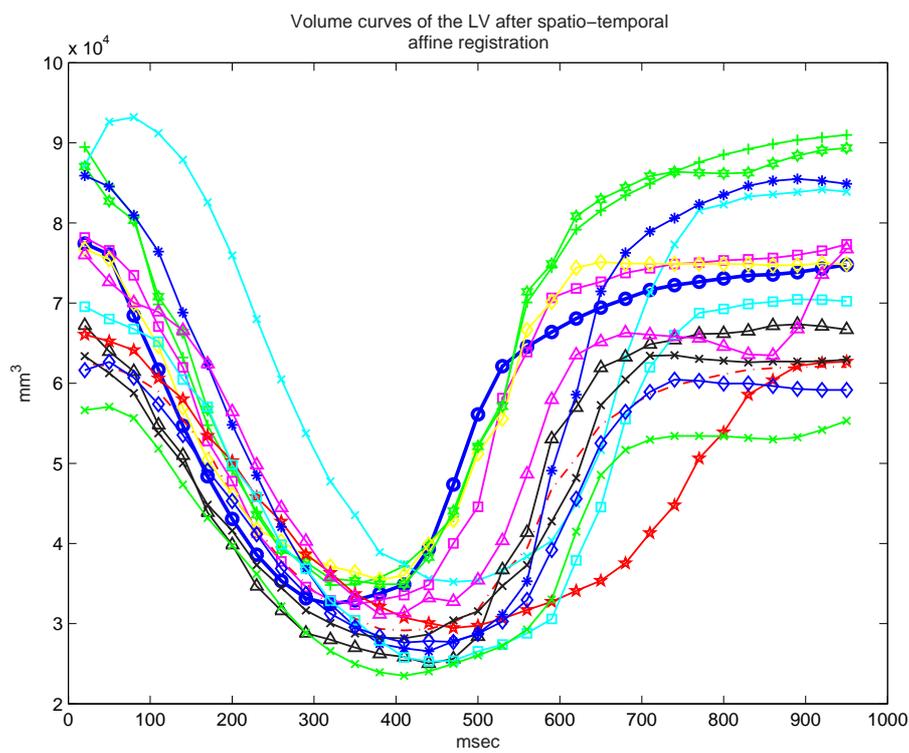
Figure 5.8: The volume curves of the left ventricle for all subjects after affine spatio-temporal registration (a) and after spatio-temporal *deformable* registration (with separate optimisation of the transformation components) (b). The spacing of spatial control points is 10*mm*. The thick line is the volume curve of the reference subject.
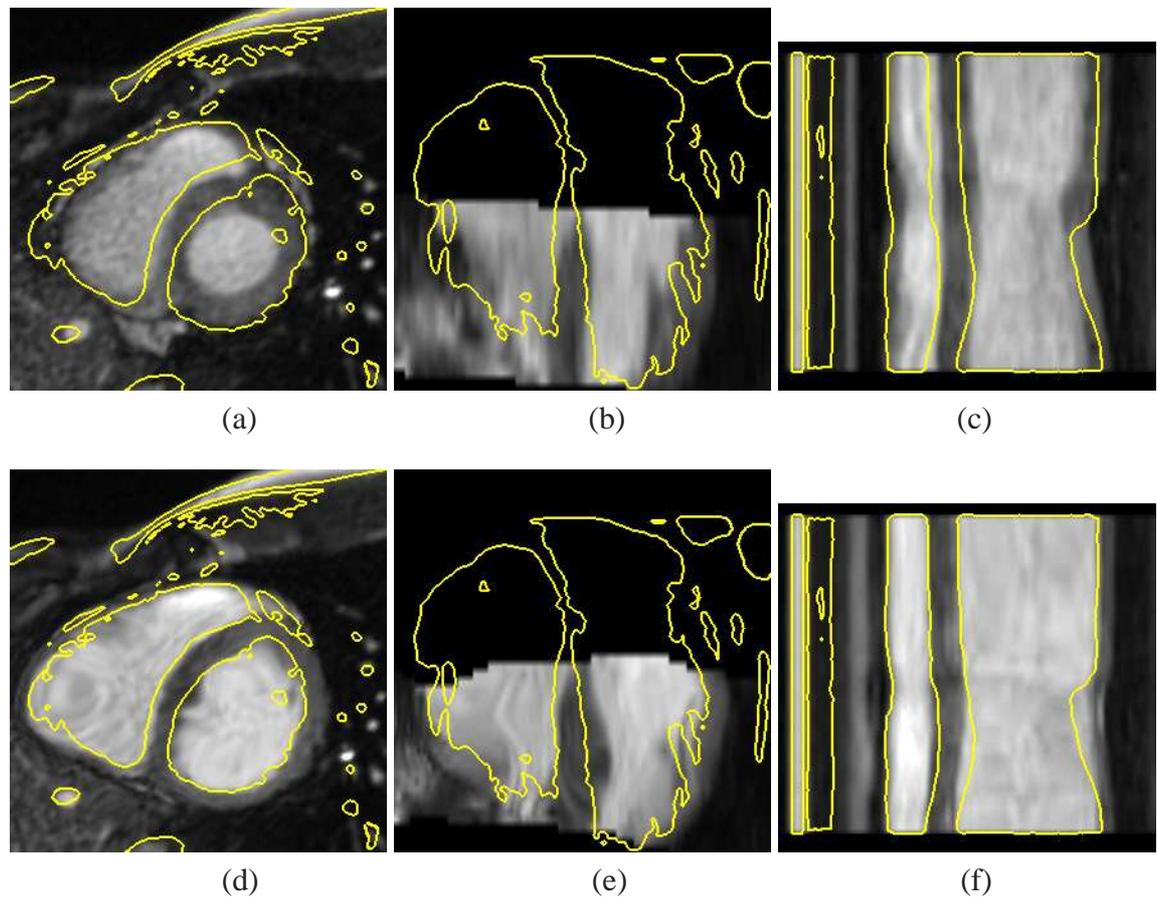
Figure 5.9: Results of the 4D cardiac MR registration algorithm. (a) The short-axis , (b) the long-axis and (c) the temporal views after the affine alignment. The corresponding (d) short-axis, (e) long-axis and (f) temporal views after the spatio-temporal free-form registration (separate optimisation of the transformation components).
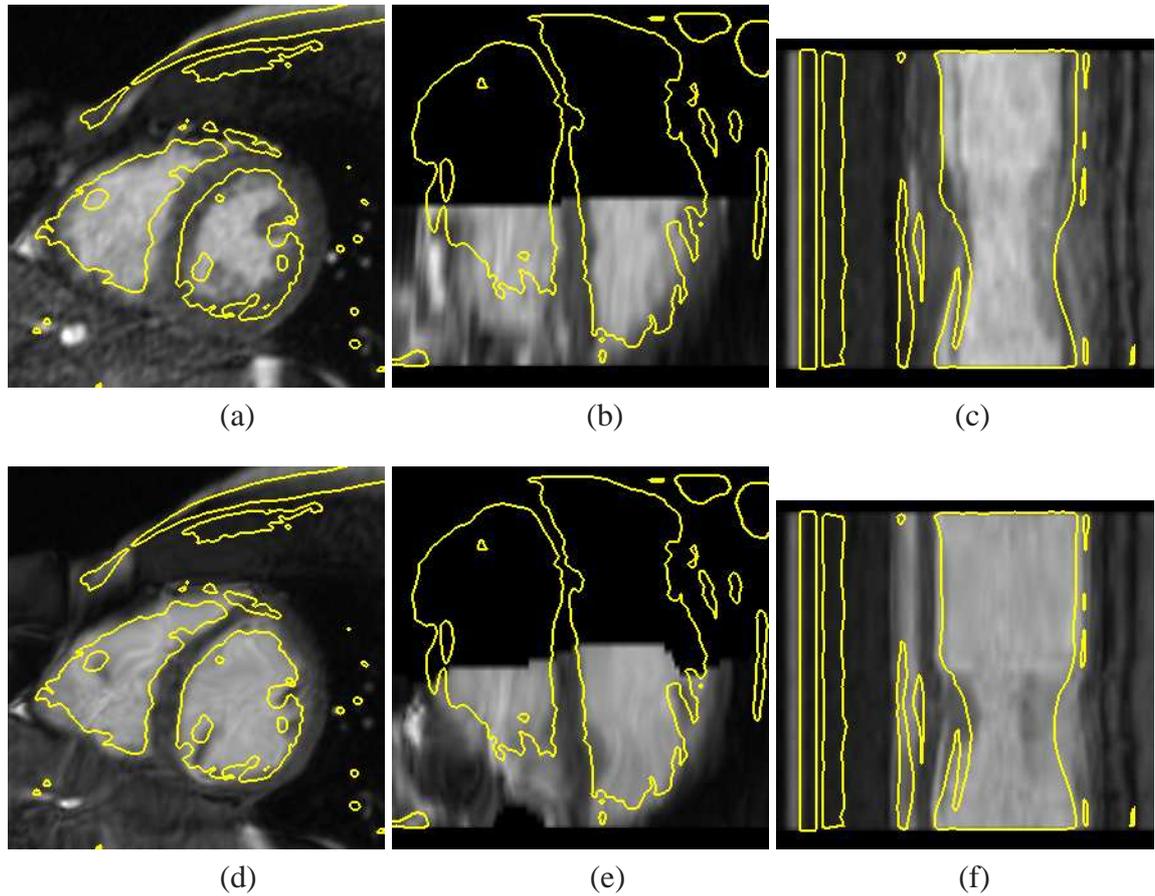
Figure 5.10: Results of the 4D cardiac MR registration algorithm. (a) The short-axis , (b) the long-axis and (c) the temporal views after the affine alignment. The corresponding (d) short-axis, (e) long-axis and (f) temporal views after the spatio-temporal free-form registration (separate optimisation of the transformation components).
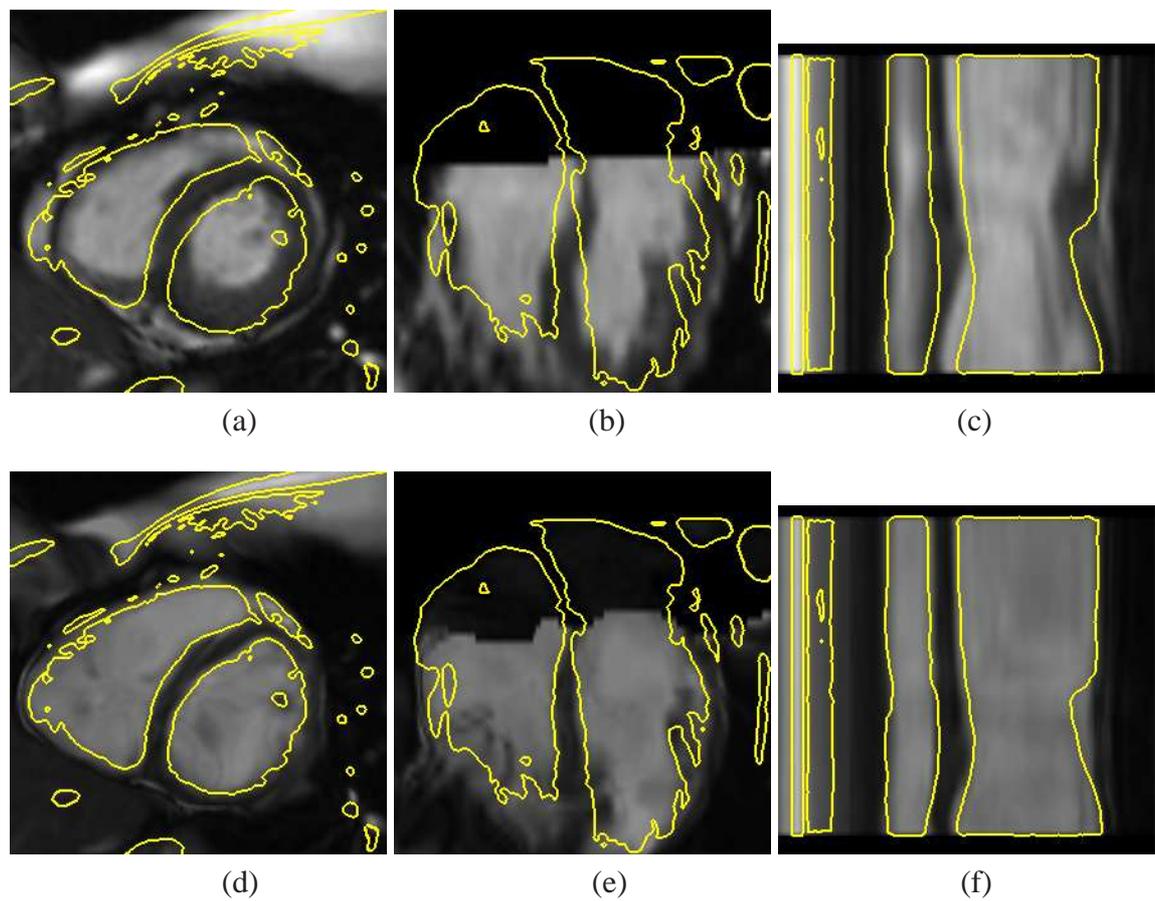
Figure 5.11: Results of the 4D cardiac MR registration algorithm. (a) The short-axis , (b) the long-axis and (c) the temporal views after the affine alignment. The corresponding (d) short-axis, (e) long-axis and (f) temporal views after the spatio-temporal free-form registration (separate optimisation of the transformation components).
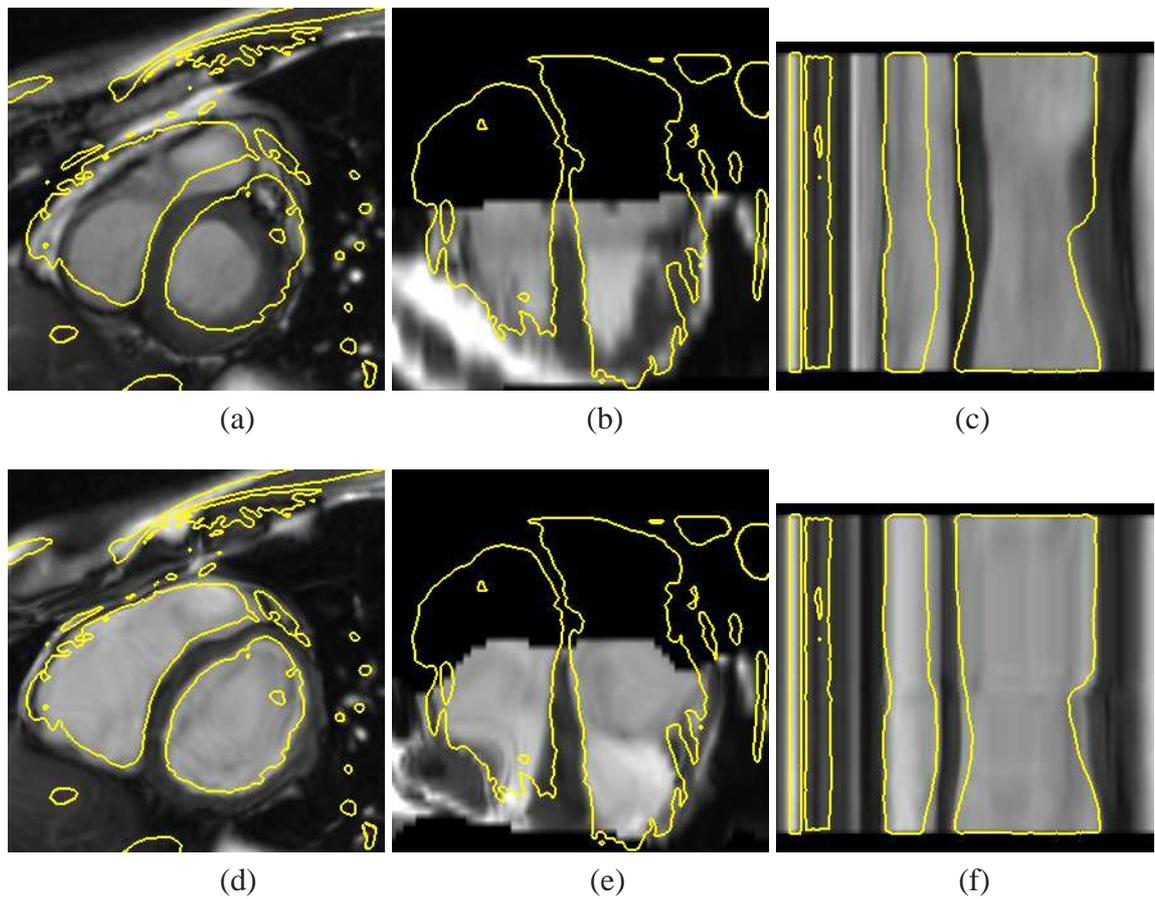
Figure 5.12: Results of the 4D cardiac MR registration algorithm. (a) The short-axis , (b) the long-axis and (c) the temporal views after the affine alignment. The corresponding (d) short-axis, (e) long-axis and (f) temporal views after the spatio-temporal free-form registration (separate optimisation of the transformation components).
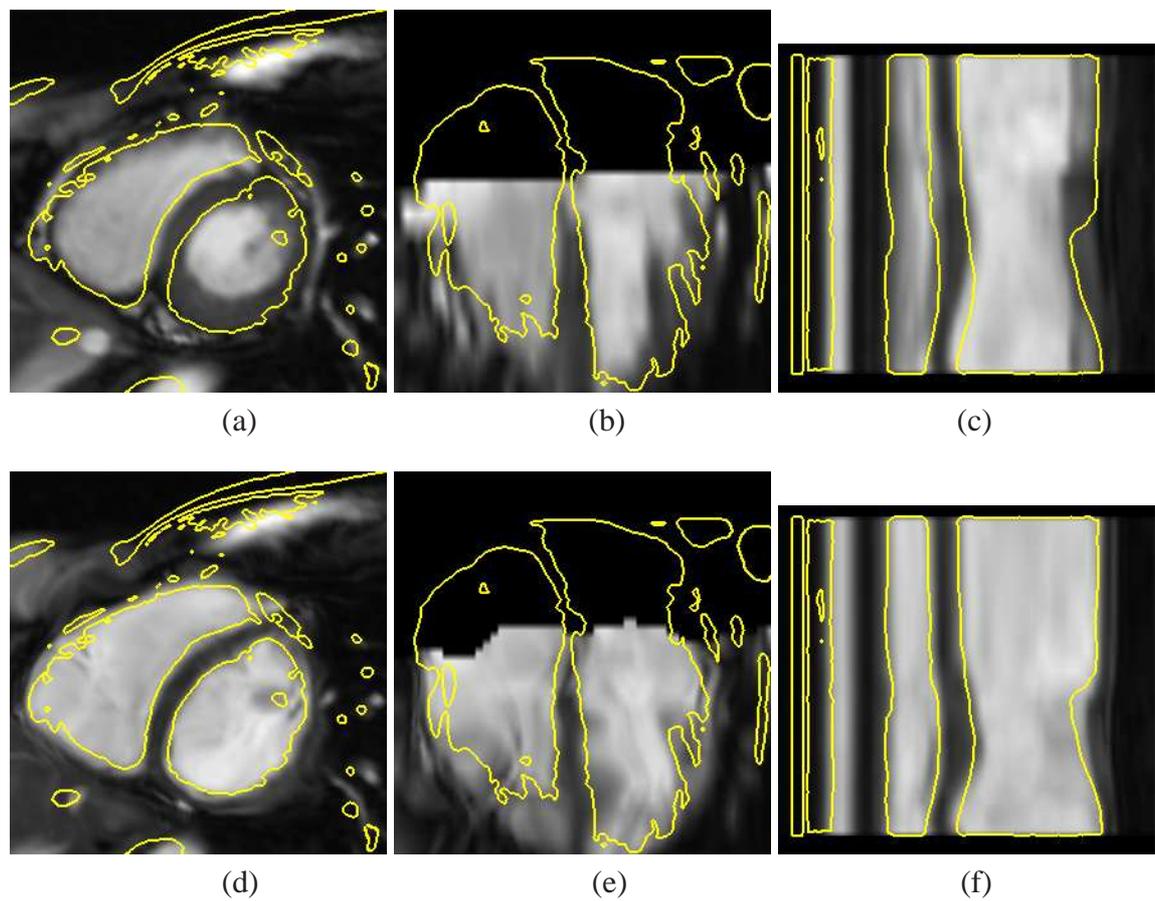
Figure 5.13: Results of the 4D cardiac MR registration algorithm. (a) The short-axis , (b) the long-axis and (c) the temporal views after the affine alignment. The corresponding (d) short-axis, (e) long-axis and (f) temporal views after the spatio-temporal free-form registration (separate optimisation of the transformation components).
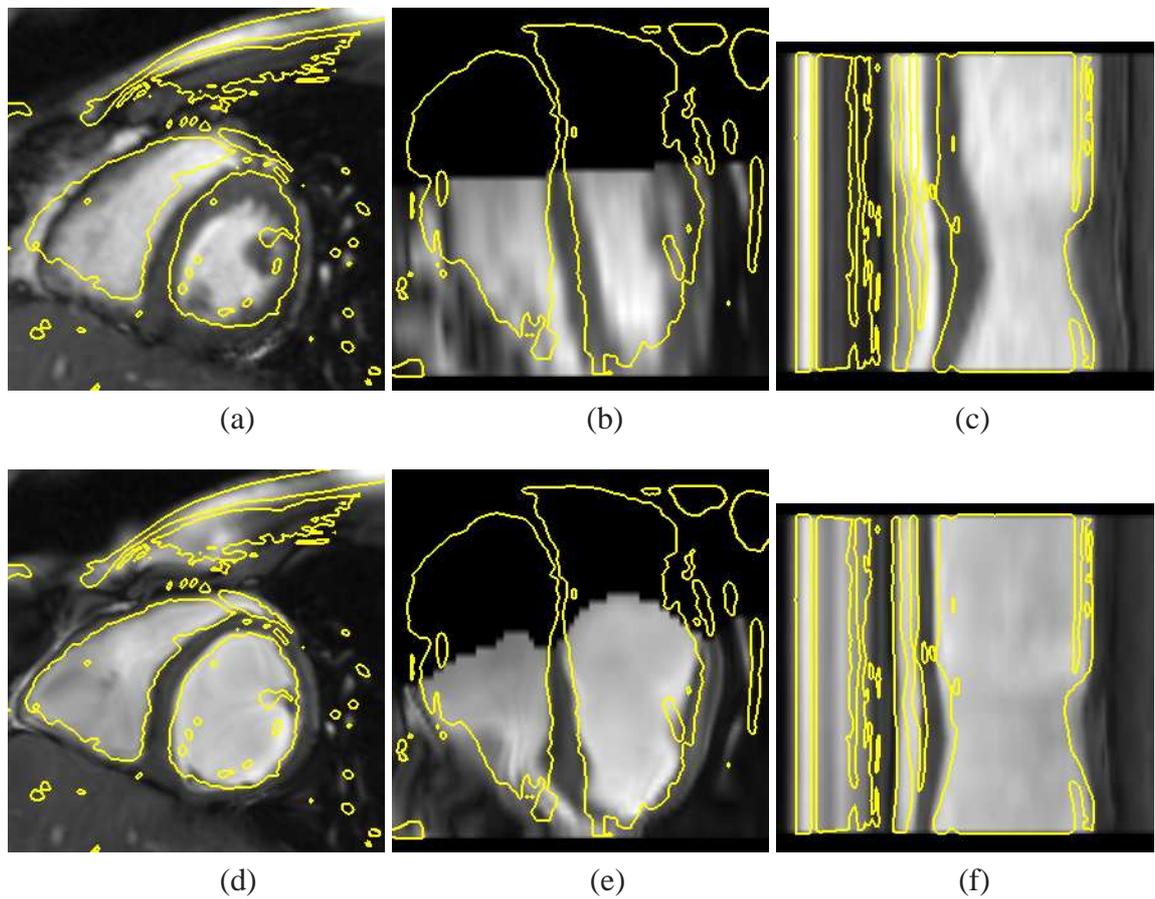
Figure 5.14: Results of the 4D cardiac MR registration algorithm. (a) The short-axis , (b) the long-axis and (c) the temporal views after the affine alignment. The corresponding (d) short-axis, (e) long-axis and (f) temporal views after the spatio-temporal free-form registration (separate optimisation of the transformation components).

5.4.1.2 Quantitative evaluation

The error in the estimation of the peak contraction and end-diastolic positions in the cardiac cycle was calculated by determining manually where these positions appear and comparing them with positions identified by the algorithm. The mean error in the detection of position where peak contraction appears is 1.2 frames while the mean error of the detection of the end-diastolic position is 0.93 frames.

Tables 5.3 and 5.4 show the mean overlap and the mean surface distance (in *mm*) (calculated over all subjects and time frames) for each anatomical region after affine spatio-temporal registration (simultaneous optimisation of the transformation components) and after spatio-temporal free-form registration (separate optimisation of the transformation components). The spacing of the control points in the spatial domain was 10*mm* (table 5.3) and 15*mm* (table 5.4). The affine spatio-temporal registration method is similar to the one described in the previous chapter [152]. It is clear from the figures that the introduction of the *deformable* models substantially improves the overlap of the anatomical features. Furthermore, the computational complexity of this optimisation approach remains low. The volume overlap of the left ventricles at two specific locations in the cardiac cycle has also been calculated. The first location is the position of peak contraction and the other location is the position where the relaxation phase of the heart ends (end-diastole) and the iso-volumetric relaxation phase starts. Table 5.5 shows the mean volume overlap of the left ventricles (calculated over all subjects) at these two positions.

In order to evaluate the quality of the temporal alignment, the mean absolute error in the temporal mapping of the same specific locations in the cardiac cycle has been measured. These temporal positions were manually determined for each image sequence. Table 5.6 shows the temporal error after affine 4D registration and after spatio-temporal *deformable* registration.

	Volume overlap		Surface distance in *mm*	
	Affine 4D	Non-rigid 4D	Affine 4D	Non-rigid 4D
Anatomical region				
Left ventricle	76.16%	82.38%	4.16	3.41
Right ventricle	77.39%	83.56%	4.95	3.93
Myocardium	70.57%	71.62%	4.77	4.21

Table 5.3: The mean volume overlap and surface distance after spatio-temporal free-form registration (separate optimisation of the transformation components) and after affine spatial-temporal registration (combined optimisation of the transformation components). The control spacing in the spatial domain is 10*mm*.

	Volume overlap		Surface distance in *mm*	
	Affine 4D	Non-rigid 4D	Affine 4D	Non-rigid 4D
Anatomical region				
Left ventricle	76.16%	82.80%	4.16	3.37
Right ventricle	77.39%	83.43%	4.95	3.73
Myocardium	70.57%	71.62%	4.77	4.98

Table 5.4: The mean volume overlap and surface distance after spatio-temporal free-form registration (separate optimisation of the transformation components) and after affine spatial-temporal registration (combined optimisation of the transformation components). The control spacing in the spatial domain is 15*mm*.

Volume overlap at specific positions in the cardiac cycle				
	ds 10*mm*		ds 15*mm*	
	Affine 4D	Non-rigid 4D	Affine 4D	Non-rigid 4D
Temporal position				
Peak contraction	74.72%	72.88%	74.72%	74.42%
End-diastole	69.27%	84.63%	69.27%	84.75%

Table 5.5: The mean volume overlap at a specific position in the cardiac cycle after affine 4D registration and after spatio-temporal *deformable* registration (separate optimisation of the transformation components).

Error in the temporal mapping of two specific positions in the cardiac cycle		
	dt 90*msec*	
Temporal position	**Affine**	**Non-rigid**
Peak contraction	73.72*msec*	46.86*msec*
End-diastole	93.57*msec*	30.352*msec*

Table 5.6: The mean absolute error in the temporal mapping of specific positions in the cardiac cycle after affine 4D registration and after spatio-temporal *deformable* registration (separate optimisation of the transformation components).

5.4.2 Combined optimisation of the transformation components

5.4.2.1 Qualitative evaluation

Figure 5.15 (a) shows the volume curves of the left ventricle after the optimisation of the spatio-temporal affine transformation, while figure 5.15 (b) shows the same volume curves and after the optimisation of the spatio-temporal *deformable* transformation. The volume of the left ventricles was calculated using the EM algorithm developed by Valdés *et al.* [110]. As with figure 5.8, it is clear that with the introduction of the *deformable* components the hearts are substantially better aligned in the temporal domain. All the hearts in figure 5.15 (b) show very similar volume curves.

Figures 5.16 - 5.21 provide examples of the spatio-temporal *deformable* registration (combined optimisation of the transformation components). The images in the top rows (a-c) are the short-axis, the long-axis and the temporal views (the temporal view is a short-axis line over time) of a frame in the middle of the image sequence after the optimisation of the global transformation components (affine spatio-temporal registration). The lines in the images represent the intensity iso-lines of the reference image sequence. The images in the bottom rows of figures 5.16 - 5.21 are the same images after spatio-temporal *deformable* registration. The introduction of the *deformable* temporal and spatial registration resulted in a substantial improvement in the alignment of the image sequences both in the spatial and in the temporal domains. The dark areas in the long-axis views are caused by the smaller field of view in the current subject.
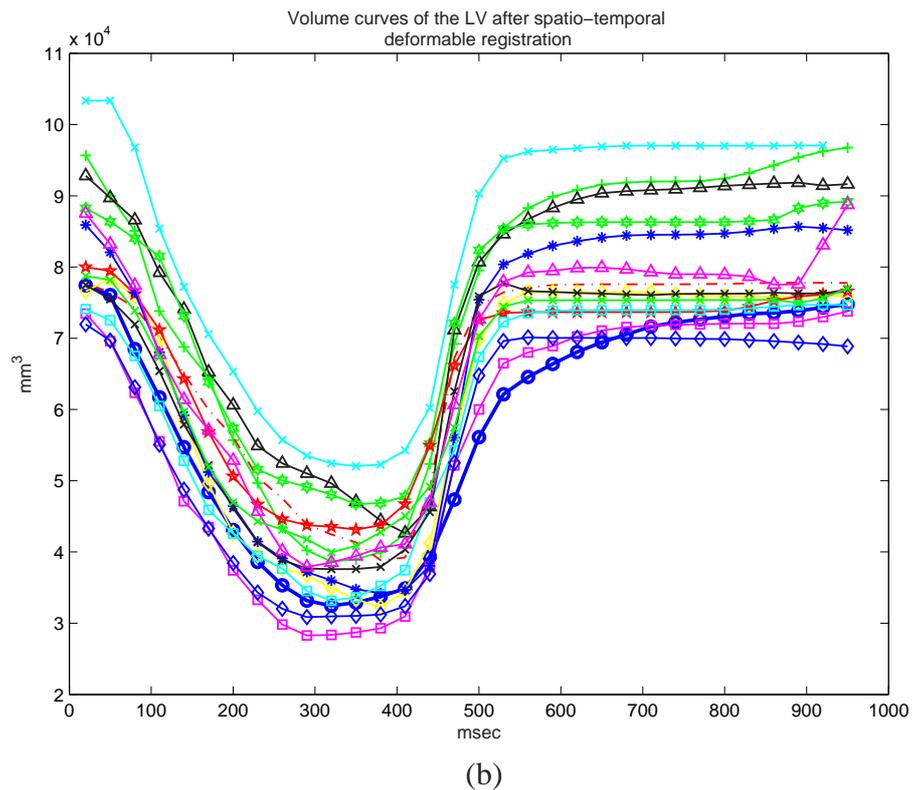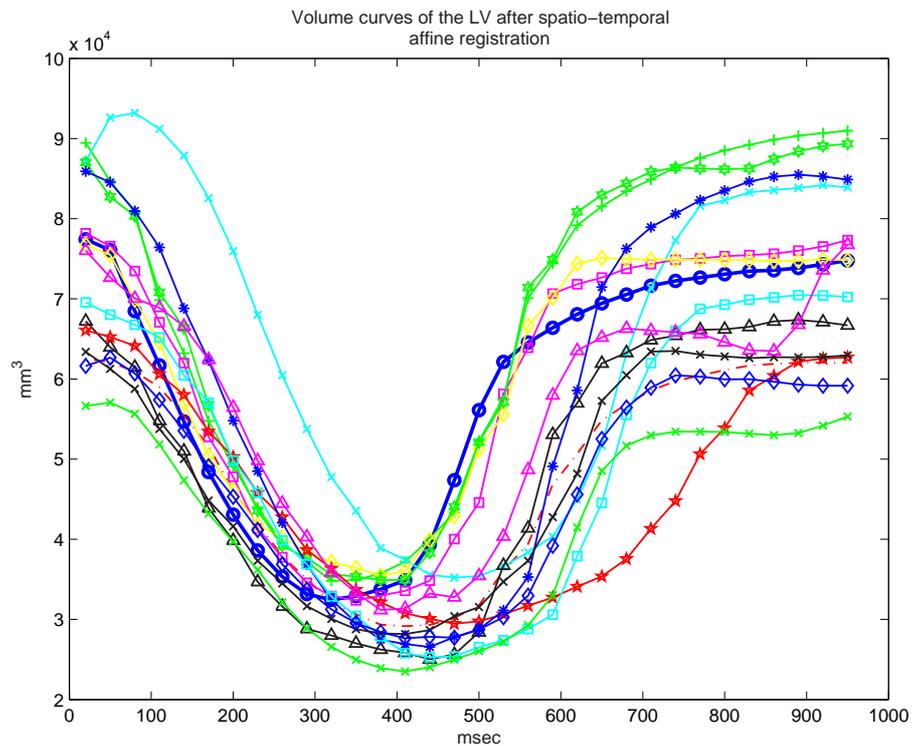
Figure 5.15: The volume curves of the left ventricle for all subjects after affine spatio-temporal registration (a) and after spatio-temporal *deformable* registration (combined optimisation of the transformation components) (b). The spacing of the spatial control point is 10*mm*. The thick line is the volume curve of the reference subject.
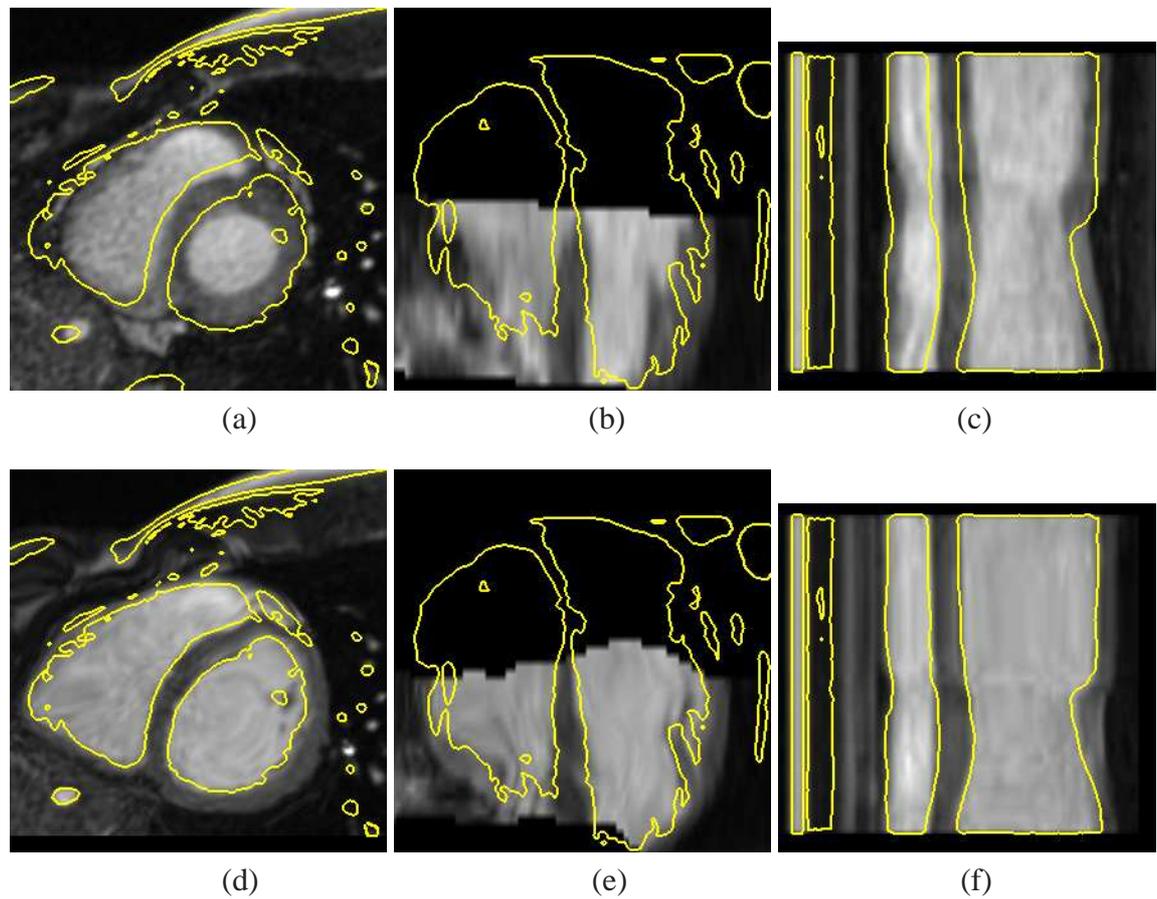
Figure 5.16: Results of the 4D cardiac MR registration algorithm. (a) The short-axis, (b) the long-axis and (c) the temporal views after the affine alignment. The corresponding (d) short-axis, (e) long-axis and (f) temporal views after the spatio-temporal free-form registration (combined optimisation of the transformation components).
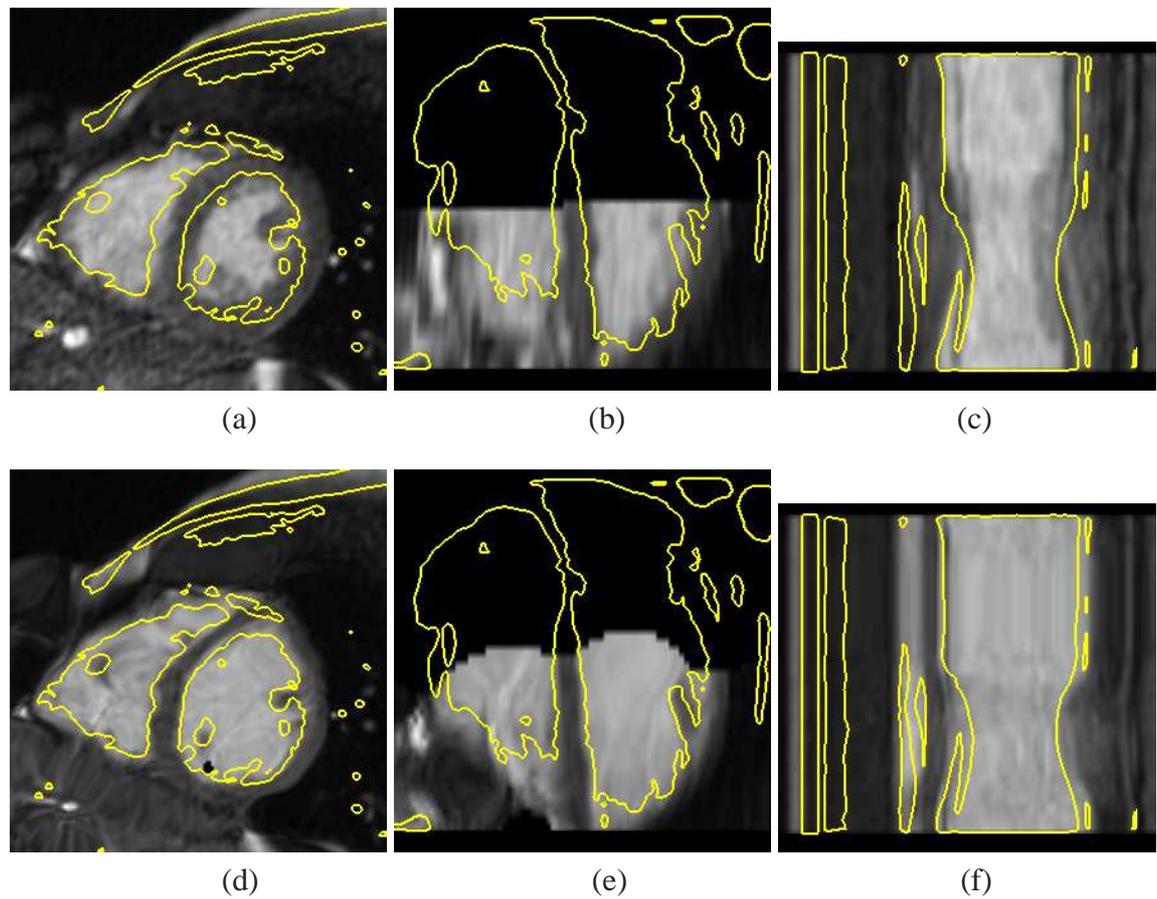
Figure 5.17: Results of the 4D cardiac MR registration algorithm. (a) The short-axis, (b) the long-axis and (c) the temporal views after the affine alignment. The corresponding (d) short-axis, (e) long-axis and (f) temporal views after the spatio-temporal free-form registration (combined optimisation of the transformation components).
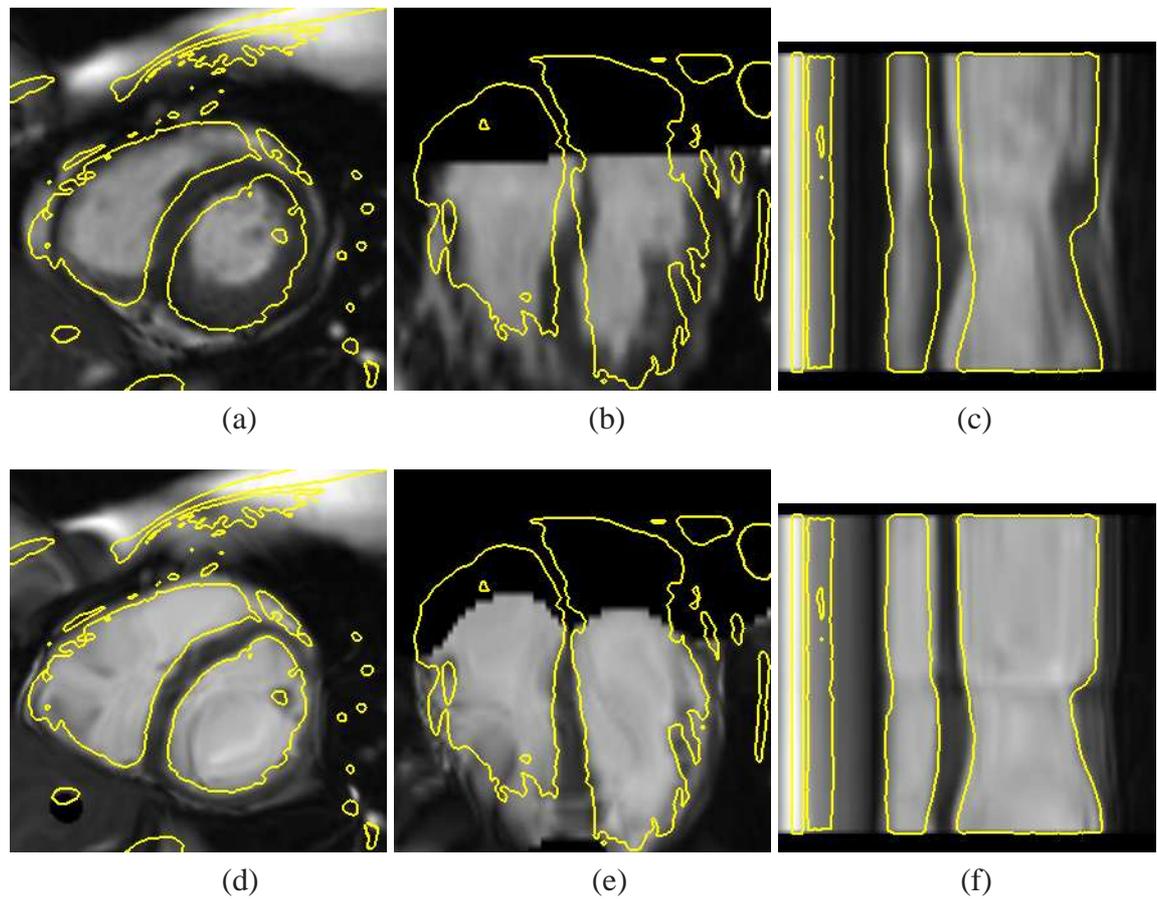
Figure 5.18: Results of the 4D cardiac MR registration algorithm. (a) The short-axis, (b) the long-axis and (c) the temporal views after the affine alignment. The corresponding (d) short-axis, (e) long-axis and (f) temporal views after the spatio-temporal free-form registration (combined optimisation of the transformation components).
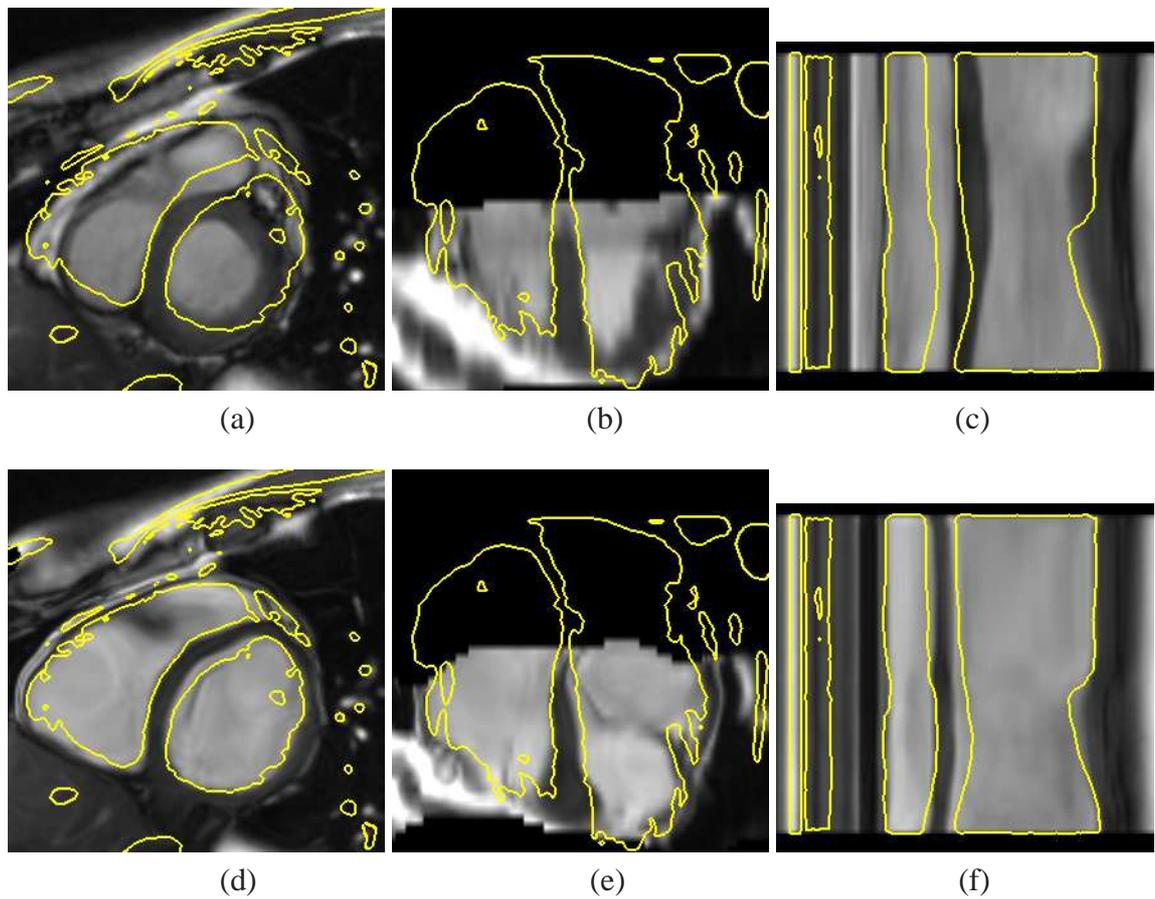
Figure 5.19: Results of the 4D cardiac MR registration algorithm. (a) The short-axis, (b) the long-axis and (c) the temporal views after the affine alignment. The corresponding (d) short-axis, (e) long-axis and (f) temporal views after the spatio-temporal free-form registration (combined optimisation of the transformation components).
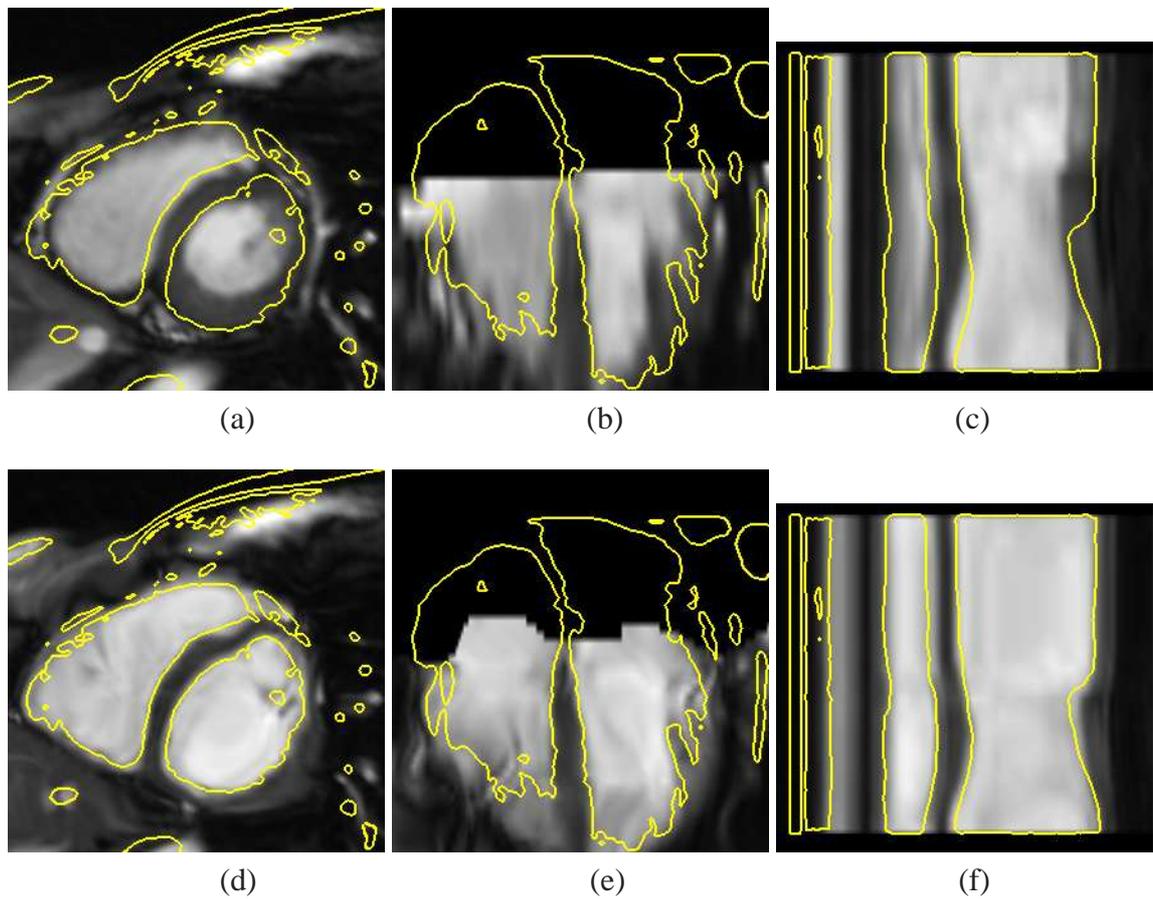
Figure 5.20: Results of the 4D cardiac MR registration algorithm. (a) The short-axis, (b) the long-axis and (c) the temporal views after the affine alignment. The corresponding (d) short-axis, (e) long-axis and (f) temporal views after the spatio-temporal free-form registration (combined optimisation of the transformation components).
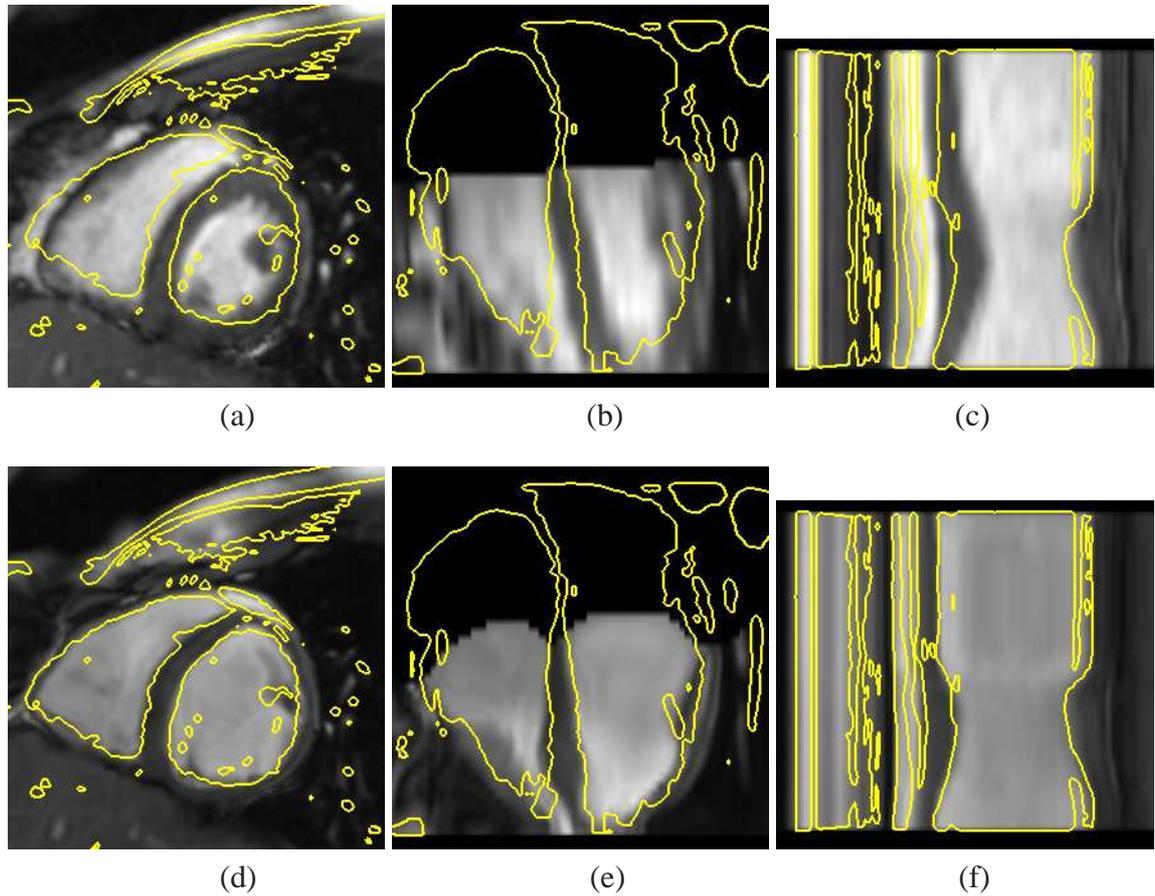
Figure 5.21: Results of the 4D cardiac MR registration algorithm. (a) The short-axis, (b) the long-axis and (c) the temporal views after the affine alignment. The corresponding (d) short-axis, (e) long-axis and (f) temporal views after the spatio-temporal free-form registration (combined optimisation of the transformation components).

5.4.2.2   Quantitative evaluation

Tables 5.7 and 5.8 show the mean volume overlap and the mean surface distance (in *mm*) (calculated for all subjects and at all time frames) for each anatomical region after spatio-temporal affine registration and after the spatio-temporal *deformable* registration. Table 5.9 shows the same overlap and surface distance measures after 3D non-rigid registration of the first frames (by matching the first and the last time frames of the image sequences). The spacing of the control points in the spatial domain was 10*mm* and 15*mm* (tables 5.7 and 5.8). In order to measure how well the images are registered in the temporal domain, the volume overlap of the left ventricles at two specific locations in the cardiac cycle has been measured. The first location is the position of peak contraction and the other location is the position where the relaxation phase of the heart ends (end-diastole) and the iso-volumetric relaxation phase starts (in which the volume of the heart remains roughly the same). The better the temporal alignment, the better the overlap will be in these two locations. Table 5.10 shows the mean volume overlap of the left ventricle (calculated over all subjects) at these two positions after spatio-temporal affine registration, non-rigid registration of the first frames and after spatio-temporal *deformable* registration. Figure 5.22 shows an example of the volume overlap of the left ventricles over time after spatio-temporal *deformable* registration and after non-rigid registration of the first frames. The volume overlap of the left ventricles remains almost constant with the introduction of the *deformable* temporal model, while without the *deformable* temporal model the image sequences become more misaligned as the hearts contract. These results indicate that the use of the *deformable* spatial and temporal parts provides a substantial improvement in the quality of the registration compared to the two methods described in the previous chapter.

Furthermore, the quality of the temporal alignment has been evaluated by calculating the mean absolute error in the temporal mapping of the two specific locations in the cardiac cycle (as previously). Table 5.11 contains the temporal error after non-rigid 3D (mapping the temporal ends) registration and after spatio-temporal *deformable* registra-

tion.

The comparison of the results of tables 5.7 vs. 5.3 and tables 5.8 vs. 5.4 suggests that the simultaneous optimisation of the transformation components provides better overlap measures than those produced by separate optimisation. One reason for this is that the temporal free-form deformation described in section 5.3.3.2 aligns a limited number of temporal positions in the cardiac cycles. On the other hand, the temporal free-form deformation described in section 5.3.3.1 provides a better temporal alignment between these temporal positions. Moreover, the optimisation of the spatial transformation described in section 5.3.3.2 is based only on the first frames of the image sequences while the combined optimisation approach is based on image information from the entire image sequences.

	Volume overlap		Surface distance in *mm*	
Anatomical region	**Affine 4D**	**Non-rigid 4D**	**Affine 4D**	**Non-rigid 4D**
Left ventricle	76.16%	85.57%	4.16	2.96
Right ventricle	77.39%	84.67%	4.95	3.60
Myocardium	70.57%	73.18%	4.77	4.16

Table 5.7: The mean volume overlap and surface distance after affine 4D registration and after spatio-temporal deformable registration (combined optimisation of the transformation components). The control spacing in the spatial domain is 10*mm* and in the temporal domain 90*msec*.

	Volume overlap		Surface distance in *mm*	
Anatomical region	**Affine 4D**	**Non-rigid 4D**	**Affine 4D**	**Non-rigid 4D**
Left ventricle	76.16%	84.95%	4.16	3.05
Right ventricle	77.39%	84.34%	4.95	3.51
Myocardium	70.57%	72.56%	4.77	4.14

Table 5.8: The mean volume overlap and surface distance after affine 4D registration and after spatio-temporal *deformable* registration (combined optimisation of the transformation components). The control spacing in the spatial domain is 15*mm* and in the temporal domain 90*msec*.

Anatomical region	Volume overlap		Surface distance in *mm*	
	ds 10*mm*	ds 15*mm*	ds 10*mm*	ds 15*mm*
	Non-rigid 3D	Non-rigid 3D	Non-rigid 3D	Non-Rigid 3D
Left ventricle	80.97%	81.66%	3.56	3.47
Right ventricle	83.38%	83.18%	3.93	3.73
Myocardium	71.63%	71.66%	4.21	4.02

Table 5.9: The mean volume overlap and surface distance after *deformable* 3D registration of the first frames (matching the temporal ends of the sequences). The control spacing in the spatial domain is 10*mm* and 15*mm*.

Volume overlap at specific positions in the cardiac cycle					
	ds 10*mm*			ds 15*mm*		
	Affine	Non-rigid		Affine	Non-rigid	
Temporal position	4D	3D	4D	4D	3D	4D
Peak contraction	74.72%	72.75%	82.33%	74.72%	73.99%	80.18%
End-diastole	69.27%	77.41%	85.62%	69.27%	79.61%	85.32%

Table 5.10: The mean volume overlap at a specific position in the cardiac cycle after affine 4D registration, after the *deformable* 3D and after spatio-temporal *deformable* registration (combined optimisation of the transformation components).

Error in the temporal mapping of two specific positions in the cardiac cycle				
	dt 90*msec*			
	ds 10*mm*		ds 15*mm*	
	Non-rigid	Non-rigid	Non-rigid	Non-rigid
Temporal position	3D	4D	3D	4D
Peak contraction	73.72*msec*	58.03*msec*	73.72*msec*	62.76*msec*
End-diastole	93.57*msec*	31.84*msec*	93.57*msec*	27.31*msec*

Table 5.11: The mean absolute error in the temporal mapping of specific positions in the cardiac cycle after non-rigid 3D registration (mapping the temporal ends) and after spatio-temporal *deformable* registration (combined optimisation of the transformation components).

5.4.2.3 Using the cross correlation based method to calculate the temporal alignment

Similar experiments with the same image sequences have been performed in order to investigate if the cross-correlation based method (described on section 5.3.3.2) can be used to calculate an initial estimate of the temporal transformation during the combined
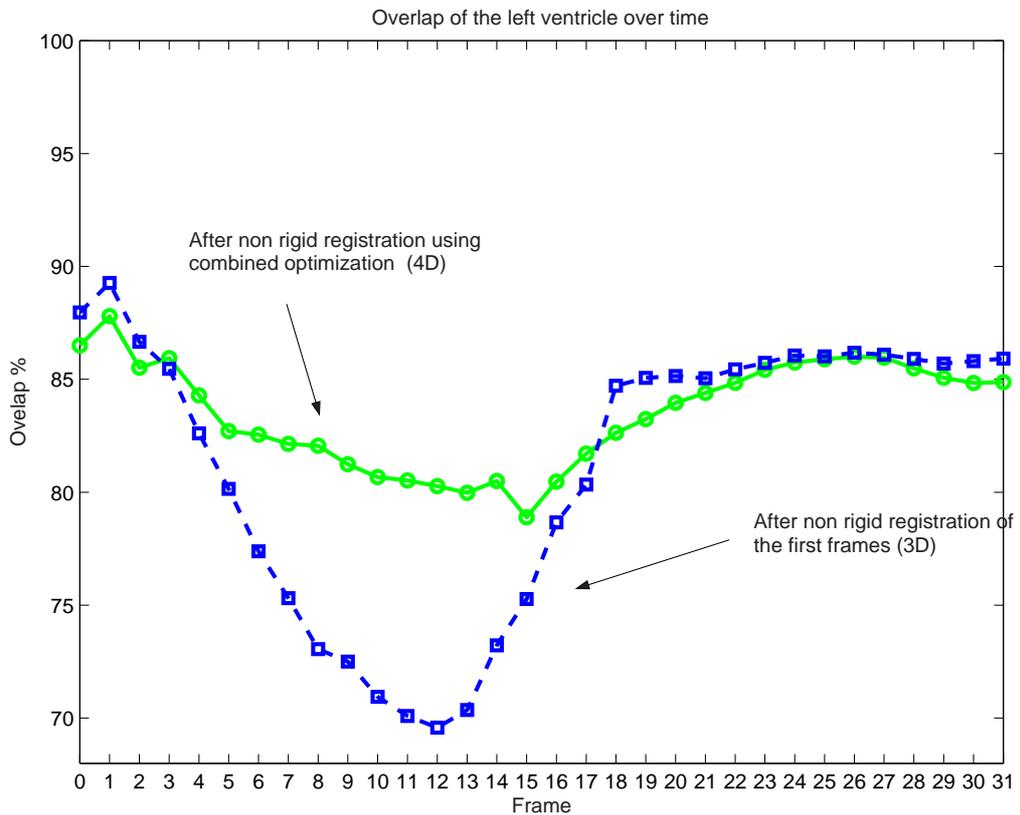
Figure 5.22: The volume overlap over time of the left ventricles of two image sequences after 4D non-rigid registration (combined optimisation of the transformation components) and after non-rigid registration of the first frames (matching the first and last frames of the image sequences).

optimisation method. More precisely the aim of the experiments is to test how the use of a better initial estimation of the temporal transformation may affect the performance of the current spatio-temporal registration method. Two experiments using the cross-correlation based method were performed:

- The aim of the first experiment was to investigate the effects of a better initial estimation of the temporal transformation on the performance of the spatio-temporal *deformable* registration algorithm which optimises the transformation components simultaneously. Hence, before optimising the *deformable* part of the spatial and temporal transformations, an initial estimate of the temporal transformation using the cross-correlation based method was calculated. This method is referred to as experiment **A**.

- The aim of the second experiment was to improve the computational efficiency of this spatio-temporal registration algorithm which optimises the transformation components simultaneously. The computational complexity of this algorithm is very high due to the large number of *degrees of freedom* of the transformation and to the large number of voxels in the image sequences. A typical combined optimisation takes more than 24 hours. The performance of the registration algorithm was investigated when the cross-correlation based method is used to calculate the temporal transformation $\mathbf{T}^{local}_{temporal}$ and only $\mathbf{T}^{local}_{spatial}$ is optimised. This method is referred to as experiment **B**.

Tables 5.12 and 5.13 provide the volume overlap and the surface distance measures for the above two experiments (with spatial control point spacing 10*mm* and 15*mm*). Figures 5.23 - 5.28 provide examples of the non-rigid registrations obtained by using the method as previously described for (experiment **A**). The results of these figure are very similar to figures 5.16 - 5.21. As expected, there is no substantial difference between the results. The algorithm converges to a similar solution whether an initial estimate of the temporal transformation has been used or not. The mean volume overlap of the left ventricles at specific temporal locations is described in table 5.14.

	Volume overlap		Surface distance in *mm*	
	Experiment A	**Experiment B**	**Experiment A**	**Experiment B**
Anatomical region	**4D**	**4D**	**4D**	**4D**
Left ventricle	85.91%	85.70%	2.97	3.16
Right ventricle	84.84%	84.79%	3.72	3.74
Myocardium	73.35%	73.26%	4.16	4.14

Table 5.12: The mean volume overlap and surface distance after STFFD registration using a temporal initial estimate and after using the cross-correlation based method to calculate the temporal part and optimising only the spatial part. The control spacing in the spatial domain is 10*mm*.

The mean absolute error in the temporal alignment of specific feature location of the cardiac cycle was also measured for experiment **A**. Table 5.15 shows the error in *msec*. A comparison of the temporal error measurements of table 5.15 to those of table 5.11 indi-

	Volume overlap		Surface distance in *mm*	
	Experiment A	Experiment B	Experiment A	Experiment B
Anatomical region	**4D**	**4D**	**4D**	**4D**
Left ventricle	85.12%	84.85%	3.06	3.13
Right ventricle	84.21%	84.29%	3.57	3.55
Myocardium	72.74%	72.45%	4.09	4.14

Table 5.13: The mean volume overlap and surface distance after STFFD registration using a temporal initial estimate and after using the cross-correlation based method to calculate the temporal part and optimising only the spatial part. The control spacing in the spatial domain is 15*mm*.

Volume overlap at specific positions in the cardiac cycle				
	ds 10*mm*		ds 15*mm*	
	Experiment A	Experiment B	Experiment A	Experiment B
Temporal position	**4D**	**4D**	**4D**	**4D**
Peak contraction	83.15%	81.92%	81.22%	79.86%
End-diastole	85.63%	86.17%	85.19%	85.79%

Table 5.14: The mean volume overlap at specific position in the cardiac cycle after STFFD registration using a temporal initial estimate and after using the cross-correlation based method to calculated temporal part and optimising only the spatial part.

Error in the temporal mapping of two specific positions in the cardiac cycle		
	dt 90*msec*	
	Experiment A	
Temporal position	ds 10*mm*	ds 15*mm*
Peak contraction	56.85*msec*	54.69*msec*
End-diastole	35.89*msec*	34.22*msec*

Table 5.15: The mean absolute error in the temporal mapping of specific positions in the cardiac cycle after the STFFD registration using a temporal initial estimate.

cates that there are no substantial differences, between the two methods, in the temporal alignment of these two specific feature positions. This also indicates that the algorithms converge to a similar solution in both cases.
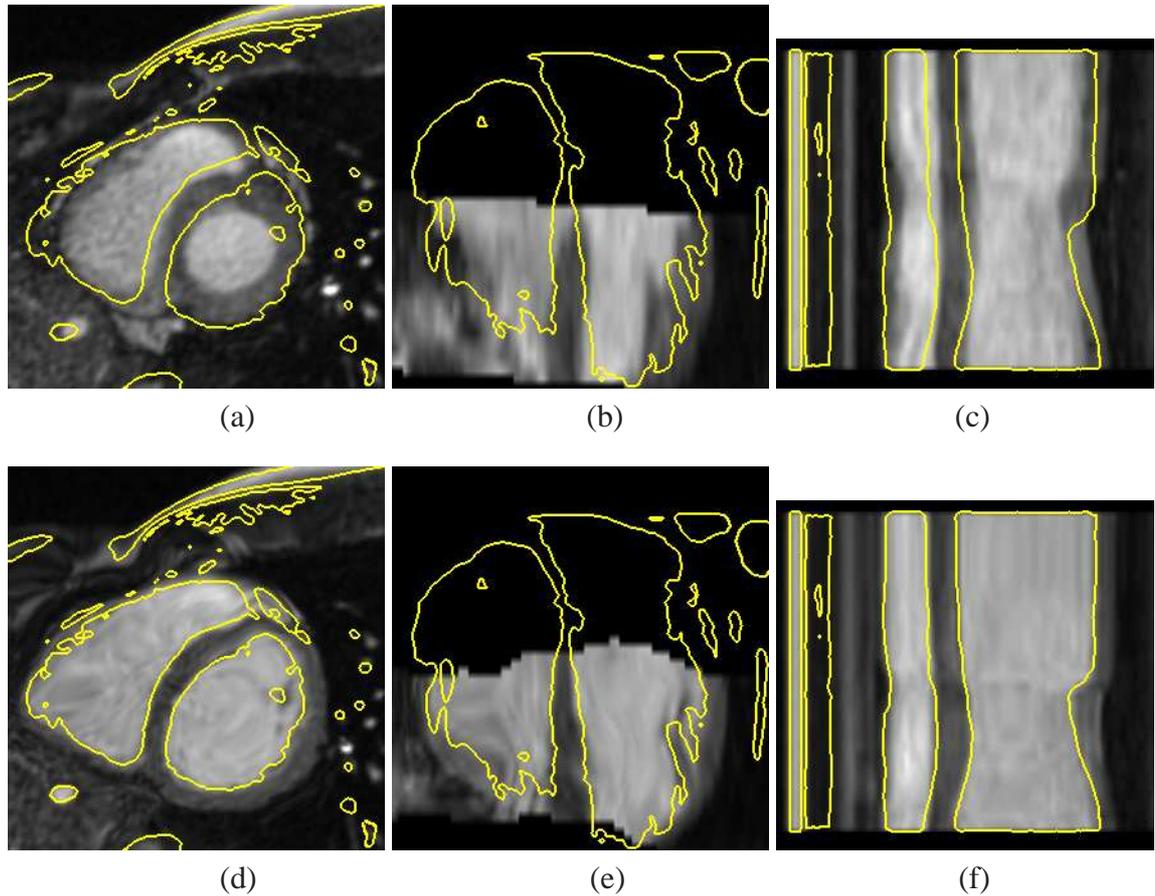
Figure 5.23: Results of the 4D cardiac MR registration algorithm. The short-axis (a), the long-axis (b) and the temporal (c) views after the affine alignment. The corresponding (d) short-axis, long-axis (e) and temporal (f) views after using the cross-correlation based method to calculate an initial estimate of the temporal part followed by combined optimisation of the transformation components (experiment **A**).

5.5   Discussion

The results reported on tables 5.6 and 5.11 and in figure 5.30 indicate that calculating the temporal transformation using either the cross-correlation based approach or the combined optimisation approach provides similar results. Figure 5.30 shows the temporal mapping of the cardiac cycles of six randomly selected image sequences and the reference image sequence. This figure presents the resulting temporal mappings $\mathbf{T}^{global}_{temporal}$ and $\mathbf{T}^{local}_{temporal}$ after combined optimisation and $\mathbf{T}^{local}_{temporal}$ after separate optimisation. The temporal mappings after separate optimisation and after combined optimisation have similar shapes. However, the temporal mapping after the combined optimisation aligns better
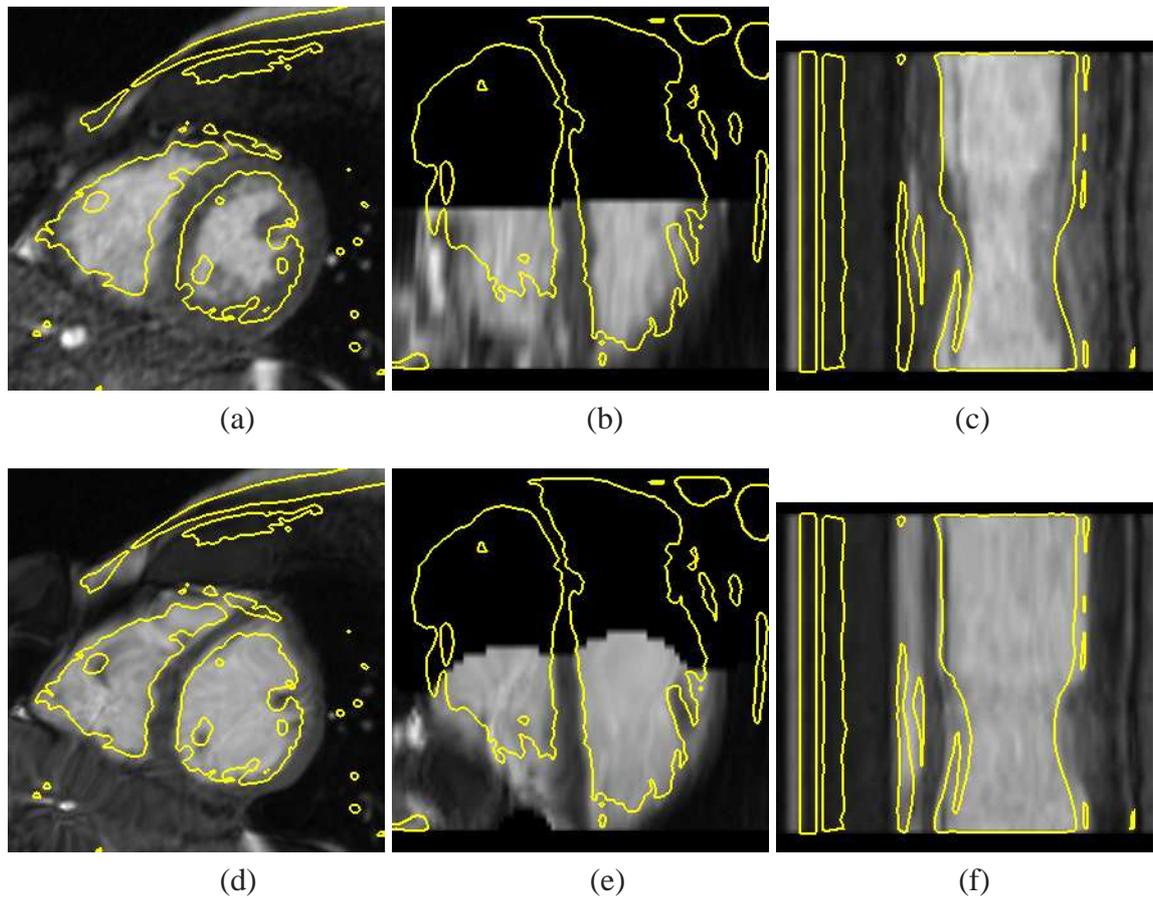
Figure 5.24: Results of the 4D cardiac MR registration algorithm. The short-axis (a), the long-axis (b) and the temporal (c) views after the affine alignment. The corresponding (d) short-axis, long-axis (e) and temporal (f) views after using the cross-correlation based method to calculate an initial estimate of the temporal part followed by combined optimisation of the transformation components (experiment **A**).

the temporal details between the two image sequences. The differences in the resulting temporal mapping towards the end of the cardiac cycle are mainly caused by the long iso-volumetric relaxation of the reference subject. As seen from the volume curves of the left ventricles in figure 5.15 the reference subject has a long iso-volumetric relaxation phase where the shape and the volume of the heart remains almost the same. This creates a small degree of ambiguity for the temporal alignment since all frames in this cardiac phase are very similar.

The volume curves in figures 5.8 and 5.15 show that the contraction phase of the hearts is better aligned than the relaxation phase. The corrected curves appear to have a long end-diastolic phase and a rapid filling phase. The main reason for this is the temporal
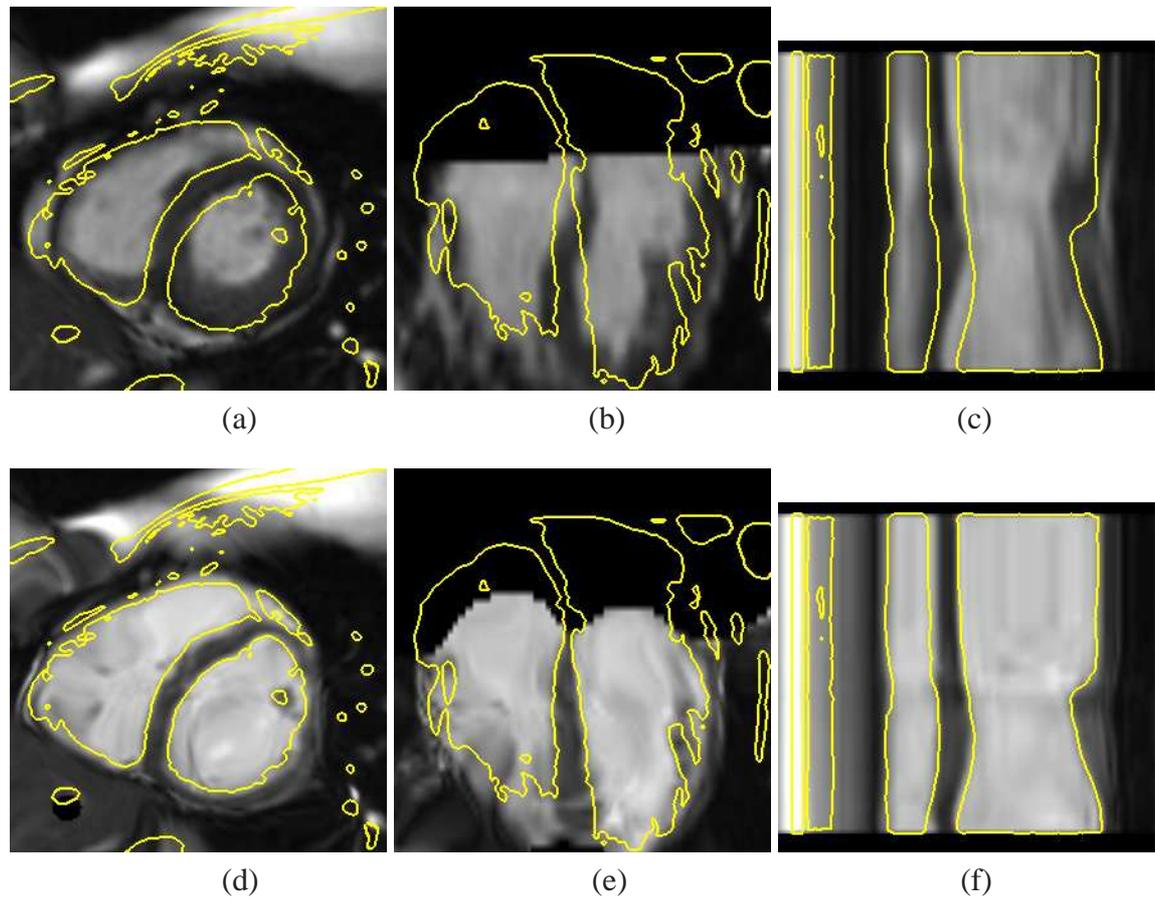
Figure 5.25: Results of the 4D cardiac MR registration algorithm. The short-axis (a), the long-axis (b) and the temporal (c) views after the affine alignment. The corresponding (d) short-axis, long-axis (e) and temporal (f) views after using the cross-correlation based method to calculate an initial estimate of the temporal part followed by combined optimisation of the transformation components (experiment **A**).

ambiguities which are generated by the long iso-volumetric relaxation of the reference subject. When the registration approach attempts to align this part of the cardiac cycles, it is forced to compress the relaxation phase.

Optimising the spatial part of the transformation based only on information from the first frame of the sequences (tables 5.3 and 5.4) does not result in as good a spatial alignment as the one based on information from the entire image sequences (tables 5.7 and 5.8). This suggests that the shape differences between two cardiac anatomies do not remain constant over the cardiac cycle. The results of table 5.5 show that the volume overlap between two cardiac anatomies decreases during the contraction phase. Therefore, the cardiac anatomies become more misaligned in the spatial domain during the contraction
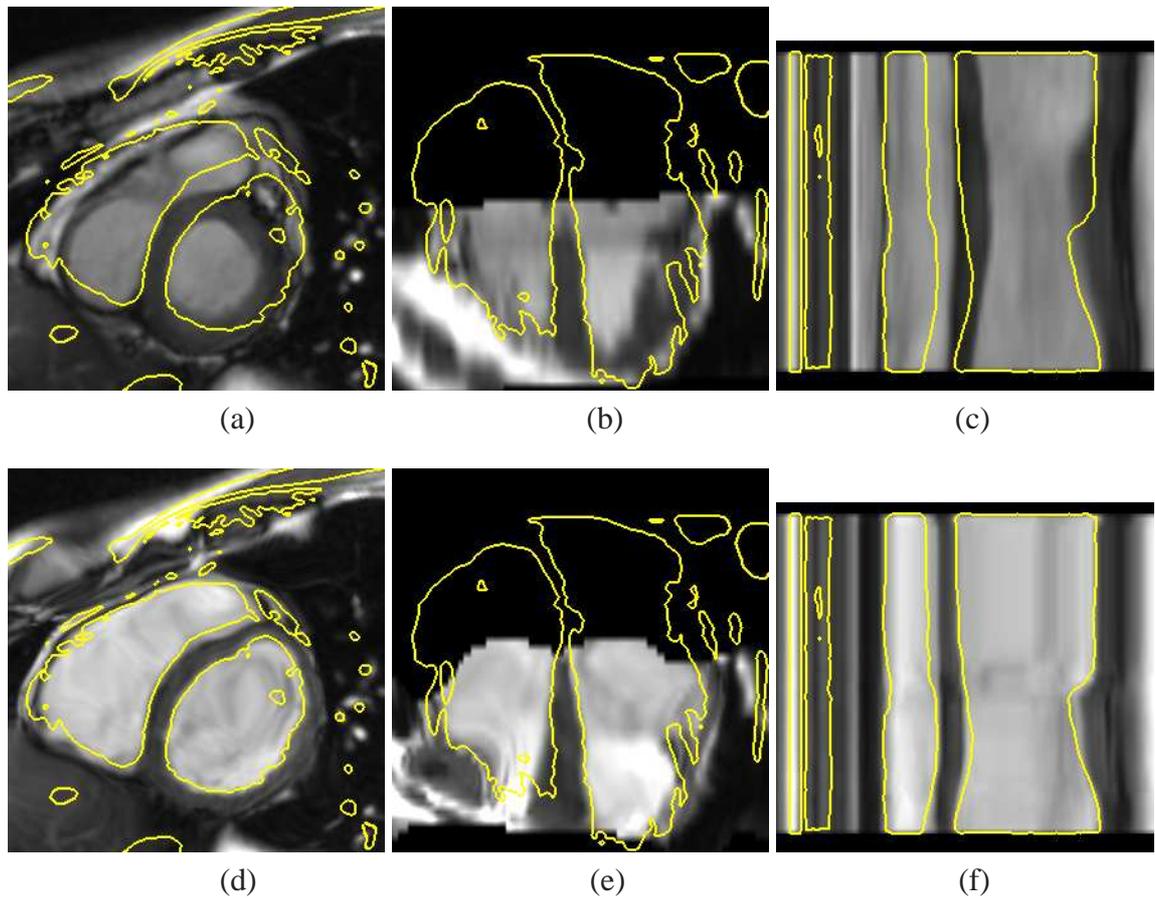
Figure 5.26: Results of the 4D cardiac MR registration algorithm. The short-axis (a), the long-axis (b) and the temporal (c) views after the affine alignment. The corresponding (d) short-axis, long-axis (e) and temporal (f) views after using the cross-correlation based method to calculate an initial estimate of the temporal part followed by combined optimisation of the transformation components (experiment **A**).

phase. At the end-diastolic position the volume overlap of the image sequences is similar for both the separate and combined optimisation of the transformation components (tables 5.5 and 5.10).

Optimising the spatial part of the transformation based on information from the entire image sequences, and not only from the first frames, takes into consideration the fact that the spatial differences of the hearts are not constant over the cardiac cycle. Hence, the mean volume overlap and surface distance measures provided in tables 5.12 - 5.14 (experiment **B**) are better than the ones in tables 5.3 - 5.5 even though the same $\mathbf{T}_{temporal}$ used in both methods.

The importance of the results of tables 5.12 - 5.14 is that using the cross-correlation

<center>(a)   (b)   (c)</center>

<center>(d)   (e)   (f)</center>
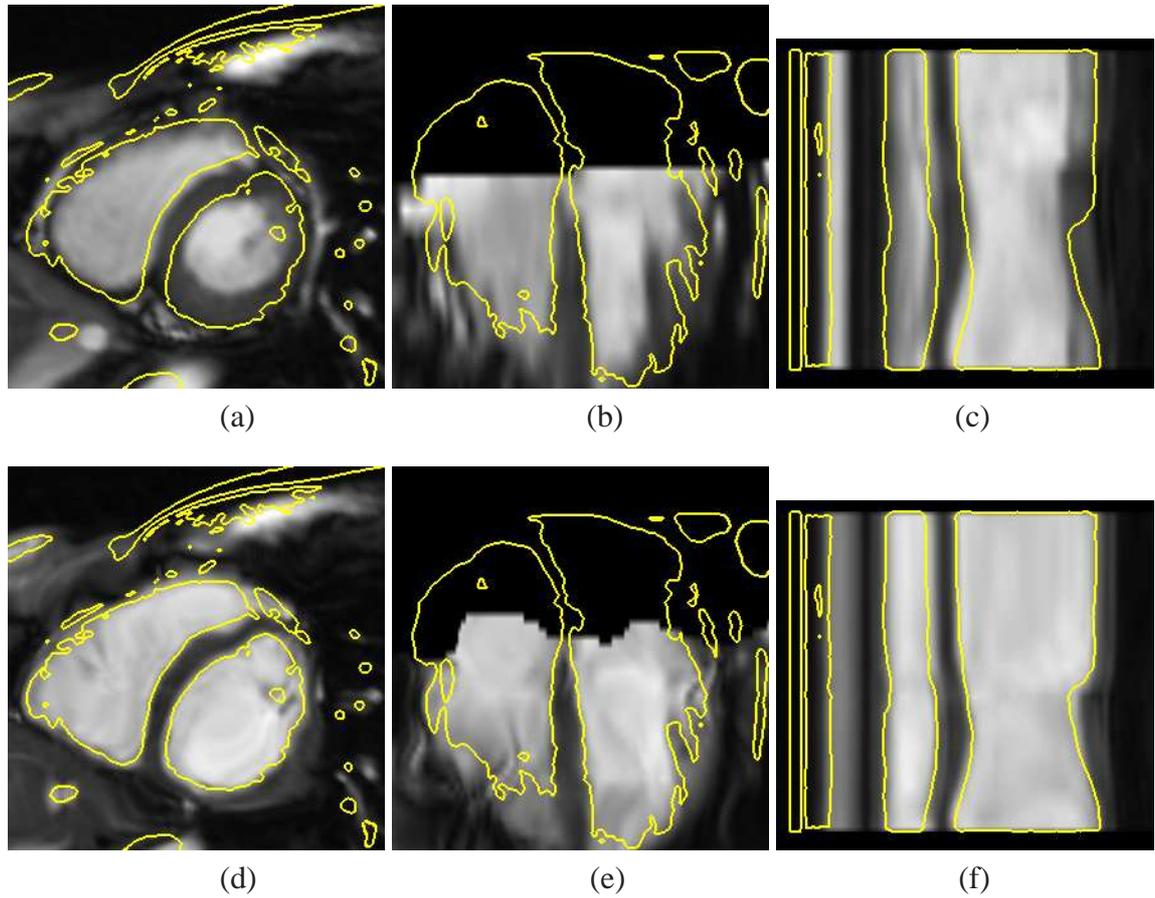
Figure 5.27: Results of the 4D cardiac MR registration algorithm. The short-axis (a), the long-axis (b) and the temporal (c) views after the affine alignment. The corresponding (d) short-axis, long-axis (e) and temporal (f) views after using the cross-correlation based method to calculate an initial estimate of the temporal part followed by combined optimisation of the transformation components (experiment **A**).

based method to calculate the temporal transformation and optimising only the spatial transformation provides similar spatio-temporal registration to the method which performs combined optimisation. However, in this case the computational complexity of the algorithm is substantially reduced since the gradient vector of the cost function (equation 5.9) with respect to the non-rigid transformation parameters Φ_t, $\nabla \mathcal{C}_t = \frac{\partial \mathcal{C}(\alpha, \beta, \Phi_t^l)}{\partial \Phi_t^l}$, is not calculated during the optimisation approach. Figures 5.31 - 5.36 provide examples of the non-rigid free-form spatio-temporal registration by using the cross-correlation method to calculate the temporal part and optimising only the spatial part (experiment **B**).

Figure 5.29 shows the overlap of the left ventricles by the methods described in experiment **A** and in experiment **B**. These results are very similar to those reported for the
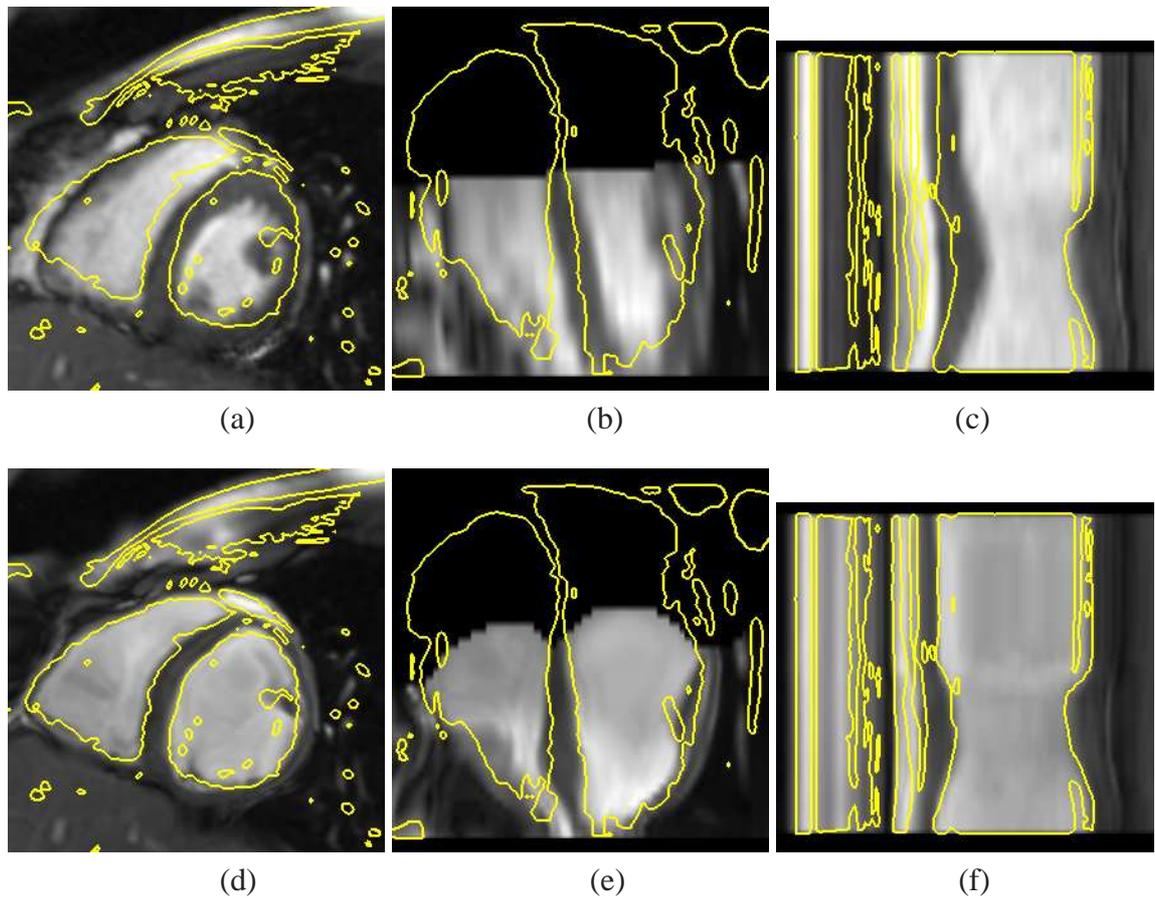
Figure 5.28: Results of the 4D cardiac MR registration algorithm. The short-axis (a), the long-axis (b) and the temporal (c) views after the affine alignment. The corresponding (d) short-axis, long-axis (e) and temporal (f) views after using the cross-correlation based method to calculate an initial estimate of the temporal part followed by combined optimisation of the transformation components (experiment **A**).

spatio-temporal free-form registration by using combined optimisation of the transformation components (figure 5.22).

Optimising the temporal and spatial *deformable* components simultaneously results in a very good spatio-temporal registration. However, the computational complexity algorithm is very high due to the large number of *degrees of freedom* in the transformation and to the large number of voxels in the image sequences. A typical combined optimisation takes more than 24 hours. Optimising only the spatial part of the transformation and using the already calculated temporal transformation as $\mathbf{T}^{local}_{temporal}$ reduces the computational complexity of the algorithm while still providing a good spatio-temporal registration.

The computational complexity can be further reduced by optimising the transforma-
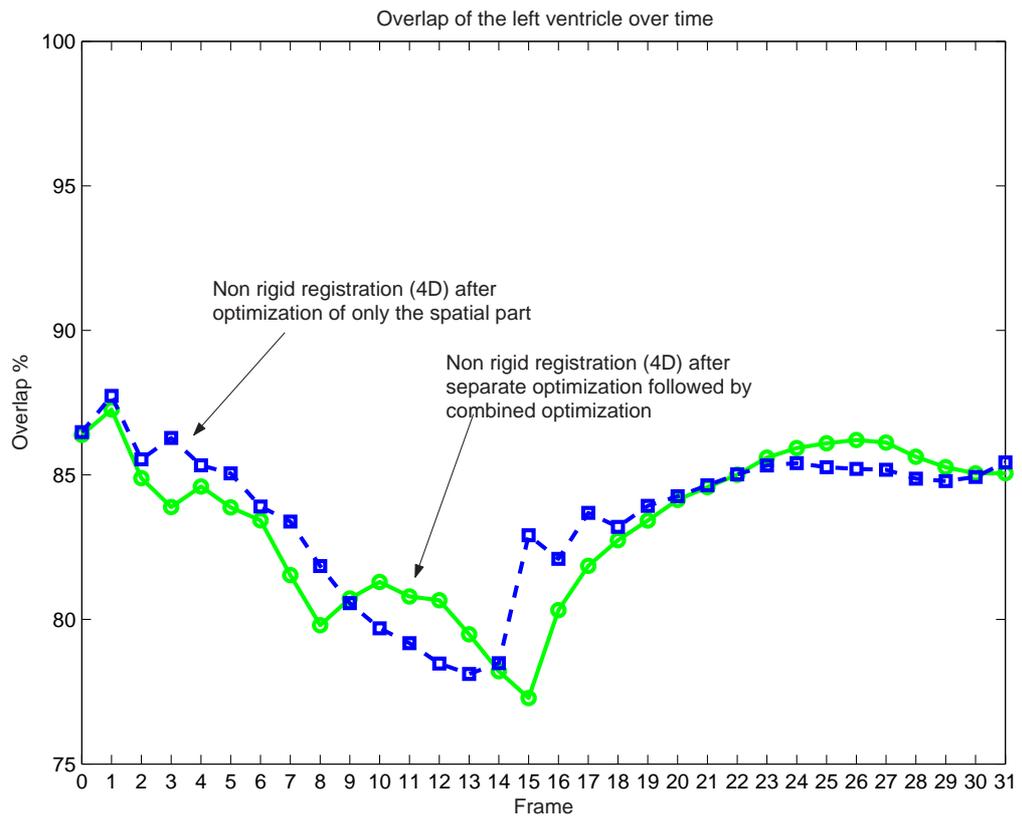
Figure 5.29: The volume overlap of the left ventricles over time after the STFFD registration using a temporal initial estimate (experiment **A**) and after using the cross-correlation based method to calculate the temporal part and optimising only the spatial part (experiment **B**).
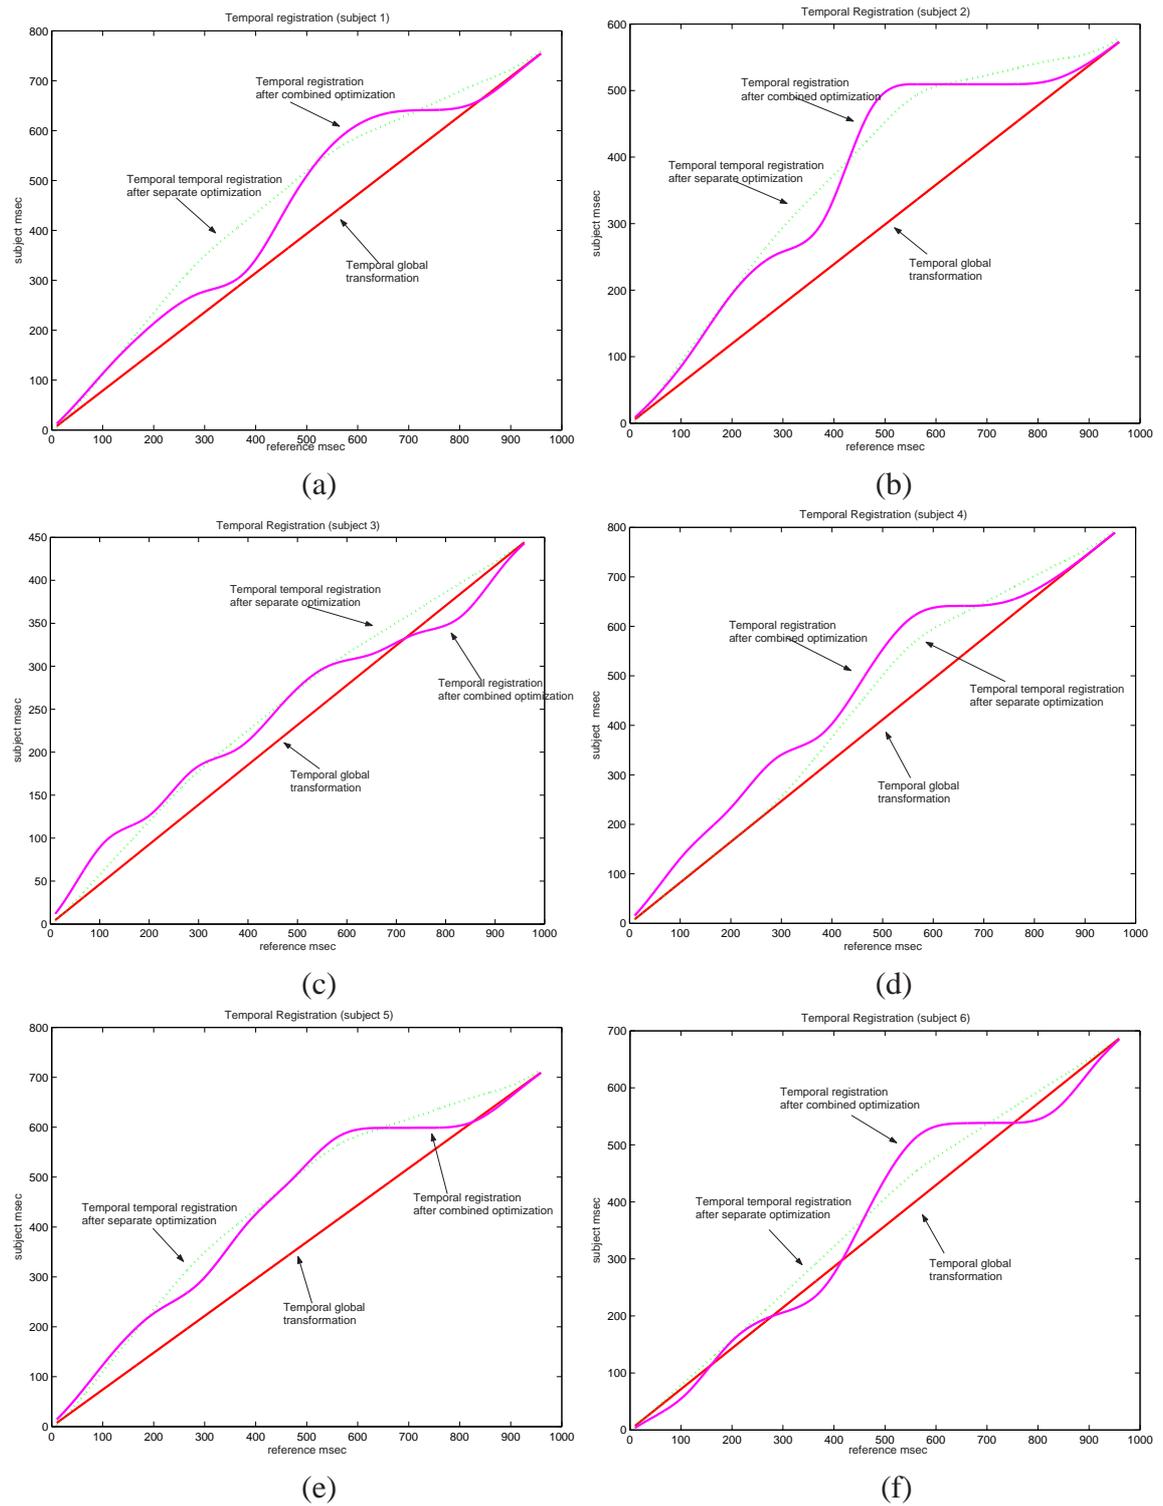
Figure 5.30: Temporal alignment between six image sequences and the reference image sequence. The plots show the temporal mapping between the image sequences after optimising the global temporal transformation component, after optimising the local temporal transformation component (using simultaneous optimisation) and after optimising the local temporal transformation component (using separate optimization)

tion components separately. However, in this case the results of the registration are not as good as those obtained by the other methods. An additional advantage of calculating the temporal transformation by using the method described in section 5.3.3.2 is that a temporal alignment between two cardiac image sequences can be generated without having to perform image registration. There is no need for the image sequences to be registered in the spatial domain in order to calculate the temporal transformation.
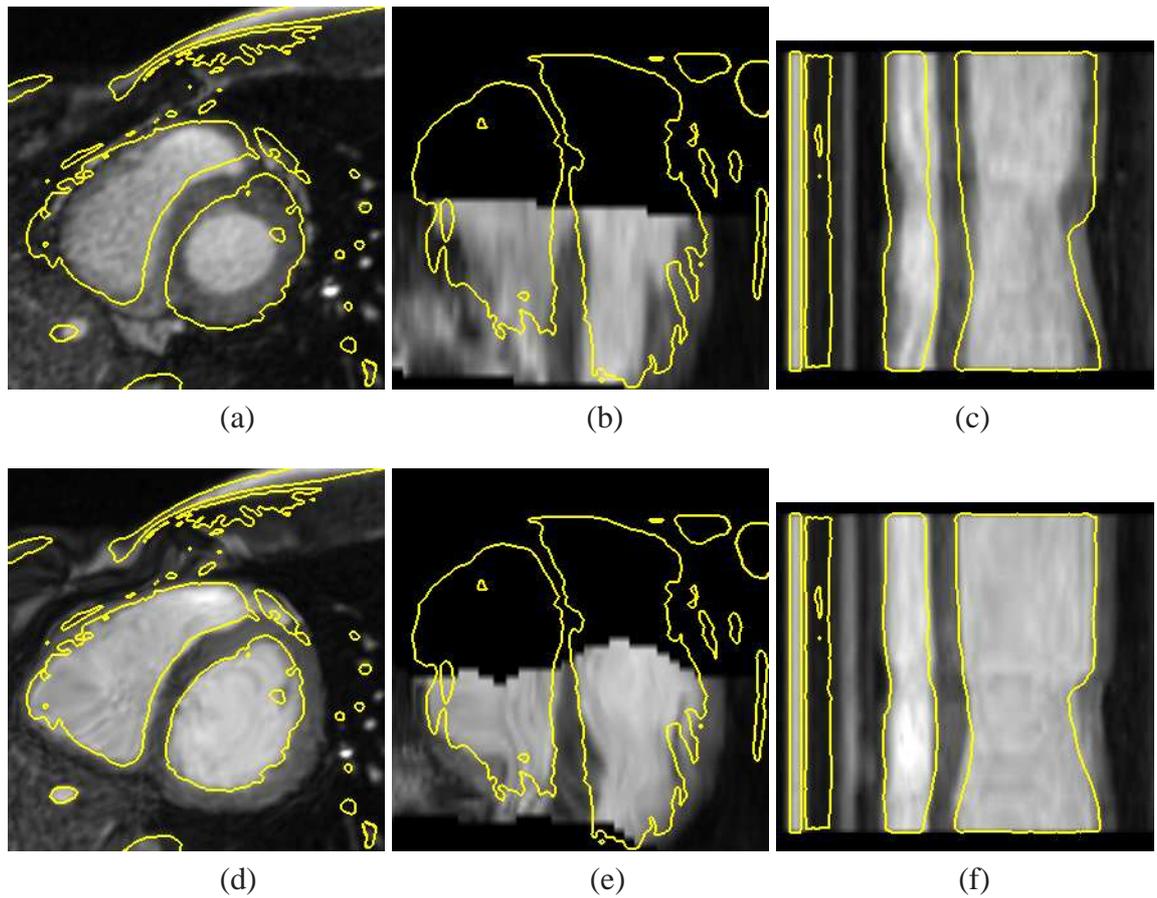
Figure 5.31: Results of the 4D cardiac MR registration algorithm. The short-axis (a), the long-axis (b) and the temporal (c) views after the affine alignment. The corresponding short-axis (d), long-axis (e) and temporal (f) views after using the cross-correlation based method to calculate the temporal part and optimising only the spatial part (experiment **B**).
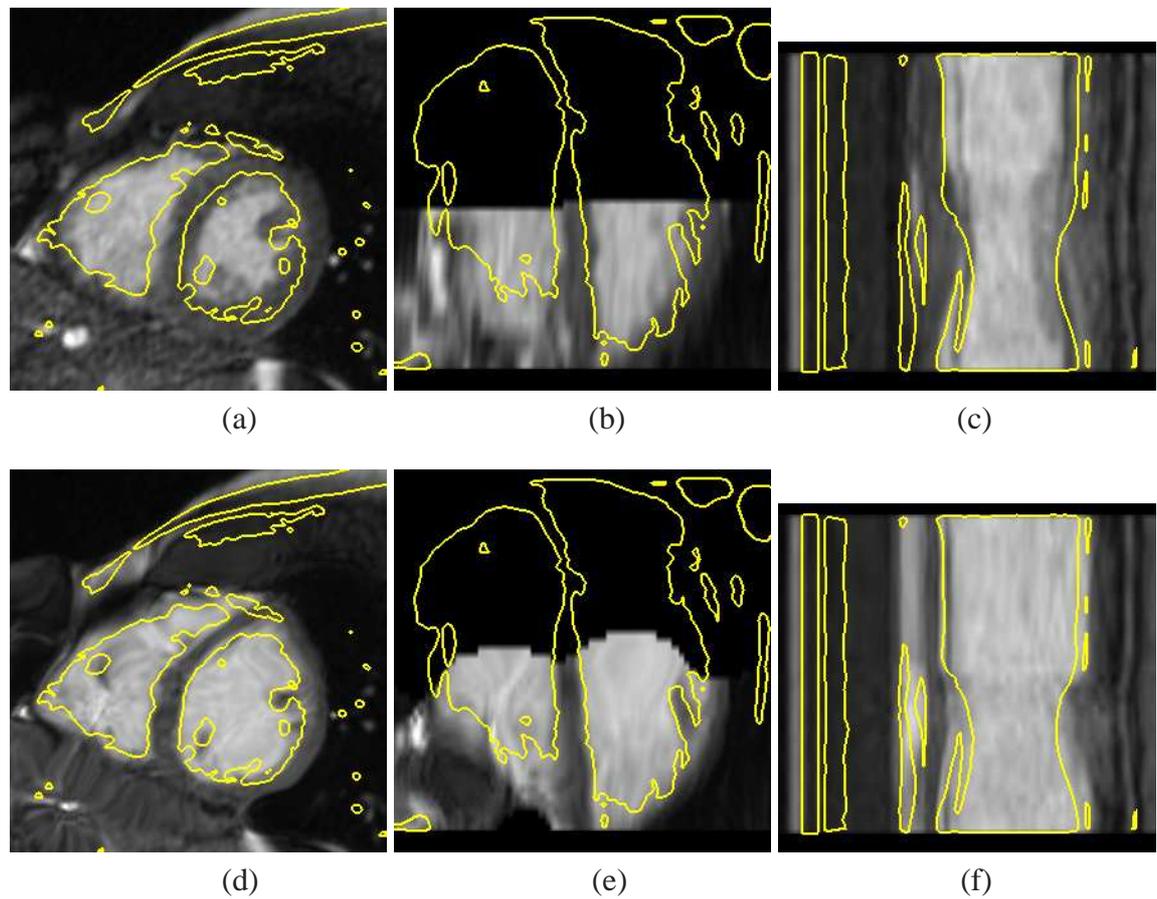
Figure 5.32: Results of the 4D cardiac MR registration algorithm. The short-axis (a), the long-axis (b) and the temporal (c) views after the affine alignment. The corresponding short-axis (d), long-axis (e) and temporal (f) views after using the cross-correlation based method to calculate the temporal part and optimising only the spatial part (experiment **B**).
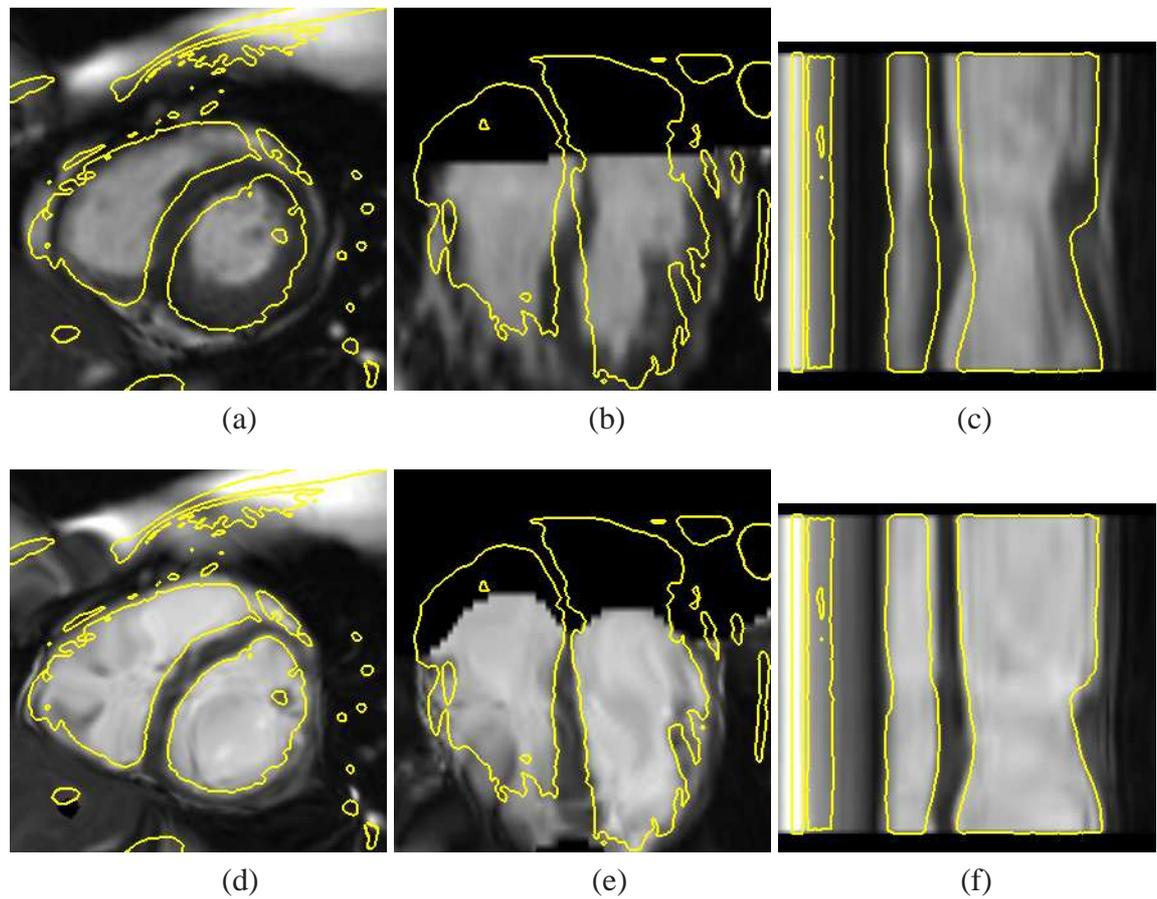
Figure 5.33: Results of the 4D cardiac MR registration algorithm. The short-axis (a), the long-axis (b) and the temporal (c) views after the affine alignment. The corresponding short-axis (d), long-axis (e) and temporal (f) views after using the cross-correlation based method to calculate the temporal part and optimising only the spatial part (experiment **B**).
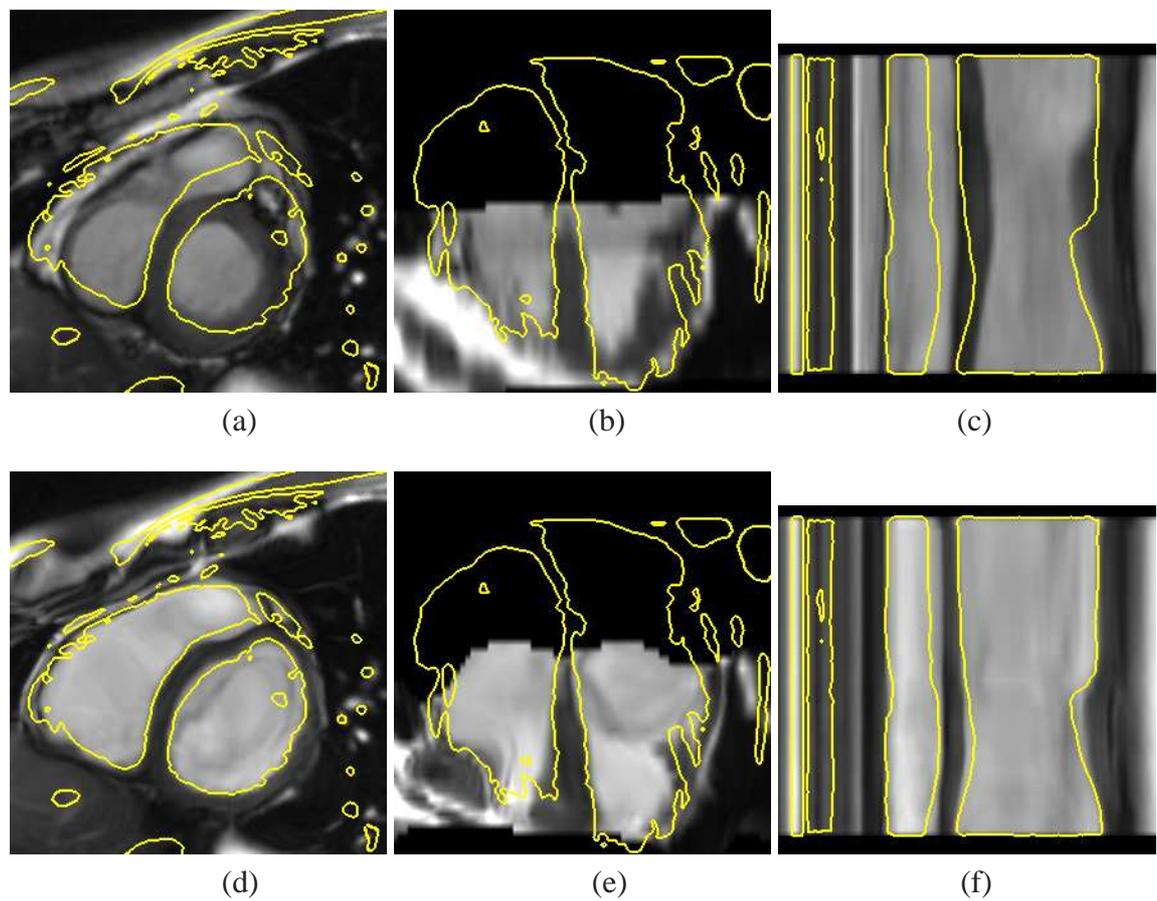
Figure 5.34: Results of the 4D cardiac MR registration algorithm. The short-axis (a), the long-axis (b) and the temporal (c) views after the affine alignment. The corresponding short-axis (d), long-axis (e) and temporal (f) views after using the cross-correlation based method to calculate the temporal part and optimising only the spatial part (experiment **B**).
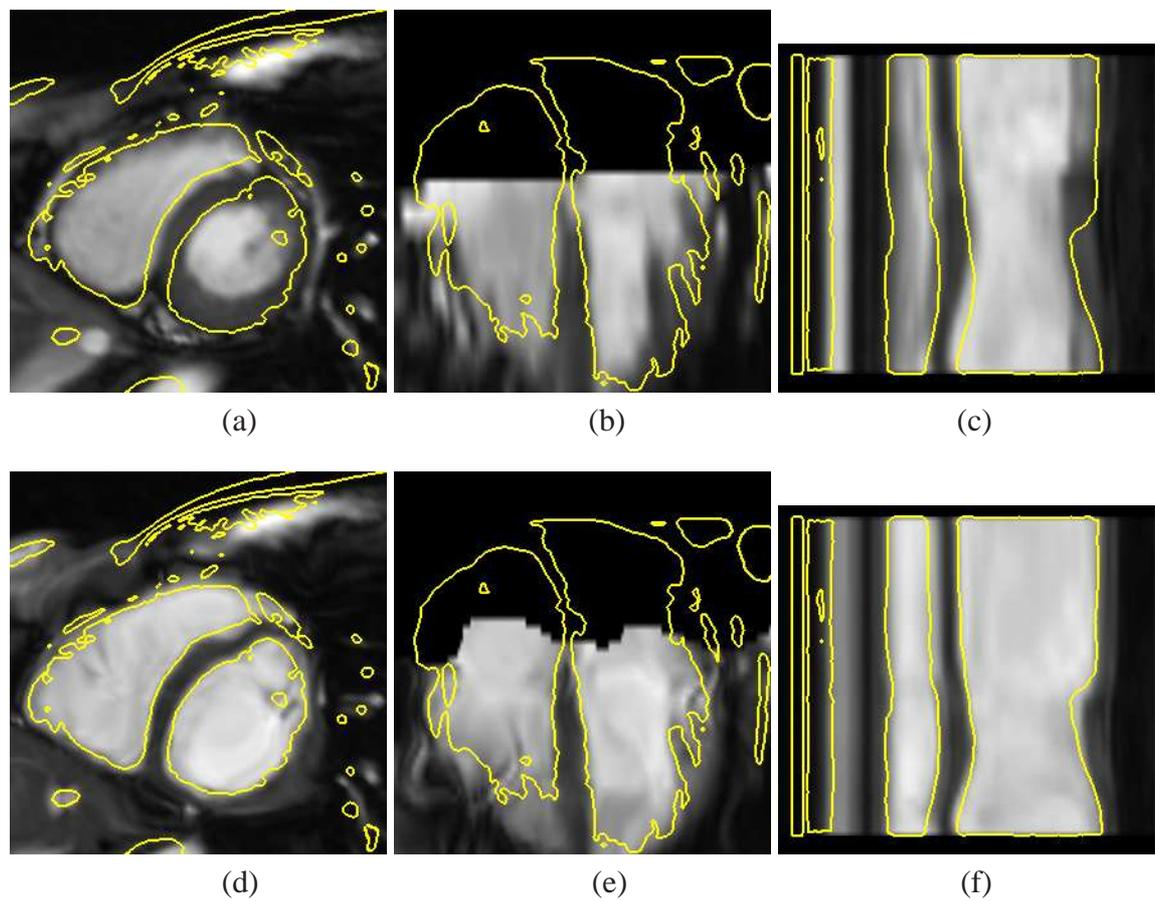
Figure 5.35: Results of the 4D cardiac MR registration algorithm. The short-axis (a), the long-axis (b) and the temporal (c) views after the affine alignment. The corresponding short-axis (d), long-axis (e) and temporal (f) views after using the cross-correlation based method to calculate the temporal part and optimising only the spatial part (experiment **B**).
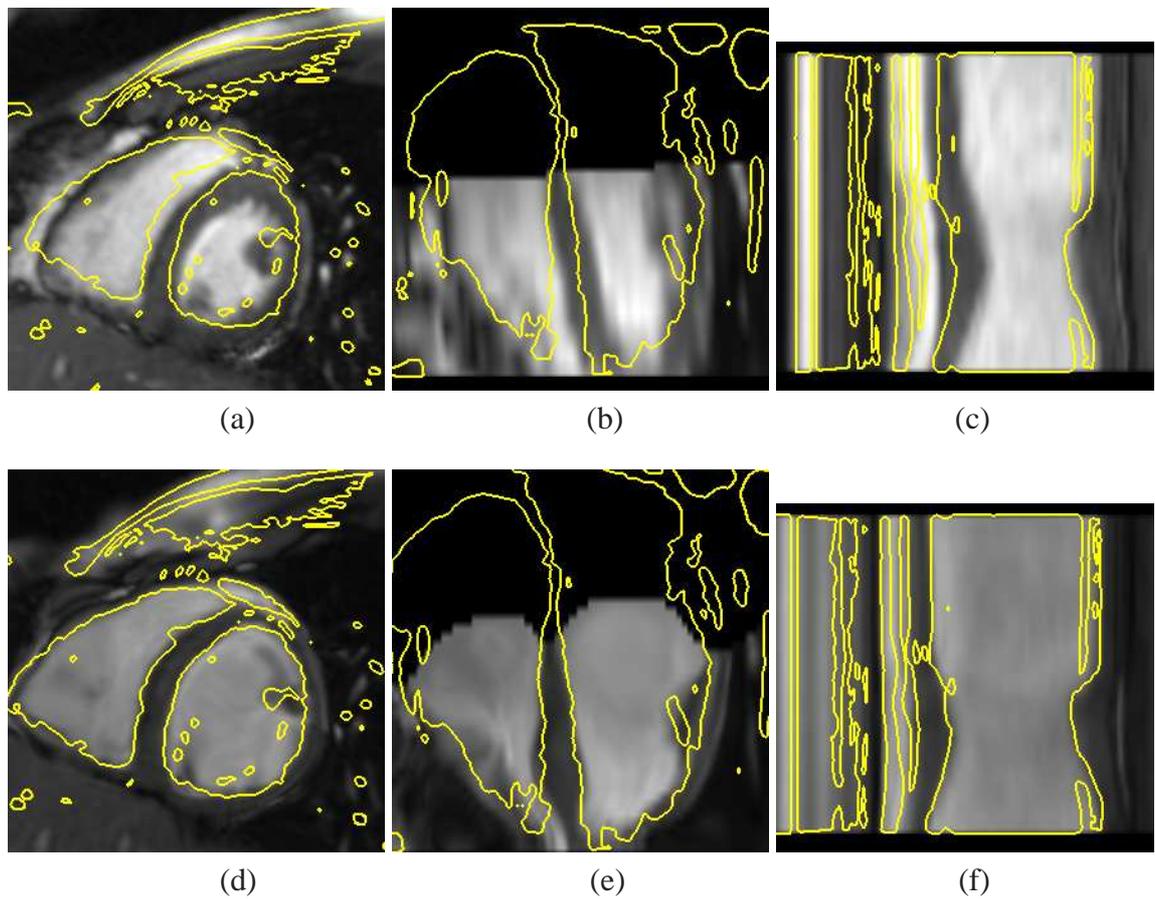
Figure 5.36: Results of the 4D cardiac MR registration algorithm. The short-axis (a), the long-axis (b) and the temporal (c) views after the affine alignment. The corresponding short-axis (d), long-axis (e) and temporal (f) views after using the cross-correlation based method to calculate the temporal part and optimising only the spatial part (experiment **B**).

5.6 Conclusions

In this chapter two spatio-temporal *deformable* registration methods for cardiac MR image sequences have been presented. The registration approaches correct spatial misalignment caused by different acquisition parameters and spatial misalignment caused by differences in the local shape of the cardiac anatomies. Furthermore, the approaches correct temporal misalignment caused by differences in the length of the cardiac cycles and temporal misalignment caused by different dynamic properties of the hearts. The first registration approach calculates the spatial and temporal components of the transformation simultaneously, while the second approach calculates first the temporal transformation and then the spatial one. In the latter, a novel approach based on normalised cross-correlation was used to calculate the temporal component of the transformation. The spatio-temporal registration algorithms have been evaluated by measuring the volume overlap and the surface distance of corresponding anatomical structures and by visual inspection. A number of experiments were performed to investigate the performance of both methods. The results indicate that the use of the spatio-temporal free-form deformation model results in a substantial improvement in the temporal and spatial alignment of the image sequences. Furthermore, performing combined optimisation of the temporal and spatial parts of the transformation results in better registration rather than when these are calculated separately. However, in this case the computational complexity is much higher. Moreover, the experiments indicated that calculating the temporal transformation using the cross-correlation based method provides a close approximation to the temporal transformation that resulted from the combined optimisation approach. Evidently, it only aligns a limited number of feature positions in the cardiac cycles, while optimising the transformation components simultaneously provides a better temporal alignment between these cardiac positions.

Chapter 6

Construction of an atlas of cardiac anatomy and function

This chapter presents two novel methods for the construction of 4D probabilistic and statistical atlases of the cardiac anatomy and function using cardiac MR imaging. The probabilistic atlas captures information regarding the cardiac anatomy and function in the form of tissue probability maps. The method for the construction of the probabilistic atlas enables the spatio-temporal modelling of tissue probability maps.

The statistical atlas describes the cardiac anatomy and how the cardiac anatomy changes during the cardiac cycle. The method for the construction of the statistical atlas divides the distribution of the cardiac shapes into two subspaces. One distribution subspace accounts for changes in cardiac shape caused by inter-subject variability. The second distribution subspace accounts for changes in cardiac shape caused by deformation during the cardiac cycle (i.e. intra-subject variability). Principal component analysis (PCA) has been performed in order to calculate the most significant modes of variation of each distribution subspace.

In order to build the atlases 30 cardiac image sequences from healthy volunteers have been used. The resulting statistical atlas has been used to differentiate between cardiac image sequences from patients with hypertrophic cardiomyopathy and normal subjects.

6.1   Introduction

A large number of approaches have been developed for the volumetric modelling of the heart. A comprehensive review of these approaches can be found by Frangi *et al.* [62]. Biomechanical models of the heart have been developed by combining surface information and motion information [183] and by using a deformation model inspired by continuum mechanics [135]. These models do not usually contain any information regarding the variation met in the population. In contrast to these biomechanical models a number of researchers have developed statistical models of the cardiac anatomy (e.g. Active Shape Models) [112] [83] and statistical models of the appearance of the heart (e.g. Active Appearance models) [128] [126]. For example, Frangi *et al.* have presented an approach for the construction of three-dimensional statistical shape models of the cardiac anatomy [64]. This approach eliminates the need for landmarking by using non-rigid registration to propagate the landmarks from an automated landmarked atlas to the rest of the images. The resulting model includes the left and right ventricles. The approach developed by Lötjönen *et al.* goes one step further: in this work, statistical shape models of the atria, ventricles and epicardium from short-axis and long-axis MR images are constructed and used for the segmentation of cardiac images [111]. In addition, a variety of methods which model shape variability have been explored including PCA [87], ICA [85], [86] and LPD. However, in both cases the statistical shape models only describe the 3D cardiac anatomy at a single time point and ignore the shape variation during the cardiac cycle. Although cardiac modelling of the anatomy is relatively well investigated, very few attempts have been made to build a computerised atlas which captures functional variability of the heart across a group of subjects. Rao *et al.* suggested a framework for building an atlas of the myocardial motion [167] by using tagged MR image sequences to calculate the cardiac motion fields. Then the calculated motion fields of different subjects are mapped into the same coordinate system using a vector field transformation technique which accounts for differences in the size, orientation and shape of the heart.

6.2   Contributions

This chapter presents a novel method for building a probabilistic atlas of the cardiac anatomy and function as well as a method for building a 4D statistical atlas of the cardiac anatomy. The probabilistic atlas contains information regarding the cardiac anatomy and function in the form of tissue probability maps. The 4D statistical atlas describes the cardiac anatomy and how the cardiac anatomy changes during the cardiac cycle. During the construction of the statistical atlas the need for manual landmarking of the cardiac images is eliminated by using a non-rigid registration algorithm to propagate a set of pseudo-landmarks from an automatically landmarked atlas to each frame of all image sequences.

The key contribution of this chapter is the extension of the statistical and probabilistic modelling of the heart in 4D. In particular, the contributions of this chapter are the following:

- The construction of a probabilistic atlas of the cardiac anatomy and function which enables spatio-temporal modelling of the cardiac anatomy.

- The construction of a 4D statistical model of the heart that subdivides the distribution space of the cardiac shape to two subspaces: one distribution space accounts for changes in cardiac shape due to deformations throughout the cardiac cycle and the other distribution space accounts for changes in the cardiac shape due to inter-subject variability.

- The use of a non-rigid registration method to propagate a set of pseudo-landmarks from an automatically landmarked atlas to all frames of the image sequences used during the construction of the statistical model. Contrary to the method presented by Frangi *et al.* [64], this approach propagates pseudo-landmarks not only to a single frame but to entire image sequences.

- The application of 4D statistical models to differentiate between cardiac image sequences from patients with hypertrophic cardiomyopathy and normal volunteers.

6.3   Building anatomical and functional atlases of the heart

Figure 6.1 summarises the steps in the construction of the atlases. The first steps in the construction of both the probabilistic and statistical atlases are identical:

- The image sequences are registered into a common spatio-temporal coordinate system. This enables direct comparison between the cardiac anatomy and function of a number of cardiac MR image sequences to be made.

- The image sequences are segmented into left and right ventricle as well as myocardium using an automated segmentation algorithm [109]. An example of cardiac MR images and their segmented tissue maps can be seen in figure 6.2.

- Shape-based interpolation [76] is used to resample the spatial domain of the segmented image sequences into isotropic voxels of size $1mm \times 1mm \times 1mm$.

- The segmented image sequences are transformed to the common spatio-temporal coordinate system.

- The transformed segmented image sequences are blurred with a Gaussian filter with $\sigma = 2mm$. The use of blurring during the construction of the atlas compensates for the low out-of-plane resolution of the images. A $\sigma = 2mm$ was chosen because it compensates for the low out-of-plane resolution without introducing substantial blurring which can destroy information regarding anatomical detail.

After these steps, the probabilistic atlas can be constructed by averaging the transformed segmented image sequences. Before calculating the average image sequence, each segmented image sequence is separated into three separate image sequences containing the left ventricle, myocardium and right ventricle. The result is a spatio-temporal map of probabilities of left ventricle, myocardium and right ventricle.

The construction of the statistical atlas requires the following additional steps:

- The surfaces are extracted from all the segmented frames of every image sequence using the *marching cubes* algorithm [107].

- A non-rigid registration method is used to automatically propagate a set of pseudo-landmarks from an automated landmarked atlas to the rest of the image sequences.

- Statistical analysis in the form of PCA is performed on the position of these landmarks.

The following sections explain these steps in more detail.

6.3.1 Spatio-temporal registration of cardiac MR image sequences

As mentioned in the previous chapters, since the heart is undergoing spatially and temporally a varying degree of motion during the cardiac cycle, 4D cardiac image registration algorithms are required when comparing two cardiac MR image sequences. Comparison of the corresponding frames (by using only spatial alignment) is not enough since these frames may correspond to different positions in the cardiac cycles (figure 4.2).

In order to register the image sequences, a spatio-temporal registration method similar to the ones presented in chapters 4 and 5 [151], [147], [153], [152] is used. In these methods the 4D transformation mapping $\mathbf{T}$ has been decoupled into a spatial $\mathbf{T}_{spatial}$ and a temporal $\mathbf{T}_{temporal}$ component ensuring causality and preventing different regions in a 3D image $S_t(x, y, z)$ from being warped differently in the temporal direction by $\mathbf{T}_{temporal}$.

The aim of the spatial part of the transformation is to relate each spatial point of a particular frame of one image sequence to a point in another particular frame of the reference image sequence. However, since the models need to explain the variability in the local shape of the hearts, the spatial transformation part needs to preserve the local shape of the hearts. It therefore only contains a global part:

$$\mathbf{T}_{spatial}(x, y, z) = \mathbf{T}_{spatial}^{global}(x, y, z) \tag{6.1}$$
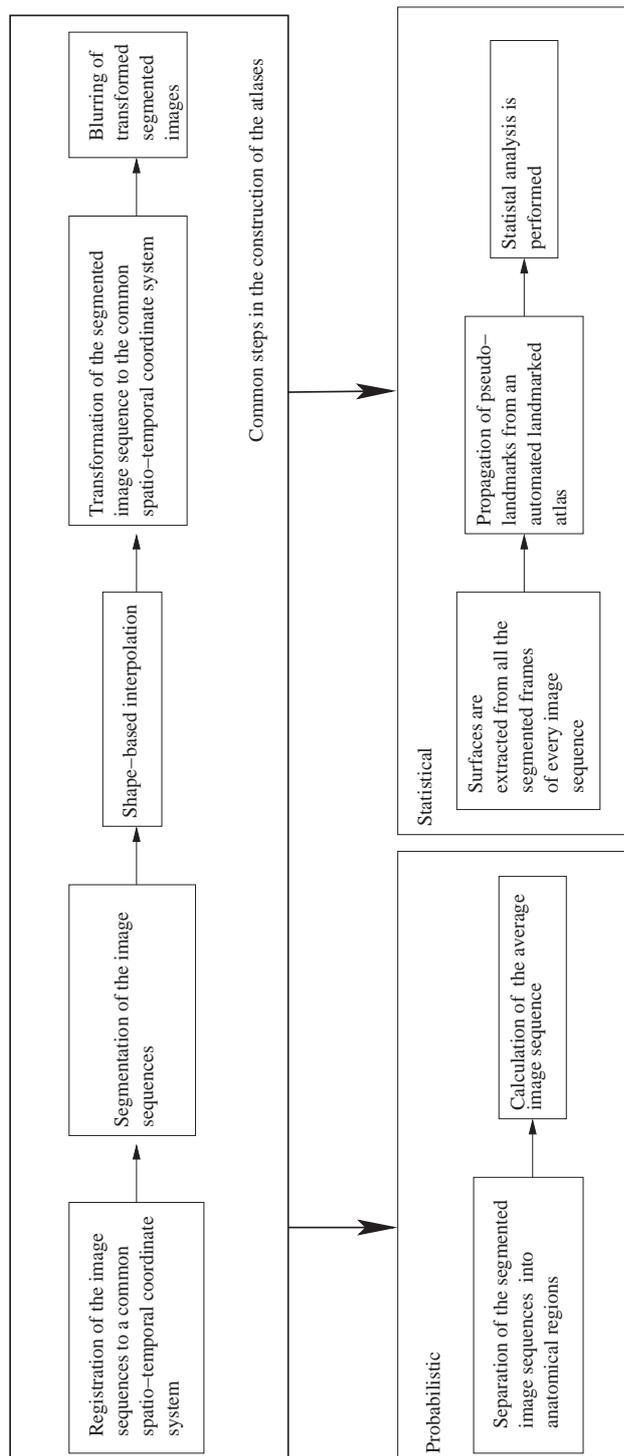
Figure 6.1: The atlas construction approach.

The global spatial transformation is an affine transformation with 9 *degrees of freedom* addressing differences in the size, orientation and alignment of the hearts (equation 4.4)

The aim of the temporal part is to relate corresponding positions in the cardiac cycle of the two hearts. As previously mentioned, in order to relate corresponding positions in the cardiac cycle of two hearts, a *deformable* temporal transformation is required. Thus, the temporal transformation consists of a global and a local part:

$$\mathbf{T}_{temporal}(t) = \mathbf{T}_{temporal}^{global}(t') + \mathbf{T}_{temporal}^{local}(t')$$

$\mathbf{T}_{temporal}^{global}$ is an affine transformation which corrects for differences in the length of the cardiac cycles and differences in the acquisition parameters (equation 5.6). $\mathbf{T}_{temporal}^{local}$ is modelled by a free-form deformation using a 1D B-Spline (equation 5.7) and corrects for temporal misalignment caused by different cardiac dynamic properties (differences in the length of contraction and relaxation phases, different motion patterns, etc). For more information on the temporal free-form deformation see section 5.3.2.

The optimal transformation is found by maximising a voxel based similarity measure, the normalised mutual information (NMI) [187] calculated directly from the joint intensity histogram of the two sequences over the spatio-temporal domain of overlap. In the first part of the optimisation procedure, NMI is optimised as a function of $\mathbf{T}_{spatial}^{global}$ and $\mathbf{T}_{temporal}^{global}$ using an iterative uphill descent algorithm. In the second part, NMI is optimised as a a function of $\mathbf{T}_{temporal}^{local}$ by also using an iterative uphill descent method.

6.3.2 Segmentation of cardiac MR image sequences

The method developed by Lorenzo-Valdés *et al.* [109] has been used to segment the image sequences. In this method the first frame of each image sequence is segmented manually and then the segmentation is propagated to the subsequent frames using a non-rigid registration algorithm. The image sequences are segmented into three anatomical structures: the left ventricle, the myocardium and the right ventricle. An example of this is shown in figure 6.2.

<div align="center">(a)   (b)   (c)   (d)</div>
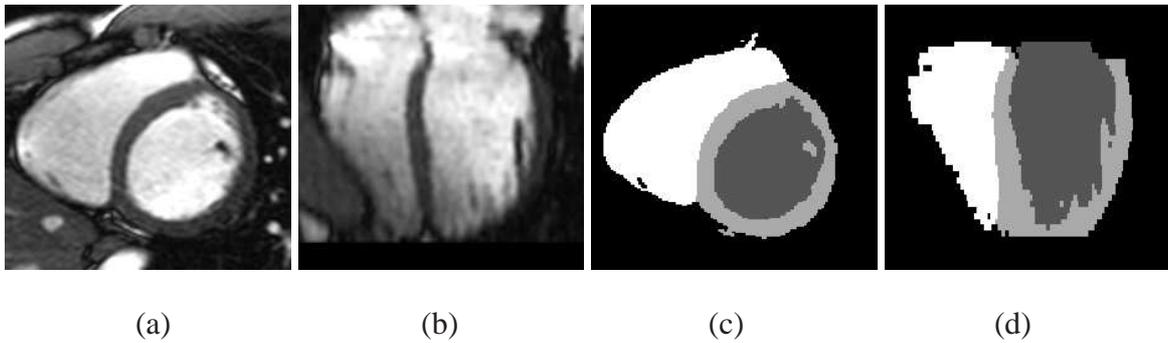
Figure 6.2: The greylevel short axis (a) and the long axis (b) views of a cardiac MR image and the corresponding segmented tissue maps (c and d).

6.3.3 Building the probabilistic atlas of the heart

After aligning all the image sequences to the same spatio-temporal coordinate system, the obtained transformations are used to map the segmented image sequences to the same spatio-temporal coordinate system. The probabilistic atlas $\mathbf{A}$ will have the same voxel dimensions and number of frames as the image sequence used as a reference during spatio-temporal registration. Each frame $\{\mathbf{A}_i; i = 0...n_f\}$ (where n_f is the number of frames in the sequences) of $\mathbf{A}$ will be formed by averaging the corresponding frames of the transformed segmented image sequences $\{\mathbf{O}_{ki}; k = 0....n_p\}$ (where n_p is the number of subjects used to construct the atlas):

$$\mathbf{A}_i(x, y, z) = \frac{1}{n_p} \sum_{k=0}^{n_p} \mathbf{O}_{ki}(x, y, z) \tag{6.2}$$

Before producing the average image sequence, each segmented image is blurred with a Gaussian kernel with $\sigma = 2$. As mentioned above, the use of blurring during the construction of the atlas is needed due to the low out-of-plane resolution which results in significant partial volume effects in the segmentation. Blurring the images with a Gaussian kernel addresses this problem by modelling this uncertainty in the tissue classification.
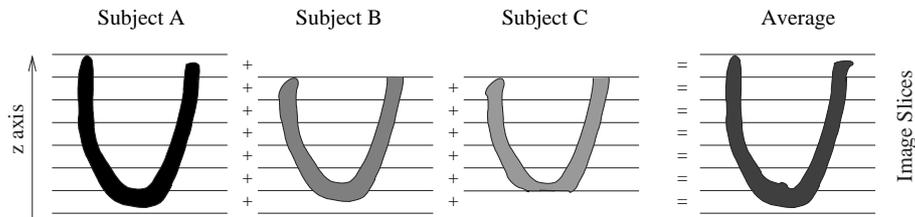
Figure 6.3: The mean image is calculated by averaging each slice separately. Calculating the average image for each slice separately reduces the effects of different coverage of the heart in the image.

A problem during the construction of the probabilistic atlas is that the image sequences may not have exactly the same coverage of the heart. For example an image may not include part of the apex or part of the base of the heart. The fact that information regarding part of the anatomy might be missing needs to be taken into account when producing the tissue probability maps. Hence, in order to reduce the effects of a different coverage of the heart in each image, the tissue probability maps have been calculated by averaging each image slice separately (figure 6.3).

6.3.4 Building a statistical atlas of the heart

As in the construction of the probabilistic atlas, during the construction of the statistical atlas the images are also blurred with a Gaussian kernel to reduce the effects of low out-of-plane resolution. In order to perform statistical analysis a set of anatomical landmarks are required to be identified in each image. Usually these landmarks are manually identified in each image. However, this is a time consuming task prone to errors. The approach presented in this thesis eliminates the need from manual landmarking by using a method similar to the one used by Frangi *et al.* [64]. In this method a set of pseudo-landmarks are propagated from an automatically landmarked atlas to all frames of each image sequence. Lötjönen *et al.* [111] have also used a similar approach for landmark propagation.

6.3.4.1   Landmark extraction and propagation

After blurring the image sequences with a Gaussian kernel with $\sigma = 2$, the *marching cubes* [107] algorithm is used to generate a dense triangulation of the boundary surface of each anatomical structure (left ventricle, myocardium and right ventricle) of all the frames of each image sequence. The vertices of the triangulation serve as *pseudo-landmarks*. In order to perform any statistical analysis, correspondence between the pseudo-landmarks of correspondent frames needs to be established. This is achieved by using a 3D surface based registration method based on B-Splines [139]. The end-diastolic frame of the image sequence used as a reference during the construction of the atlases is also used as the reference surface in these registrations. The extracted surfaces from each frame of all image sequences are registered to the reference surface using a non-rigid registration method. The main difference between this approach the method proposed by Frangi *et al.* [64] is that the former propagates landmarks over the entire image sequences rather the first frames only. Furthermore, this approach registers corresponding extracted surfaces while the approach proposed by Frangi *et al.* [64] registers the corresponding segmented images. After registering all surfaces, the obtained transformations are used to propagate the *pseudo-landmarks* of the reference surface to each frame (figure 6.4).

6.3.4.2   Modelling shape variability

Once landmark correspondence has been established between all shapes, modelling techniques, such as principal components analysis can be used to analyse shape variability. Let $\{\mathbf{q_{ik}}; i = 0...n_p; k = 0...n_f\}$ denote n shapes (n_p subjects with n_f frames each). Each shape consists of m 3D landmarks, $\mathbf{p_j} = (p_{1j}, p_{2j}, p_{3j}; j = 1....m)$. Each vector $\mathbf{q_{ik}}$ consists of the landmarks $p_{11}, p_{21}, p_{31}, p_{12}, p_{22}, p_{32}, ..., p_{1m}, p_{2m}, p_{3m}$. The aim of the statistical analysis is to approximate the distribution of the landmarks with a linear model of the form:

$$\mathbf{q} = \bar{\mathbf{q}} + \Phi \mathbf{b} \qquad (6.3)$$

where $\bar{\mathbf{q}}$ is the average landmark vector, $\mathbf{b}$ is the shape parameter vector of the model, and Φ is a matrix of eigenvectors. The matrix Φ is obtained by performing principal component analysis (PCA) [87] to the covariance matrix $\mathbf{C}$. During the principal component analysis, the principal components of $\mathbf{C}$ are calculated as its eigenvectors ϕ_i and the corresponding eigenvalues λ_i are also calculated (such that $\lambda_i < \lambda_{i+1}$).

The aim of this statistical analysis is to identify what changes in the cardiac anatomy occur due to the cardiac cycle and what changes occur due to inter-subject variation. Therefore, the aim of this statistical analysis is to approximate the distribution of the landmarks with two linear models similar to equation 6.3:

$$\mathbf{q}_{within} = \bar{\mathbf{q}} + \Phi_{within} \mathbf{b}_{within} \tag{6.4}$$

$$\mathbf{q}_{between} = \bar{\mathbf{q}} + \Phi_{between} \mathbf{b}_{between} \tag{6.5}$$

Principal component analysis (PCA) will be used to find the estimate of two subspaces of the overall distribution. In order to achieve this, two separate principal component analyses are performed. The covariance matrix for the total shape distribution is given by:

$$\mathbf{C}_{total} = \frac{1}{n_f n_p} \sum_{i=1}^{n_p} \sum_{k=1}^{n_f} (\mathbf{q_{ik}} - \bar{\mathbf{q}})(\mathbf{q_{ik}} - \bar{\mathbf{q}})^T \tag{6.6}$$

where n_f is the number of frames of each image sequence, n_p is the number of image sequences and $\bar{\mathbf{q}}$ is the mean shape.

The covariance matrix of the shape differences occurring across the cardiac cycle (intra-subject) is given by:

$$\mathbf{C}_{within} = \frac{1}{n_f n_p} \sum_{i=1}^{n_p} \sum_{k=1}^{n_f} (\mathbf{q_{ik}} - \bar{\mathbf{q_i}})(\mathbf{q_{ik}} - \bar{\mathbf{q_i}})^T \tag{6.7}$$

where $\bar{\mathbf{q_i}}$ is the mean for the subject i (the image sequences contain the same number of frames since they are registered in the temporal domain) and $\mathbf{q_{ik}}$ is the shape of frame k of subject i.

The covariance matrix which describes the shape differences occurring across the population (inter-subject) is given by:

$$\mathbf{C}_{between} = \frac{1}{n_p} \sum_{i=1}^{n_p} (\bar{\mathbf{q}}_{\mathbf{i}} - \bar{\mathbf{q}})(\bar{\mathbf{q}}_{\mathbf{i}} - \bar{\mathbf{q}})^T \tag{6.8}$$

where, as in eq. 6.7, n_p is the number of image sequences and $\bar{\mathbf{q}}$ is the total mean.

Figure 6.5 explains which changes in the cardiac shape are explained by the covariance matrix $\mathbf{C}_{between}$ and which by the covariance matrix $\mathbf{C}_{within}$. In order to find the principal components of each subspace the eigenvalues and eigenvectors of each covariance matrix (eq. 6.7 and 6.8) are calculated. The eigenvectors of $\mathbf{C}_{between}$ are used to form $\Phi_{between}$, while the eigenvectors of $\mathbf{C}_{within}$ are used to form Φ_{within}. A similar decomposition of the total distribution space to subspaces has been used by Costen *et al.* for the automatic extraction of the face identity-subspace [41]. New shape examples can be generated by varying the parameters $\mathbf{b}_{within}$ and $\mathbf{b}_{between}$ of equations 6.4 and 6.5. Assuming that the distribution of the data follows a multidimensional Gaussian distribution, the variance of the ith parameter of $\mathbf{b}$ across the training set is given by λ_i. If limits in the variation of b_i are applied such that $b_i \leq \pm 3\sqrt{\lambda_i}$, then it is ensured that the generated shape is similar to those contained in the training class.
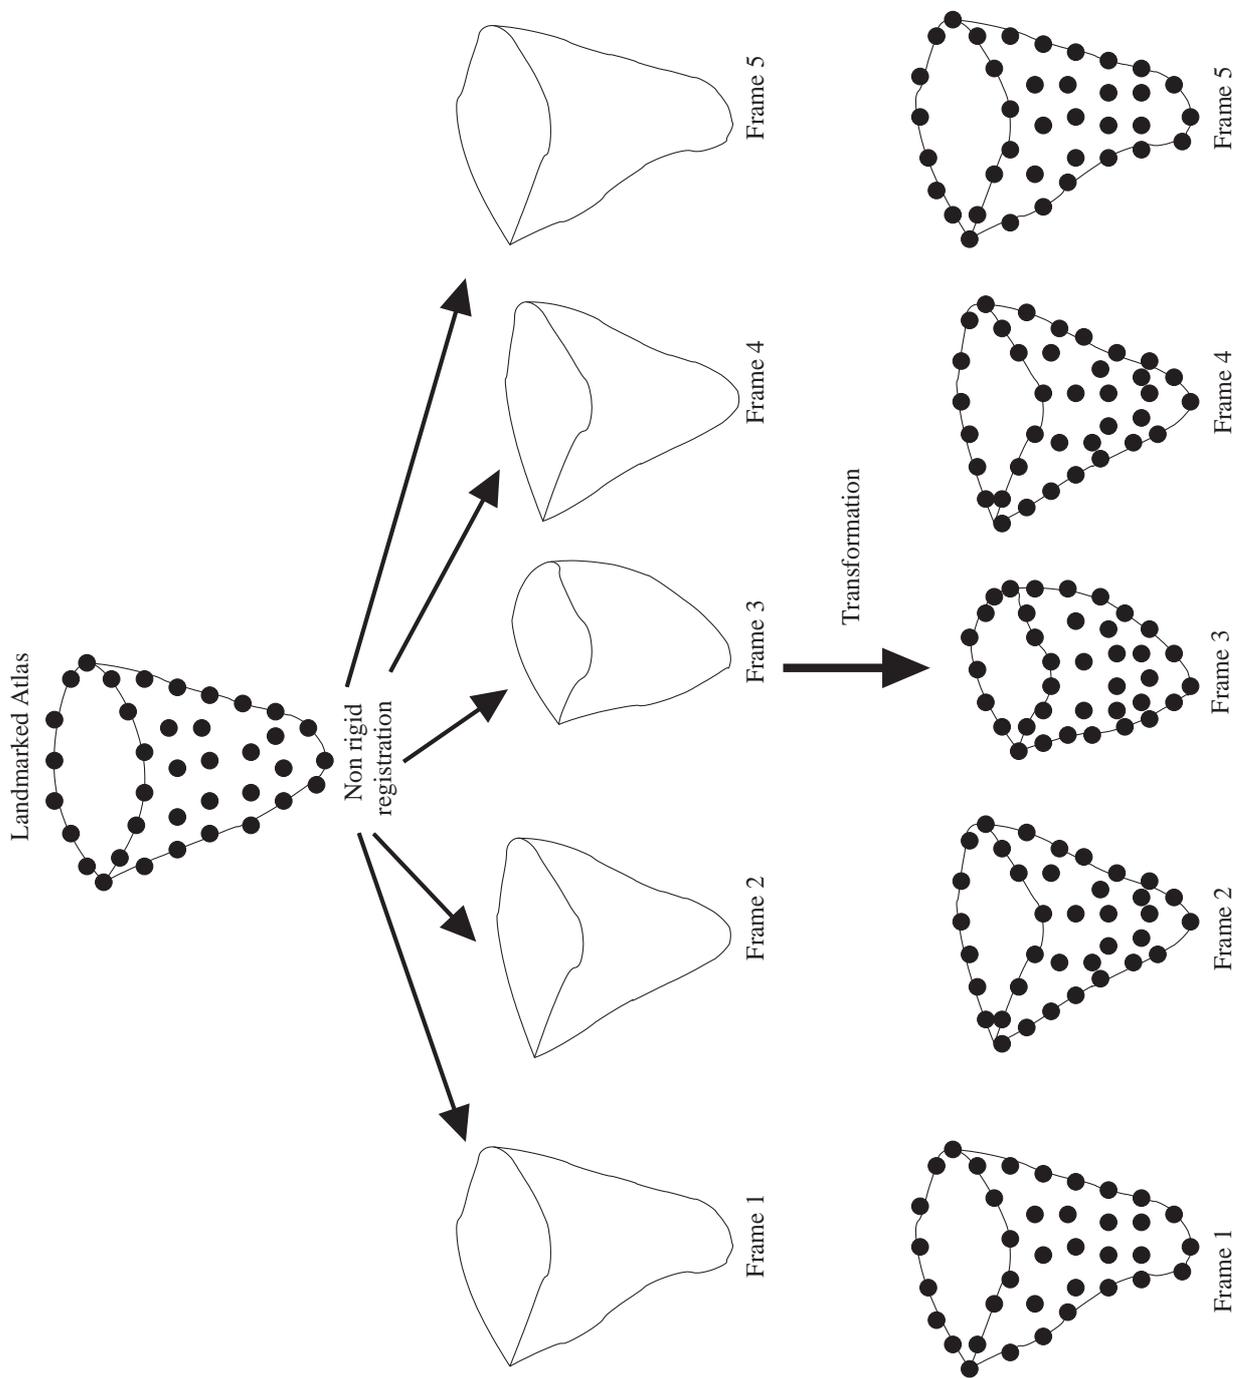
Figure 6.4: An example of propagating the pseudo-landmarks from the automatically landmarked atlas to a subject.
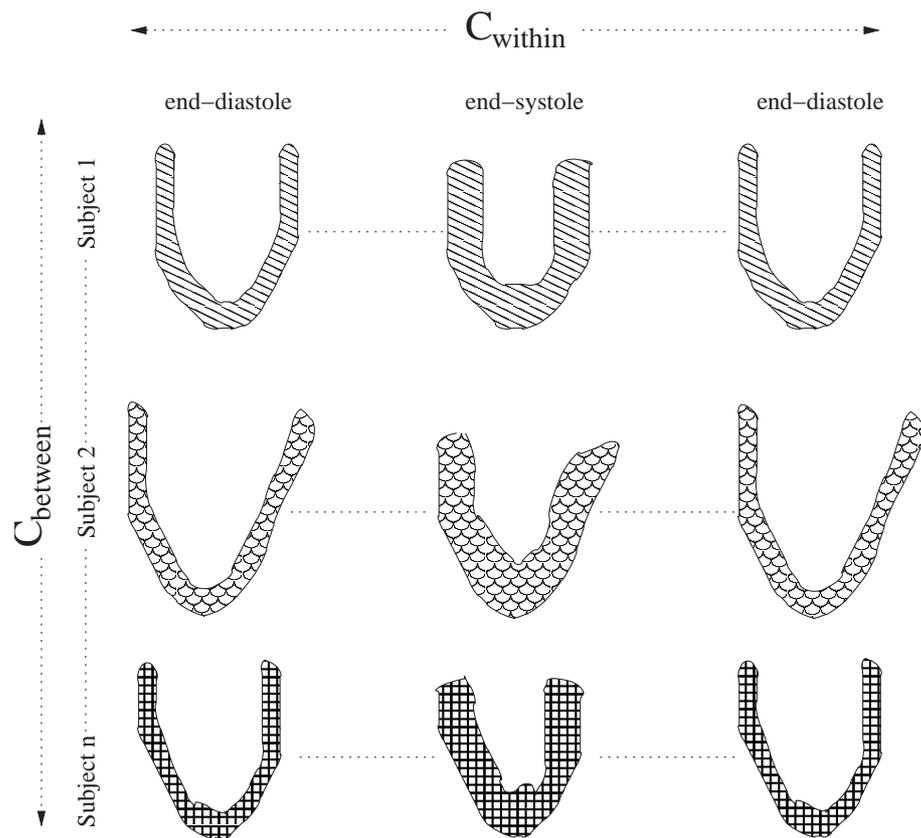
Figure 6.5: The changes in the cardiac shape described by the covariance matrices C_{within} and $C_{between}$.

6.4   Results

6.4.1   Materials

In order to produce the probabilistic and statistical 4D atlases, 30 untagged MR image sequences from healthy volunteers have been acquired. The image sequences have been acquired using a Siemens Sonata 1.5T scanner using TrueFisp pulse sequence in a form of a series short-axis images. Seven of the image sequences were acquired using retrospective gating while the rest were acquired using prospective gating. One of the image sequences was selected to be the reference subject for the spatio-temporal registration. Care was taken by visual inspection to ensure that the reference subject was a normal representative of the population. The image sequence of the reference subject had 18 frames with in-plane resolution of 192×256 and pixel size of $1.48mm \times 1.48mm$. Furthermore, each frame of the reference subject contained 10 slices covering the heart from the apex to the base. The thickness of each slice was $10mm$ (which is the typical pixel size used in these acquisitions). The length of the cardiac cycle of the subject was $792.5msec$. The global temporal transformation was calculated in order to compensate for the differences in the length of the cardiac cycles of the subjects (by matching the first and the last time frames of the image sequences). This is a fair assumption since all image sequences contained almost entire cardiac cycles. Furthermore, an initial estimate of the global spatial transformation was also provided due to the large variety in the position and orientation of the hearts. Most of the images covered the entire left ventricle from base to apex, while a limited number of images did not include the apex of the heart. One image sequence was excluded from the data set due to poor registration with the reference image sequence.

6.4.2   Probabilistic atlas

During the construction of the probabilistic atlas, all the image sequences were included independent of whether the images covered the entire heart or not. The reason for including images with different coverage of the heart is that, at the current stage, it is considered

very important to include as many subjects as possible in the atlas. In the future, where the population would be substantially increased (for example more than 100 subjects), it would be possible to ensure that all the images have exactly the same coverage. As explained in section 6.3.3, in order to minimise the effects of different image coverage the tissue probability maps are calculated by averaging each slice separately.

Figures 6.6 - 6.8 provide examples of the probabilistic atlases of the left ventricle, the myocardium and the right ventricle. In the probabilistic atlases, the smaller the intensity values, the smaller is the probability of a particular voxel belonging to a particular structure (the intensity of 255 corresponds to a probability of 1, while the intensity 0 corresponds to a probability of 0). Figure 6.7, shows that the papillary muscles are more blurred that the rest of the myocardium. Furthermore, the probability of a voxel to belong to a certain anatomical structure is smaller towards the edges of the anatomical structures (figures 6.6 - 6.8). This indicates that there is a large variation in the position and the size of the papillary muscles and also in the local shape of the hearts. The registration approach, described in section 6.3.1, uses only an affine spatial transformation model. This transformation model addresses spatial difference caused by translation, orientation and scaling. It cannot correct differences in the shape of the structures which results in a high degree of blurring on the edges of the structures.

Volume renderings of the tissue probability maps have also been produced. The volume renderings were produced using the *Visualisation Toolkit* package (www.vtk.org). Figures 6.9 - 6.11 show the volume rendering of the atlases of the left ventricle, the myocardium and the right ventricle from two different views. In the volume renderings, the smaller the probability of a voxel to belong to a structure the more transparent the voxel is rendered. Similar to figures 6.6 - 6.8, the voxels near the sides of the anatomical structures in figures 6.9 - 6.11 are rendered more transparent than the voxels in the centre due to their smaller probability values.
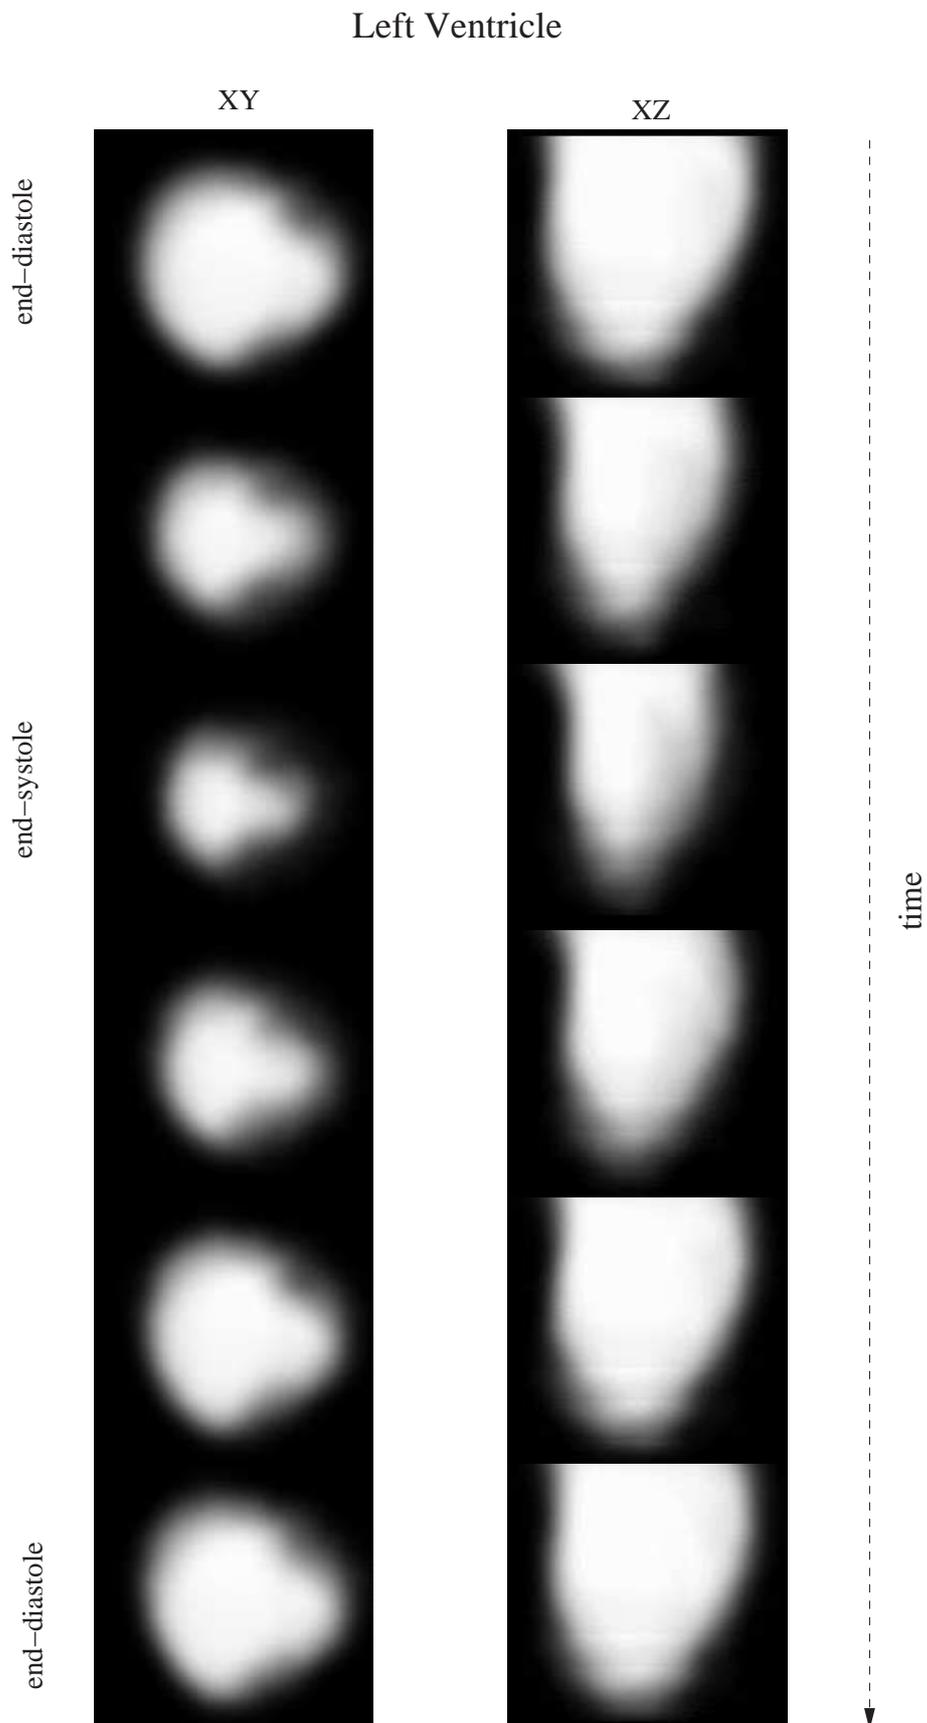
Left Ventricle

Figure 6.6: The probabilistic atlas of the left ventricle.

Myocardium

XY   XZ
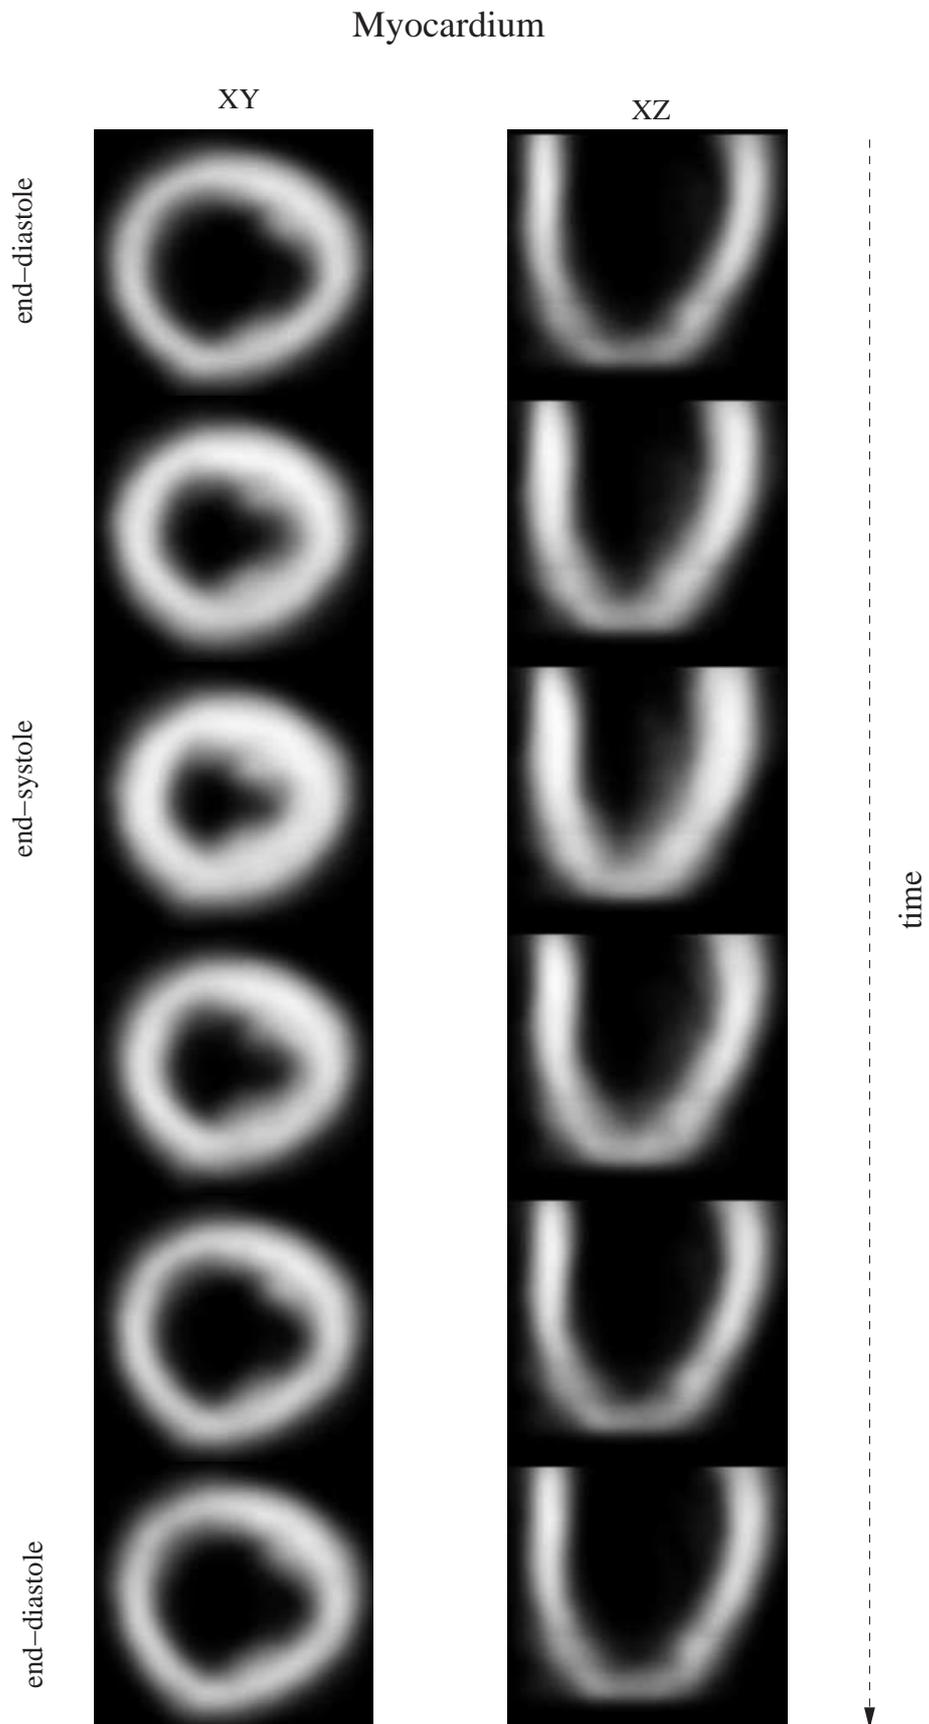
Figure 6.7: The probabilistic atlas of the myocardium.
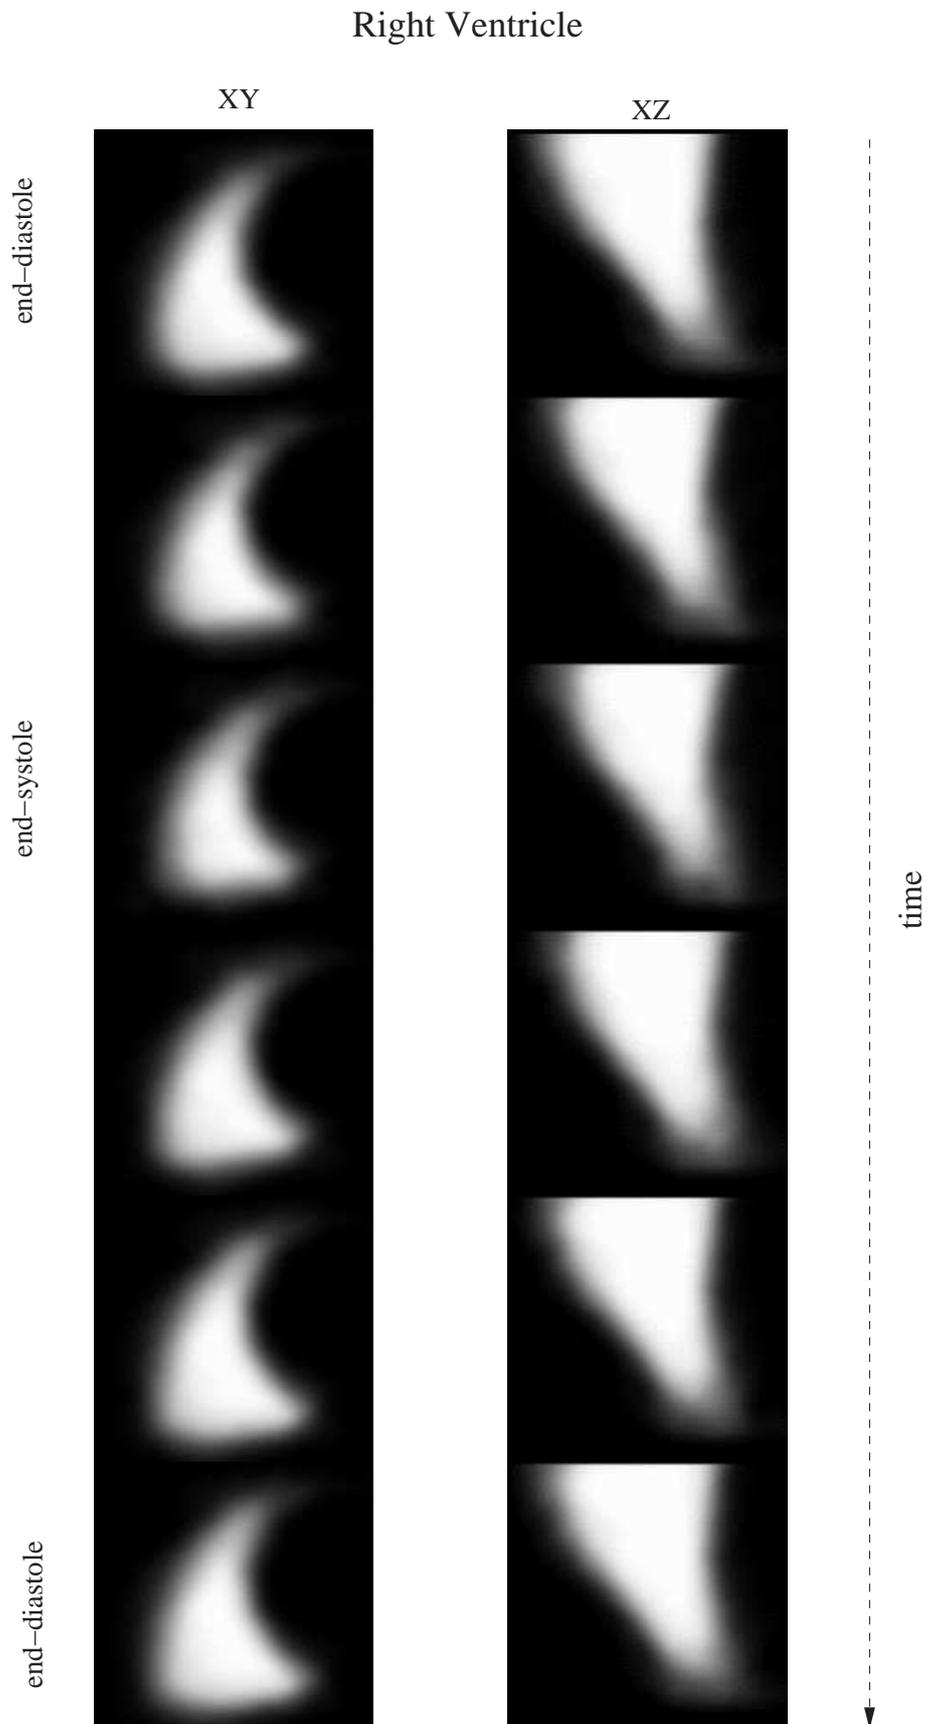
Right Ventricle

Figure 6.8: The probabilistic atlas of the right ventricle.
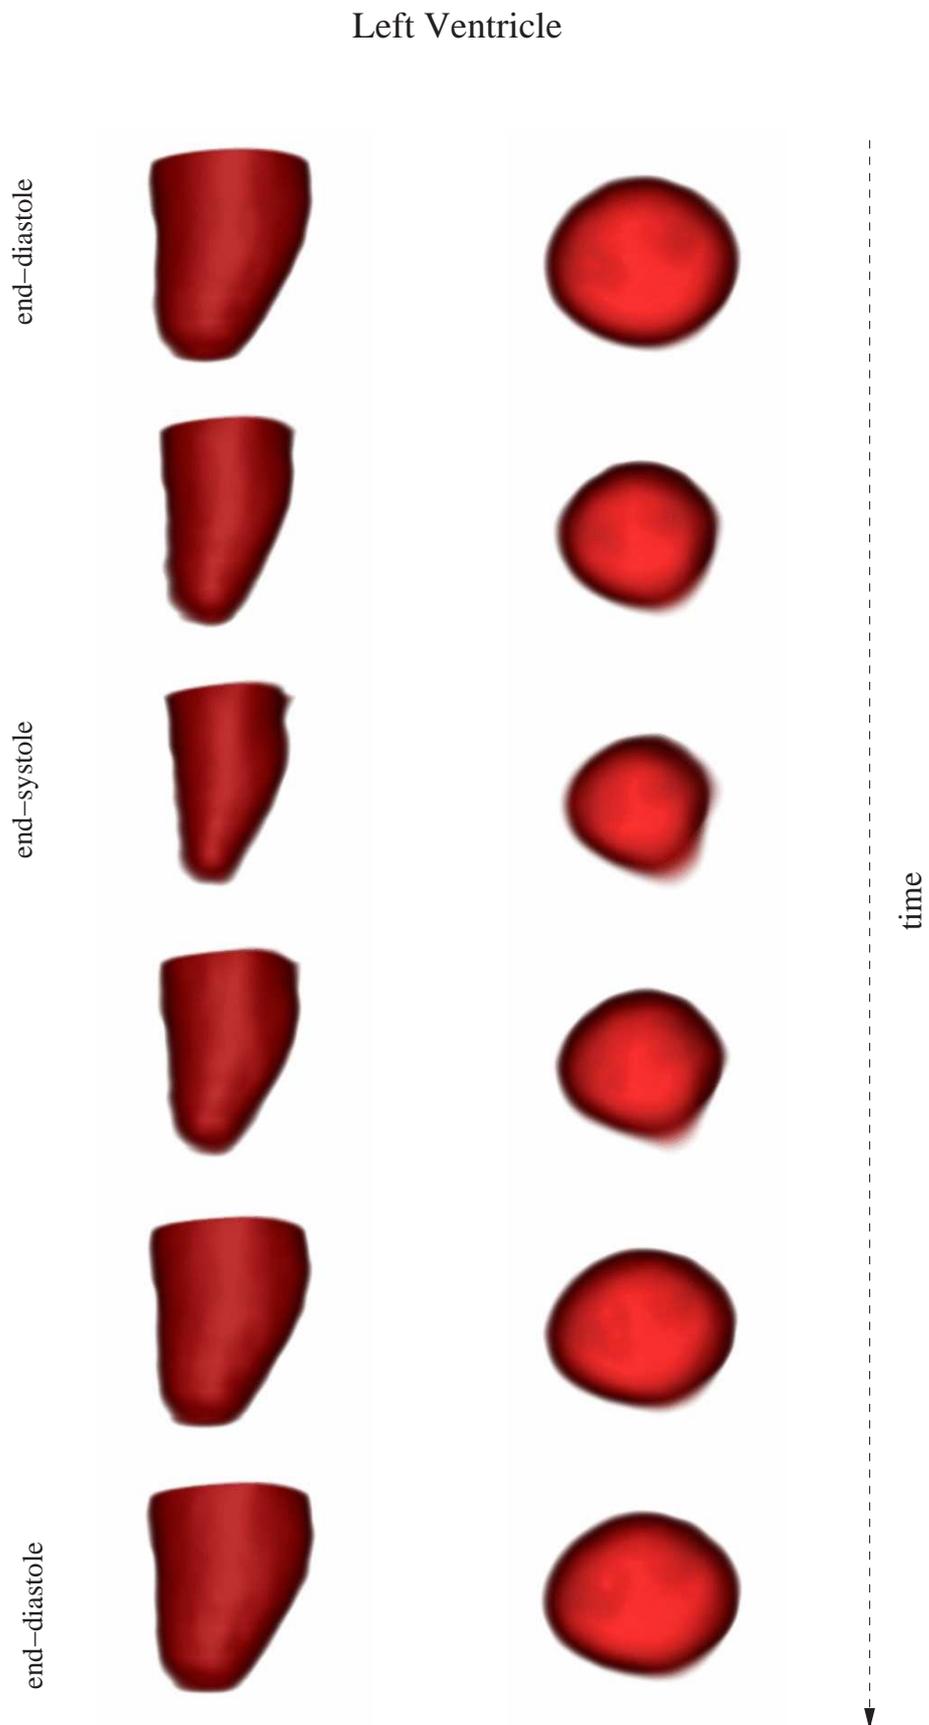
Left Ventricle

Figure 6.9: Examples of the volume renderings of the left ventricle viewed from two different positions.
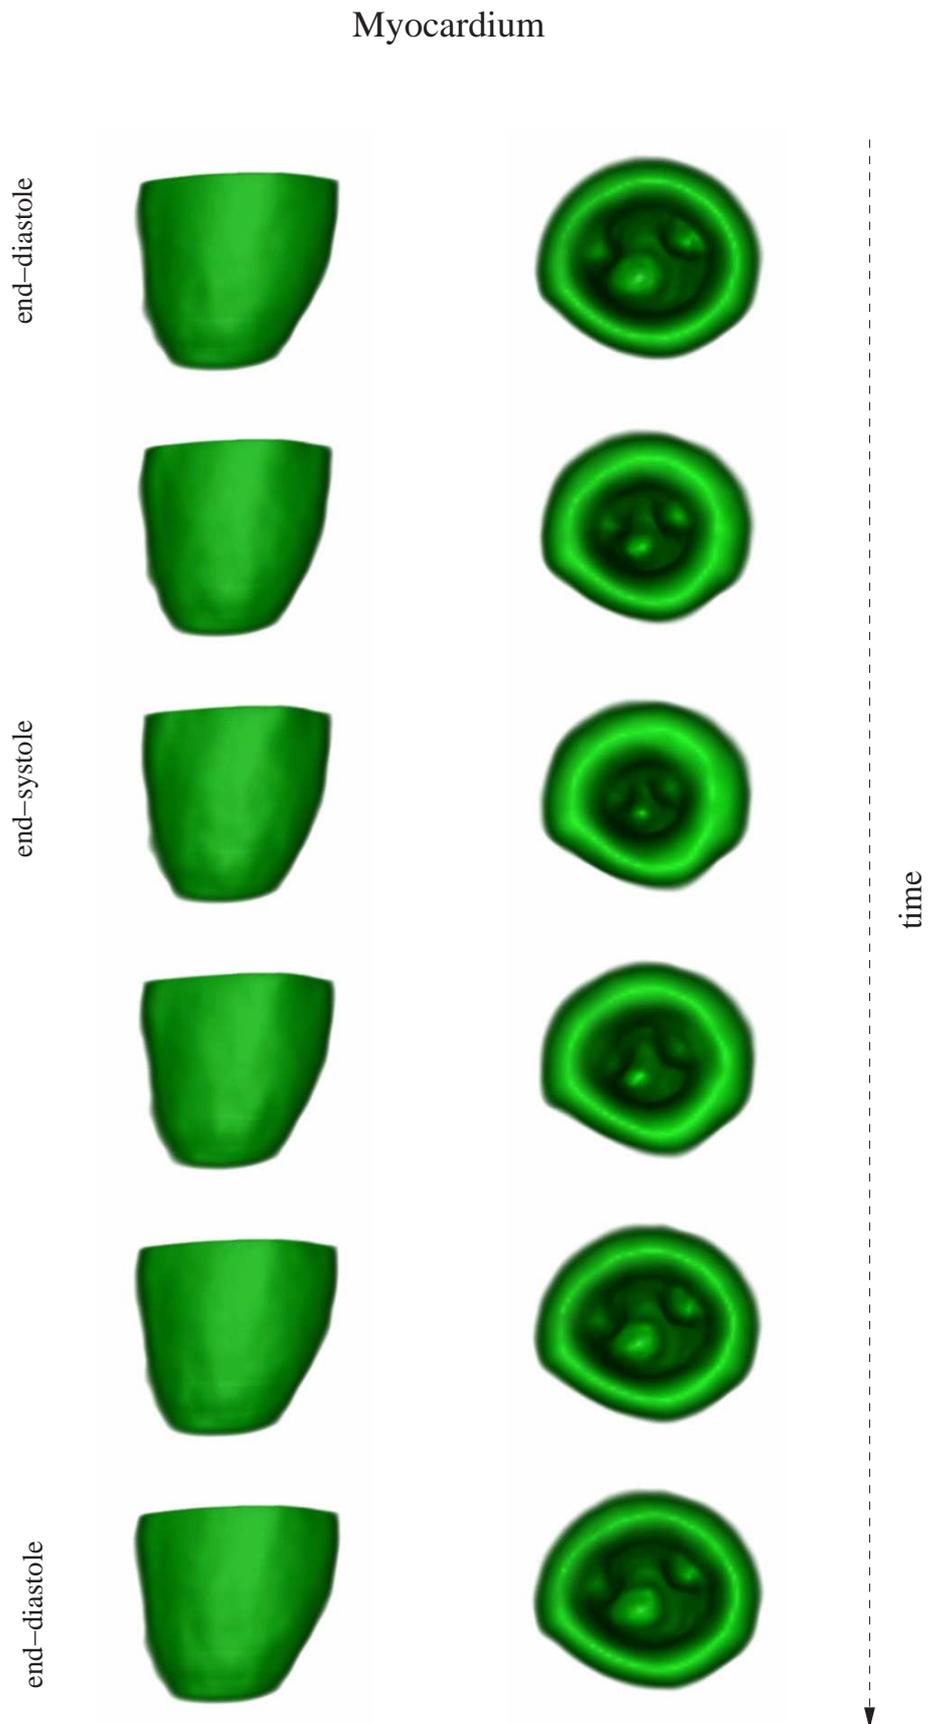
Myocardium

Figure 6.10: Examples of the volume renderings of the myocardium viewed from two different positions.
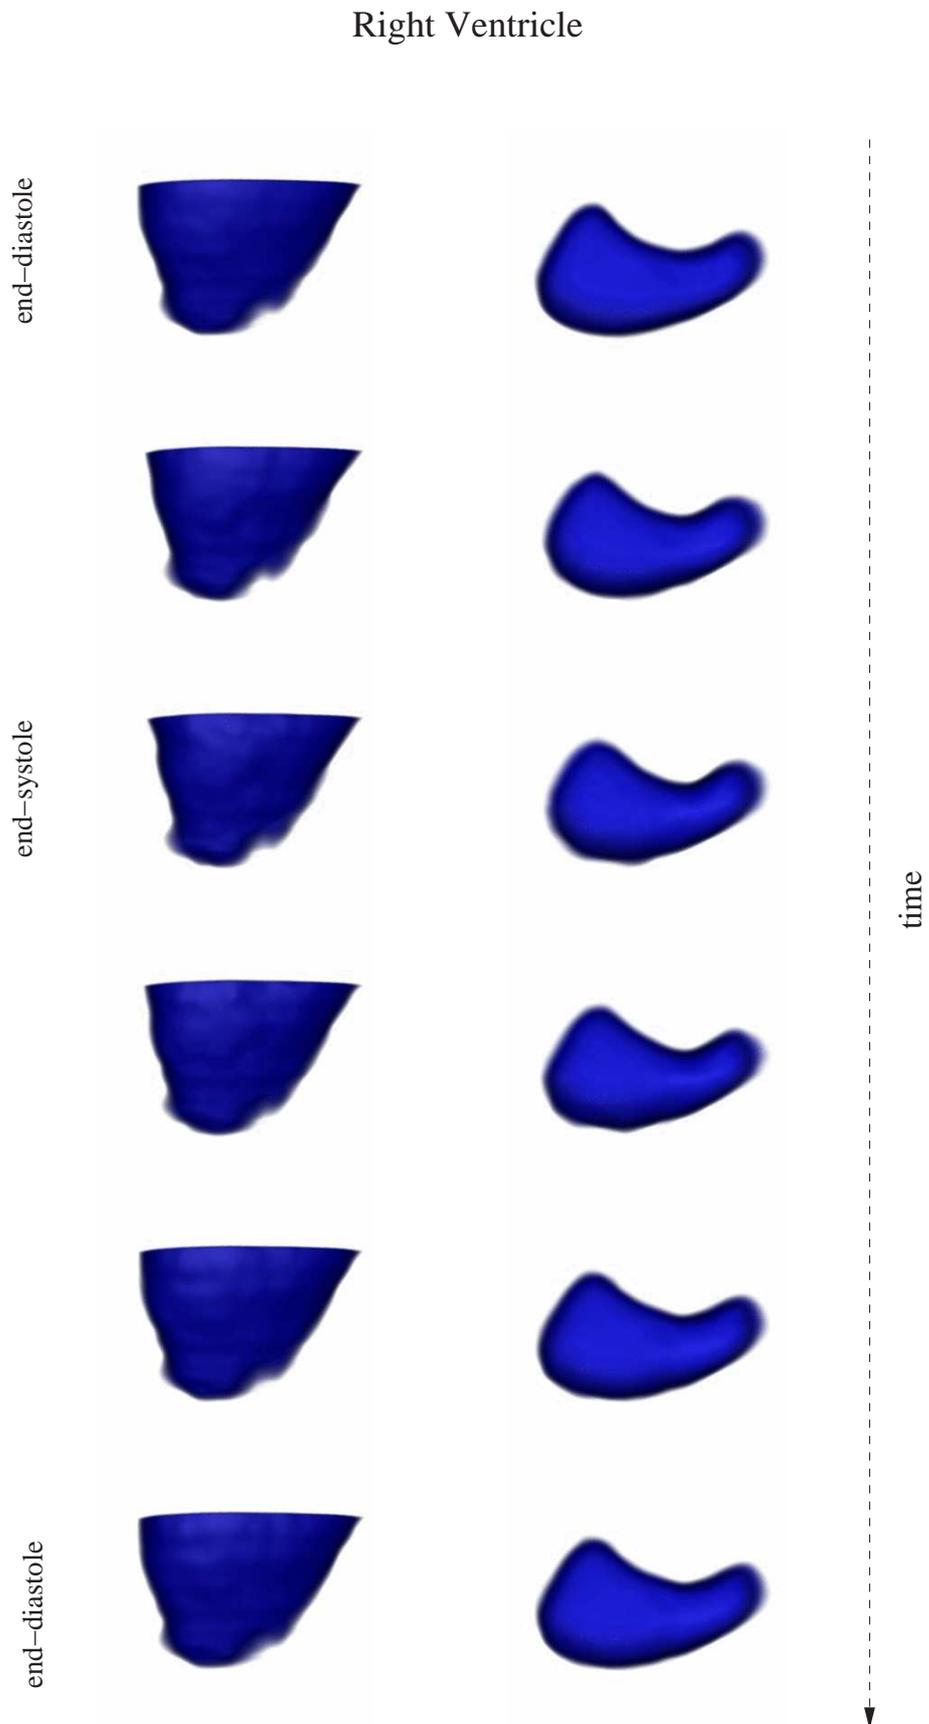
Right Ventricle

Figure 6.11: Examples of the volume renderings of the right ventricle viewed from two different positions.

6.4.3   Statistical atlas

In the construction of the statistical atlas it is very important for all the images to include the same area of the anatomy. Otherwise the most significant modes of shape variation will describe the missing areas of some images. As mentioned above, during the construction of the probabilistic atlas, the tissue probability maps are calculated by averaging each image slice separately. In contrast, in the construction of the statistical atlas, PCA analysis is performed on the surfaces extracted from the entire image rather than on each separate slice. Three image sequences were further excluded from the data due to missing a significant part of the hearts towards the apex. Hence, the statistical atlases were constructed using cardiac image sequence from 26 healthy volunteers.

6.4.3.1   Statistical model of the inter-subject variability

These models describe the significant changes in the shape of the left ventricle, the myocardium and the right ventricle due to inter-subject variability. Figures 6.12 - 6.14 show the three most significant modes of variation for the left ventricle, the myocardium and the right ventricle. For the left ventricle, the three most significant modes of shape variation describe the differences in the size of the left ventricle (mode 1), the variation of the position of the apex of the heart (mode 2) and the elongation of the apex of the heart (mode 3). For the myocardium, the three most significant modes of variation describe the size of the myocardium in the long-axis direction (mode 1), the size and thickness of the myocardium (mode 2) and the direction of the myocardium long-axis (mode 3). For the right ventricle the three most significant modes of variation describe the elongation of the apex of the right ventricle (mode 1), the size of the right ventricle (mode 2) and the shape of the right ventricle (mode 3). Table 6.1 provides a description of the three most significant modes of shape variation.
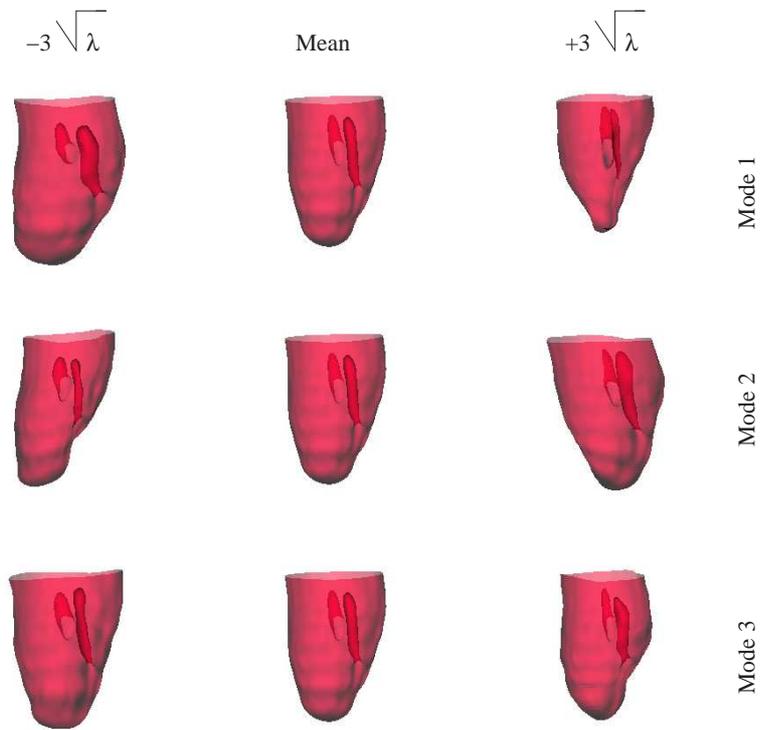
Figure 6.12: The significant modes of shape variation of the left ventricle due to intersubject variability.
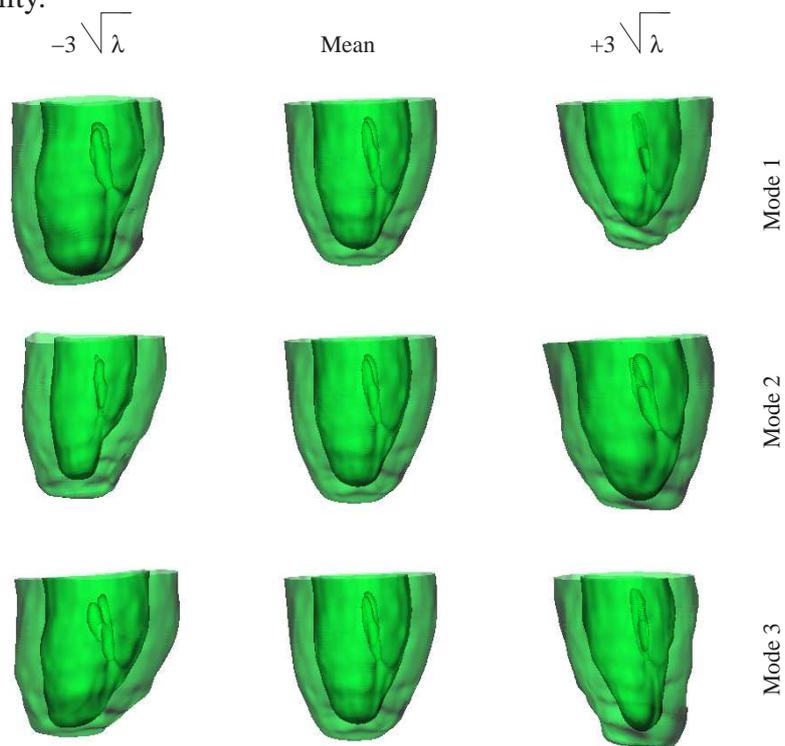
Figure 6.13: The significant modes of shape variation of the myocardium due to intersubject variability.
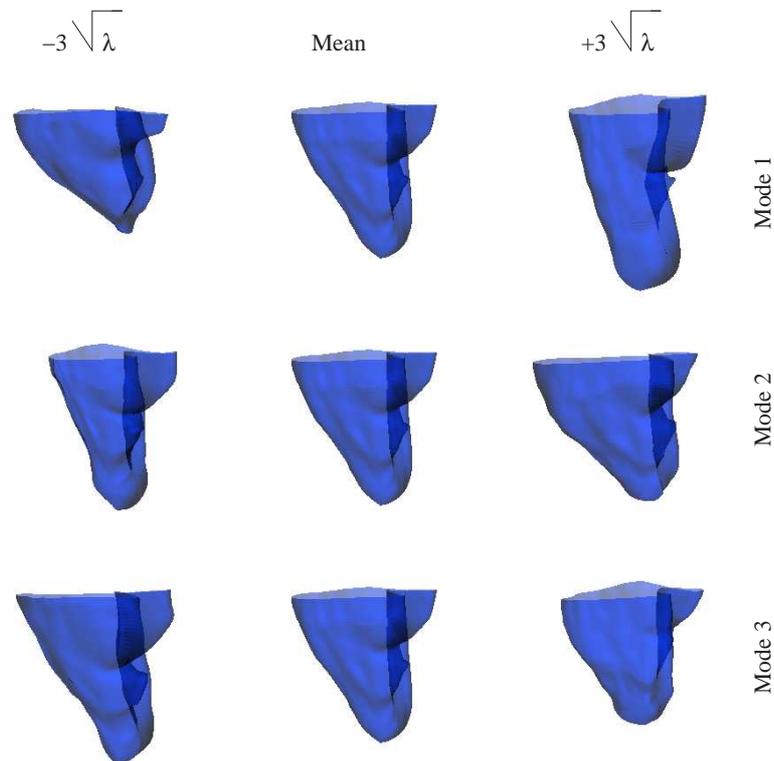
Figure 6.14: The significant modes of shape variation of the right ventricle due to inter-subject variability.

6.4.3.2 Statistical model of the intra-subject variability

These models describe the most significant changes in the shape of the left ventricle and the myocardium which occur across the cardiac cycle (i.e. intra-subject variability). Figures 6.15 - 6.17 show the three most significant modes of variation of the left ventricle, the myocardium and the right ventricle. For the left ventricle, the three most significant modes of variation describe the differences in the volume of the left ventricle during the cardiac cycle (mode 1), the twisting of the heart during the contraction phase (mode 2) and the changes in the position of the apex of the left ventricle as well as the position of the papillary muscles (mode 3). For the myocardium, the three most significant modes of variation describe the changes in the size of the left ventricle and the thickening of the myocardium (mode 1), the twisting of the myocardium during the contraction phase (mode 2) and the changes in the size of the left ventricle in the direction of the long-axis and the movement of the cardiac wall (mode 3). For the right ventricle, the first two most

significant modes of variation describe the changes in the volume of the right ventricle (mode 1) and the twisting of the right ventricle as well as the changes in the position of the right ventricle's apex (modes 2). Finally, the third most significant mode of variation of the right ventricle's shape also describes twisting of the right ventricle (mode 3). Table 6.2 provides a description of the three most significant modes of shape variation.
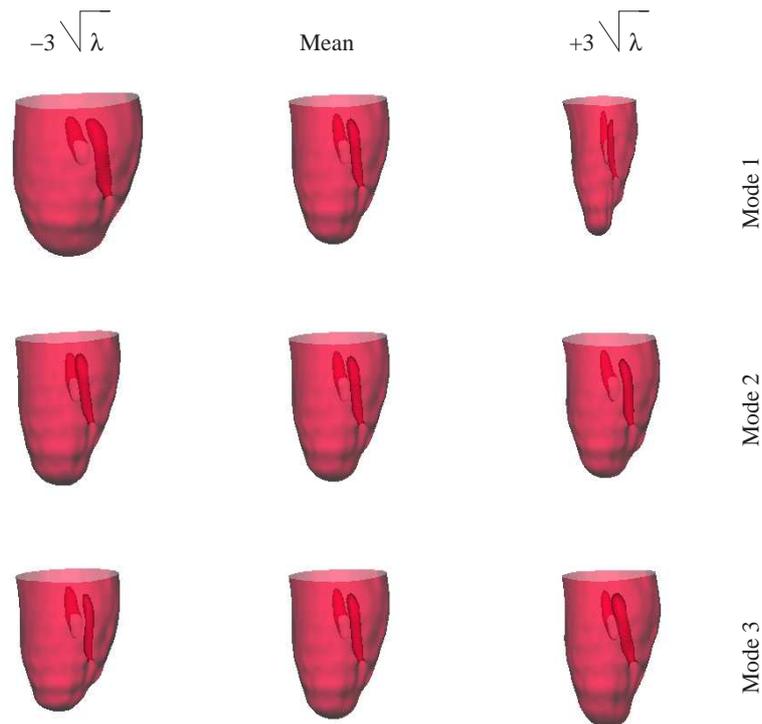
Figure 6.15: The significant modes of shape variation of the left ventricle due to intra-subject variability.
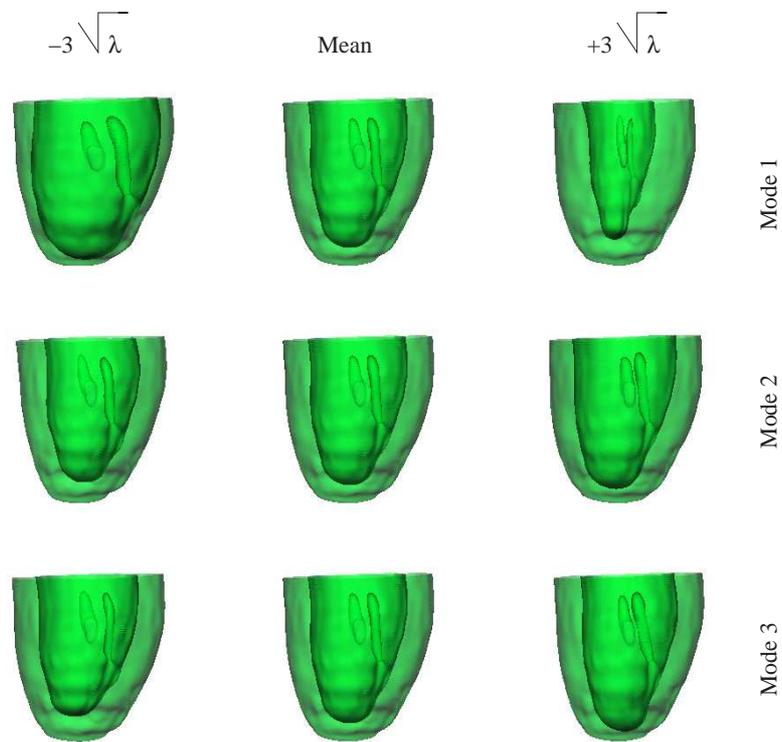
Figure 6.16: The significant modes of shape variation of the myocardium due to intra-subject variability.
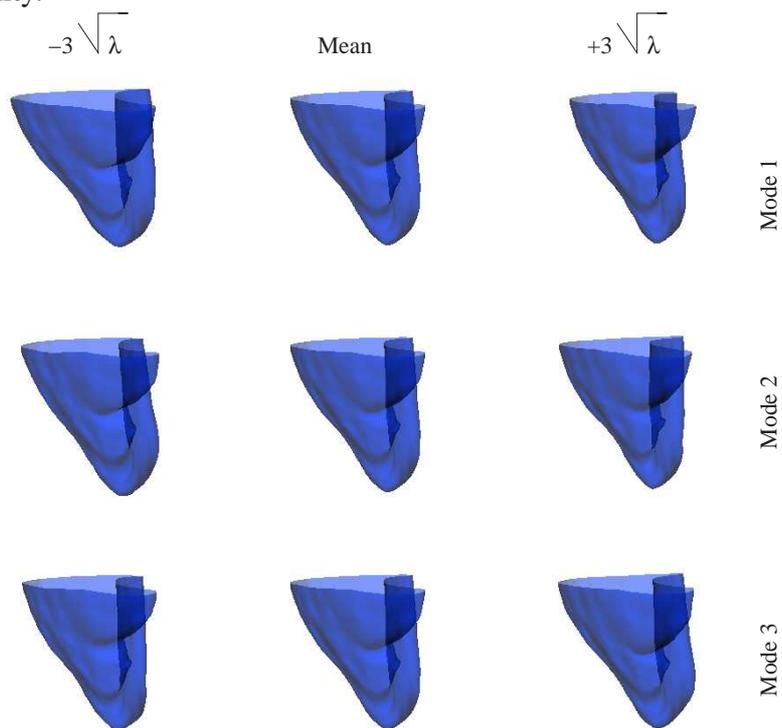
Figure 6.17: The significant modes of shape variation of the right ventricle due to intra-subject variability.

Left Ventricle		
Mode	**Shape Variance**	**Description**
Mode 1	25.07%	Size of the left ventricle
Mode 2	17.34%	Position of the apex of the heart
Mode 3	11.63%	Elongation of the apex of the heart
Myocardium		
Mode	**Shape Variance**	**Description**
Mode 1	20.31%	Size of the myocardium in the long-axis direction
Mode 2	17.06%	Size and thickness of the myocardium
Mode 3	12.54%	Direction of the myocardium long-axis
Right Ventricle		
Mode	**Shape Variance**	**Description**
Mode 1	25.86%	Elongation of the apex of the right ventricle
Mode 2	23.11%	Size of the right ventricle
Mode 3	13.42%	Elongation of the apex and shape of the right ventricle

Table 6.1: The three most significant modes of variation for the left ventricle, the myocardium and the right ventricle due to inter-subject variability.

Left Ventricle		
Mode	**Shape Variance**	**Description**
Mode 1	71.48%	Volume of the left ventricle during the contraction
Mode 2	3.87%	Twisting of the heart during the contraction
Mode 3	2.30%	Changes in the position of the apex and the papillary muscles
Myocardium		
Mode	**Shape Variance**	**Description**
Mode 1	65.87%	Size of the left ventricle and thickening of the myocardium
Mode 2	4.69%	Twisting of the myocardium during the contraction phase
Mode 3	2.13%	Changes in the size of the left ventricle in the direction of the long-axis and movement of the cardiac wall
Right Ventricle		
Mode	**Shape Variance**	**Description**
Mode 1	42.45%	Volume of the right ventricle during the contraction
Mode 2	8.93%	Twisting of the right ventricle and changes in the position of the apex
Mode 3	3.90%	Twisting of the right ventricle

Table 6.2: The three most significant modes of variation for the left ventricle, the myocardium and the right ventricle due to intra-subject variability.

6.4.3.3 Discussion

Figure 6.18 shows the amount of shape variance explained by various modes of the two shape distribution subspaces (inter-subject (a) and intra-subject (b)) for the left ventricle, the myocardium and the right ventricle. In order to describe 90% of the left ventricle's inter-subject shape variability, 13 modes (out of 26) of shape variation are required. Furthermore, 16 (of 468) modes of shape variation are required to describe 90% of the left ventricle's intra-subject shape variability. Similarly, for the myocardium 13 modes are required to describe 90% of its inter-subject shape variability and 28 modes of shape variation are required to describe 90% of its intra-subject shape variability. In order to describe 90% of the right ventricle's inter-subject shape variability, 12 modes (out of 26) of shape variation are required. Finally, 42 modes of shape variation are required to describe 90% of the right ventricle's intra-subject shape variability. The inter-subject statistical model has 26 modes while the intra-subject statistical model has 468. Therefore, in order to enable direct comparison between the curves in figure 6.18 (a) and 6.18 (b) more subjects need to be included in the statistical atlas. The inter-subject shape variability curves are steeper than the intra-subject curves which indicates that the shape variability of the cardiac anatomy is larger across different subjects than across the cardiac cycle.

During the construction of the atlases it is assumed that the distribution of the cardiac shapes in both the inter-subject and intra-subject models is Gaussian. Unfortunately, this assumption has some limitations as seen on figures 6.20 and 6.21. This limitation may produce a difficulty in explaining a shape with a few modes of shape deviation. This will lead to a suboptimal ability of the model to represent unseen instances of cardiac shapes and to generate new instances of cardiac shapes that are similar to those in the training set.

In order to construct the atlases the image sequences need to be registered both in spatial domain and temporal domains. As previously mentioned, the shape of the hearts has to be preserved while the temporal characteristics of each heart have to be deformed so that every corresponding frame of the image sequences corresponds to the same position
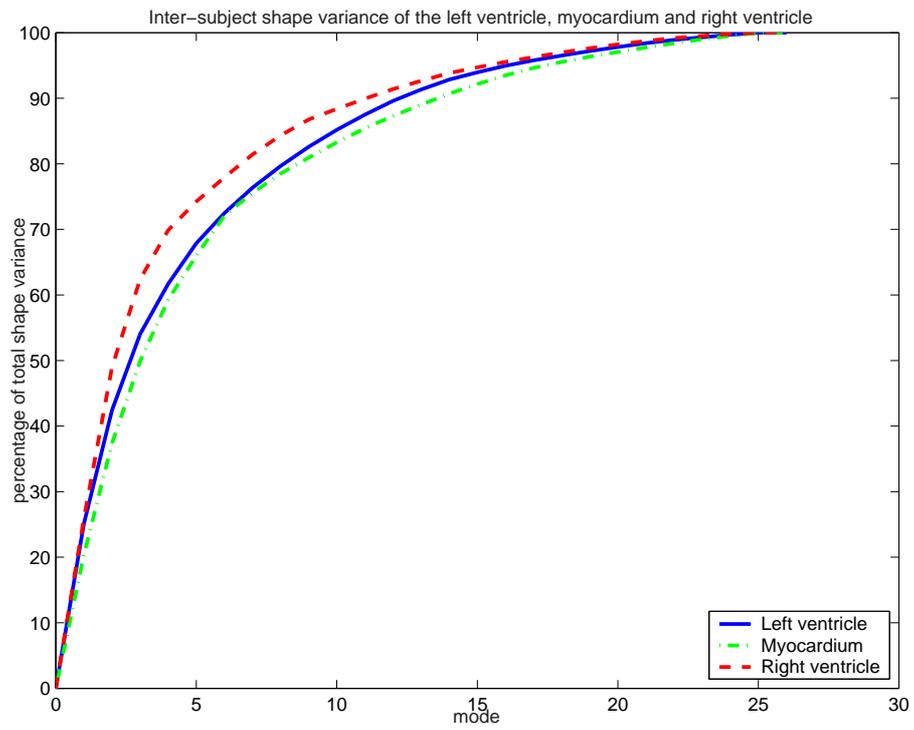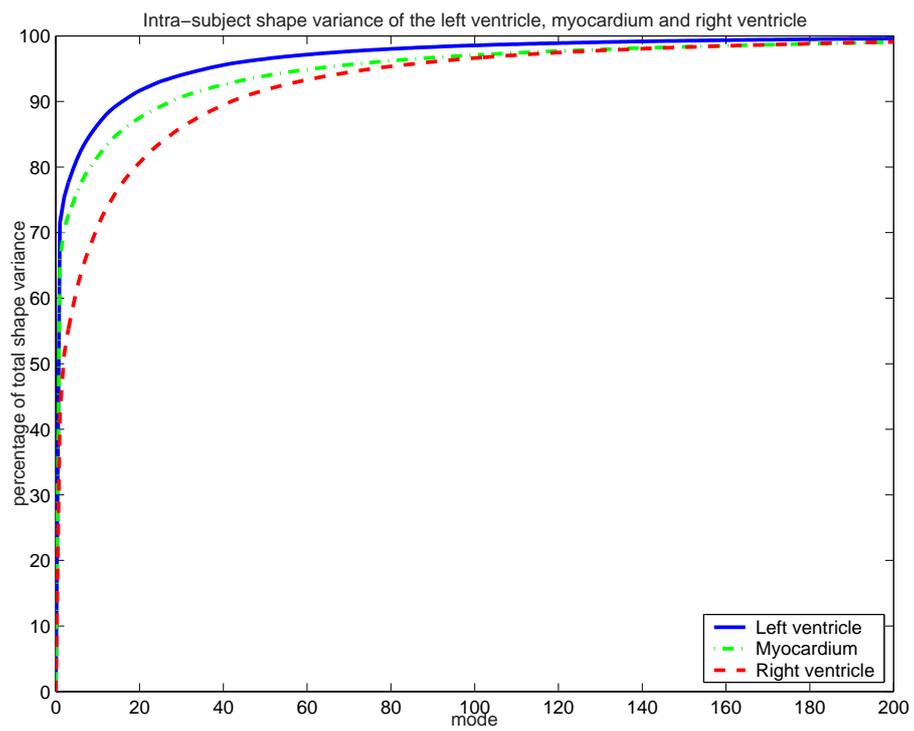
Figure 6.18: The cumulative variance of the left ventricle, the myocardium and the right ventricle: (a) inter-subject and (b) intra-subject.

in the cardiac cycle of the hearts. Figure 6.19 provides the volumes of the left-ventricle of the subjects over the cardiac cycle and shows that after the spatio-temporal registration the hearts follow a similar motion pattern (the volume curves have been separated into two parts: figure 6.19 (a) and 6.19 (b)) . However, it is not expected that the volume curves will completely match since only affine differences are corrected in the spatial domain.

6.5 Differences between the probabilistic and the statistical atlases

The probabilistic atlases contain information about the degree of variability at every voxel of the atlas, i.e. each voxel contains a probability of belonging to a certain structure. However, the probabilistic atlases cannot provide information regarding the type of variability. Furthermore, information regarding shape variability occurring across the population and due to the cardiac cycle is combined to a single probability value. The statistical atlases provide additional information regarding the variability of the cardiac shape. In particular, performing statistical analysis in each shape distribution subspace allows the visualisation of the average shape of the cardiac anatomy and also the identification of the most significant modes of variation in the cardiac shape due to inter-subject and intra-subject shape variability.

As mentioned above, the first steps in the construction of both atlases are identical. However, the probabilistic atlas is created directly from the segmented image sequences (after mapping them to a common spatio-temporal domain), while during the construction of the statistical atlases surfaces need to be extracted from the segmented images and also exact correspondence between landmarks needs to be established.
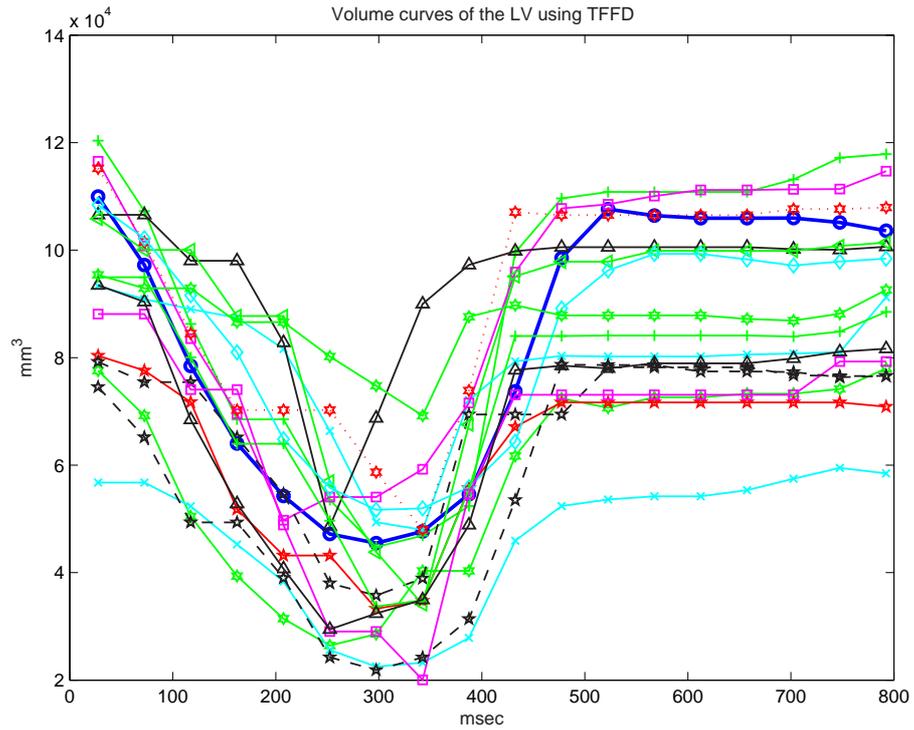
(a)
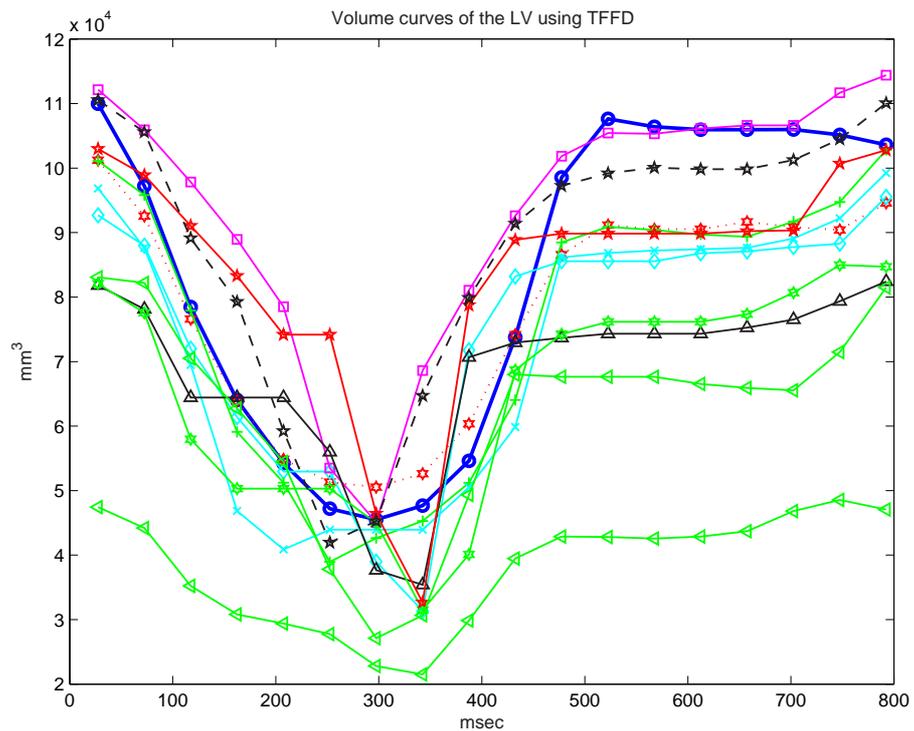
(b)

Figure 6.19: The volume curves of the left ventricle of the subjects over cycle.
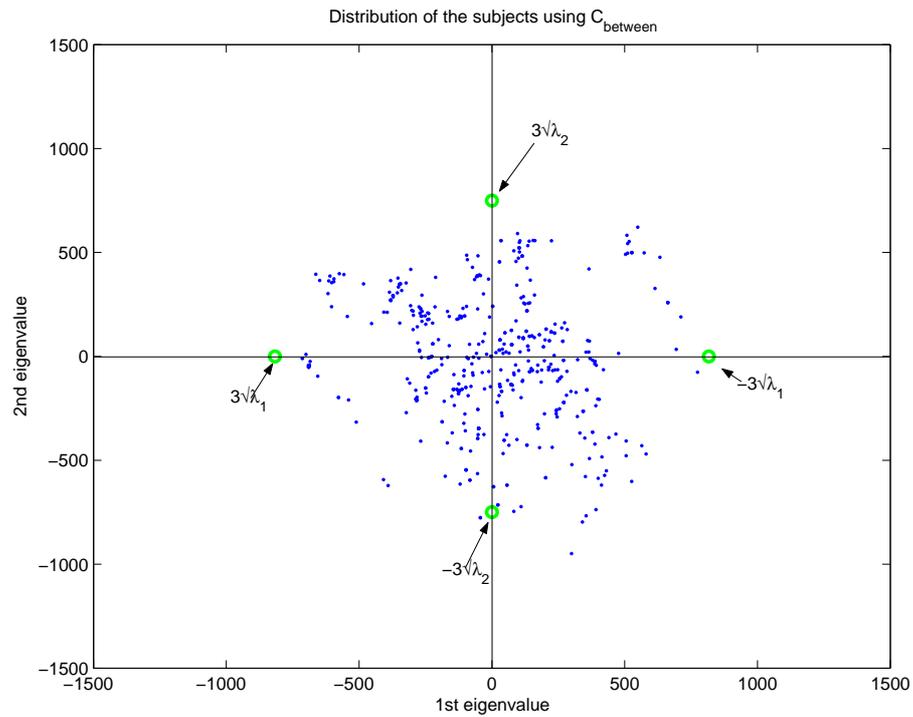
Figure 6.20: The distribution of all subjects' myocardium using the inter-subject covariance matrix.
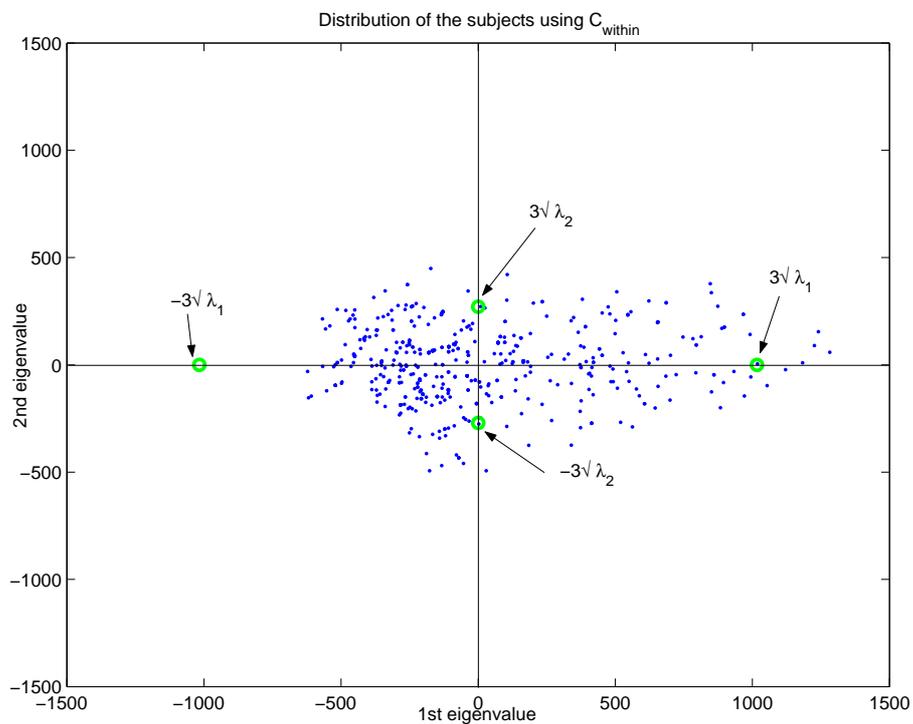
Figure 6.21: The distribution of all subjects' myocardium using the intra-subject covariance matrix.

6.6   Classification of cardiac MR image sequences using the statistical models

In this section a possible use of the statistical atlases for the classification of cardiac data is demonstrated. The aim of this section is not to provide a detailed classification approach but to demonstrate how the statistical models can be used for classification of cardiac MR image sequences. The above statistical models have been used to classify cardiac data from normal volunteers and patients with hypertrophic cardiomyopathy (a condition in which the myocardium has an excessive thickening). In order to perform this classification, six normal subjects have been excluded from the model (i.e. the model has been constructed from only 20 healthy subjects) and cardiac MR image sequences from 10 patients with hypertrophic cardiomyopathy have been acquired. The subjects with hypertrophic cardiomyopathy were not age and sex matched. In a more detailed study of the disease it would be more appropriate to use subjects from the same sex and age group. This forms a set of 16 image sequences to be classified. The same processing steps for the registration and pseudo-landmark extraction and propagation were used for these image sequences as those used for the construction of the statistical models (see section 6.3). Then, for each image sequence, the mean surface (over the cardiac cycle) was calculated. These mean surfaces were projected to the space of the statistical models. Figure 6.22 shows the projections of the subjects' myocardium to the space of the inter-subject population atlas (a) and intra-subject atlas (b) . It is clear from the distribution of the data that a simple classifier should enable the correct differentiation between normal and hypertrophic subjects.

In order to classify the data a *k-weighted NN-classifier* has been used. After projecting the subject to be classified to the space of the statistical models, the classifier calculates its distance from the 5 nearest subjects. Then two weighted sums of distances are calculated. The first is the weighted sum of the subject's distance to each of the normal subjects. The second is the weighted sum of the subject's distance to each of the subjects with hypertrophic cardiomyopathy. The subject is classified depending on the smallest of the

weighted distance sums. Two subjects were excluded from the experiments due to poor registration with the reference subject. A leave one out experiment was performed using the rest 14 subjects (the algorithm is described in listing 4).

Listing 4 The leave one out experiment using *k-weighted NN-classifier* algorithm.

 1: **repeat**
 2:     **select** a subject, s_i, to be used for classification
 3:     **describe** s_i using the model's eigenvectors
 4:     **describe** the rest of the subjects using the model's eigenvectors
 5:     **calculate** the distance, d_i, between s_i and each subject i
 6:     **sort** the subjects such that $d_i \leq d_{i+1}$
 7:     **calculate** the weight factor for the first 5 nearest subjects ($N = 5$): $w = \sum_{i=0}^{i=N} d_i$
 8:     **calculate** the total distance of normal subjects:

$$D_n = \sum_{i=0}^{4} \frac{d_i}{w}, \text{ if } i \in normals$$

 9:     **calculate** the total distance of hypertrophic subjects:

$$D_h = \sum_{i=0}^{4} \frac{d_i}{w}, \text{ if } i \in hypertrophics$$

10:     **if** $D_n \leq D_h$ **then**
11:         s_i is normal
12:     **end if**
13:     **if** $D_h \leq D_n$ **then**
14:         s_i has hypertrophic cardiomyopathy
15:     **end if**
16: **until** all subjects have been classified

The first 4 principal components were employed when using the statistical atlas describing the intra-subject cardiac shape variability, while the second and the third principal components were employed when using the statistical atlas describing the inter-subject shape variability. In this case, the first principal component was not used in the classification since it describes the size of the myocardium in the long-axis direction. Furthermore, the combination of these modes was also used for data classification. The classification results are reported in table 6.3.

All the subjects with hypertrophic cardiomyopathy and 83% of the normal subjects were classified correctly using the statistical model of the myocardium describing changes in the cardiac anatomy due to intra-subject variation (model A). All the normal subjects

	Model A	Model B	Both
Normal	83%	100%	100%
Hypertrophic cardiomyopathy	100%	87.5%	87.5%

Table 6.3: Accuracy of the classification using the statistical atlas of the myocardium describing changes due to intra-subject shape variation (model A), inter-subject shape variation (model B) and a combination of both.

and 87.5% of the subjects with hypertrophic cardiomyopathy were classified correctly using the statistical model of the myocardium describing changes in the cardiac anatomy due to inter-subject variation (model B). The combination of the features of the two statistical models does not improve the results. This might be due to the limited size in the two data sets. The fact the hypertrophic cardiomyopathy is a progressive disease could affect the classification approach. This is especially true for subjects at the early stages of the disease.

6.7   Conclusions

This chapter presented methods for building a 4D probabilistic atlas of the cardiac anatomy and function and a 4D statistical atlas of the cardiac anatomy. The probabilistic atlas captures information regarding the cardiac anatomy and function in the form of tissue probability maps. However, the probabilistic framework cannot provide information regarding the type of variability. The 4D statistical atlas describes the cardiac anatomy and how the cardiac anatomy changes during the cardiac cycle. Contrary to probabilistic atlases, the statistical atlases provide not only information regarding how much variability exists in the data but also what the variability is. The method for building statistical atlases separates the distribution space of the cardiac shape into two subspaces. One distribution subspace accounts for the changes in cardiac shape caused by inter-subject variability. The second distribution subspace accounts for the changes in cardiac shape caused by deformation in the cardiac cycle (i.e. intra-subject variability). Principal component analysis (PCA) has been performed in order to calculate the most significant modes of variation of each distribution subspace. Moreover, this method eliminates the need for manual land-
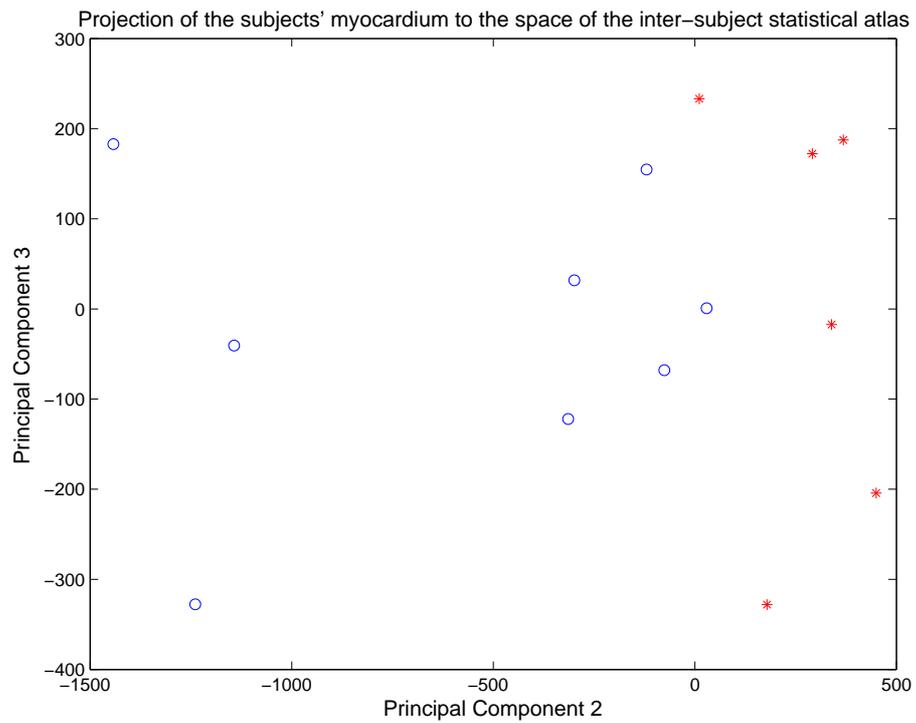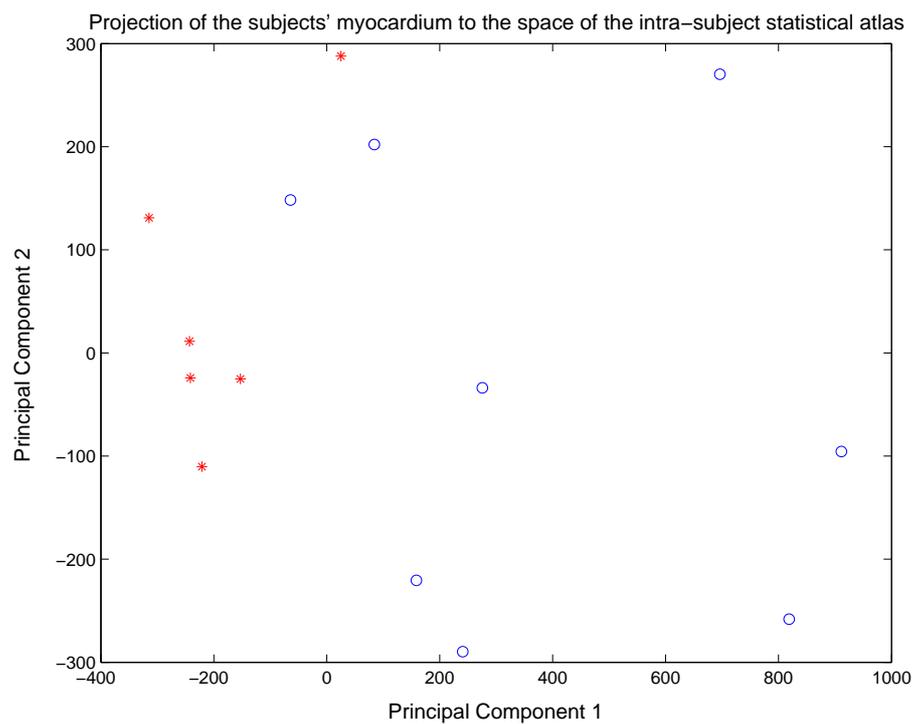
Figure 6.22: Projection of the subjects' myocardium to the space of the (a) inter-subject atlas and (b) the intra-subject atlas (the circles on the graphs represent the subjects with hypertrophic cardiomyopathy while the stars represent the normal subjects).

marking of the cardiac images by using a non-rigid registration algorithm to propagate landmarks from an automatically landmarked atlas to each frame of the image sequences. Both atlases of the myocardium have been used to classify cardiac image sequences from six healthy volunteers and ten patients with hypertrophic cardiomyopathy. The results showed that the use of statistical atlas describing shape changes due to intra-subject variability enables 83% of the normal subjects and 100% of the hypertrophic subjects to be classified correctly. The use of the statistical atlas describing shape changes due to inter-subject variability enables 100% of the normal subjects and 87.5% of the hypertrophic subjects to be classified correctly.

Chapter 7

Conclusions

Spatio-temporal registration and modelling of the cardiac anatomy and function using cardiac MR imaging are challenging tasks. In this chapter, the principal contributions of this thesis are analysed, the general limitations of the presented work are discussed and a description of possibilities for future research is presented. Animation examples (of spatio-temporal registration and the atlases) can be found at:

http://www.doc.ic.ac.uk/~dp1/Research/Thesis/

7.1   Contributions

Chapters 4 and 5 describe novel registration methods which enable the spatial and temporal registration of cardiac MR image sequences. Furthermore, chapter 6 describes methods for modelling the cardiac anatomy and function. In contrast to published methods for building probabilistic and statistical atlases of the cardiac anatomy, the work presented in chapter 6 aims to build probabilistic and statistical atlases which will describe not only the cardiac anatomy but also how it changes over the cardiac cycle.

In chapter 4, a new method for the simultaneous spatial and temporal alignment of cardiac MR image sequences to the same coordinate system is presented. The presented registration algorithm has the ability to correct spatial misalignment of affine nature between the image sequences and also temporal misalignment which could be the result

of differences in the length of the cardiac cycles of the subjects and in the temporal acquisition parameters. The registration method has been extended by the introduction of a *deformable* spatial transformation model which not only corrects global spatial shape differences but also local differences in the cardiac anatomy.

In chapter 5, the registration method for the simultaneous spatial and temporal alignment of cardiac MR image sequences is further extended by the introduction of a *deformable* temporal transformation part. Therefore, this spatio-temporal registration method has the ability to correct spatial misalignment between the images caused by global and local shape differences. Furthermore, it has the ability to correct temporal misalignment caused by differences in the length of the cardiac cycles and in the dynamic properties of the hearts. An alternative method for the *deformable* spatio-temporal alignment of cardiac MR image sequences is also presented. The major difference compared to the previous method is that this method optimises the temporal and spatial components separately. In this method, the temporal registration of the image sequences is based on a normalised cross-correlation measure, while the spatial mapping is based on image information from only the first frames of the sequences. This registration method is significantly faster than the previous methods. Moreover, this registration method can be used to enable only the temporal alignment of cardiac MR image sequences, since it does not require the cardiac image sequences to be aligned in the spatial domain in order to find their temporal correspondence.

In chapter 6, a probabilistic atlas of the cardiac anatomy and function (in terms of how the anatomy changes during the cardiac cycle) is presented. Modelling the cardiac anatomy and function addresses the limitations of current probabilistic atlases of the heart which are only limited to cardiac anatomy and not to the cardiac function. A statistical atlas containing statistical information regarding the anatomy of a healthy heart and how the anatomy of the heart changes during the cardiac cycle is also presented. This statistical cardiac atlas addresses the limitations of current statistical cardiac atlases which describe either the cardiac anatomy or the cardiac function. In order to build the statistical atlases the distribution space of the cardiac shape has been subdivided into two separate

subspaces. The first distribution space describes the changes in the cardiac shape caused by inter-subject variability. The second distribution space describes the changes in the cardiac shape caused by the cardiac cycle. Two separate *principal component analyses* (PCA) have been performed in order to calculate the most significant modes of variation of each subspace.

In chapter 6, a possible use of the statistical atlases for the classification of cardiac data is demonstrated. The statistical atlases of the myocardium are used for the classification of image sequences from normal subjects and subjects with hypertrophic cardiomyopathy.

7.2 Conclusions

7.2.1 Spatio-temporal registration

The registration methods presented in chapters 4 and 5 enable the spatial and temporal comparison of cardiac MR image sequences. The resolution of the image sequences and the spacing of the control points of the B-Spline transformation models used in this thesis, enable spatio-temporal alignment of cardiac MR image sequences with sufficient accuracy for building statistical and probabilistic models of the heart.

A linear spatial transformation model and a linear temporal transformation model are not sufficient to provide complete mapping between the cardiac anatomy and function of two cardiac MR image sequences. Several applications require a more detailed mapping both in the spatial and temporal domain. For example, the construction of 4D models of the cardiac anatomy and function requires the use of a *deformable* temporal transformation model since a mapping between the cardiac anatomies at corresponding time points in the cardiac cycles needs to be established. The evaluation of the registration methods showed that the introduction of the *deformable* spatial part, in the registration method described in chapter 4, significantly improves the spatial alignment of the image sequences. Furthermore, the *deformable* temporal part (chapter 5) addresses the limitations of the affine temporal model. The affine model can only scale and translate the entire temporal

domain of the image sequences. However, the *deformable* model can deform parts of the temporal domain differently. Its use improves significantly the temporal mapping of the image sequences by deforming the motion characteristics of each image sequence to map the reference image sequence.

The calculation of the temporal transformation using the simultaneous optimisation approach provides similar results as the cross-correlation based approach. The temporal mapping after simultaneous optimisation aligns better the temporal details between the two image sequences (figure 5.30). However, the cross-correlation based approach is significantly faster and does not require the image sequences to be registered in the spatial domain.

The results of chapters 5.4.1 and 5.4.2 showed that optimising the spatial part of the transformation based on information from the entire image sequences and not only from the first frames, results in a better spatial alignment (figure 5.22). This is due to the fact that the combined optimisation approach takes into consideration that the spatial differences of the hearts are not constant over the cardiac cycle and also can resolve ambiguities.

The simultaneous optimisation of the transformation components provides better spatio-temporal registration than optimising the transformation components separately. Using the cross-correlation based method to calculate the temporal mapping and optimising only the spatial transformation reduces the computational complexity of the combined optimisation approach without affecting its performance significantly. The different approaches for *deformable* spatio-temporal alignment of cardiac MR image sequences can be used depending on the type of application (e.g. how fast the spatio-temporal mapping should be calculated and how accurate it should be).

7.2.2 Models of the cardiac anatomy and function

The techniques presented in this thesis allow the building of 4D atlases containing information regarding the anatomy and function (in terms of how the anatomy changes during

the cardiac cycle) of a healthy heart. The probabilistic atlas contains information in the form of tissue probability maps. However, the probabilistic framework cannot provide information regarding the type of variability. The statistical atlases provide not only information regarding how much variability exists in the data but also what the variability is. The method for the construction of the statistical atlas allows the decomposition of the distribution space of the cardiac shape into two subspaces. This enables the construction of two separate models: one describing changes in the cardiac shape caused by inter-subject variability, and the other by intra-subject variability. The most significant modes of shape variation of these models describe meaningful variations of the cardiac shape.

7.3 General limitations

The approaches proposed in this dissertation have a number of limitations. There are limitations due to the underlying assumptions of the spatio-temporal registration. The type of mapping, $\mathbf{T}(x, y, z, t) = (x'(x, y, z), y'(x, y, z), z'(x, y, z), t'(t))$, used in the spatio-temporal registration methods assumes that the spatial differences between two hearts are constant over the cardiac cycle. However, in chapter 5 it has been shown that this assumption has some limitations. This can been seen from the overlap measures (section 5.4.1.2) of the registration method which calculates the optimal spatial transformation using only information from the first frames of the sequences (separate optimisation). Furthermore, figures 5.3 and 5.4 show that the hearts become more misaligned in the spatial domain during the contraction phase of the cardiac cycle. Finally, the alignment of the volume curves of figure 5.8 is not as good as the volume curves of figure 5.15 because the optimisation of the spatial transformation is based only on the first frames and not the entire sequences. Although the proposed registration methods cannot provide 100% spatio-temporal alignment between different image sequences, they still provide a good approximation to the problem. The simultaneous optimisation of the temporal and spatial components is not able to completely separate the spatial differences from the temporal differences. Therefore, the spatial mapping will compensate for temporal differences to a

certain extent. Similarly, the temporal mapping will compensate for spatial differences to a certain extent.

There are also limitations due to image acquisition and quality which could affect the results. The quality of the resulting temporal mapping of the proposed methods will depend on the temporal resolution of the image sequence. The more frames a sequence has (i.e. temporal information), the better the temporal alignment will be. Furthermore, the quality of the image sequences will also affect the results since the similarity measure (normalised mutual information) used during registration is based only on image information. However, the use of image blurring and the choice of normalised mutual information as a similarity measure reduces the effect of image noise on the quality of the resulting spatio-temporal registration.

The identification of features in the cardiac cycle using the cross-correlation based method will also depend on the quality of the images and the temporal resolution of the image sequences. Calculating the normalised cross-correlation between frames with low signal to noise ratio may affect the resulting similarity values which could influence the identification of the temporal features.

The low resolution in the through-plane direction (typical slice thickness is 8-10mm) affects the registration and the segmentation of the apex and the base of the heart. This introduces artifacts in the probabilistic and statistical atlases making them less accurate in the areas of the apex and base of the heart.

7.4   Future research

There are two main directions for future research. The first involves the development of new methods for the spatio-temporal alignment of image sequences and for modelling the cardiac anatomy and function. The use of a transformation mapping of the following form could be investigated:

$$\mathbf{T}(x, y, z, t) = (x'(x, y, z, t'), y'(x, y, z, t'), z'(x, y, z, t'), t'(t)) \qquad (7.1)$$

This type of 4D mapping does not assume that the spatial differences between two hearts are constant over the cardiac cycle. It can provide different spatial alignment for different frames of the image sequences. A 4D B-Spline transformation can be used as $\mathbf{T}_{spatial}$. In order to define a 4D B-Spline transformation, let Φ_{4D} denote an n_x, n_y, n_z, n_t mesh of control points $\phi_{i,j,k,t}$ with uniform spacing δ in the spatial domain and δ_t in the temporal domain. Then the 4D spatial mapping can be written as:

$$\mathbf{T}^{local}_{spatial}(x, y, z, t) = \sum_{s=0}^{3} \sum_{l=0}^{3} \sum_{m=0}^{3} \sum_{n=0}^{3} B_s(o)B_l(u)B_m(v)B_n(w)\phi_{t+s,i+l,j+m,k+n} \quad (7.2)$$

As mentioned in chapter 4, finding the optimal transformation of such type will be significantly more computationally expensive than the type of transformation used in this thesis. For example, a transformation consisting of a 3D B-Spline with $10 \times 10 \times 10$ control points and a 1D temporal B-Spline with 10 control points will have 3010 *degrees of freedom*. On the other hand, a transformation consisting of a 4D B-Spline with $10 \times 10 \times 10 \times 10$ control points and a 1D temporal B-Spline with 10 control points will have 30010 *degrees of freedom*. Increasing the spacing of the control points in the temporal domain, δ_t, will reduce the number of *degrees of freedom* in the 4D transformation. It might be appropriate to calculate δ_t such that it provides spatial alignment for a limited number of temporal feature positions (e.g. for the beginning and ends of the cardiac cycles and for the feature position of peak contraction). In this case, the 4D B-Spline based transformation for the spatial alignment of the image sequences will have $10 \times 10 \times 10 \times 3$ control points and 9000 *degrees of freedom*.

A statistical deformation atlas of the cardiac anatomy and function could be constructed. As mentioned in chapter 3, the key difference of statistical deformation atlases is that statistical analysis is performed directly on the deformation fields which describe a dense correspondence between the anatomies and also cardiac function of two image sequences. Statistical deformation atlases have mainly been developed for the human brain [173]. The construction of a statistical deformation model from cardiac MR image sequences will provide statistical information regarding the deformation required to

map the anatomies of several hearts. It will also provide information regarding the temporal deformation required to map the motion characteristics of several hearts. In order to obtain these deformation fields (spatial and temporal), a *deformable* spatio-temporal registration method will be used.

The second direction for future research is the development of new applications which use the proposed methods. As mentioned in chapter 3, the construction of an atlas should be seen as an ever evolving process. The more data an atlas has, the more accurate representation of the population it will be. Hence, it is very important to add more cardiac MR image sequences to the atlas to make it more representative of the population.

Different statistical and probabilistic atlases can be constructed for different groups of the population. For example, statistical and probabilistic atlases can be constructed for specific cardiovascular diseases. These atlases can be used in a large number of tasks including for the segmentation of cardiac MR images. They can also be used to enable the better understanding of the cardiovascular diseases and also the classification of image sequences. Furthermore, the comparison of different atlases can enable clinicians to assess qualitative and quantitative differences between groups (e.g. normal subjects and subjects with hypertrophic cardiomyopathy).

The spatio-temporal registration methods can be used for the assessment of pharmacological and surgical intervention. The effect of pharmacological or surgical intervention will be studied by comparing images prior and after to the intervention. Mapping a number of image sequences to a common spatio-temporal coordinate system (this coordinate system could be defined by an atlas) will enable to study the effect of the intervention to the cardiac physiology. It may also be possible to use the same methodology to study how the cardiac anatomy and function changes due to aging.

The methods presented in this dissertation enable the spatio-temporal mapping of image sequences with minimal user interaction. The methods (except an initialisation step) are automatic and require no segmentation of anatomical features and also identification of temporal features in the cardiac cycles. These methods are used to build atlases describing the cardiac anatomy and how the anatomy changes over the cardiac cycle. The use

of atlases is expected to play a significant role in the interpretation of cardiac MR images since they collect anatomical and functional information from a large set of the population to a single model. The work presented in this dissertation provides the foundations for the spatio-temporal mapping of cardiac MR image sequences and also for modelling the cardiac anatomy and function. This will hopefully encourage many important new developments still to come.

Publications

Publications in journals

- D. Perperidis, R. Mohiaddin, and D. Rueckert. Spatio-temporal free-form registration of cardiac MR image sequences. *In Medical Image Analysis*, 9(5):441-456, 2005.

Publications in conference proceedings

- D. Perperidis, R. Mohiaddin, and D. Rueckert. Construction of a 4D statistical atlas of the cardiac anatomy and its use in classification. *In Eight Int. Conf. on Medical Image Computing and Computer- Assisted Intervention (MICCAI '05)*, Lecture Notes in Computer Science, 2005.

- D. Perperidis, R. Mohiaddin, and D. Rueckert. Fast Spatio-temporal registration of cardiac MR Image Sequences. *In Functional Imaging and Modeling of the Heart (FIMH'05)*, Lecture Notes in Computer Science, LNCS 3504, pages 414-424, 2005.

- D. Perperidis, R. Mohiaddin, and D. Rueckert. Spatio-temporal free-form registration of cardiac MR image sequences. *In Seventh Int. Conf. on Medical Image Computing and Computer- Assisted Intervention (MICCAI '04)*, Lecture Notes in Computer Science, LNCS 3217, pages 911-919, 2004.

- D. Perperidis, R. Mohiaddin, and D. Rueckert. Spatio-temporal free-form registra-

tion of cardiac MR image sequences. *In Proc. of Medical Image Understanding and Analysis '04*, pages 157- 160, 2004.

- D. Perperidis, M. Lorenzo-Valdés, R. Chandrashekara, A. Rao, R. Mohiaddin, G. I. Sanchez- Ortiz, and D. Rueckert. Building a 4D atlas of the cardiac anatomy and motion using MR imaging. *In IEEE International Symposium on Biomedical Imaging*, pages 412-415, 2004.

- D. Perperidis, M. Lorenzo-Valdés, R. Chandrashekara, A. Rao, R. Mohiaddin, G. I. Sanchez- Ortiz, and D. Rueckert. Building a 4D atlas of the cardiac anatomy and motion using MR imaging. *In Proc. of Medical Image Understanding and Analysis '04*, pages 9-12, 2004.

- D. Perperidis, A. Rao, R. Mohiaddin, and D. Rueckert. Non-rigid spatio-temporal alignment of 4D cardiac MR images. *In Second International Workshop on Biomedical Image Registration (WBIR '03)*, Lecture Notes in Computer Science, LNCS 2717, pages 191-200, 2003.

- D. Perperidis, A. Rao, M. Lorenzo-Valdés, R. Mohiaddin, and D. Rueckert. Spatio-temporal alignment of 4D cardiac MR images. *In Functional Imaging and Modeling of the Heart (FIMH '03)*, Lecture Notes in Computer Science, LNCS 2674, pages 205-214, 2003.

Bibliography

[1] N. A. Ablitt, J. Gao, J. Keegan, L. Stegger, D. N. Firmin, and G. Z. Yang. Predictive cardiac motion modeling and correction with partial least squares regression. *IEEE Transactions on Medical Imaging*, 23(10):1315–1324, 2004.

[2] American Heart Association. Heart disease and stroke statistics - 2005 update. http://www.americanheart.org/, 2005.

[3] K. S. Arun, T. S. Huang, and S. D. Blonstein. Least square fitting of two 3-D point sets. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(5):698–700, 1987.

[4] L. Axel. Biomechanical dynamics of the heart with MRI. *Annual Review of Biomedical Engineering*, 4:321–347, 2002.

[5] L. Axel. Tagged MRI-based studies of cardiac function. In *Functional Imaging and Modeling of the Heart, FIMH'03*, number 2674 in Lecture Notes in Computer Science, pages 1–7, Lyon, France, 2003. Springer.

[6] L. Axel and L. Dougherty. Heart wall motion: improved method of spatial modulation of magnetization for MR imaging. *Radiology*, 172(2):349–350, 1989.

[7] L. Axel and L. Dougherty. MR imaging of motion with spatial modulation of magnetization. *Radiology*, 171(3):841–849, 1989.

[8] R. Bajcsy and S. Kovacic. Multiresolution elastic matching. *Computer Vision, Graphics and Image Processing*, 46(1):1–21, 1989.

[9] R. S. Balaban, S. Chesnick, K. Hedges, F. Samaha, and F. W. Heineman. Magnetization transfer contrast in MR imaging of the heart. *Radiology*, 180(3):671–675, 1991.

[10] E. Bardinet, L. D. Cohen, and N. Ayache. Tracking and motion analysis of the left ventricle with deformable superquadrics. *Medical Image Analysis*, 1(2):129–149, 1996.

[11] E. Bardinet, L. D. Cohen, and N. Ayache. A parametric deformable model to fit unstructured 3D data. *Computer Vision and Image Understanding*, 71(1):39–54, 1998.

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

[13] K. K. Bhatia, J. V. Hajnal, B. K. Puri, A. D. Edwards, and D. Rueckert. Consistent groupwise non-rigid registration for atlas construction. In *2004 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, 2004.

[14] L. M. Bidaut and J. P. Vallée. Automated registration of dynamic MR images for the quantification of myocardial perfusion. *Journal of Magnetic Resonance Imaging*, 13(4):648–655, 2001.

[15] F. L. Bookstein. Principal warps: thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(6):567–585, 1989.

[16] F. L. Bookstein. Thin-plate splines and the atlas problem for biomedical images. In *Information Processing in Medical Imaging: Proc. 12th International Conference (IPMI'91)*, number 511 in Lecture Notes in Computer Science, pages 326–342, 1991.

[17] J. G. Bosch, S. C. Mitchell, B. P .F. Lelieveldt, F. Nijland, O. Kamp, M. Sonka, and J. H. C. Reiber. Automatic segmentation of echocardiographic sequences by active

appearance motion models. *IEEE Transactions on Medical Imaging*, 21(11):1374–1383, 2002.

[18] J. J. Bray, A. Cragg, A. D. C. Macknight, R. G. Mills, and D. W Taylor. *Lecture Notes in Human Physiology*. Blackwell Science, third edition, 1994.

[19] M. Bro-Nielsen and C. Gramkow. Fast fluid registration of medical images. In *Proc. Visualization in Biomedical Computing*, number 1131 in Lecture Notes in Computer Science, pages 267–276, 1996.

[20] C. Broit. *Optimal registration of deformed images*. PhD thesis, University of Pennsylvania, USA, 1981.

[21] L. G. Brown. A survey of image registration techniques. *ACM Computing Surveys*, 24(4):325–276, 1992.

[22] T. M. Buzug, J. Weese, C. Fassnacht, and C. Lorenz. Image registration: Convex weighting functions for histogram-based similarity measures. In *Proceedings of Computer Vision, Virtual Reality and Robotics in Medicine (CVRMed/MRCAS)*, number 1205 in Lecture Notes in Computer Science, pages 203–212, Grenoble, France, 1997.

[23] J. F. Cardoso. Higher order contrast for independent component analysis. *Neural Computation*, 11(1):157–192, 1999.

[24] A. Carrington and A. D. McLachlan. *Introduction to magnetic resonance*. Chapman and Hall, 1967.

[25] Y. Caspi and M. Irani. Spatio-temporal alignment of sequences. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(11):1409–1424, 2002.

[26] J. C. Cauvin, J. Y. Boire, M. Zanca, J. M. Bonny, J. Maublant, and A. Veyre. 3D modelling in myocardial ^{210}TL SPECT. *Computerized Medical Imaging and Graphics*, 17(4-5):345–350, 1993.

[27] M. D. Cerqueira, N. J. Weissman, V. Dilsizian, A. K. Jacobs, S. Kaul, W. K. Laskey, D. J. Pennell, J. A. Rumberger, T. J. Ryan, and M. S. Verani. Standardized myocardial segmentation and nomenclature for tomographic imaging of the heart. *Journal of Cardiovascular Magnetic Resonance*, 4(2):203–210, 2002.

[28] R. Chandrashekara, R. Mohiaddin, and D. Rueckert. Analysis of myocardial motion in tagged MR images using nonrigid image registration. In M. Sonka and J. Michael Fitzpatrick, editors, *Proceedings of the SPIE International Symposium on Medical Imaging*, pages 1168–1179, San Diego, California USA, 24-28 February 2002.

[29] R. Chandrashekara, R. Mohiaddin, and D. Rueckert. Analysis of 3D myocardial motion in tagged MR images using nonrigid image registration. *IEEE Transactions on Medical Imaging*, 3(10):1245–1250, 2004.

[30] R. Chandrashekara, A. Rao, G. I. Sanchez-Ortiz, R. Mohiaddin, and D. Rueckert. Construction of a statistical model for cardiac motion analysis using non-rigid image registration. In *Information Processing in Medical Imaging: Proc. 18th International Conference (IPMI'03)*, number 2732 in Lecture Notes in Computer Science, 2003.

[31] Z. Chen, F. S. Prato, and C. A. McKenzie. T_1 fast acquisition relaxation mapping (T_1-FARM): an optimized reconstruction. *IEEE Transactions on Medical Imaging*, 17(2):155–160, 1998.

[32] G. E. Christensen, R. D. Rabbitt, and M. I. Miller. Deformable templates using large deformation kinematics. *IEEE Transactions on Image Processing*, 5(10):1435–1447, 1996.

[33] A. Collignon, F. Maes, D. Delaere, D. Vandermeulen, P. Suetens, and G. Marchal. Automated multimodality image registration using information theory. In *Information Processing in Medical Imaging: Proc. 14th International Conference (IPMI'95)*, pages 263–274, 1995.

[34] D. L. Collins, A. C. Evans, C. Holmes, and T. M. Peters. Automatic 3D segmentation of neuro-anatomical structures from MRI. In *Information Processing in Medical Imaging: Proc. 14th International Conference (IPMI'95)*, pages 139–152, 1995.

[35] D. L. Collins, P. Neelin, T. M. Peters, and A. C. Evans. Automatic 3D intersubject registration of MR volumetric data in standardized talairach space. *Journal of Computer Assisted Tomography*, 18(2):192–205, 1994.

[36] D. L. Collins, T. M. Peters, and A. C. Evans. An automated 3D nonlinear deformation procedure for determination of gross morphometric variability in human brain. In *Proc. SPIE, International Conference Visualization in Biomedical Computing*, volume 2359, pages 180–190, 1994.

[37] T. F. Cootes, G. J. Edwards, and Taylor C. J. Active appearance models. In H.Burkhardt and B. Neumann, editors, *European Conference on Computer Vision 1998*, volume 2, pages 484–498. Springer, 1998.

[38] T. F. Cootes, A. Hill, C. J. Taylor, and J. Haslam. The use of active shape models for locating structures in medical images. *Image Vision Computing*, 12(6):355–366, 1994.

[39] T. F. Cootes and C. J. Taylor. Statistical models of appearance for medical image analysis and computer vision. In *Proceedings SPIE Medical Imaging 2001*, San Diego, CA, 2001.

[40] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models - their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, 1995.

[41] N. P. Costen, T .F. Cootes, G. J. Edwards, and C. J. Taylor. Automatic extraction of the face identity-subspace. *Image and Vision Computing*, 20(5-6):319–329, 2002.

[42] J. G. Cseransky, S. Joshi, L. Wang, J. W. Haller, M Gado, J. P. Miller, U. Grenander, and M. I. Miller. Hippocampal morphometry in schizophrenia by high-dimensional brain mapping. In *Proceedings of the National Academy of Sciences of the United States of America*, volume 95, pages 11406–11411, 1998.

[43] C. Davatzikos. Spatial transformation and registration of brain images using elastically deformed models. *Computer Vision and Image Understanding*, 66(2):207–222, 1997.

[44] C. Davatzikos, M. Vaillant, S. M. Resnick, J. L. Prince, S. Letovsky, and R. N. Bryan. A computerized approach for morphological analysis of the corpus callosum. *Journal of Computer Assisted Tomography*, 20(1):88–97, 1996.

[45] M. H. Davis, A. Khotanzad, D. P. Flaming, and S. E. Harms. A physics-based coordinate transformation for 3D image matching. *IEEE Transactions on Medical Imaging*, 16(3):317–328, 1997.

[46] J. Declerck, N. Ayache, and E. R. McVeigh. Use of a 4D planispheric transformation for the tracking and analysis of LV motion with tagged MR images. In *SPIE Medical Imaging*, volume 3660, pages 69–80, 1999.

[47] J. Declerck, J. Feldmar, M. L. Goris, and F. Betting. Automatic registration and alignment on a template of cardiac stress and rest reoriented SPECT images. *IEEE Transactions on Medical Imaging*, 16(6):727–737, 1997.

[48] E. R. E. Denton, L. I. Sonoda, D. Rueckert, S. Rankin, C. Hayes, M. O. Leach, D. L. G. Hill, and D. J. Hawkes. Comparison and evaluation of rigid, affine, and nonrigid registration of breast MR images. *Journal of Computer Assisted Tomography*, 23(5):800–805, 1999.

[49] R. J. DeQuardo, M. S. Keshavan, F. L. Bookstein, W. W. Bagwell, W. D. Green, J. A. Sweeny, G. L. Haas, R. Tandon, N. R. Schooler, and J. W. Pettegrew.

Landmark-based morphometric analysis of first-episode schizophrenia. *Biological Psychiatry*, 45(10):1321–1328, 1999.

[50] I. D. Dinov, M. S. Mega, P. M. Thompson, L. Lee, R. P. Woods, C. J. Holmes, W. L. de Sumners, and A. W. Toga. Analyzing functional brain images in a probabilistic atlas: A validation of subvolume thresholding. *Journal of Computer Assisted Tomography*, 24(1):128–138, 2000.

[51] R. O. Duba and P. E. Hart. *Pattern Classification and Scene Analysis*. John Wiley, New York, 1973.

[52] J. Duchon. Interpolation des functions de deux variables suivant les principes de la flexion des plaques minces. *RAIO Analyse Numérique*, 10(5):5–12, 1976.

[53] M. C. Dulce, G. H. Mostbeck, K. K. Friese, G. R. Caputo, and C. B. Higgings. Quantification of the left ventricular volumes and function with cine MR imaging: Comparison of geometric models with three-dimensional data. *Radiology*, 188(2):371–376, 1993.

[54] J. Duncan, P. Shi, T. Constable, and A. Sinusas. Physical and geometrical modeling for image-based recovery of left ventricular deformation. *Progress in Biophysics and Molecular Biology*, 69(2-3):333–351, 1998.

[55] T. L. Faber, R. W. McColl, R. M. Opperman, J. R. Corbett, and R. M. Peshock. Spatial and temporal registration of cardiac SPECT and MR images: Methods and evaluation. *Radiology*, 179(3):857–861, 1991.

[56] T. L. Faber, E. M. Stokely, R. M. Peshock, and J. R. Corbett. A model-based four-dimensional left ventricular surface detector. *IEEE Transactions on Medical Imaging*, 10(3), 1991.

[57] A. Fenster, D. B. Downey, and H. N. Cardinal. Three-dimensional ultrasound imaging. *Physics in Medicine and Biology*, 46(5):67–99, 2001.

[58] S. E. Fischer, G. C. McKinnon, S. E. Maier, and P. Boesiger. Improved myocardial tagging contrast. *Magnetic Resonance in Medicine*, 30(2):191–200, 1993.

[59] J. M. Fitzpatrick, D. L. G. Hill, and C. R. Maurer, Jr. *Handbook of Medical Imaging*, chapter Image Registration. SPIE Press, 2000.

[60] International Consortium for Brain Mapping. ICBM probabilistic atlases. http://www.loni.ucla.edu/ICBM/, 2002.

[61] A. Frangi, D. Rueckert, J. Schnabel, and W. Niessen. Automatic construction of biventricular shape models. In *Functional Imaging and Modeling of the Heart, FIMH'03*, number 2674 in Lecture Notes in Computer Science. Springer, 2003.

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

[63] A. F. Frangi, D. Rueckert, J. A. Schnabel, and W. J. Niessen. Automatic construction of multiple-object three-dimensional statistical shape models: Application to cardiac modelling. In *Information Processing in Medical Imaging: Proc. 17th International Conference (IPMI'01)*, number 2082 in Lecture Notes in Computer Science, pages 78–91, Davis, CA, 2001.

[64] A. F. Frangi, D. Rueckert, J. A. Schnabel, and W. J. Niessen. Automatic construction of multiple-object three-dimensional statistical shape models: Application to cardiac modeling. *IEEE Transaction on Medical Imaging*, 21(9):1151–1165, 2002.

[65] P. A. Freeborough and N. C. Fox. Modeling brain deformations in Alzheimer disease by fluid registration of serial 3D MR images. *Journal of Computer Assisted Tomography*, 22(5):838–843, 1998.

[66] P. A. Freeborough, R. P. Woods, and N. C. Fox. Accurate registration of serial 3D MR brain images and its application to visualizing change in neurodegenerative disorders. *Journal of Computer Assisted Tomography*, 20(6):1012–1022, 1996.

[67] J. C. Gee and R. K. Bajcsy. Elastic matching: Continuum mechanical and probabilistic analysis. In A. W. Toga, editor, *Brain Warping*, pages 183–197. Academic Press, 1999.

[68] A. Goshtasby. Registration of images with geometric distortions. *IEEE Transactions on Geoscience and Remote Sensing*, 26(1):60–64, 1988.

[69] J. C. Gower. Generalized procrustes analysis. *Phychometrika*, 40:33–50, 1975.

[70] G. Grenander and M. I. Miller. Computational anatomy: An emerging discipline. *Quarterly of Applied Mathematics*, 56(4):617–694, 1998.

[71] A. Guimond, J. Meunier, and J. P. Thirion. Average brain models: A convergence study. *Computer Vision and Image Understanding*, 77(9):192–210, 2000.

[72] T. Gustavsson, R. Paschel, and K. Caidahl. Model based dynamic 3D reconstruction and display of the left ventricle from 2D cross-sectional echocardiograms. *Computerized Medical Imaging and Graphics*, 17(4-5):273–278, 1993.

[73] E. Haber, D. Metaxas, and L. Axel. Three-dimensional geometric modeling of cardiac right and left ventricles. In *Biomedical Engineering Society Annual Meeting*, San Diego, CA, 1997.

[74] J. V. Hajnal, D. L. G. Hill, and D. J. Hawkes. *Medical Image Registration*. CRC Press, 2001.

[75] G. Hamarneh and T. Gustavsson. Deformable spatio-temporal shape models: Extending ASM to 2D+time. In *12th British Machine Vision Conference*, volume 1, pages 13–22, 2001.

[76] G. T. Herman, J. Zheng, and C. A. Bucholtz. Shape-based interpolation. *IEEE Computer Graphics and Applications*, 12(3):69–79, 1992.

[77] D. L. G. Hill, P. G. Batchelor, M. Holden, and D. J. Hawkes. Medical image registration. *Physics in Medicine and Biology*, 46(3):R1–R45, 2001.

[78] D. L. G. Hill, J. V. Hajnal, S. M. Smith, T. Hartkens, and K. McLeish. A dynamic brain atlas. In *Fifth Int. Conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI '02)*, number 2488 in Lecture Notes in Computer Science, pages 532–539, Tokyo, Japan, 2002.

[79] D. L. G. Hill and D. J. Hawkes. Medical image registration using knowledge of adjacency of anatomical structures. *Image and Vision Computing*, 12(3):173–178, 1994.

[80] C. K. Hoh, M. Dahlbom, G. Harris, Y. Choi, R. A. Hawkins, M. E. Philips, and J. Maddahi. Automated iterative three-dimensional registration of positron emission tomography images. *Journal of Nuclear Medicine*, 34(11):2009–2018, 1993.

[81] M. Holden, D. L. G. Hill, E. R. E. Denton, J. M. Jarosz, T. C. S. Cox, T. Rohlfing, J. Goodey, and D. J. Hawkes. Voxel similarity measures for 3D serial MR brain image registration. *IEEE Transactions on Medical Imaging*, 19(2):94–102, February 2000.

[82] P. Horkaew. *Statistical modelling of complex topological shapes with application to cardiovascular imaging*. PhD thesis, Imperial College London, 2004.

[83] P. Horkaew and G. Z. Yang. Optimal deformable surface models for 3D medical image analysis. In *Information Processing in Medical Imaging: Proc. 18th International Conference (IPMI'03)*, number 2732 in Lecture Notes in Computer Science, pages 13–24. Springer, 2003.

[84] P. Horkaew and G. Z. Yang. Construction of 3D dynamic statistical models for complex topological shapes. In *Seventh Int. Conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI '04)*, number 3216 in Lecture Notes in Computer Science, pages 217–224. Springer, 2004.

[85] A. Hyvärinen, J. Karhunen, and E. Oja. *Independent Component Analysis*. John Wiley and Sons, 2001.

[86] A. Hyvärinen and E. Oja. A fast fixed algorithm for independent component analysis. *Neural Computation*, 9(7):1483–1492, 1997.

[87] R. A. Johnson and D. W. Wichern. *Multivariate Statistics: A Practical Approach*. Chapman and Hall, 1988.

[88] S. C. Joshi. *Large deformation diffeomorphisms and Gaussian random fields for statistical characterization of brain sub-manifolds*. PhD thesis, Washington University, 1998.

[89] M. Kachelriess, D. A. Sennst, W. Maxlmoser, and W. A. Kalender. Kymogram detection and kymogram-correlated image reconstruction from subsecond spiral computed tomography of the heart. *Medical Physics*, 29(7):1489–1555, 2002.

[90] J. Kittler and J. Illingworth. Relaxation labeling algorithms - a review. *Image and vision computing*, 3(4):206–216, 1986.

[91] R. E. Klabunde. *Cardiovascular Physiology Concepts*. Lippincott Williams and Wilkins, 2004.

[92] G. J. Klein, B. W. Reutter, and R. H. Huesman. 4D affine registration models for respiratory-gated PET. In *IEEE Nuclear Science Symposium Conference Record 2000*, Lyon France, 2000.

[93] G. J. Klein, B. W. Reutter, and R. H. Huesman. Four dimensional affine registration models for respiratory-gated PET. *IEEE Transactions on Nuclear Science*, 48(3):756–760, 2001.

[94] R. M. Lapp. Four dimensional cardiac segmentation. Master's thesis, Imperial College London, UK, Sep 2003.

[95] R. M. Lapp, M. Lorezo-Valdés, and D. Rueckert. 3D/4D cardiac segmentation using active appearance models, non-rigid registration, and the insight toolkit. In

Seventh Int. Conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI '04), number 3216 in Lecture Notes in Computer Science, pages 419–426. Springer, 2004.

[96] A. C. Larson, R. D. White, G. Laub, E. R. McVeigh, D. Li, and O. P. Simonetti. Self-gated cardiac cine MRI. *Magnetic Resonance in Medicine*, 51(1):93–102, 2004.

[97] P. C. Lauterbur. Image formation by induced local interactions: Examples employing nuclear magnetic resonance. *Nature*, 242:190–191, 1973.

[98] M. J. Ledesma-Carbayo, J. Kybic, M. Sühling, P. Hunziker, M. Desco, A. Santos, and M. Unser. Cardiac ultrasound motion detection by elastic registration exploiting temporal coherence. In *2004 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, 2004.

[99] S. Lee, G. Wolberg, K.Y. Chwa, and S. Y. Shin. Image metamorphosis with scattered feature constraints. *IEEE Transactions on Visualization and Computer Graphics*, 2(4):337–354, 1996.

[100] S. Lee, G. Wolberg, and S. Y. Shin. Scattered data interpolation with multilevel B-splines. *IEEE Transactions on Visualization and Computer Graphics*, 3(3):228–244, 1997.

[101] T. M. Lehmann, C. Gönner, and K. Spitzer. Survey: Interpolation methods in medical image processing. *IEEE Transactions on Medical Imaging*, 18(11):1049–1075, 1999.

[102] B. P. F. Lelieveldt, S. C. Mitchell, J. G. Bosch, R. J. van der Geest, M. Sonka, and J. H. C. Reiber. Time-continuous segmentation of cardiac image sequences using active appearance motion models. In *Information Processing in Medical Imaging: Proc. 17th International Conference (IPMI'01)*, number 2082 in Lecture Notes in Computer Science, pages 446–452. Springer, 2001.

[103] L. Lemieux, U. C. Wieshmann, N. F. Moran, D. R. Fish, and S. D. Shorvon. The detection and significance of subtle changes in mixed-signal brain lesions by serial MRI scan matching and spatial normalization. *Medical Image Analysis*, 2(3):227–242, 1998.

[104] H. Lester and S. Arridge. A survey of hierarchical non-linear medical image registration. *Pattern Recognition*, 32(1):129–149, 1999.

[105] D. N. Levin, C. A. Pelizzari, G. T. Chen, C. T. Chen, and M. D. Cooper. Retrospective geometric correlation of MR, CT and PET images. *Radiology*, 169:817–823, 1988.

[106] D. Li, P. Dhawale, P. J. Rubin, E. M. Haacke, and R. J. Gropler. Myocardial signal response to dipyridamole and dobutamine: demonstration of the BOLD effect using a double-echo gradient-echo sequence. *Magnetic Resonance in Medicine*, 36(1):16–20, 1996.

[107] W. E. Lorensen and H. E. Cline. Marching cubes: a high resolution 3D surface reconstruction algorithm. In *Computer Graphics: SIGGRAPH'87 Conference Proceedings*, volume 21, pages 163–169, 1987.

[108] M. Lorenzo-Valdés, G. I. Sanchez-Ortiz, A. Elkington, R. Mohiaddin, and D. Rueckert. Segmentation of 4D cardiac MR images using a probabilistic atlas and the EM algorithm. *Medical Image Analysis*, 8(3):255–265, 2004.

[109] M. Lorenzo-Valdés, G. I. Sanchez-Ortiz, R. Mohiaddin, and D. Rueckert. Atlas based segmentation and tracking of 3D cardiac MR images using non rigid registration. In *Fifth Int. Conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI '02)*, number 2488 in Lecture Notes in Computer Science, pages 642–650. Springer, 2002.

[110] M. Lorenzo-Valdés, G. I. Sanchez-Ortiz, R. Mohiaddin, and D. Rueckert. Segmentation of 4D cardiac MR images using a probabilistic atlas and the EM algorithm.

In *Sixth Int. Conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI '03)*, number 2878 in Lecture Notes in Computer Science. Springer, 2003.

[111] J. Lötjönen, S. Kivistö, J. Koikkalainen, D. Smutek, and K. Lauerma. Statistical shape model of atria, ventricles and epicardium from short- and long-axis MR images. *Medical Image Analysis*, 8(3):371–386, 2004.

[112] J. Lötjönen, J. Koikkalainen, D. Smutek, S. Kivistö, and K. Lauerma. Four-chamber 3D statistical shape model from cardiac short-axis MR images. In *Sixth Int. Conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI '03)*, number 2878 in Lecture Notes in Computer Science, pages 491–498. Springer, 2003.

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

[114] J. B. A. Maintz and M. A. Viergever. A survey of medical image registration. *Medical Image Analysis*, 2(1):1–37, 1998.

[115] T. Mäkelä, P. Clarysse, J. Lötjönen, O. Sipilä, K. Lauerma, H. Hännimen, E. P. Pyökkimies, J. Nenonen, J. Knuuti, T. Katila, and I. E. Magnin. A new method for the registration of cardiac PET and MR images using deformable model based segmentation of the main thorax structures. In *Fourth Int. Conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI '01)*, number 2208 in Lecture Notes in Computer Science. Springer, 2001.

[116] T. Mäkelä, P. Clarysse, O. Sipilä, N. Pauna, and Q. C. Pham. A review of cardiac image registration methods. *IEEE Transactions on Medical Imaging*, 21(9):1011–1021, 2002.

[117] W. J. Manning and D. J. Pennel. *Cardiovascular Magnetic Resonance*. Churchill Livingstone, 2002.

[118] P. Mansfield and I. L. Pykett. Biological and medical imaging by NMR. *Journal of Magnetic Resonance*, 29:355–373, 1978.

[119] J. Mazziotta, A. Toga, A. Evans, P. Fox, J. Lancaster, R. Zilles, K. Woods, T. Paus, G. Simpson, B. Pike, C. Holmes, L. Collins, P. Thompson, D. MacDonald, M. Iacoboni, T. Schormann, K. Amunts, N. Palomero-Gallagher, S. Geyer, L. Parson, K. Narr, N. Kabani, G. Le Goualher, J. Feidler, K. Smith, D. Boomsma, P. H. Hulshoff, T. Cannon, R. Kawashima, and B. Mazoyer. A four dimensional probabilistic atlas of the human brain. *Journal of the American Medical Informatics Association*, 8(5), 2001.

[120] J. Mazziotta, A. Toga, A. Evans, P. Fox, J. Lancaster, R. Zilles, K. Woods, T. Paus, G. Simpson, B. Pike, C. Holmes, L. Collins, P. Thompson, D. MacDonald, M. Iacoboni, T. Schormann, K. Amunts, N. Palomero-Gallagher, S. Geyer, L. Parson, K. Narr, N. Kabani, G. Le Goualher, D. Boomsma, Cannon T., R. Kawashima, and B. Mazoyer. A probabilistic atlas and reference system for the human brain: International consortium for brain mapping (ICBM). *The Royal Society*, 356:1293–1322, 2001.

[121] J. C. Mazziotta, A. W. Toga, A. Evans, P. Fox, and J. Lancaster. A probabilistic atlas of the human brain: Theory and rationale for its development. The international consortium for brain mapping. *Neuroimage*, 2:89–101, 1995.

[122] K. McLeish. *Combining data acquisition and post-processing techniques for magnetic resonance of moving objects*. PhD thesis, University of London, April 2003.

[123] K. McLeish, D. L. G. Hill, D. Atkinson, J. M. Blackall, and R. Razavi. A study of the motion and deformation of the heart due to respiration. *IEEE Transactions on Medical Imaging*, 21(9):1141–1150, 2002.

[124] J. Meinguet. Multivariate interpolation at arbitrary point made simple. *Zeitschrift für angewandte Mathematik und Physik*, 30:292–304, 1979.

[125] D. Metaxas and D. Terzopoulos. Shape and nonrigid motion estimation through physics-based synthesis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(3):580–591, 1993.

[126] S. C. Mitchell, J. G. Bosch, B. P. F. Lelieveldt, R. J. van der Geest, J. H. C. Reiber, and M. Sonka. 3D active appearance models: Segmentation of cardiac MR and ultrasound images. *IEEE Transactions on Medical Imaging*, 21(9):1167–1178, 2002.

[127] S. C. Mitchell, B. P. F. Lelieveldt, R. J. van der Geest, J. G. Bosch, J. H. C. Reiber, and M. Sonka. Time-continuous segmentation of cardiac image sequences using active appearance motion models. In *SPIE Medical Imaging*, volume 4322, pages 249–256, 2001.

[128] S. C. Mitchell, B. P. F. Lilieveldt, R. J. van der Geest, H. G. Bosch, J. H. C. Reiber, and M. Sonka. Multistage hybrid active appearance model matching: Segmentation of left and right ventricles in cardiac MR images. *IEEE Transactions on Medical Imaging*, 20(5):415–523, 2001.

[129] T. J. Mosher and M. B. Smith. A DANTE tagging sequence for the evaluation of translational sample motion. *Magnetic Resonance in Medicine*, 15(2):334–339, 1990.

[130] N. M. I. Noble, D. L. G. Hill, M. Beeuwer, J. A. Schnabel, D. J. Hawkes, F. A. Gerritsen, and R. Razavi. Myocardial delineation via registration in a polar coordinate system. In *Fifth Int. Conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI '02)*, number 2488 in Lecture Notes in Computer Science, pages 651–658. Springer, 2002.

[131] W. L. Nowinski and A. Thirunavuukarasuu. Atlas-assisted localization analysis of functional images. *Medical Image Analysis*, 5(3):201–220, 2001.

[132] W. G. O'Dell and A. D. McCulloch. Imaging three dimensional cardiac function. *Annual Review of Biomedical Engineering*, 2:431–456, 2000.

[133] World Health Organisation. The world health report 2003. http://www.who.int/whr/2003/en/, 2003.

[134] S. Pallotta, M. C. Gillardi, V. Bettinardi, G. Rizzo, C. Landoni, G. Striano, R. Masi, and F. Fazio. Application of surface matching image registration technique to the correlation of cardiac studies in positron emission tomography by transmission images. *Physics in Medicine and Biology*, 40:1695–1708, 1995.

[135] X. Papademetris. *Estimation of 3D Left Ventricular Deformation from Medical Images Using Biomechanical Models*. PhD thesis, Yale University, 2000.

[136] X. Papademetris, P. Shi, D. P. Dione, A. J. Sinusas, R. T. Constable, and J. S. Duncan. Recovery of soft tissue object deformation from 3D image sequences using biomechanical models. In *Information Processing in Medical Imaging: Proc. 16th International Conference (IPMI'99)*, number 1613 in Lecture Notes in Computer Science, pages 352–57. Springer-Verlag, 1999.

[137] X. Papademetris, A. J. Sinusas, D. P. Dione, T. R. Constable, and J. S. Duncan. Estimation of 3D left ventricular deformation from medical images using biomechanical models. *IEEE Transactions on Medical Imaging*, 21(7):786–800, 2002.

[138] X. Papademetris, A. J. Sinusas, D. P. Dione, and J. S. Duncan. 3D cardiac deformation from ultrasound images. In *Second Int. Conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI '99)*, number 1679 in Lecture Notes in Computer Science, pages 420–429. Springer, 1999.

[139] T. Papatheodorou and D. Rueckert. Evaluation of automatic 4D face recognition

using surface and texture registration. In *In Sixth IEEE International Conference on Automatic Face and Gesture Recognition*, pages 321–326, 2004.

[140] H. Park, P. H. Bland, and C. R. Meyer. Construction of an abdominal probabilistic atlas and its application in segmentation. *IEEE Transactions on Medical Imaging*, 22(4):483–492, 2003.

[141] J. Park, D. N. Metaxas, A. A. Young, and L. Axel. Deformable models with parameter functions for cardiac motion analysis from tagged MRI data. *IEEE Transactions on Medical Imaging*, 15(3):278–289, 1996.

[142] N. J. Pelc, R. J. Herfkens, A. Shimakawa, and D. Enzmann. Phase contrast cine magnetic resonance imaging. *Magnetic Resonance Quarterly*, 7(4):229–254, 1991.

[143] C. A. Pelizzari, G. T. Chen, D. R. Spelbring, R. R. Weichselbaum, and C. T. Chen. Accurate three dimensional registration of CT, PET, and/or MR images of the brain. *Journal of Computer Assisted Tomography*, 13(1):20–26, 1989.

[144] G. P. Penney, J. Weese, J. A. Little, P. Desmedt, D. L. G. Hill, and D. J. Hawkes. A comparison of similarity measures for use in 2D-3D medical image registration. *IEEE Transactions on Medical Imaging*, 17(4):586–595, 1998.

[145] D. Perperidis, M. Lorenzo-Valdés, R. Chandrashekara, A. Rao, R. Mohiaddin, G. I. Sanchez-Ortiz, and D. Rueckert. Building a 4D atlas of the cardiac anatomy and motion using MR imaging. In *2004 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, 2004.

[146] D. Perperidis, M. Lorenzo-Valdés, R. Chandrashekara, A. Rao, R. Mohiaddin, G. I. Sanchez-Ortiz, and D. Rueckert. Building a 4D atlas of the cardiac anatomy and motion using MR imaging. In *Medical Image Understanding and Analysis (MIUA 2004)*, pages 9–12, 2004.

[147] D. Perperidis, R. Mohiaddin, and D. Rueckert. Spatio-temporal free-form registration of cardiac MR image sequences. In *Seventh Int. Conf. on Medical Image*

Computing and Computer-Assisted Intervention (MICCAI '04), number 3216 in Lecture Notes in Computer Science, pages 911–919. Springer, 2004.

[148] D. Perperidis, R. Mohiaddin, and D. Rueckert. Spatio-temporal free-form registration of cardiac MR image sequences. In *Medical Image Understanding and Analysis (MIUA 2004)*, pages 157–160, 2004.

[149] D. Perperidis, R. Mohiaddin, and D. Rueckert. Construction of a 4D statistical atlas of the cardiac anatomy and its use in classification. In *Eighth Int. Conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI '05)*, Lecture Notes in Computer Science, 2005.

[150] D. Perperidis, R. Mohiaddin, and D. Rueckert. Fast spatio-temporal registration of cardiac MR image sequences. In *Functional Imaging and Modeling of the Heart,(FIMH'05)*, number 3504 in Lecture Notes in Computer Science, pages 414–424. Springer, 2005.

[151] D. Perperidis, R. Mohiaddin, and D. Rueckert. Spatio-temporal free-form registration of cardiac MR image sequences. *Medical Image Analysis*, 9(5):441–456, 2005.

[152] D. Perperidis, A. Rao, M. Lorenzo-Valdés, R. Mohiaddin, and D. Rueckert. Spatio-temporal alignment of 4D cardiac MR images. In *Functional Imaging and Modeling of the Heart, FIMH'03*, number 2674 in Lecture Notes in Computer Science, Lyon, France, June 5-6, 2003. Springer.

[153] D. Perperidis, A. Rao, R. Mohiaddin, and D. Rueckert. Non-rigid spatio-temporal alignment of 4D cardiac MR images. In *Second International Workshop on Biomedical Image Registration (WBIR '03)*, number 2717 in Lecture Notes in Computer Science. Springer, 2003.

[154] C. Petitjean, N. Rougon, and F. Prêteux. Building and using a statistical 3D motion atlas for analyzing myocardial contraction in MRI. In *Proceedings SPIE Confer-*

ence on Image Processing - SPIE International Symposium Medical Imaging'04, San Diego, CA, volume 5370, 14-19 February 2004.

[155] S. Plein, T. N. Bloomer, J. P. Ridgway, T. R. Jones, G. J. Bainbridge, and M. U. Sivananthan. Steady-state free precession magnetic resonance imaging of the heart: comparison with segmented k-space gradient-echo imaging. *Journal of Magnetic Resonance Imaging*, 14(3):230–236, 2001.

[156] S. Plein, W. H. T. Smith, P. J. Ridgway, A. Kassner, D. J Beacock, T. N. Bloomer, and M. U. Sivananthan. Measurements of left ventricular dimensions using real time acquisition in cardiac magnetic resonance imaging: comparison with conventional gradient echo imaging. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 13(2):101–108, 2001.

[157] J. Pluim. *Mutual information based registration of medical images*. PhD thesis, Utrecht University, 2000.

[158] J. P. W. Pluim, J. B. A. Maintz, and M. A. Viergever. Image registration by maximization of combined mutual information and gradient information. *IEEE Transactions on Medical Imaging*, 19(8):809–814, 2000.

[159] J. P. W. Pluim, J. B. A Maintz, and M. A. Viergever. Interpolation artifacts in mutual information-based image registration. *Computer Vision and Image Understanding*, 77(2):211–232, 2000.

[160] J. P. W. Pluim, J. B. A. Maintz, and M. A. Viergever. Mutual-information-based registration of medical images:a survey. *IEEE Transactions on Medical Imaging*, 22(8):986–1004, 2003.

[161] P. V. Prasad, D. Burstein, and R. R. Edelman. MRI evaluation of myocardial perfusion without a contrast agent using magnetization transfer. *Magnetic Resonance in Medicine*, 30(2):267–270, 1993.

[162] W. K. Pratt. Correlation techniques of image registration. *IEEE Transactions on Aerospace and Electronic Systems*, 10(3):353–358, 1974.

[163] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, Cambridge, England, 2 edition, 1992.

[164] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger. SENSE: Sensitivity encoding for fast MRI. *Magnetic Resonance in Medicine*, 42(5):952–962, 1999.

[165] A. Rao, R. Chandrashekara, G. I. Sanchez-Ortiz, M. Lorenzo-Valdés, R. Mohiaddin, and D. Rueckert. Construction of a cardiac motion atlas from MR using non-rigid registration. In *Functional Imaging and Modeling of the Heart, FIMH'03*, number 2674 in Lecture Notes in Computer Science. Springer, 2003.

[166] A. Rao, R. Chandrashekara, G. I. Sanchez-Ortiz, R. Mohiaddin, P. Aljabar, J. V. Hanjnal, B. K. Puri, and D. Rueckert. Spatial transformation of motion and deformation fields using non-rigid registration. *IEEE Transactions on Medical Imaging*, 23(9):1065–1076, 2004.

[167] A. Rao, G. I. Sanchez-Ortiz, R. Chandrashekara, M. Lorenzo-Valdés, R. Mohiaddin, and D. Rueckert. Comparison of cardiac motion across subjects using non-rigid registration. In *Fifth Int. Conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI '02)*, number 2488 in Lecture Notes in Computer Science, pages 722–729. Springer, 2002.

[168] P. E. Rasser, P. J. Johnston, P. B. Ward, and P. M. Thompson. A deformable brodmann areas atlas. In *2004 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, 2004.

[169] D. Rey, G. Subsol, H. Delingette, and N. Ayache. Automatic detection and segmentation of evolving processes in 3D medical images: Application to multiple

sclerosis. In *Information Processing in Medical Imaging: Proc. 16th International Conference (IPMI'99)*, number 1613 in Lecture Notes in Computer Science, pages 154–167, 1999.

[170] G. Rizzo, P. Pasquali, and M. Gilardi. Multimodal biomedical image registration: Use of a cross correlation technique. In *International Conference of the IEEE Engineering in Medicine and Biology Society*, volume 13, pages 219–220, 1991.

[171] D. Rueckert, M. J. Clarkson, D. L. G. Hill, and D. J. Hawkes. Non-rigid registration using high order mutual information. In *In Proc. SPIE Medical Imaging 2000: Image Processing*, pages 438–447, San Diego,CA, 2000.

[172] D. Rueckert, A. F. Frangi, and J. A. Schnabel. Automatic construction of 3D statistical deformation models of the brain using non-rigid registration. In *Fourth Int. Conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI '01)*, number 2208 in Lecture Notes in Computer Science, pages 74–84. Springer, 2001.

[173] D. Rueckert, A. F. Frangi, and J. A. Schnabel. Automatic construction of 3D statistical deformation models of the brain using non-rigid registration. *IEEE Transactions on Medical Imaging*, 22(8):1014–1025, 2003.

[174] D. Rueckert, L. I. Sonoda, C. Hayes, D. L. G. Hill, M. O. Leach, and D. J. Hawkes. Non-rigid registration using free-form deformations: Application to breast MR images. *IEEE Transactions on Medical Imaging*, 18(8):712–721, 1999.

[175] D. B. Russakoff, T. Rohling, A. Ho, D. H. Kim, R. Shahidi, J.R. Adler, and Jr Maurer, C. R. Evaluation of intensity-based 2D-3D spine image registration using clinical gold-standard data. In *Proc. 2nd International Workshop on Biomedical Image Registration (WBIR '03)*, number 2717 in Lecture Notes in Computer Science, pages 151–160, 2003.

[176] G. Schaltenbrand and P. Bailey. *Introduction to Stereotaxis with an atlas of the human brain*. Georg Thieme Verlag, 1959.

[177] J. A. Schnabel, D. Rueckert, M. Quist, J. M. Blackall, A. D. Castellano Smith, T. Hartkens, G. P. Penney, W. A. Hall, H. Liu, C. L. Truwit, F. A. Gerritsen, D. L. G. Hill, and D. J. Hawkes. A generic framework for non-rigid registration based on non-uniform multi-level free form deformations. In *Fourth Int. Conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI '01)*, number 2208 in Lecture Notes in Computer Science, pages 573–581. Springer, 2001.

[178] W. J. Schroeder, J. A. Zarge, and W. E. Lorensen. Decimation of triangle meshes. In *19th Annual Conference on Computer Graphics and Interactive Techniques*, volume 26, pages 65–70, 1992.

[179] T. W. Sederberg and S. R. Parry. Free-form deformation of solid geometric models. In *13th Annual Conference on Computer Graphics and Interactive Techniques*, pages 151–160. ACM Press, 1986.

[180] M. Sermesant, C. Forest, X. Pennec, H. Delingette, and N. Ayache. Biomechanical model construction from different modalities: Application to cardiac images. In *Fifth Int. Conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI '02)*, number 2488 in Lecture Notes In Computer Science, pages 714–721. Springer, 2002.

[181] M. Sermesant, C. Forest, X. Pennec, H. Delingette, and N. Ayache. Deformable biomechanical models: Application to 4D cardiac image analysis. *Medical Image Analysis*, 7(4):475–488, 2003.

[182] P. C. Shi, A. J. Sinusas, R. T. Constable, E. Ritman, and J. S. Duncan. Point-tracked quantitative analysis of left ventricular surface motion from 3D image sequences: Algorithms and validation. *IEEE Transactions on Medical Imaging*, 19(1):36–50, 2000.

[183] P. C. Shi, A. J. Sinusas, T. Constable, and J. S. Duncan. Volumetric deformation analysis using mechanics-based data fusion: Applications in cardiac motion recovery. *International Journal of Computer Vision*, 35(1):87–107, 1999.

[184] D. K. Sodickson and W. J. Manning. Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays. *Magnetic Resonance in Medicine*, 38(4):591–603, 1997.

[185] M. Sonka, B. P. F. Lelieveldt, S. C. Mitchell, J. G. Bosch, R. J. van der Geest, and J. H. Reiber. Active appearance motion model segmentation. In *2nd International Workshop on Digital and Computational Video*, pages 64–68, 2001.

[186] C. Studholme. *Measures of 3D Medical Image Alignment*. PhD thesis, United Medical and Dental Schools of Guy's and St. Thomas' Hospitals,, 1997.

[187] C. Studholme, D. L. G. Hill, and D. J. Hawkes. An overlap invariant entropy measure of 3D medical image alignment. *Pattern Recognition*, 32(1):71–86, 1999.

[188] P. Suetens. *Fundamentals of Medical Imaging*. Cambridge University Press, 2002.

[189] J. Talairach and G. Szikla. *Atlas d' Anatomie stereotaxique du Telencephalic: Etudes Anatomo-Radiologiques*. Masson and Cie, Paris, 1967.

[190] J. Talairach and P. Tournoux. *Co-Planar Stereotactic atlas of the human brain: 3-Dimensional proportional system: An approach to Cerebral imaging*. Georg Thieme Verlag, 1988.

[191] P. M. Thompson and A. W. Toga. Detection, visualization and animation of abnormal anatomic structure with a deformable probabilistic atlas based on random vector field transformations. *Medical Image Analysis*, 1(4):271–294, 1997.

[192] A. W. Toga and P. M. Thompson. The role of image registration in brain mapping. *Image and Vision Computing*, 19:3–24 Sp. Iss., 2001.

[193] H. K. Tu, A. Matheny, D. B. Goldgof, and H. Bunke. Left ventricular boundary detection from spatio-temporal volumetric computed tomography images. *Computerized Medical Imaging and Graphics*, 19(1):27–46, 1995.

[194] T. G. Turkington, T. R. DeGradom, M. W. Hanson, and E. R. Coleman. Alignment of dynamic cardiac PET images for correction of motion. *IEEE Transactions on Nuclear Science*, 44(2):235–242, 1997.

[195] R. Underwood and D. Firmin. *Magnetic resonance of the cardiovascular system.* Blackwell Scientific Publications, 1991.

[196] M. Unser. Splines: A perfect fit for signal and image processing. *IEEE Signal Processing magazine*, 16(6):22–38, 1999.

[197] M. Üzümcü, A. F. Frangi, M. Sonka, J. H. C. Reiber, and B. P. F. Lelieveldt. ICA vs. PCA active appearance models: Application to cardiac MR segmentation. In *Sixth Int. Conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI '03)*, number 2878 in Lecture Notes in Computer Science, pages 451–458. Springer, 2003.

[198] P. Viola and W. M. Wells. Alignment by maximisation of mutual information. *International Journal of Computer Vision*, 24(2):137–154, 1997.

[199] P. A. Viola. *Alignment by maximization of mutual information*. PhD thesis, Massachusetts Institute of Technology, 1995.

[200] C. Vuille and A. E. Wayman. *Principles and practice of echocardiography*, chapter Left ventricle I: General considerations, assessment of chamber size and function. Lea and Fabiger, Philadelphia, PA, 2nd edition, 1994.

[201] G. Wahba. Spline models for observational data. *Society for Industrial and Applied Mathematics*, 1990.

[202] L. Wang, S. C. Joshi, M. I. Miller, and J. G. Cseransky. Statistical analysis of hippocampal asymmetry in schizophrenia. *NeuroImage*, 14(3):531–545, 2001.

[203] S. Webb. *The Physics of Medical Imaging*. Bristol; Philadelphia : Hilger, 1988.

[204] W. M. Wells, P. Viola, H. Atsumi, and S. Nakajima. Multi-modal volume registration by maximization of mutual information. *Medical Image Analysis*, 1(1):33–51, 1996.

[205] D. S. Williams, J. A. Detre, J. S. Leigh, and A. P. Koretsky. Magnetic resonance imaging of perfusion using spin inversion of arterial water. In *Proceedings of the National Academy of Sciences of the United States of America*, volume 89, pages 212–126, 1992.

[206] R. P. Woods. *Handbook of Medical Imaging*, pages 465–497. Academic Press, 2000.

[207] E. A. Zerhouni, D. M. Parish, W. J. Rogers, A. Yang, and E. P. Shapiro. Human heart:tagging with MR imaging - a method for noninvasive assessment of myocardial motion. *Radiology*, 169:59, 1988.

[208] B. Zitová and J. Flusser. Image registration methods: a survey. *Image and Vision Computing*, 21:977–1000, 2003.