Abstract

Myocardial perfusion is an indicator of heart health that can be used to locate damaged areas of the heart. It can be measured by an imaging modality such as magnetic resonance imaging (MRI).

When measuring myocardial perfusion in a living human we must align images over time to correct for motion caused by breathing. In this thesis we will introduce a version of the normalized gradient fields similarity method, modified to improve noise robustness and decrease computational complexity. Furthermore we will use a segmentation in the reference time-frame as *a priori* information.

The results of this thesis is the development of a registration method that aligns the myocardium in perfusion MRI images. The myocardial alignment, measured by the Dice coefficient, increased from 0.79 ± 0.11 before registration to 0.89 ± 0.028 after registration. The resulting perfusion curves looks like we would expect from previous studies of myocardial perfusion.

Acknowledgements

Firstly i would like to thank my supervisor Helen Soneson for her help and patience. I would also like to thank Robert Jablonowski for performing manual segmentations of the left ventricle. Finally i would like to thank everyone else at the Cardiac MR group at Lund's university who contributed through discussions.

Contents

1 Introduction **5**

2 Background **7**
2.1 The human heart . . . 7
2.1.1 Myocardial Perfusion . . . 7
2.2 Magnetic resonance imaging . . . 8
2.2.1 Contrast enhanced MRI . . . 8
2.3 Registration . . . 8
2.3.1 Feature and intensity based methods . . . 10
2.4 The software Segment . . . 10

3 Aim **11**

4 Theory **13**
4.1 Image transformations . . . 13
4.1.1 Affine transformations . . . 14
4.1.2 Thin plate spline . . . 14
4.2 Cost functions . . . 15
4.2.1 Sum squared difference . . . 15
4.2.2 Normalized gradient fields . . . 15
4.2.3 Noise estimation . . . 16
4.2.4 Regularization . . . 17
4.2.5 Weights . . . 17
4.3 Scale spaces . . . 17
4.4 Optimization . . . 18
4.4.1 Downhill simplex . . . 18
4.4.2 Simulated annealing . . . 19
4.5 Validation . . . 19
4.6 Consistency . . . 21
4.7 Statistical analysis . . . 21

5 Results **23**
5.1 Validation . . . 23
5.2 Consistency . . . 23
5.3 Effects of segmentation . . . 23
5.3.1 Cost-function validation . . . 26
5.3.2 Improvement from initial alignment . . . 26
5.4 Perfusion . . . 26

5.5 Cost function comparisons 26

6 Discussion and Conclusions **31**
6.1 Registration and segmentation 31
6.2 Limitations 31
6.3 Noise Estimation 33
6.4 Implementation 33

Bibliography **35**

Chapter 1

Introduction

In the process of medical imaging, images and data of the interior of the human body are created by one of the various imaging modalities, such as ultrasound, magnetic resonance imaging (MRI) and radiography i.e x-ray. This data can then be used for clinical diagnosis as well as medical research. Depending on the application and imaging modality, the data can be in 2, 3 or 4 dimensions. When imaging a living person the subject should preferably hold still during the process. This is especially important if we want to track changes over time in time-resolved data.

In this thesis we will study the problem that arises when movement cannot be avoided in one particular imaging application, myocardial perfusion imaging by MRI. In this case we want to track intensity changes in images the of heart in a living human over several minutes. The problem to solve is correcting for the breathing motions that will cause movement and deformation of the heart. A solution to this is to align the image in every time frame to a common coordinate system where the heart does not move. To achieve this we need a model to describe how the heart moves and a method to estimate these movements.

Chapter 2

Background

2.1 The human heart

The heart is the organ that provides circulation of the blood in the body. Figure 2.1 shows an image of a human heart imaged with MRI. The heart is divided into four chambers, the right atrium and ventricle, and the left atrium and ventricle. De-oxygenated blood enters through the right atrium and is then pumped to the lungs via the right ventricle. The arterial (oxygenated) blood from the lungs then enters the left atrium and flows back to the body via the left ventricle. The muscle surrounding the chambers is called the myocardium. The innermost layer of the myocardium is called endocardium while the outside is called epicardium.

Figure 2.1: MRI image of a human heart where the myocardium of the left ventricle (LV) has be segmented. The right ventricle (RV) is also marked.

2.1.1 Myocardial Perfusion

Perfusion is the process where arterial blood is delivered to tissues though the smallest blood vessels of the body, the capillaries. In this thesis we are specifically interested in myocardial perfusion of the left ventricle. Myocardial perfusion is an important indicator of heart health and function. A region with decreased perfusion is called a perfusion defect. Perfusion defects is caused by an occlusion of the coronary arteries or a congenial heart defect.

2.2 Magnetic resonance imaging

Magnetic resonance imaging (MRI) is a technique commonly used for medical imaging i.e. to produce images of the inside of the human body. The first MRI image was published in 1974 by Paul Lauterbur. Since then, MRI has become an important tool in medicine and research. To acquire a MR image, the patient is placed inside the MR scanners powerful magnetic field, usually 1.5 or 3 Tesla. The human body consist mostly of water in different concentrations in different tissue types. A portion of the hydrogen atoms in the water molecules aligns their magnetization in the direction of the magnetic field. An electromagnetic field is then turned of and on with a resonance frequency that causes protons to adsorb energy from the field and flip its spin when turned on. When the field turns off again the proton will realign with the field i.e relax and emit an electromagnetic signal which can be detected by the scanner. Protons in different tissue types have distinct relaxation times allowing us to differentiate among them. The image gathered is a cross section slice of the body, most often a short-axis image, as in Figure 2.1. The short-axis image plane is perpendicular to the long-axis of the left ventricle. By acquiring several slices, a 3-D image stack covering the whole heart can be generated.

2.2.1 Contrast enhanced MRI

To quantitatively study perfusion of the heart a contrast agent, usually gadolinium, is injected intravenously before using the MR scanner. The agent decreases relaxation time where it is present, thus increases the image intensity in proportion to its concentration. Then several cross-sectional short-axis images is gathered every heartbeat for 2-3 minutes. To minimize movement due to the heart beating, cardiac gating is used. Cardiac gating is a technique that synchronize images over the heart cycle by ECG triggering. By identifying regions in the left myocardium across all relevant time-frames, perfusion can be measured by tracking the changes in intensity over time in each region. To locate perfusion defects, the myocardium is divided into segments, and perfusion is measured separately in each. The American Hearth Association (AHA) recommends a standardized 17 segment model[1] based on three cross-sectional images, so that results can be compared across different studies. The division of the left ventricle into segments are illustrated in Figure 2.2, and Figure 2.3 illustrates the corresponding three slices by contrast enhanced MRI. To identify these segments we must identify the borders of the left ventricle myocardium, the epicardium and endocardium. This can be achieved by segmenting every time-frame separately either automatically or manually. Alternatively segmentation is only performed in one reference time-frame and overlayed to all other frames by registration. The next section will present the later approach.

2.3 Registration

Registration is the process of aligning several data sets or images to a common coordinate system. Often one image is used as a reference, and thereafter the other images are aligned to the reference by applying a transformation. For registration to be successful we want the transformed image to be equivalent to the reference in some sense. At the same time we want to keep information that genuinely differs between the images. In this thesis we align the myocardium in the images in different

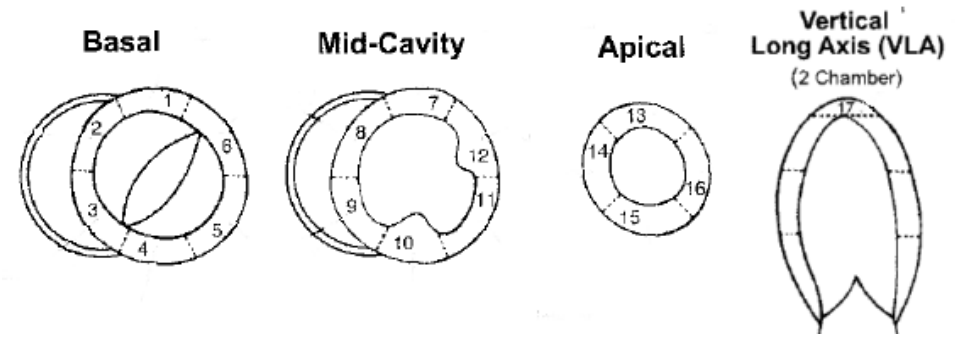

Figure 2.2: Illustration of the AHA segmentation model, in three short-axis slices and a long-axis slice.

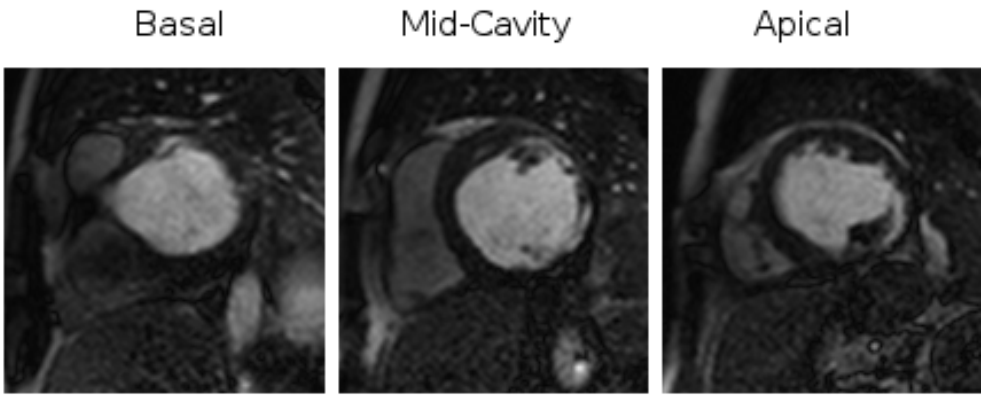

Figure 2.3: Three contrast enhanced MRI short-axis image slices from a patient.

time-frames to a reference time-frame. The reference and transformed image becomes similar in the sense that the myocardium is in the same position in the images. Once registered, a segmentation in the reference time frame will give a segmentation in every time frame, i.e the same segmentation as in the reference frame. The reverse is not true, a segmentation in every time-frame will not immediately yield a registration. For example; imagine the image and segmentation in Figure 2.1 rotated an unknown amount, we would know the location of the myocardium, but not how a part of it corresponds to the original image.

2.3.1 Feature and intensity based methods

There are two main types of image registration methods, feature based and intensity based. In the original feature based method, Scale Invariant Feature Transform (SIFT)[2] a feature vector is calculated for each image coordinate. The points with the most unique features in each image is selected. By matching these points feature vectors from one image to another we can find matches between the coordinates of the points. Using these coordinate matches and assuming a linear transformation between the images we get a system of linear equations. The equation system can be solved algebraically, yielding the parameters of the linear transformation.

This approach is most useful when we have different images of the same scene captured by the same sensor from different perspectives. If we transform the images to a common coordinate system they become the same in the sense that the matching points are in the same location. However some parts of the scene may be visible in one image but not in another.

While feature based methods can be used for deformation models, the number of good matches will limit the number of parameters that can be used in the transformation. If the system is under-determined we cannot uniquely solve it, and there is no guarantee of how many feature points matches we can find in a set of images.

Another option is intensity based methods. In these method a cost function is calculated based on the intensity differences between the images. The parameters of the transformation model can be estimated by numerically minimizing the differences between the reference and transformed image. After transformation the images will be similar in a sense defined by the cost function, which can be an advantage if we choose a cost function appropriate for our purposes.

2.4 The software Segment

Segment is a software package for Medical image analysis. Segment is developed by Medviso AB together with Lund Cardiac MR Group at Lund University. The registration algorithm developed in this thesis is implemented in Segment by using existing functions and add new features for studying MR perfusion images.

Chapter 3

Aim

The aim of this thesis is to develop a registration method suitable for time-resolved contrast enhanced MRI images used in perfusion analysis. The method should be fully automatic after giving a reference segmentation of the left ventricle in one time-frame as *a priori* information.

Chapter 4

Theory

In this chapter we will go though the theory needed for the image registration method. The following notation will be used,

- $\mathbf{x}$ bold lower case used for coordinate vectors
- $\mathbf{A}$ bold upper case used for matrices
- $I(\mathbf{x})$ used for images defined on $\Omega \in \mathcal{R}^2$

To register a whole set of MRI images, each short-axis slice will be processed separately. One time-frame will be chosen in advance as the reference. Thereafter, we can either register all other time-frames to the reference directly, or do it step by step and register each time-frame to the previously registered one, starting from the reference. When choosing a transformation and cost function for registration we need to keep the objective of aligning the myocardium in the reference and transformed image. Myocardial segmentations is only performed in the reference time-frame so we can not directly compare the overlap of reference and transformed image. Cost functions must instead rely on related image properties. Figure 4.1 shows three time-frames from one slice in a perfusion study, with segmentation only in the first image.

4.1 Image transformations

To register images we need an image transformation that aligns the base image frame, I_b, to the coordinate system of the reference image frame, I_r. The registered image is denoted as the target, I_t. The image transformation

$$T(I_b) = I_t : \Omega \rightarrow \Omega$$

is defined by the coordinate transformation function $f(\mathbf{x})$

$$I_t(f(\mathbf{x})) = I_b(\mathbf{x})$$

The coordinate transformation function must have sufficient degrees of freedom (DoF) to allow us to model changes in perspective, i.e movement of heart relative to the camera as well as deformation of the heart caused by the patient breathing, while remaining smooth.

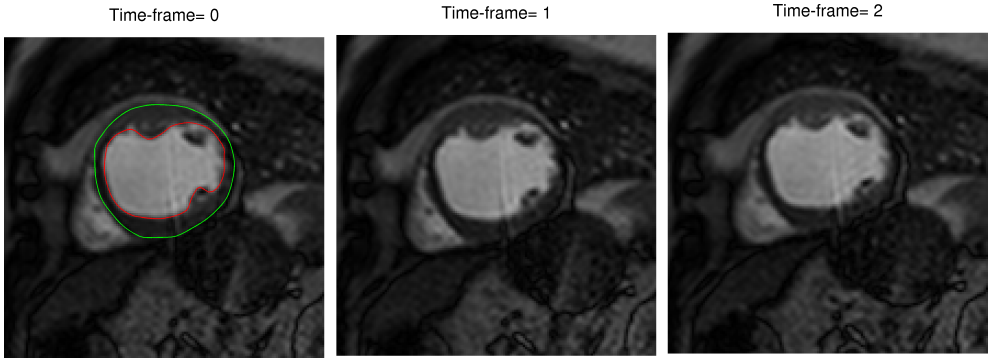

Figure 4.1: Three time-frame MR images from one patient, with segmentation in the first time-frame image.

4.1.1 Affine transformations

An affine transformation is a combination of a 2D linear transform and a translation, and is used in image registration to model small changes in perspective. The affine transformation is given by

$$f(\mathbf{x}) = \mathbf{A}\mathbf{x} + \begin{pmatrix} t_x \\ t_y \end{pmatrix}$$

Where $\mathbf{A}$ is a 2×2 linear transformation matrix and t_x and t_y represent the translation in x and y direction respectively. By using homogeneous coordinates the transform is expressed by

$$\begin{pmatrix} f(\mathbf{x}) \\ 1 \end{pmatrix} = \begin{pmatrix} \mathbf{A} & & t_x \\ & & t_y \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ 1 \end{pmatrix}$$

Affine transformations have six DoF, these are shearing i.e rotation of the xy-plane around the x axis, rotation around the z axis, and scaling and translation in the x and y directions respectively.

4.1.2 Thin plate spline

For deformation modeling the most commonly used transform is the Thin Plate Spline (TPS). The name is derived from the analogy of bending a thin sheet of metal, the sheet or image plane is bent in the z direction causing displacements in the x and y direction around a set of N control points $\mathbf{w_1}, ..., \mathbf{w_i}, ..., \mathbf{w_N}$, with corresponding weights $\mathbf{c_1}, ..., \mathbf{c_i}, ..., \mathbf{c_N}$. The TPS transform is given by

$$f(\mathbf{x}) = \sum_{i=1}^{K} \mathbf{c_i}\rho(||\mathbf{x} - \mathbf{w_i}||)$$

With the kernel $\rho(r) = r^2 log(r)$, where r is the euclidean distance between an image point and a control point. The transformation used in this thesis is a combination

between TPS and an affine transform. By the combination we get a transformation that can correct for both movement and deformation of the myocardium. The transformation is given by

$$\begin{pmatrix} f(\mathbf{x}) \\ 1 \end{pmatrix} = \mathbf{A} \begin{pmatrix} \mathbf{x} \\ 1 \end{pmatrix} + \mathbf{Kc}$$

Where $K_{(i,j)} = \rho(||\mathbf{x_i} - \mathbf{w_j}||)$

4.2 Cost functions

In some literature the term similarity measure is used analogous to cost function, where registration is accomplished by maximizing similarity between images. However, in practice it is generally the differences that are minimized. In this thesis we will use the term cost function, C, to describe differences between images. To register the base image, I_b, to the reference image, I_r, we estimate the parameters for a transformation T, by

$$T = \underset{T}{\text{Min}} \left(C(T(I_b), I_r) + \lambda E(T) \right) \tag{4.1}$$

where E(T) is a regularizing function to keep the transformation smooth, and λ is the smoothing parameter.

4.2.1 Sum squared difference

A simple cost function is sum squared difference (SSD)

$$\mathbf{C}_{SSD}(I_r, I_t) := \int_{x \in \Omega} (I_r(x) - I_t(x))^2 d\mathbf{x}$$

The main advantage of usage of C_{SSD} is the low computational complexity. The disadvantages with comparing intensity directly is related to the nature of the contrast enhanced MR images we want to study. Differences in intensity within an image due to contrast agent introduce an undesired weighting. In the early time-frames the intensity will be higher in the ventricles, see Figure 4.1, and then slowly increase in the myocardium over time. The resulting transformation will register the high intensity regions well while the low intensity regions will be ignored.

4.2.2 Normalized gradient fields

Cost functions based on normalized gradient were originally proposed by Haber, Eldad and Modersitzk[3] to deal with registration of images between different modalities. Normalized Gradient Fields (NGF) has also been used for contrast enhanced MRI[4]. Using gradient information is intuitive for our problem as the myocardial boundaries we want to align should have strong gradients. The normalized gradient field $\mathbf{n}_\epsilon$ of an image $\mathbf{I}$ is defined by

$$\mathbf{n}_\epsilon(I, x) = \frac{\Delta I(\mathbf{x})}{\sqrt{|\Delta I(\mathbf{x})|^2 + \epsilon^2}}$$

Where ϵ is a soft thresholding of the gradient function. Regions where ϵ is much greater than $\Delta I(\mathbf{x})$ contribute very little as $n_\epsilon(I, x)$ is almost zero.

The choice of ϵ proposed in [3] is

$$\epsilon = \eta \frac{\int_\Omega |\Delta I(\mathbf{x}) d\mathbf{x}|}{\int_\Omega d\mathbf{x}}$$

where η is the noise level in the image. Noise level estimation will be described on the next section. Several cost functions based on NGF have been proposed that uses directional information for example scalar product.

$$\mathbf{C}_{NFG*}(I_r, I_t) := \int_{\mathbf{x}\in\Omega} ||\mathbf{n}_\epsilon(I_r, \mathbf{x}) * \mathbf{n}_\epsilon(I_t, \mathbf{x})||^2 d\mathbf{x}$$

For our application the location of strong gradients i.e the endocardium and epicardium is more important than gradient direction. Therefore we will only use the amplitude of $\mathbf{n}_\epsilon$ in our cost function;

$$\mathbf{C}_{|NFG|}(I_r, I_t) := \int_{\mathbf{x}\in\Omega} (|\mathbf{n}_\epsilon(I_r, \mathbf{x})| - T(|\mathbf{n}_\epsilon(I_b, \mathbf{x})|))^2 d\mathbf{x}$$

This will also reduce computational complexity when evaluating the cost function, as $\mathbf{n}_\epsilon$ only has to be calculated once per image.

4.2.3 Noise estimation

In the original NGF implementation, noise is estimated globally by the standard deviation of the gradient amplitude $|\Delta(I\mathbf{x})|$. This estimation will only be unbiased for an uniform image with added Gaussian noise. In contrast enhanced MR images we have boundary gradients that will be averaged with the noise. Perhaps more importantly the contrast agent will increase noise as well as the desired signal roughly in proportion to its concentration, which will give the noise a nonuniform spatial distribution. Shen-chuan Tai and Shih-ming Yang [5] suggest using edge detection to exclude edge points during noise estimation. In its simplest form it can be done by thresholding $|\Delta I(\mathbf{x})|$, excluding the highest portion as outliers. To account for the spatial distribution, noise can be made a local feature by estimating it in a point $\mathbf{x}$ by the standard deviation in a neighborhood $N_\mathbf{x}$ around $\mathbf{x}$

$$\eta^2(\mathbf{x}) = \frac{\int_{N_\mathbf{x}} (\mu - |\Delta I(\mathbf{x})|)^2 d\mathbf{x}}{\int_{N_\mathbf{x}} d\mathbf{x}}$$

$$\mu = \frac{\int_{N_\mathbf{x}} |\Delta I(\mathbf{x})| d\mathbf{x}}{\int_{N_\mathbf{x}} d\mathbf{x}}$$

With a global noise estimation, gradients in regions with low intensity that are strong compared to their neighborhood, may be suppressed due to the higher gradients and noise in regions with higher contrast agent concentration. By making noise a local feature we also deal with the undesired weighing caused by differences in contrast concentration. Gradients are normalized with respect to their neighborhood and thus we can fit the transformation to structures in low intensity regions as well as high. Figure 4.2a show a contrast enhanced MRI image where the contrast agent is concentrated in the left ventricle. With global noise estimation, the gradient amplitudes of the endocardium are clearly visible, while the epicardal ones are suppressed, as illustrated by Figure 4.2b. With local noise estimation, this problem is remedied to

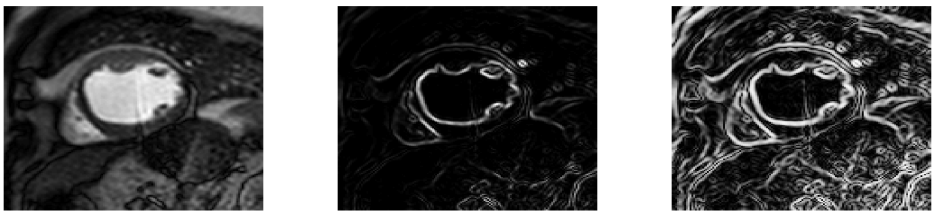

(a) Original image (b) NGF amplitude with global noise estimation (c) NGF estimation with local noise estimation

Figure 4.2: Comparison of NGF amplitudes between global or local noise estimation.

an extent, as illustrated by Figure 4.2c. In our registration method a square neighborhood with a side about half the width of the myocardium was used. This allows us to normalize the epicardium and endocardium separately.

4.2.4 Regularization

To increase the smoothness of the transformation we add a Laplacian regularization to our cost function.

$$E(T) := \int_{x\in\Omega} ||\nabla f||^2 d\mathbf{x}$$

4.2.5 Weights

Images such as Figure 4.1 contains background objects such as the lungs as well as the heart. Since we are only interested in how the myocardium moves, we will add a weighting function that emphasize the region of the image we want to register. The cost function in Equation 4.1 becomes

$$\mathbf{C}_{|NFG|}(I_r, I_t) := \int_{x\in\Omega} (|\mathbf{n}_\epsilon(I_r, x)| - T(|\mathbf{n}_\epsilon(I_b, \mathbf{x})|))^2 w(\mathbf{x}) d\mathbf{x} + \lambda \int_{x\in\Omega} ||\nabla f||^2 w(\mathbf{x}) d\mathbf{x}$$

$$w(\mathbf{x}) = \begin{cases} 1 & \text{if } \Upsilon(\mathbf{x}) \leq R \\ 0 & \text{if } \Upsilon(\mathbf{x}) > R \end{cases}$$

where the function $\Upsilon(\mathbf{x})$ is the minimum euclidean distance between the point $\mathbf{x}$ and the epicardium or endocardium in the reference frame and R is a constant. With this weighting function we will use knowledge of the reference frame segmentation to fit the transformation specifically for the myocardium. Its binary nature is especially useful if we make the assumption that the moments of the myocardium are small so that we can set the constant R small. The cost function does not need to be evaluated in the region with zero weights.

4.3 Scale spaces

In the previous section we have suppressed noise by amplitude. Another question is how small structures we should consider when registering images. The scale space rep-

resentation of an image, I, is defined as the convolution of the image with a Gaussian kernel

$$S(\mathbf{x}, t) = \frac{1}{2\pi t} e^{-(x^2+y^2)/2t} * I(\mathbf{x})$$

Where t is the scale parameter. Structures of size smaller than $\sqrt(t)$ will be smoothed out in the scale space at scale t. Fitting a transformation to the scale-space representation of the reference and base images at a suitable scale will suppress small changes while taking the larger ones into account.

4.4 Optimization

Finding the optimal parameters for the transformation is a minimization problem in N = 6 + 2n dimensions. Where n is the number of control points in the TPS transform, and the additional six are the affine parameters. Since the cost function is always positive, a finite global minimum exists. There are no guaranties that the global minimum is unique, i.e that the function is globally convex. An optimization method explores the parameter space in a number if steps, attempting to locate the lowest point.

Optimization methods for convex problems, so called greedy methods, will only take downhill steps. These methods will terminate when a local minimum is reached, which is determined by either the step size or change in function value is below a predetermined threshold. Using such a method on a non-convex problem will still converge to a local minimum. However, we can not know if the minimum is the global minimum.

Identifying the global minimum can be difficult in large parameter spaces, instead we will focus on finding a good enough solution. The optimization method of choice is a combination of the downhill simplex and simulated annealing methods, which will be presented in the following sections.

4.4.1 Downhill simplex

The Downhill simplex or amoeba method is an algorithm for nonlinear optimization in multiple dimensions. A *simplex* is a geometric figure. In N dimensions it consists of N+1 points and the line segments connecting each point to every other point. In 2 dimensions it is a triangle, in 3 a tetrahedron etc. In the downhill simplex we begin from an initial simplex, for example the set of points

$$\mathbf{P_i} = \mathbf{P_0} + \lambda \mathbf{e_i}$$

Where $\mathbf{P_0}$ is the initial guess and the first point in the simplex, λ is a constant and $\mathbf{e_i}$ is the unit vector of the N dimensional space. The algorithm then moves the simplex in a series of steps. The list below and Figure 4.3 shows the basics steps of the method.

- Reflections moves the point with the highest function value through the lowest face of the simplex, the step size is designed to conserve the volume of the simplex.
- Expansions, if the reflected point is the lowest in the simplex the reflection is a downhill step, the simplex will expand in the same direction, taking larger steps.

- Contractions, conversely if the reflected point become the highest point we have reached a valley floor and now need to contract the simplex and take smaller steps.

4.4.2 Simulated annealing

The downhill simplex method will generally only take downhill steps. For a method to be effective in the presence of multiple local minima, we need to be able to sometimes take uphill steps to transition from one local minimum to another potentially lower one. One way of accomplishing this is to sometimes take random steps. Simulated annealing is derived from thermodynamics, in analogy of how metals cool and anneal. At high temperatures ions can move more freely within the metal, as the temperature lowers, this mobility is lost. If a metal is cooled slowly the ions line up in a pure crystal, with the lowest energy state. If the metal instead is cooled quickly it will not reach this state.

In simulated annealing we introduce a temperature parameter T to the optimization method. The method sometimes takes an uphill step, with a probability dependent on T. The temperature is then gradually lowered, and the method will accept fewer and fewer uphill steps.

In Numerical recipes in C [6] this idea is applied to the downhill simplex method, with the same basic steps as described in the previous section. In this simulated annealing version, a logarithmically distributed random variable, proportional to the temperature T is added to the function value of each point in the simplex, and a similar random variable is subtracted from each point tried as a replacement point. This way the method will always accept a true downhill step if available. But if the method is stuck in a local minimum or on a flat surface it will nevertheless accept a random step. The temperature is gradually reduced making random steps fewer and smaller. As $T \to 0$ the method will revert to the original downhill simplex.

Instead of a termination criteria we will budget m moves at each temperature. Starting at an initial temperature T_0 the temperature will be reduced to $T_i = T_0(1-\alpha)^i$ every m moves, where $0 < \alpha < 1$, until T_i reaches a preset temperature T_{end}. If the temperature is reduced slowly enough we are likely to end up in the vicinity of the global minimum.

The speed that the temperature is reduced by is proportional to α/m, and an appropriate value can for a specific problem be determined by experiment.

4.5 Validation

Ground truth myocardial segmentations can be used to validate how well a transformation registers the myocardium. To evaluate the relative overlap between the transformed myocardium, M_t, and reference myocardium , M_r, we use Dice coefficient s, calculated by

$$s(M_r, M_t) = \frac{2|M_r \cap M_t|}{|M_r| + |M_t|}$$

The norm |M| is the area of the myocardium or intersection. Dice gives a value between one and zero where one corresponds to perfect overlap between transformed and reference myocardium, and zero corresponds to no overlap at all.

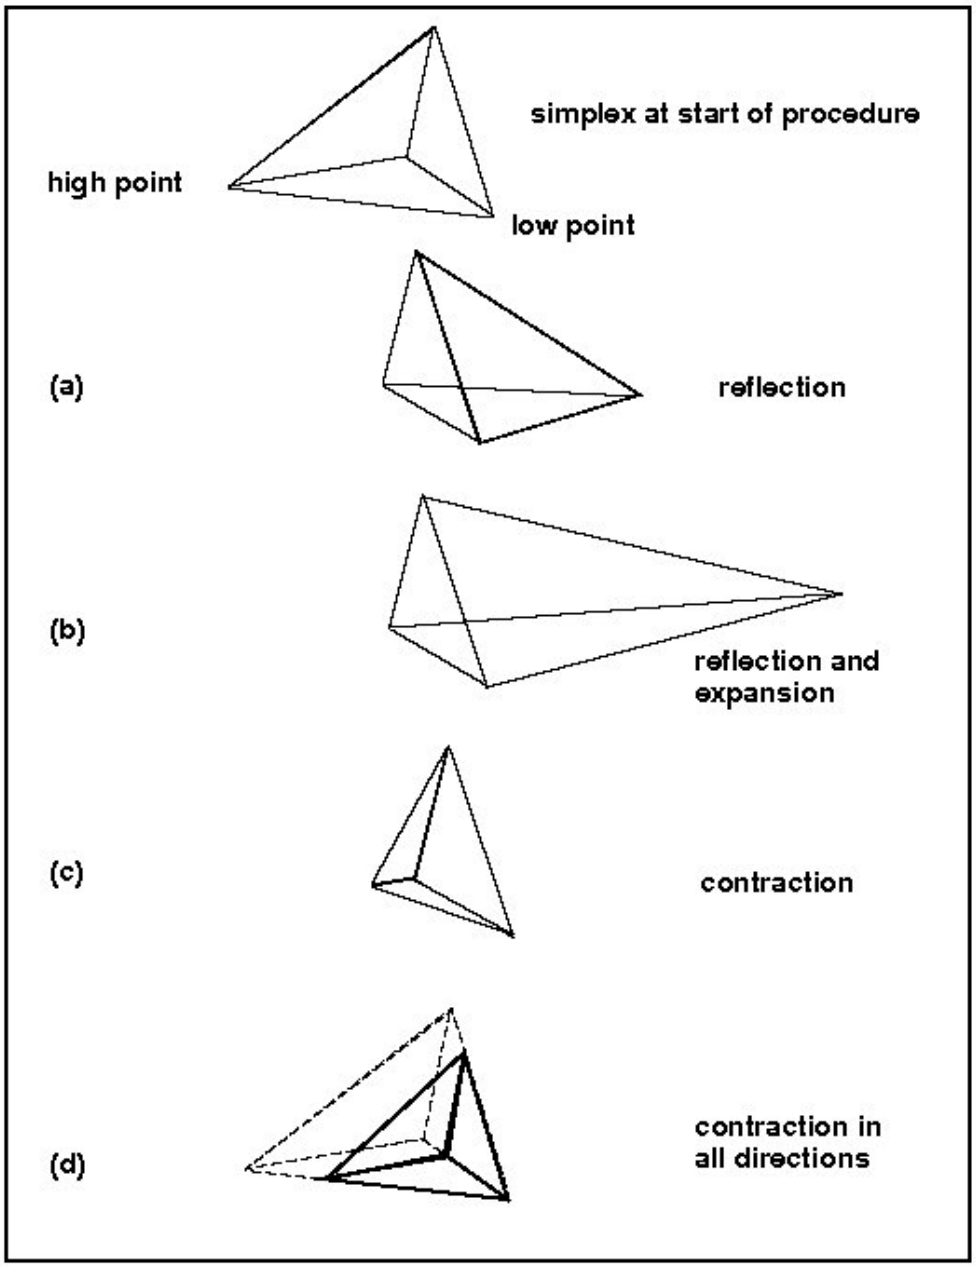

Figure 4.3: Basic moves in the downhill simplex method

4.6 Consistency

Simulated annealing is a stochastic method and will converge to different parameter values each run. Because of this we need to compare the consistency of the algorithm if we run it several times with the same configuration and data sets. Every time the algorithm is run we get a segmentation for each image in the series. By comparing the Dice similarity between the myocardium in the same image in different iterations, the consistency of the algorithm can be determined. Running the algorithm 10 times for each validation set, the Dice similarity $s(M_{ti}, M_{tj})$ were calculated, where i and j are the iteration numbers.

4.7 Statistical analysis

Values are presented as mean $\pm$ standard deviation. Differences in bias were analyzed by the t-test, and differences in the standard deviation were analyzed by the F-test. A p-value < 0.05 was considered statistically significant.

Chapter 5

Results

5.1 Validation

To validate the method we used image stacks from four different patients with manual segmentation of the left ventricle as ground truth. Two patients have segmentations in five time-frames each, one patient have segmentations in 8 time-frames and one patient has segmentation in 15 time-frames.

The first segmented image in every stack was used as the reference, the rest of the segmented images were the registered to the reference. It was only the segmentation in the reference time-frame image that was used in the registration. The remaining segmentations were used in the validation to calculate the Dice coefficient. Running the algorithm 10 times for each validation set, the average Dice coefficient for all the images was 0.89 ± 0.028. The average Dice coefficient before registration was 0.79 ± 0.11. Dice coefficient improved after registration in every image except for one. The results from the registration is illustrated for one patient in Figure 5.1.

5.2 Consistency

The average consistency Dice coefficient, $s(M_{ti}, M_{tj})$, for $i \neq j$ was 0.96 ± 0.016. This corresponds to an average inconsistency in area of $98mm^2$ or an average movement of the epicardium and endocardium of $0.2mm$. In the MRI images $0.2mm$, corresponds to around $1/7$ pixels, where each image is 146x146 pixels.

5.3 Effects of segmentation

How we segment the images in the validation sets has an effect on the results we measure. For example if we segmented something completely different in one time-frame the Dice coefficient would be (close to) zero regardless of how good the registration is. Segment has a function that refines segmentation based on gradients. In Figure 5.2 the refined segmentation and the registered images resulting from it can be seen. Running the algorithm 10 times for each validation set results in a average Dice coefficient of 0.93 ± 0.023. The consistency Dice coefficient was 0.97 ± 0.014.

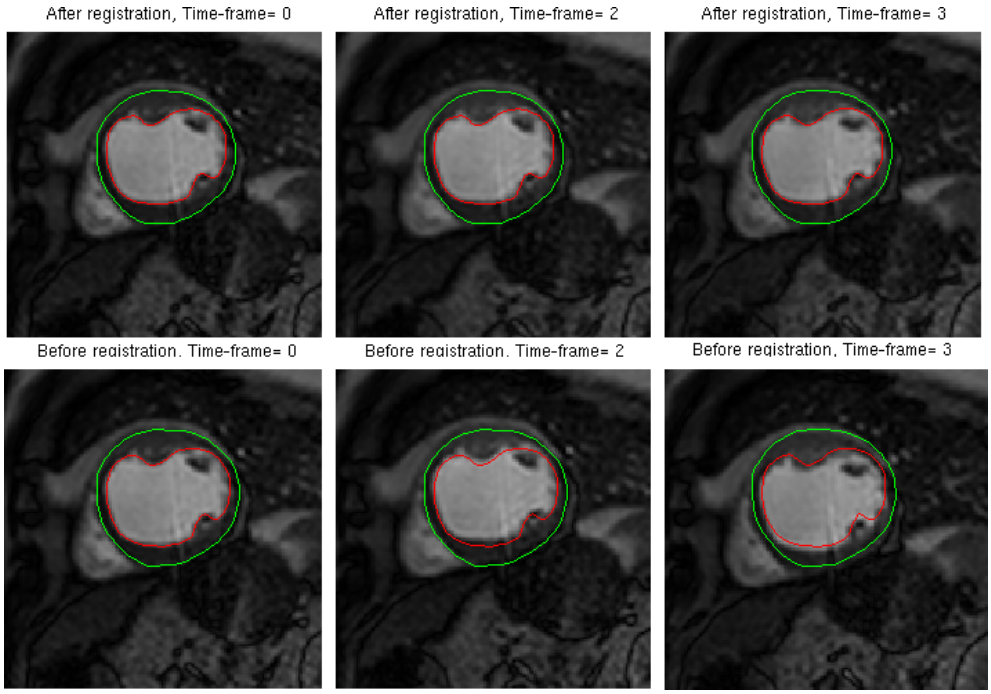

Figure 5.1: Comparison between registered and unregistered images for one patient. The green and red lines illustrate the epicardium and endocardium respectively, with the reference segmentation from the first time-frame image overlayed to all images.

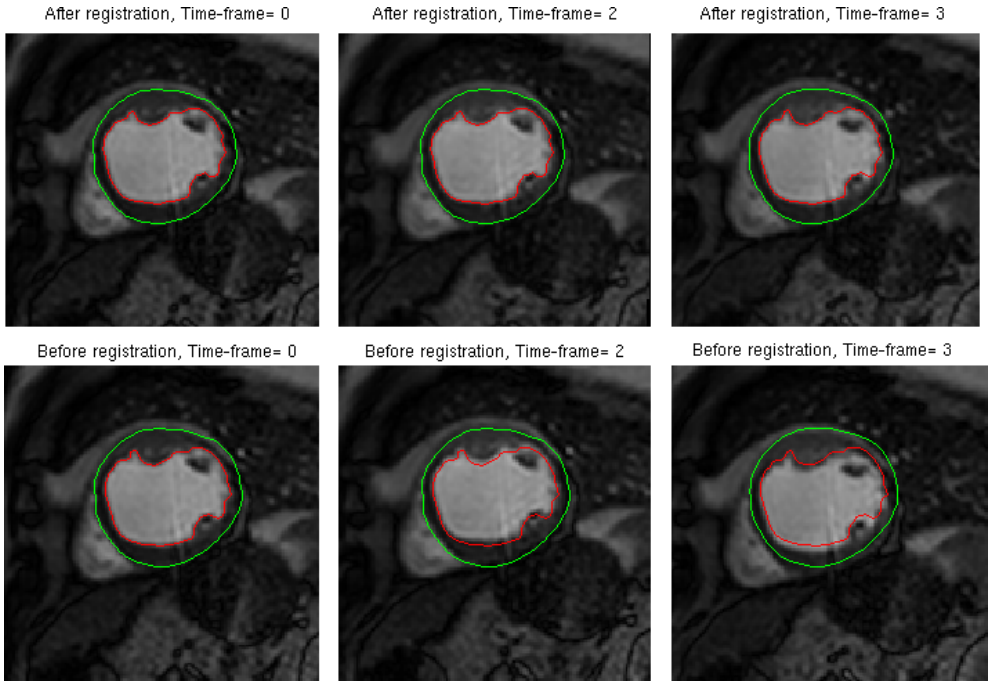

Figure 5.2: Comparison between registered and unregistered images with refined segmentation, with same images as in Figure 5.2

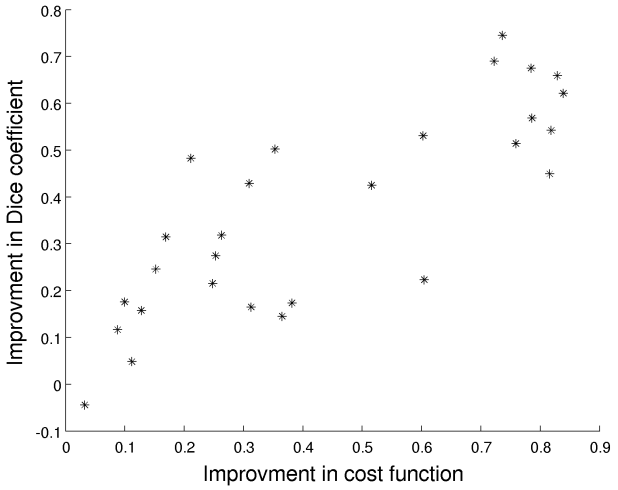

Figure 5.3: Comparison between improvements in cost function and Dice coefficient.

5.3.1 Cost-function validation

For our cost function to performs well, we would expect improvements in the cost function to correlate to improvement of the validation measurement. Figure 5.3 shows the dependence of Dice coefficient improvement on the cost function improvement. The values are correlated with $\mathbf{R}^2 = 0.52$ (p-value < 0.001) and covariance 0.94.

5.3.2 Improvement from initial alignment

As we see in images like the ones in Figure 5.2 some images are more misaligned before registration than others. Ideally the alignment after registration should not depend on the initial alignment. Figure 5.4 shows a weak dependence of the final Dice coefficient upon the initial one.

5.4 Perfusion

Myocardial perfusion is proportional to the increase in intensity over time in the contrast enhanced MRI image. Figure 5.5 show how the image intensity in the left ventricle myocardium changes over time for one patient after the contrast agent was injected. Intensity in the bloodpool, interior of the left ventricle, is also shown.

5.5 Cost function comparisons

After registration using the SSD cost function the average Dice coefficient was 0.81 ± 0.073. A feature introduced in this thesis for myocardial perfusion studies is local noise estimation. To evaluate if this improves the algorithm we will also use the NGF cost function without it. With global noise estimation, the average Dice coefficient was 0.86 ± 0.043.

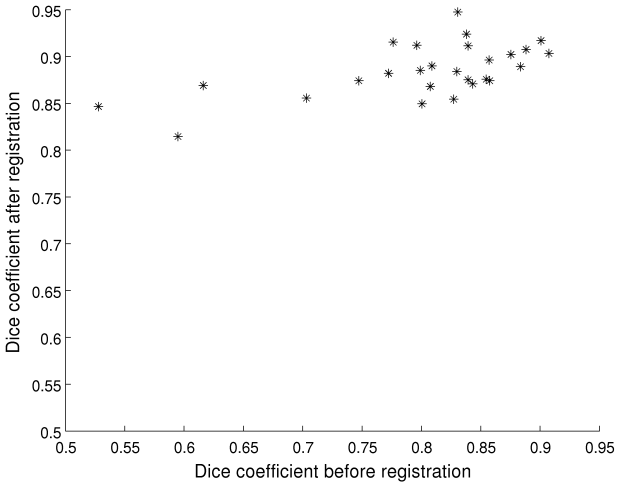

Figure 5.4: Comparison between initial and final Dice coefficient

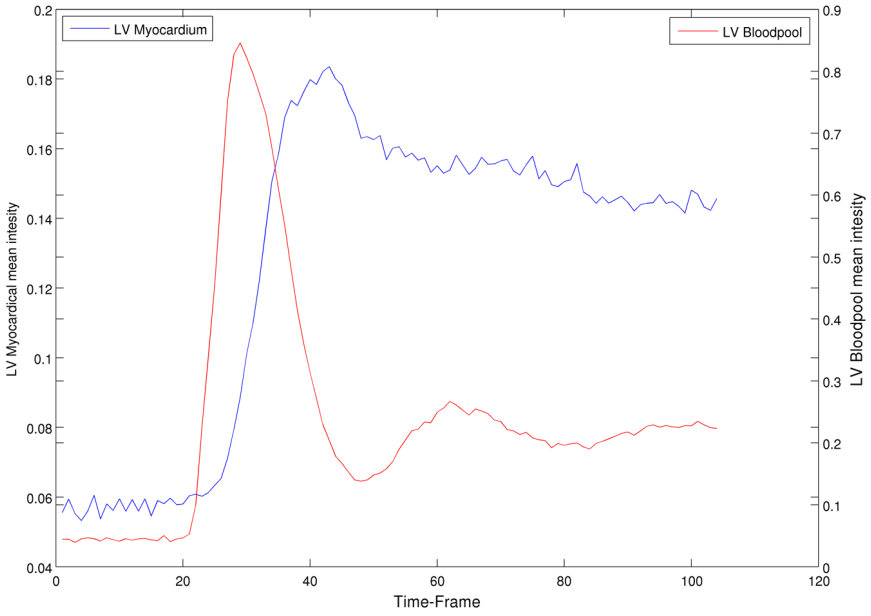

Figure 5.5: Mean intensity in the LV myocardium and bloodpool, the interior of the left ventricle.

Both of these cost function performed worse on average than NGF with local noise estimation. More importantly they result in higher variability with very poor results in some images, especially in longer time series. NGF with local noise estimation performed significantly better than SSD in 89 percent of the images, and significantly better than NGF with global noise estimation in 79 percent of the images. Figure 5.7 shows an example where SSD performs poorly, and Figure 5.6 shows the same images registered using NGF with global noise estimation. For comparison see Figure 5.8 where the same images have been registered using the NGF method with local noise estimation.

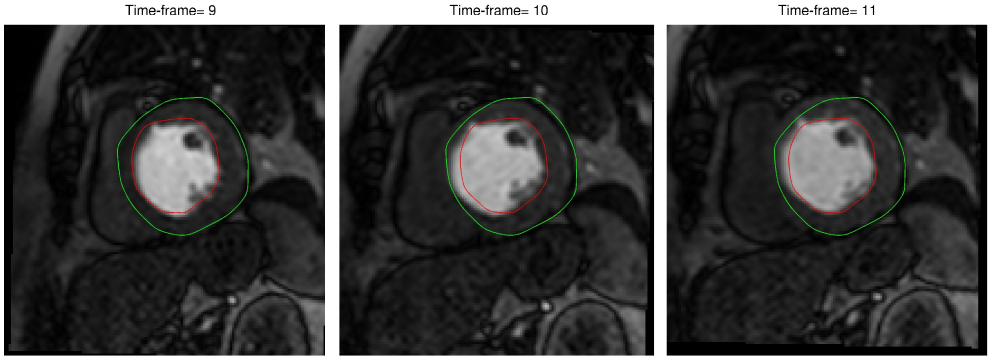

Figure 5.6: Registered images using NGF with global noise estimation.

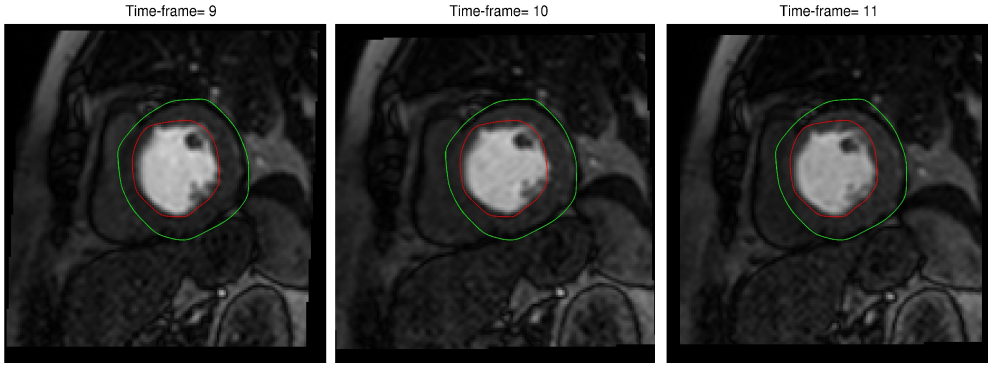

Figure 5.7: Registered images using SSD cost function.

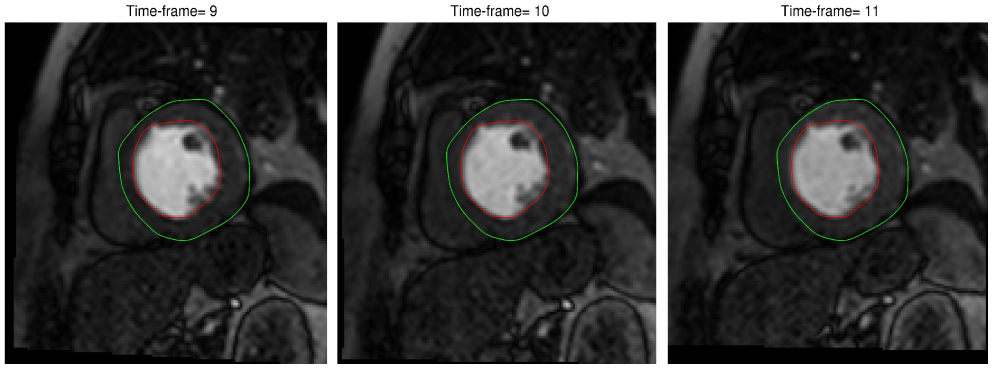

Figure 5.8: Registered images using NGF with local noise estimation.

Chapter 6

Discussion and Conclusions

The results of this thesis is the development of a registration method for perfusion MRI images that allow usage of the same left ventricular segmentation in all time-frames, in the perfusion analysis. Overall the method performed well, the Dice coefficient improved on average from 0.79 ± 0.11 before registration to 0.89 ± 0.028 after registration. The Dice coefficient improved the most in images with the lowest coefficient prior to registration. The Dice coefficient improved in all but one image, this image was already aligned with the reference with Dice coefficient 0.91 before registration and decreased to 0.90 after registration i.e still quite well aligned but slightly lower. Visually the overlayed left ventricular segmentations after registration looked reasonable in the images studied. The perfusion curves produced when analyzing the intensity in the registered images look like what we would expect from previous studies of myocardial perfusion. Despite using a stochastic method, variability between iterations was low with a consistency of 0.96.

6.1 Registration and segmentation

As seen in the previous chapter how we segment images effect the overall performance. In Figure 6.1 we see a manual segmentation in the first time-frame that does not follow the edges in the image perfectly. Hence the segmentation in the registered images will follow the edges in the image in similar fashion. This is frequent in the segmentations but is often a deliberate judgment rather then a human error. How to manually segment images are beyond the scope of this thesis, but we can note that any error made in segmentation of the reference time-frame image will effect all time-frames.

6.2 Limitations

While myocardial alignment improved in almost all images in the validation sets, the Dice coefficient never reaches one. Part of this residual error may be be due to small inconsistencies in segmentation. A more fundamental limitation is representing a 3-D object such as the heart with a very limited resolution in one dimension, the three slices along the long axis. In the affine model, movement along the long axis of the left ventricle, i.e perpendicular to the image plane is modeled by scaling. In reality this kind of movement will result in imaging a different slice along the long axis. This

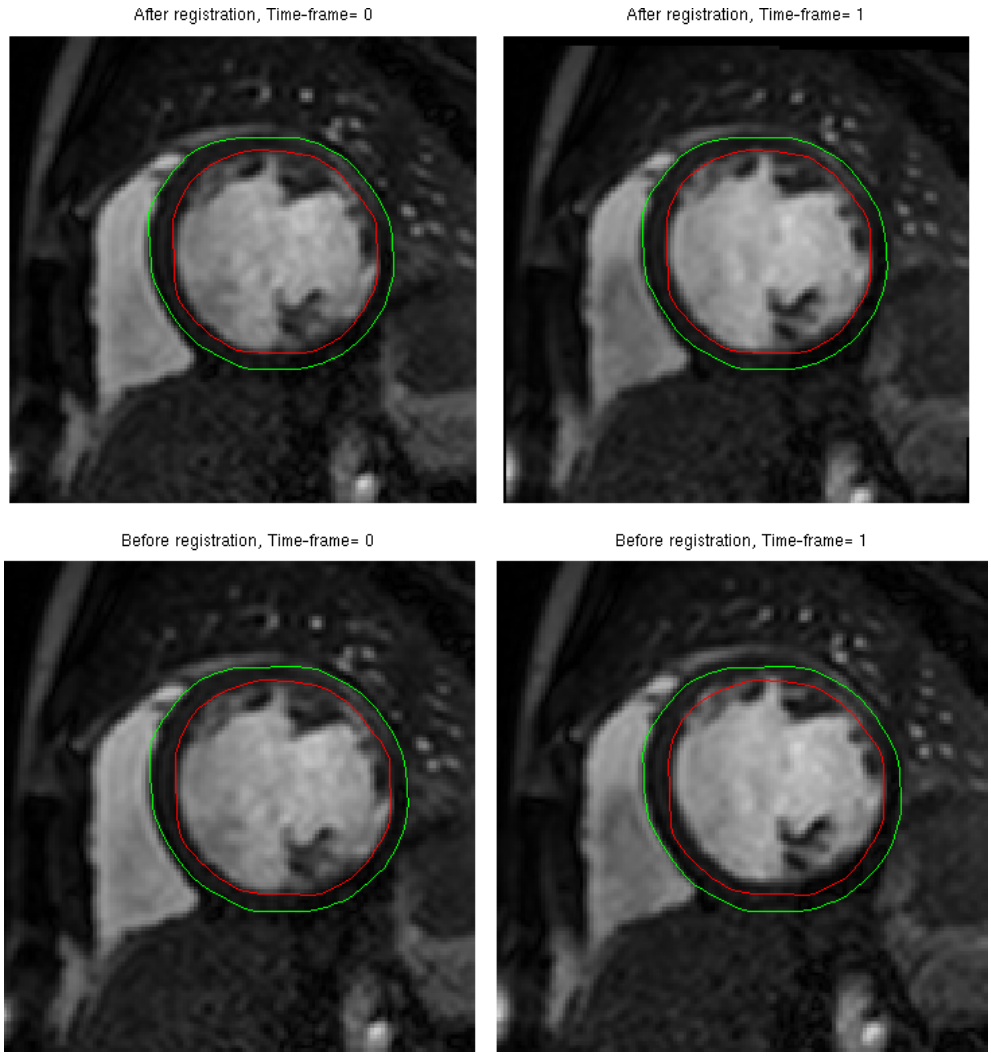

Figure 6.1: Segmentation in registered images will follow the image edges in a similarly to the segmentations in the reference frame.

can to some extent be modeled by the TPS deformation. However the theory for the registration algorithm could be extended to 3 dimensions, if we want to study perfusion in true 3-D in the future.

6.3 Noise Estimation

Local noise estimation significantly improved the robustness of the algorithm, increasing the Dice coefficient to 0.89 ± 0.028 from 0.86 ± 0.046 compared to global noise estimation. The improvement and reduction of variability using local noise estimation was due to improved results in a few images where NGF with global noise estimation performed poorly. Proving it to be useful to deal with the problems that arise due to shifting contrast levels in different time-frame images.

In our implementation we used simple square neighborhoods, the shape and size of the neighborhoods is something that could be studied further.

6.4 Implementation

In its current implementation registration of one image takes about 4 seconds, this is highly dependent on the hardware we use. The bulk of the computations used in the method are the matrix multiplications performed when calculating the transformation in a cost function evaluation. These computations are highly parallelizable and would be ideal for a graphical processing unit (GPU) implementation.

Bibliography

[1] A. H. A. W. G. on Myocardial Segmentation, R. for Cardiac Imaging:, M. D. Cerqueira, N. J. Weissman, V. Dilsizian, A. K. Jacobs, S. Kaul, W. K. Laskey, D. J. Pennell, J. A. Rumberger, T. Ryan, and M. S. Verani, "Standardized myocardial segmentation and nomenclature for tomographic imaging of the heart," *Circulation*, vol. 105, no. 4, pp. 539–542, 2002.

[2] D. Lowe, "Object recognition from local scale-invariant features," *Computer Vision, 1999. The Proceedings of the Seventh IEEE International Conference on*, vol. 2, pp. 1150 – 1157, 1999.

[3] E. Haber and J. Modersitzki, "Beyond Mutual Information: A Simple and Robust Alternative," in *Bildverarbeitung für die Medizin 2005* (Brauer, H.-P. Meinzer, H. Handels, A. Horsch, and T. Tolxdorff, eds.), Informatik aktuell, ch. 72, pp. 350–354–354, Berlin/Heidelberg: Springer Berlin Heidelberg, 2005.

[4] G. Wollny, M. J. Ledesma-Carbayo, P. Kellman, and A. Santos, "A new similarity measure for non-rigid breathing motion compensation of myocardial perfusion MRI," *Engineering in Medicine and Biology Society, 2008. EMBS 2008. 30th Annual International Conference of the IEEE*, pp. 3389–3392, Oct. 2008.

[5] S. chuan Tai and S. ming Yang, "A fast method for image noise estimation using laplacian operator and adaptive edge detection," *Communications, Control and Signal Processing, 2008. ISCCSP 2008. 3rd International Symposium on*, pp. 1077 – 1081, 2008.

[6] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, "Numerical recipes in C (2nd ed.): the art of scientific computing," 1992.