



HAL
open science

Reduced representation of segmentation and tracking in cardiac images for group-wise longitudinal analysis

Marc-Michel Rohé

► **To cite this version:**

Marc-Michel Rohé. Reduced representation of segmentation and tracking in cardiac images for group-wise longitudinal analysis. Human health and pathology. COMUE Université Côte d'Azur (2015 - 2019), 2017. English. NNT: 2017AZUR4051 . tel-01575292v2

HAL Id: tel-01575292

<https://inria.hal.science/tel-01575292v2>

Submitted on 4 Oct 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

UNIVERSITY OF NICE - SOPHIA ANTIPOLIS
DOCTORAL SCHOOL STIC
SCIENCES ET TECHNOLOGIES DE L'INFORMATION
ET DE LA COMMUNICATION

PHD THESIS

to obtain the title of

PhD of Science

of the University of Nice - Sophia Antipolis

Specialty : COMPUTER SCIENCE

Defended by

Marc-Michel ROHÉ

Reduced Representation of Segmentation and Tracking in Cardiac Images for Group-wise Longitudinal Analysis

Thesis Advisor: Xavier PENNEC

Thesis Co-Advisor: Maxime SERMESANT

prepared at INRIA Sophia Antipolis, ASCLEPIOS Team

to be defended on July 3rd, 2017

Jury :

<i>President :</i>	Patrick CLARYSSE	-	CREATIS, CNRS, Inserm, Lyon
<i>Reviewers :</i>	Julia SCHNABEL	-	King's College London
	Alistair YOUNG	-	University of Auckland
<i>Advisor :</i>	Xavier PENNEC	-	Inria Sophia-Antipolis
<i>Co-Advisor :</i>	Maxime SERMESANT	-	Inria Sophia-Antipolis
<i>Examinator :</i>	Andrew TAYLOR	-	University College London

Reduced Representation of Segmentation and Tracking in Cardiac Images for Group-Wise Longitudinal Analysis

Abstract This thesis presents image-based methods for the analysis of cardiac motion to enable group-wise statistics, automatic diagnosis and longitudinal study. This is achieved by combining advanced medical image processing with machine learning methods and statistical modelling.

The first axis of this work is to define an automatic method for the segmentation of the myocardium. We develop a very-fast registration method based on convolutional neural networks that is trained to learn inter-subject heart registration. Then, we embed this registration method into a multi-atlas segmentation pipeline.

The second axis of this work is focused on the improvement of cardiac motion tracking methods in order to define relevant low-dimensional representations. Two different methods are developed, one relying on Barycentric Subspaces built on references frames of the sequence, and another based on a reduced order representation of the motion from polyaffine transformations.

Finally, in the last axis, we apply the previously defined representation to the problem of diagnosis and longitudinal analysis. We show that these representations encode relevant features allowing the diagnosis of infarcted patients and Tetralogy of Fallot versus controls and the analysis of the evolution through time of the cardiac motion of patients with either cardiomyopathies or obesity.

These three axes form an end to end framework for the study of cardiac motion starting from the acquisition of the medical images to their automatic analysis. Such a framework could be used for diagnosis and therapy planning in order to improve the clinical decision making with a more personalised computer-aided medicine.

Keywords: Medical image analysis, Non-rigid registration, Deep learning, Statistical model reduction, Longitudinal analysis

Représentation réduite de la segmentation et du suivi des images cardiaques pour l'analyse longitudinale de groupe

Résumé Cette thèse présente des méthodes d'imagerie pour l'analyse du mouvement cardiaque afin de permettre des statistiques groupées, un diagnostic automatique et une étude longitudinale. Ceci est réalisé en combinant des méthodes d'apprentissage et de modélisation statistique.

En premier lieu, une méthode automatique de segmentation du myocarde est définie. Pour ce faire, nous développons une méthode de recalage très rapide basée sur des réseaux neuronaux convolutifs qui sont entraînés à apprendre le recalage cardiaque inter-sujet. Ensuite, nous intégrons cette méthode de recalage dans une pipeline de segmentation multi-atlas.

Ensuite, nous améliorons des méthodes de suivi du mouvement cardiaque afin de définir des représentations à faible dimension. Deux méthodes différentes sont développées, l'une s'appuyant sur des sous-espaces barycentriques construits sur des frames de référence de la séquence et une autre basée sur une représentation d'ordre réduit du mouvement avec des transformations polyaffine.

Enfin, nous appliquons la représentation précédemment définie au problème du diagnostic et de l'analyse longitudinale. Nous montrons que ces représentations encodent des caractéristiques pertinentes permettant le diagnostic des patients atteint d'infarct et de Tétralogie de Fallot ainsi que l'analyse de l'évolution dans le temps du mouvement cardiaque des patients atteints de cardiomyopathies ou d'obésité.

Ces trois axes forment un cadre pour l'étude du mouvement cardiaque de bout en bout de l'acquisition des images médicales jusqu'à leur analyse automatique afin d'améliorer la prise de décision clinique grâce à un traitement personnalisé assisté par ordinateur.

Mots Clefs: Analyse d'images médicales, Recalage non-rigide, Apprentissage profond, Réduction statistique de modèle, Analyse longitudinale

Acknowledgments

First of all, I would like to warmly thank my two supervisors, Xavier Pennec and Maxime Sermesant for their guidance and support all along the three years of my Ph.D. You were always here to answer my questions and to help me during my research work and I have learned a lot from your deep and complementary scientific knowledge. Xavier, your passion for theoretical mathematics was highly communicative and you never refused to share your time in order to help me when I needed. I could come with one simple question and get drag into hours of interesting discussions about completely different topics and ideas of research. Maxime, you helped me to turn innovative mathematical concepts into practical sound applications and you always managed to come out with ideas to close the gap between theory and applications. It was a pleasure to work under both of your guidance, not only on a professional level but also on a personal level with the multiple trips abroad we had the chance to do together! I would also like to thank Nicholas Ayache for accepting me in the Asclepios team where I enjoyed outstanding working conditions and a friendly research environment for the last three years.

I am extremely grateful to Prof. Julia Schnabel and Prof. Alistair Young for accepting to be my reviewers and for spending their precious time to read and correct this manuscript. Their encouraging compliments and constructive comments on my work are invaluable. I am also indebted to Andrew Taylor and Patrick Clarysse who accepted to be member of my jury and to come at my defense. Dear committee, thank you, it has been a great honor for me to have such an outstanding jury.

During this PhD, I had the chance to have such a great time so that I always enjoyed working even when the pressure was high. Special thanks to the people who worked in my office: my ginger friend Nina who introduced me to the team and made me feel welcomed when I arrived, Krissy, Loïc D., Sofia and Yann! To Thomas with whom I probably drank thousands of coffee and had thousands of laughs. To Sophie who managed to cancel the journal club just a month after arriving in the lab. To Roch with whom I had the chance to travel to multiple places for meetings of the MD-Paedegree project. To Rocio for inviting us at her incredible wedding in Mexico. Thanks to all the other students for the good time we have spent together in and also outside the lab. Those who are not in the lab anymore but are not forgotten nonetheless: Cloclo, Nicolas C., Vikash, Milkymat, Bisheh, Loïc... And those who still need to work a little bit to reach our level and get the almighty PhD: Rafifou, Shuman, Pawel, Julian, Qiao, Wen, Luigi, Nicoco and probably many more by the time your read this!

I finally would like to thank my whole family for their love, support, encouragement and kindness: my mother of course, my father, my sister, my grand-mother and my cousins. You always believed in me and supported me during my long years of education and I would never have succeeded without that. Thank you for always pushing me to do what makes me happy, no matter the obstacles or challenges. This thesis is dedicated to you.

I have probably forgotten to thank many people so thank you all!

Contents

1	Clinical and Technical Context	1
1.1	Introduction	1
1.2	Healthy Cardiac Structure and Function	2
1.2.1	Cardiac Structure	3
1.2.2	Cardiac Cycle	4
1.3	From Healthy to Pathological Heart	6
1.3.1	MD-Paedigree Project	7
1.3.2	Cardiomyopathies	9
1.3.3	Obesity	11
1.4	Cardiac imaging	12
1.4.1	Echocardiography	13
1.4.2	Magnetic Resonance Imaging	14
1.4.3	Computed Tomography	17
1.5	Manuscript Organization and Objectives	17
1.6	Publications	23
1.6.1	Journal Articles	23
1.6.2	Selective Peer-Reviewed Conference Paper	24
1.6.3	Awards	24
I	FROM MEDICAL IMAGES TO 3D SHAPE	25
2	Registration in Medical Imaging	27
2.1	Introduction	27
2.2	Registration Algorithms	28
2.2.1	Similarity Measures	30
2.2.2	Transformation Spaces and Regularization	31
2.2.3	Optimization Methods	32
2.3	LCC Log-Domain Diffeomorphic Demons	33
2.3.1	Optimization Method: Alternate Optimization	33
2.3.2	Deformation Parametrization: Stationary Velocity Field	34
2.3.3	Similarity Metric: Local Correlation Coefficient	35
2.4	Framework of Currents for Shape Registration	36
2.4.1	The Formalism of Currents for Surface Representation	38
2.4.2	Surface Registration Using Currents and LDDMM	40
3	SVF-Net: Learning Deformable Image Registration Using Shape Matching	43
3.1	Motivations	43
3.2	Reference Deformations Computation	46

3.2.1	Currents for Shape matching	47
3.2.2	Elastic Body Spline Interpolation	47
3.2.3	Stationary Velocity Fields for Diffeomorphism Parametrization	48
3.3	SVF-Net: Fully Convolutional Neural Network Architecture	49
3.4	Experiments	50
3.4.1	Training	51
3.4.2	Evaluation	51
3.5	Conclusions	53
4	An Automatic Multi-Atlas Myocardium Segmentation	55
4.1	Chapter Overview	56
4.2	Motivations	57
4.3	Segmentation in Clinical Practice	57
4.4	Segmentation Methods In Medical Imaging	59
4.5	Multi-Atlas Segmentation: Overview of the Method	60
4.6	Rigid Alignment by Landmarks Detection	61
4.6.1	Method for Rigid Alignment	62
4.6.2	CNNs and Heatmap for Landmarks Detection	63
4.6.3	Training and Results	63
4.6.4	Alignment of the Images and Cropping	65
4.7	Registration of Images and Fusion of Segmentation	66
4.7.1	Points Propagation	66
4.7.2	Points Combination	66
4.7.3	Supervised Learning of Local Distance	67
4.7.4	Computation of Weights Based on Local Distance	68
4.8	Conclusions	70
II LOW-DIMENSIONAL REPRESENTATION OF CARDIAC MOTION		73
5	Barycentric Subspace for Cardiac Motion Tracking	75
5.1	Chapter Overview	75
5.2	Background: Cardiac Motion Tracking	76
5.2.1	Cardiac Motion Tracking Algorithms	77
5.2.2	Evaluation of Tracking Methods	79
5.3	Motivations: Low-Dimensional Representation of the Cardiac Motion	79
5.4	Methodology: Barycentric Subspaces	82
5.4.1	Definition of the Subspace	82
5.4.2	Projection of an Image to the Subspace	83
5.4.3	Computation of the optimal References of a sequence of Images	84
5.4.4	Computation of an Image within the Subspace	84
5.5	Using Barycentric Subspaces as a prior on the Registration	85
5.5.1	Barycentric Log-Demons Algorithm	88

5.5.2	Evaluation using a Synthetic Sequence	88
5.5.3	Towards Symmetric Transitive Registration	89
5.6	Conclusion	90
5.A	Projection of an image to the Barycentric Subspace	91
5.B	Frame-To-Frame Barycentric Registration Formulation	92
6	Highly Reduced Model of the Cardiac Function for Fast Simulation and Personalization	93
6.1	Chapter Overview	93
6.2	Motivations	94
6.3	Reduced-Polyaffine Projection for Compact Cardiac Motion Representation	95
6.4	Biophysical Model of the Heart Simulation Database	96
6.5	Parameters Mapping through PLS Regression	97
6.6	Applications	99
6.6.1	Direct Highly Reduced Cardiac Function Model	99
6.6.2	Personalization of Model Parameters	100
6.7	Conclusion	100
III	APPLICATION TO DIAGNOSIS AND LONGITUDINAL ANALYSIS	101
7	Polyaffine Transformations for the Automatic Diagnostic of LV Infarct	103
7.1	Motivations	104
7.2	Database of Asymptomatic Subjects and Patients with Myocardial Infarction	105
7.3	Extraction of features of interest through shape and motion dimensionality reduction	106
7.3.1	Polyaffine projection	106
7.3.2	Thickness parameters	108
7.4	Dimensionality reduction of the parameters and classification	110
7.4.1	Learnt Dimensionality Reduction	110
7.4.2	Cross-Validation on Training Set: Classifier Selection	112
7.5	Results and Validation on Testing Set	113
7.6	Conclusion	114
8	Cardiac Motion Signature Using Barycentric Subspaces	117
8.1	Chapter Overview	117
8.2	Cardiac Motion Signature from Low-Dimensional Representation	118
8.2.1	Data	118
8.2.2	Methods	118
8.2.3	Optimal References Frames and Barycentric Curves	119
8.2.4	Group-Wise Analysis of Differences	120

8.3	Reconstruction of Cardiac Sequences	121
8.3.1	Qualitative results	122
8.3.2	Quantitative results	123
8.4	Conclusion	124
9	Longitudinal Analysis of the Cardiac Motion	125
9.1	Chapter Overview	125
9.2	Motivations	126
9.3	Cardiac Motion Features Extraction	127
9.4	Cardiomyopathies	129
9.4.1	Data	129
9.4.2	Mean Motion Model	130
9.4.3	Longitudinal Motion Analysis of CMP patients	131
9.5	Obesity	132
9.5.1	Data: Obesity Patients (CVD)	132
9.5.2	Mean Motion Model Parametrized by BMI	133
9.5.3	Longitudinal Motion Analysis of CVD patients	134
9.6	Conclusion	136
10	Conclusions and Perspectives	139
10.1	Summary of the Main Contributions	139
10.2	Perspectives and Future Applications	141
10.3	Virtual Patient in the Age of Artificial Intelligence: the Future of Medicine	144
	Bibliography	147

Clinical and Technical Context

Contents

1.1	Introduction	1
1.2	Healthy Cardiac Structure and Function	2
1.2.1	Cardiac Structure	3
1.2.2	Cardiac Cycle	4
1.3	From Healthy to Pathological Heart	6
1.3.1	MD-Paedigree Project	7
1.3.2	Cardiomyopathies	9
1.3.3	Obesity	11
1.4	Cardiac imaging	12
1.4.1	Echocardiography	13
1.4.2	Magnetic Resonance Imaging	14
1.4.3	Computed Tomography	17
1.5	Manuscript Organization and Objectives	17
1.6	Publications	23
1.6.1	Journal Articles	23
1.6.2	Selective Peer-Reviewed Conference Paper	24
1.6.3	Awards	24

The purpose of this chapter is to give the clinical and technical context, building the basis for the following parts of this manuscript. We start from the highest level - the patient - to give the clinical motivation of our work. We describe what a healthy cardiac motion is and the differences seen in two medical conditions: cardiomyopathy and obesity. Then, we introduce some background on imaging techniques used in cardiology. This will bring us from a patient to a representation of its cardiac motion using medical images. Finally, we present the objectives: *Using sequences of images of cardiac motion, how can we automatically analyze and derive relevant information to help diagnosis, prognosis and therapy planning?*

1.1 Introduction

Clinical practice is usually divided in three steps. First, the diagnosis, where a potential disease is identified by the clinician using the information available about

the patient and his symptoms. Then the prognosis, where a set of likely outcomes of the evolution of the patient's disease is predicted: it usually includes the expected duration, the impact on main functions of the body, and potential symptoms of the disease. Finally, therapy planning, where all previous information are put together in order to plan one or several possible treatments, which have the best chances of curing the patient or reducing the symptoms of the disease.

In many cases, medical images are a very important source of information that can help clinicians in the process of these three steps and they are getting increasingly used in clinical practice. In the context of cardiac imaging, it gives clinicians knowledge about the shape of the heart (when acquiring a single image) and its function (when acquiring a series of images or a video of the beating heart with multiple frames). To compare those acquisitions over several time periods (or longitudinal study) is essential to keep track of the progression of a disease, and will help with decision making and therapy planning.

Automatic analysis of medical images, using computational methods, can give clinicians additional insight on what cannot be seen by the human eye or calculated directly. A simple example of how these tools can help clinical workflow is the computation of important quantitative parameters, such as ejection fraction (derived from the size of the ventricles) or strain values. Algorithms can also help process big amount of data, which a single clinician could not. This can, for example, allow the comparison of data of one patient with a large database of subjects with known evolution and therapy in order to find similar cases. They can also replace the clinician by carrying out tasks that could be done manually, but are too time-consuming within a very time-constrained clinical workflow.

The last decades have seen spectacular advances and progress in computational methods for the automatic analysis of medical images using computer vision techniques, especially in cardiology. Among the most important developments led by medical image analysis, there is the automatic extraction of the geometry of the the heart (both left and right ventricles), which is called *Segmentation* [Heller 2002, Zheng 2008, Ecabert 2008], and the automatic evaluation of cardiac motion or *Cardiac Motion Tracking* [Zerhouni 1988, Tobon-Gomez 2013], using *Registration* techniques. These methods have already been applied to many clinical problems [Ferre 1999, Zhang 2004, Goshtasby 2005, Norouzi 2014]. However, they have yet to be disseminated widely in clinical routines, due to their lack of robustness and automation. The purpose of this thesis is therefore to improve those methods, in the context of cardiac motion analysis. To do so, we propose contributions to both the technical methodologies and their applications to multiple different diseases with a focus on cardiomyopathies and obese children.

1.2 Healthy Cardiac Structure and Function

The main function of the heart is to pump blood through the body. Even though the concept is straightforward, the biological machinery involved in this process is




Figure 1.1: Percentage of the population (both male and female) reporting heart problems in major European countries. Taken from [Townsend 2016].

very complicated, especially considering the amount of blood and the frequency at which it has to function. A normal heart rate at rest ranges from 60 to 100 beats per minute in adults with a stroke volume - the amount of blood pumped at each heartbeat - of around 70 mL for each ventricle. This gives an idea of the important constraints that the heart has to endure during a human life. The consequences of a heart failure are often dramatic: shortness of breath, excessive tiredness, leg swelling and possibly death. It affects between 2% and 10% of the whole population (see Fig. 1.1) and it is the number one cause of death in many countries, accounting for approximately 25% of deaths. Therefore, it is of crucial importance to monitor heart function and to understand what separates an efficient heart from a deficient one. We present here the main characteristics of a typical heart structure and cycle. Then, we show how it differs in a pathological condition.

1.2.1 Cardiac Structure

In this section, we describe the heart and its anatomy. A detailed review of the cardiac structure can be found in [Anderson 2004]. The heart is located slightly to the left of the chest between the lungs. It is cone-shaped (see Fig. 1.2) with the base (the top of the heart) positioned upwards and tapering down towards the apex (the bottom of the heart). An average adult heart has a mass of 250 to 350 grams [Robb 1942], it is typically of the size of a fist: 12 cm long, 8 cm wide, and 6 cm thick. Its size can vary between individuals, especially for those doing sports and exercises who tend to have larger hearts with a more efficient pumping mechanism.

The heart has two sides (left and right), both divided into two parts: atrium and ventricle making up four chambers in total. The atria are smaller than the ventricles and have thinner, less muscular walls than the ventricles. They are connected to veins, which role is to transport blood from the heart to the rest of the body, and act as chambers where blood transits. The ventricles are larger and stronger pumping chambers that send blood out of the heart, and are connected to arteries that carry blood away from the heart (Fig. 1.2).

The wall of the heart is made of four layers: the *epicardium/pericardium*, the *myocardium* and the *endocardium*. The epicardium is the outermost layer and is a thin layer of serous membrane. It helps lubricate and protect the outside of the heart. An additional layer, the *pericardium* (a thin fibrous sac that does not contract), further protects the heart and isolates it from other organs. The second layer, the myocardium, is the muscular layer of the heart: the largest one and the one that makes up for the majority of its mass. It is the most important layer, it consists of cardiac muscle tissue responsible for the contraction of the heart and therefore the pumping of the blood through the organism. Finally, the last innermost layer, the endocardium keeps the blood from sticking to the inside of the heart. It is a simple squamous endothelium layer that lies at the border with the blood pool.

Right and left ventricles have different functions. The left side of the heart sends blood all the way to the extremities of the body in the systemic circulatory loop while the right side is responsible for the pulmonary circulation to the nearby lungs. This difference in function results in a difference in muscle thickness, with the left ventricle being 15-mm thick - in order to have enough strength to push blood to the whole body - and the right ventricle being only 5-mm thick, thus making it more difficult to identify and delimitate in medical images.

1.2.2 Cardiac Cycle

The heart is a non-stopping pumping machine, repeating over and over again the same cardiac cycle: the period that starts the beginning of a heartneat and ends at to the beginning of the next. Its frequency is described by the heart rate, which is typically expressed in beats per minute. A single cycle of cardiac activity (or heartbeat) can be divided into two basic phases: diastole (roughly one third of the whole cycle) and systole. These two phases describe an alternating sequence of contraction and relaxation of the myocardium, in which blood gets delivered into both the systemic and pulmonary circulations [Boron 2012].

These two phases are further divided into sub-phases (see Fig. 1.4):

- **Isovolumic Contraction:** ventricular depolarization causes the ventricles to contract and pressure inside the ventricles to increase rapidly. Immediately after the start of the contraction, pressure in the ventricles exceeds that of in the atria, causing the atrioventricular valves to close. Pressure in the ventricles is lower than in the aorta and the pulmonary artery. Therefore, all the valves are closed and no blood can be ejected. During this phase, the ventricular




Figure 1.2: Anatomy of the heart. Taken from <https://www.utdlab.com/>.

volume remains unchanged. In the electrocardiogram (ECG), it is identified by the so-called QRS complex, which is the most visible part of the ECG signal and corresponds to the depolarization of the left and right ventricles.

- Ejection: pressure in the left ventricle exceeds that of the aorta. Similarly, pressure in the right ventricle goes beyond that of the pulmonary trunk, causing the pulmonary valves to open. Both the left and right ventricles eject blood to the aorta and pulmonary arteries while the atrioventricular valves are still closed. Ejection starts rapidly and slows down as systole progresses. It corresponds to the ST segment of the ECG: first the ventricles are completely depolarized, then the T wave appears with the repolarization in the second half of this phase. This is the end of the systole and the beginning of the relaxation with the diastole.
- Isovolumic Relaxation: the ventricles relax, causing a rapid drop in ventricular pressure. Shortly after, pressure in the ventricles goes below that of the pul-




Figure 1.3: Two complete cardiac cycles, together with their main events, as well as their corresponding pressure and volume curves. *Taken from <https://www.fastbleep.com/biology-notes/1/573>.*

monary artery and the aorta, and blood returns toward the ventricles, causing the semilunar valves to close while the atrioventricular valves are still closed. It corresponds to the end of the T wave of the ECG.

- **Passive ventricular filling:** the atrioventricular valves open and the ventricles are filled with blood coming from the atria. This passive phase accounts for most of ventricular filling. Pressure in ventricles does not change significantly, and volume increase is compensated by ventricular relaxation. No electrical activity can be found on the ECG.
- **Active ventricular filling:** the atria contract and complete ventricular filling. Only a small amount of blood entered the ventricles during this phase. Pressure in both ventricles is close to zero. Electrical activity is non-existent at the beginning of the phase, then depolarization of the atrial can be detected with the T wave, resulting in atrial contraction.

1.3 From Healthy to Pathological Heart

The work of this thesis was conducted within the MD-Paedigree project which is presented in this section. Then, we describe the clinical context centered on two




Figure 1.4: Schema representing a cardiac cycle with the active and passive systolic and diastolic phases. Taken from <http://highered.mheducation.com/>.

medical conditions affecting children and young adolescents: cardiomyopathies and obesity, and we show how these conditions affect the cardiac structure and function.

1.3.1 MD-Paedigree Project

MD-Paedigree aims at providing decision support to clinicians when treating their young patients in four areas: cardiomyopathies, obesity-related cardiovascular disease, juvenile idiopathic arthritis and neurological neuromuscular diseases. It is strongly embedded in The Virtual Physiological Human (VPH) [Hunter 2010, Hunter 2003, Ayache 2006] framework, an European initiative which focuses on building methodology and technology that enable collaborative investigation of the human body as a single complex system. This clinically-driven project regroups 7 world-renowned clinical centers of excellence: Ospedale Pediatrico Bambino Gesù (Italy), University College of London (United Kingdom), Istituto Gianna Gaslini (Italy), Deutsche Herzzentrum Berlin (Germany), Katholieke Universiteit Leuven & University Hospital Leuven (Belgium), Stichting Vu-VUmc (Netherlands) and Universtair Medisch Centrum Utruchet (Netherlands). This European project brings together these clinical centers with multiple technical centers (among which Inria and Siemens). This collaborative work between clinicians and computer scientists aims at improving interpretability of pediatric bio-medical information, data and knowledge by developing together a set of reusable and adaptable multi-scale

models for individualized, more predictive, effective and safer pediatric healthcare.

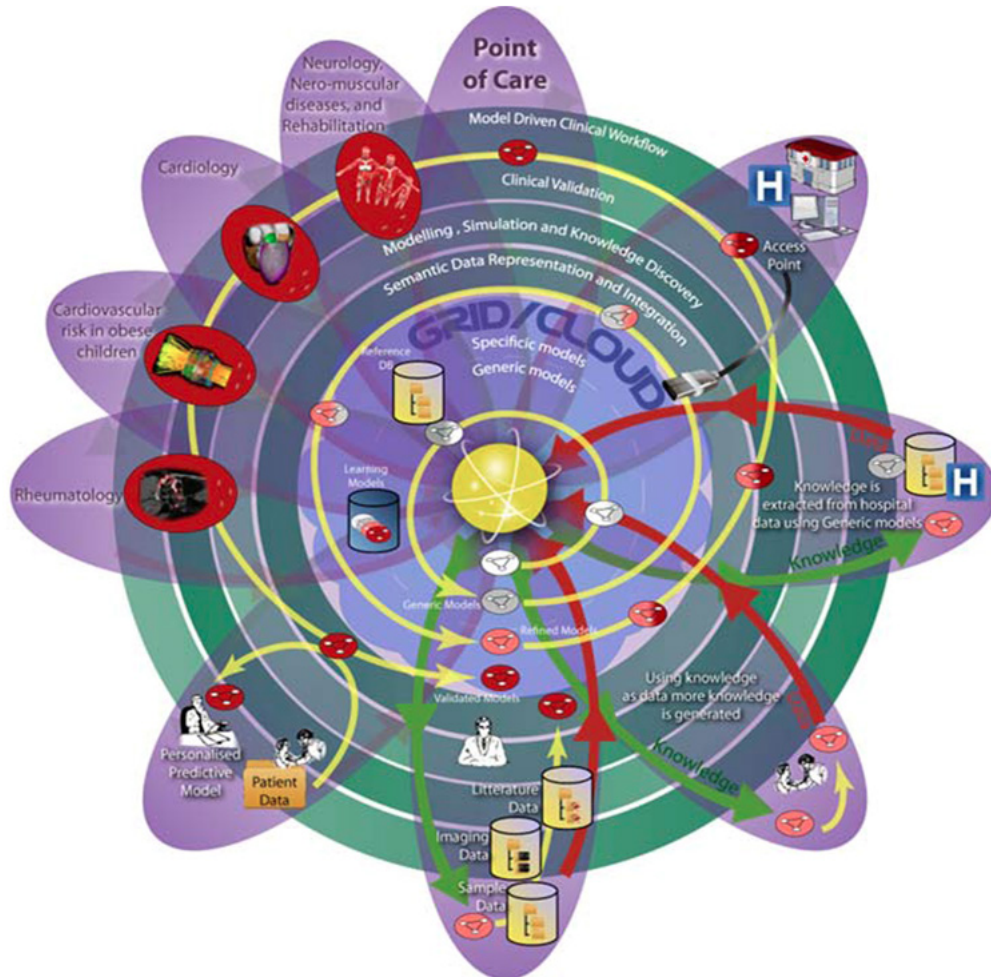


Figure 1.5: Diagram of the SOKU-Vision connecting patients and clinical centers to the models and the technical centers. On top-left the four medical conditions studied in the project (this thesis focuses on two: Cardiomyopathies and Obesity). The yellow spiral representing the models goes from the center - the simplest models - and turns around to become more and more specific and efficient. As they improve, models are validated in the different centers and integrated in the clinical workflow. Taken from <http://www.md-paedegree.eu/>.

MD-Paedegree validates and brings to maturity patient-specific computer-based predictive models of various paediatric diseases. It aims at increasing their potential acceptance in the clinical and biomedical research environment by making them readily available not only in the form of sustainable models and simulations, but also as newly-defined workflows for personalised predictive medicine at the point of care. These tools can be accessed and used through an innovative model-driven




Figure 1.6: (Left): Summary of ESC 2008 classification [Elliott 2007]. HCM: hypertrophic cardiomyopathy; DCM: dilated cardiomyopathy; ARVC: arrhythmogenic right ventricular cardiomyopathy; RCM: restrictive cardiomyopathy. Taken from [Elliott 2012]. (Right) : Schematic representation of the different types of cardiomyopathies as proposed in [Marcus 1982] and their effect on anatomy. Taken from Wikipedia.

infrastructure powered by an established digital repository solution able to integrate multimodal health data.

MD-Paedigree implements the so-called SOKU vision which is illustrated in Fig. 1.5. The aim of this framework is to make the design and development of innovative predictive models simpler to reuse and to integrate into a clinical context. It all starts with the basic models (in the center of the diagram) which are incubated in the system with progressive semantic enrichment and model modifications in a circular way represented by the yellow spiral. During this improvement, there is an increasing intervention of both automated database-guided learning and knowledge experts validation. Then, as these models reach the status when they can be used in clinical practice, they are applied and validated by the different clinical centers and clinical researchers. In this project, four different areas of clinical pathologies were studied by the partners: cardiomyopathies, cardiovascular risk in obese children and adolescents, juvenile idiopathic arthritis and neurological and neuromuscular disease. In this thesis, we work with the first two medical conditions and we present them in the following section.

1.3.2 Cardiomyopathies

Cardiomyopathies are a group of diseases impacting the heart leading to negative effects on the heart muscle size, shape, and function. It results in a deficiency of the heart to provide enough oxygenated blood to the rest of the body and remove carbon dioxide and other waste products. As the disease worsens and the heart weakens, classical signs and symptoms of heart failure usually occur which include: shortness of breath or trouble breathing, increased fatigue (tiredness), swelling in the ankles, feet, legs, abdomen, and veins in the neck.

There has been different ways to classify cardiomyopathies. One of the first proposed classification can be found in [Marcus 1982], in which cardiomyopathies are defined as *"a heart muscle disease of unknown cause"*. They are classified according to their pathophysiological phenotype into dilated cardiomyopathy, hypertrophic cardiomyopathy, or restrictive cardiomyopathy (Fig. 1.6 right). Since this first classification, more advanced classifications have been proposed to incorporate new types of diseases, such as arrhythmogenic right ventricular dysplasia (ARVD) [Elliott 2007, Elliott 2012], and to incorporate genetic mutation testing within the framework of classification (see Fig. 1.6 left).

The differences between the four main types of cardiomyopathies according to the European Society of Cardiology (ESC) [Elliott 2007] (see Fig. 1.6) can be summarized as follows [Lung 2016]:

- Hypertrophic cardiomyopathy: it is very common and it affects men and women equally of any age (1 out of every 500 people). It is caused by an enlargement and thickening of the heart muscle without any obvious cause. The ventricles, the lower chambers of the hearts and the septum thicken creating narrowing or blockages in the ventricles. The efficiency of the heart to pump blood to the body is decreased. It can also cause stiffness of the ventricles, changes in the mitral valve, and cellular changes in the heart tissue. While the cause is not always known, hypertrophic cardiomyopathy is often inherited genetically.
- Dilated cardiomyopathy: it affects 1 out of 2,500 persons. It is associated with left ventricular remodeling, which manifests as increases in left ventricular end-diastolic and end-systolic volumes, wall thinning, and a change in chamber shape to something more spherical and less elongated. Weakened chambers of the heart no longer pump efficiently, causing the heart muscle to work harder. Possible consequences are heart failure, heart valve disease, irregular heart rate, and blood clots in the heart. The causes may be alcohol, heavy metals, coronary heart disease, cocaine use, and viral infections and it can also be inherited from a person's parents.
- Restrictive cardiomyopathy: in opposition to other types, walls of the heart do not thicken but ventricles become stiff and rigid. It causes ventricles not to relax and not to fill with the normal blood volume. With the progression of the disease, the ventricles do not pump as well and the heart muscle weakens. It can lead to heart failure and problems with the heart valves. Possible causes of this disease are amyloidosis, hemochromatosis, and some cancer treatments.
- Arrhythmogenic right ventricular dysplasia: it is a rare type of cardiomyopathy and it occurs when the muscle tissue in the right ventricle is replaced with fatty or fibrous tissue, possibly leading to disruptions in the heart's electrical signals and arrhythmias. It usually affects teenagers and can cause sudden cardiac arrest in young athletes.




Figure 1.7: (Left-Top): The International Classification of adult underweight, overweight and obesity according to BMI as defined by the World Health Organization (WHO). Taken from [Organization 1987]. (Left-Bottom): list of medical complications due to obesity. (Right): Prevalence rates of obesity in child of major developed countries for boys and girls. OECD data.

1.3.3 Obesity

Obesity is a medical condition in which abnormal or excessive fat has been accumulated to the point where it can have a negative impact on health. Assessment of the obesity of a person is done using the Body Mass Index (BMI). The BMI is defined as the body mass divided by the square of the height and it is universally expressed in units of kg/m^2 . The value of the BMI defines a classification of a person from underweight to obese using a benchmark proposed by the World Health Organization (WHO) (see Fig. 1.7 left).

Obesity is a growing public health concerns in major developed countries. In 2008, approximately 35% of adults aged 20+ were overweight (BMI superior to 25 kg/m^2) and 10% of men and 14% of women in the world were obese (BMI superior to 30 kg/m^2). The worldwide prevalence of obesity worldwide has nearly doubled between 1980 and 2008 with almost 1 billion persons according to data from the Organisation de coopération et de développement économiques (OECD). Relating to children, 43 millions were estimated to be overweight and obese in 2010 (WHO data) while 92 millions were at risk. The worldwide prevalence of childhood overweight and obesity has been increasing from 4.2% in 1990 to 6.7% in 2010 and is expected to reach 9.1% in 2020 (Fig. 1.7).




Figure 1.8: Relative risk of death over 10 years by BMI for white men (left) and women (right) who have never smoked in the United States. *Taken from [Berrington de Gonzalez 2010]*

Obesity is one of the main preventable causes of death worldwide. More than 2.8 million people die each year as a result of being overweight or obese. It is associated with various diseases, particularly cardiovascular diseases, diabetes mellitus type 2, obstructive sleep apnea, certain types of cancer, osteoarthritis and asthma and cause multiple medical complications (Fig. 1.7 left-bottom). Studies [Berrington de Gonzalez 2010] have found a strong dependency between mortality risk and BMI (see Fig. 1.8). It is lowest at a BMI of 18-25 kg/m^2 , with risk increasing with lower and higher BMI. A high BMI has a strong negative impact on the life expectancy and it increases the mortality risk significantly: pre-obesity (BMI of 25-30 kg/m^2) reduces life expectancy by six to seven years, obesity class I (BMI of 30-35 kg/m^2) by two to four years, while severe obesity (BMI ≥ 40 kg/m^2) by ten years.

1.4 Cardiac imaging

To classify an individual's condition into separate and distinct categories that allow medical decisions about treatment and prognosis to be made, clinicians have data concerning the patient's medical condition. These data can take multiple forms: the symptoms reported by the patients, his medical history, his etiological data and the physical examination of the patient. When a cardiac problem is suspected, it is of crucial importance to analyze the myocardium motion in order to detect functional abnormalities and lesions. To do so, multiple non-invasive imaging techniques have been developed giving clinicians a vision of the heart anatomy and motion. These medical images give clinicians additional data to improve their decisions. We present the three main imaging techniques used in cardiology: Magnetic Resonance Imaging (MRI), Echocardiography (Echo), and Computed Tomography (CT) with a particular focus on MRI which will be the main modality used in this manuscript.




Figure 1.9: 2D and 3D echo images of the heart with the 3 different views.

1.4.1 Echocardiography

An echocardiogram (see Fig. 1.9) is an imaging technique based on the application of ultrasound (standard two-dimensional, three-dimensional, and Doppler ultrasound). Ultrasounds are sound waves whose frequencies are way higher than those audible to humans ($>20,000$ Hz) and ultrasonic images (also known as sonograms) are constructed by sending ultrasound pulse with a probe. The sound echoes off with varying degrees depending on the tissue. Then, these echoes are recorded and displayed as an image to the operator. Because of the difference in the reflection of the sound, the ultrasound image is able to differentiate between internal body structures such as tendons, muscles, joints, vessels and internal organs.

Echocardiography is one of the most widely used diagnostic tool in cardiology because of its ease of use, low-cost and rapidity [Belohlavek 1993]. It is used as a very efficient and fast tool to assess the cardiac anatomy and function of a patient [Schiller 1989]. The Echo acquisition is done with the patient shirtless, lying down on the left side. The clinician places the ultrasound probe using a gel in order to favor the transmission of the waves through the skin. Additionally, ECG electrodes can be used in order to gate the acquisition with the cardiac rhythm. Then, multiple 2D temporal acquisitions are done on different locations. An acquisition takes from 10 to 30 minutes and is completely painless and safe which is a major quality of Echocardiography.




Figure 1.10: 12 short-axis slices of a cardiac MRI from apex (top-left) to base (bottom-right).

A recent progress in echocardiography is the development of $3D$ temporal echocardiography. Traditional approach relies on multiple $2D$ acquisitions synchronized together using the ECG gating in order to reconstruct a $3D$ volume. However, a more advanced technique permits real-time three-dimensional echocardiography where the $3D$ acquisition is done in just one heartbeat. This technique has been applied successfully to identify structural abnormalities related to monomorphic ventricular tachycardia [Goland 2008] or to assess aortic valve area in aortic stenosis by continuity equation [Poh 2008].

The main advantages of the Echo is that it is easy to use, low-cost and completely non-invasive making it the principal imaging modality in cardiology. These qualities come at the cost of a relatively high noise to signal ratio, the presence of artifacts, and a low tissue contrast which might impair the diagnosis and decrease the quality of the assessment of the cardiac function. Therefore, other imaging techniques are often used as a complement to the Echo acquisition.

1.4.2 Magnetic Resonance Imaging

Magnetic resonance imaging (MRI) (see Fig. 1.10) is a relatively recent medical imaging technique. It was developed in the 70's notably by Paul C. Lauterbur in 1971 [Lauterbur 1973] and Sir Peter Mansfield [Mansfield 1977]. In 2003, they both received the Nobel Prize in Medicine for their "*discoveries concerning magnetic resonance imaging*", although differences in tissue relaxation time values were already known from the 50's [Odeblad 1955]. MRI allows the acquisition of images of the anatomy of the body in a non-invasive manner. Its use in cardiology has been increasing over the last decades for disease detection, diagnosis, and treatment monitoring, for example to assess congenital heart disease [Razavi 2003].

The technology relies on strong magnetic fields, radio waves, and field gradients to compute images of the inside of the body. More precisely, a strong and uniform magnetic field (most of the scanners used in clinics operate at 1.5 or 3 Teslas but systems with power ranging from 0.2 to 7 T are available, mostly for research purpose) aligns the spin of the hydrogen atoms. Then, an additional magnetic field,




Figure 1.11: (Left): A 1.5T MRI scanner: entering such a machine can be unpleasant for those who are claustrophobic or otherwise uncomfortable with the imaging device surrounding them. Taken from <https://www.med-ed.virginia.edu>. (Right): Principle of ECG-gated acquisition. A R-R interval on ECG, representing 1 cardiac cycle, is divided into N (here 8) frames of equal duration (A). Image data from each frame are acquired over multiple cardiac cycles and stored separately (B). When all data are added together, each frame represents a specific time of the cardiac cycle. These images are processed, and clinical indices such as the volume curve can be extracted. (C). Taken from [Paul 2004].

the pulse, is overlaid to re-orient the aligned spins. This field is disabled and the time it takes for the spins to realign with the magnetic field is detected. This time depends on the environment and the chemical nature of the molecules which allows the scanner to reconstruct an image showing different contrasts on the biological structures. Further details can be found in [Liang 1999].

Traditional MRI acquisition is modified for the acquisitions of cardiac images in several ways. The adaptation of MRI to cardiac images is often referred to as Cardiovascular magnetic resonance imaging (CMR) [Pennell 2004, Bogaert 2005]. The main difference of cardiac MRI over classical MRI applied to other organs is the necessity to acquire a whole temporal sequence of images spanning a cardiac cycle. Such an acquisition faces multiple challenges. There is a necessity to have a common starting point, for the sake of comparison between different patients. It is also needed when one wants to stack multiple 2D slice acquisition in one 3D sequence (see Fig. 1.10). Also, because such an acquisition of multiple 2D slices can take time, the motion of the lungs during the respiratory cycle can produce motion artifacts in the image, thereby making the alignment of the 2D slices challenging.

To overcome these problems, CMR uses Electrocardiography (ECG) to synchronize the image acquisition [Nacif 2012]. ECG records the electrical activity of the




Figure 1.12: Different views of the heart acquired using Cardiac MRI and their relations with either a focus on the ventricles or the valves. Usual procedure to reconstruct a 3D cardiac sequence of the beating heart is to stack multiple multiple Short-Axis (SAX) 2D slices together.

heart using electrodes placed on the skin. The very small electrical changes on the skin coming from the depolarization during each heartbeat is detected, and similar electrical patterns correspond to similar points in the cardiac cycle. The R wave of the ECG (the most prominent wave of the QRS complex and the most easily detected) is used as a reference point corresponding to the end of the diastole and the start of the systole. Data acquisition is initiated after a given delay following the R wave and images are created from data collected over a series of cardiac cycles (R to R intervals) [Paul 2004]. ECG gating allows for stop motion imaging by acquiring data only during a specified portion of the cardiac cycle, typically during diastole when the heart is not moving. The patient is usually asked to hold his breath during imaging in order to alleviate respiratory motion. This can be difficult, especially for children. Therefore many artifacts such as slices misalignment are present after the

acquisition and they have to be corrected as a post-processing step.

MRI is widely used in hospitals and clinics for medical diagnosis, staging of disease and follow-up without exposing the body to ionizing radiation. MRI does not involve X-rays, as opposed to computed tomography (CT scan). While the hazards of X-rays are now well-controlled in most medical contexts, MRI can still be seen as superior to CT in this regard. Compared with CT, the defaults of MRI scans are: the acquisition takes more time, the machines are louder, and they require the subject to go into a narrow tube which can be a source of discomfort. In addition, people with medical implants or other non-removable metal inside the body may be unable to safely undergo an MRI examination.

1.4.3 Computed Tomography

Cardiac Computed Tomography [Schoenhagen 2005, Topol 2007] (also known as CT scan) is a non-invasive imaging modality that uses X-rays to take many detailed 2D pictures of the heart. Multiple X-rays images are combined to produce cross-sectional views of the body. Cardiac CT can be used to visualize the heart anatomy, coronary circulation, and great vessels (including the aorta, pulmonary veins, and arteries). In some cases, the patient is injected with an intravenous dye (iodine) before the acquisition. This contrast medium is a chemical substance that reveals what is happening inside the hollow parts of the body [Ropers 2003, Morin 2003] (such as the blood vessels, the stomach, bowel or even the fluid around the spinal cord) on the images.

The scan takes between 15 minutes and 1 hour to complete, including preparation time. Similarly to MRI sequence of cardiac motion, multiple temporal frames of CT images are acquired and ECG gated to reconstruct the whole sequence. CT images benefit from better spatial resolution (see Fig. 1.13) than MRI and echo images with isotropic voxel size in each direction. Therefore, there is no slice gap as with MRI images and no slice misalignment. With multi-slice CT scans [Kachelrieß 2000, McCollough 1999, Pan 2004], up to 256 slices can be acquired at the same time and the heart can be imaged during an entire cardiac cycle. But this higher quality comes at the cost of large amount of X-rays irradiation, which can be dangerous and carcinogenic [Donnelly 2001, Pearce 2012, Schulze 2014]. This is why CT acquisitions are rarely done with children and were not acquired for our project.

1.5 Manuscript Organization and Objectives

In this thesis we present a whole technical pipeline of cardiac motion analysis from the instant a patient enters the hospital to an image based group-wise statistical study of its cardiac motion. In the introduction, we have described the different diseases that will be the focus of this work. We presented the different cardiac imaging modalities which will be the source of our imaging data. The rest of the manuscript



Figure 1.13: 3D CT image with the 3 different views. Voxel resolution is better compared to MRI images and is isotropic in each direction.

presents a whole process starting from the images to a statistical study of the cardiac motion with a focus on applications to clinical problems. The organization of the manuscript is schematically presented in Fig. 1.14. We now describe the 3 main parts of this manuscript, each of which will be driven by a key question.

Part - I

From Medical Images to 3D Shape

The accurate segmentation of the myocardium is of great importance in clinical practice. This is usually the first step of medical image analysis. The most commonly used clinical indices (e.g. the ejection fraction) to evaluate a cardiac disease are derived from the information given by the segmentation of the myocardium during a cardiac motion. Therefore, clinicians often perform manual segmentation of the contours of the myocardium as the first step in the analysis of cardiac images. However, considering the typical work-flow that clinicians are facing, manual segmentation is often too time-consuming and not a viable option. Furthermore, experts can have different opinions as to how to perform the segmentation (inclusion/exclusion of papillary muscles and trabeculations for example) leading to high inter-rater variability of manual segmentations. This makes the comparison between different patients not segmented by the same expert often difficult. This expert variability can be decreased by developing automatic methods which require reduced or no user input. For all these reasons, there is a need to develop a completely automated, fast and robust segmentation method.

Going a little bit outside the direct clinical practice, most of the methods of cardiac function analysis require a segmentation of the myocardium. Cardiac motion tracking algorithms, that will be the topic of Part II, usually require an initial segmentation of the myocardium at the first frame. This segmentation can be used to track the motion and monitor the main clinical indices during the cardiac cycle. It can also be used to give a local parametrization and estimation of the motion (with the AHA regions for example, as will be done in Chapter 7) for group-wise local comparison between different patients. Finally, bio-mechanical and electrophysiological simulations also rely on tetrahedral mesh segmentation of the whole




Figure 1.14: The global organization of the manuscript with a pipeline for cardiac motion analysis. First, from the images, a tetrahedral segmentation is extracted in Part I. Then, cardiac motion tracking is performed on the sequence of images in Part II. Finally, we show applications of this representation in diagnosis and longitudinal analysis in Part III.

myocardium. These simulations can be used for therapy planning and for quantitative understanding of the functioning of the heart in healthy and pathological subjects.

This non-exhaustive list of applications of segmentation leads to our first objective with the question driving part I of this manuscript:

- *Can we develop a fast, automatic and accurate segmentation of the myocardium represented by tetrahedral mesh for shape analysis, cardiac motion tracking, and bio-mechanical simulations?*

The main objective is to compute the segmentation of the myocardium. The method we use to get the segmentation rely heavily on registration methods. We actually show that the two main topic of medical imaging: registration and segmentation are intrinsically related and can be used together to perform one task or the other. In Chapter 2, we present the background of registration in medical imaging with the focus on two different algorithms that are used extensively in the remaining of the manuscript: *Log-Demons* for image registration and *LDDM on Currents* for shape registration. In Chapter 3, we use shape segmentation matching to train a predictive registration algorithm and we develop a new fast and robust algorithm for image registration. Finally, in Chapter 4, we go the other way around and use image registration to compute shape segmentation by developing an automatic, robust and fast multi-atlas segmentation pipeline. Results show that our method has good accuracy and robustness while being faster than traditional multi-atlas segmentation algorithms.

The main contributions developed in this part are:

- A method for computing reference transformations between pair of images, using mesh segmentations which are registered in the space of currents. These transformations can be used efficiently to train a learning-based registration algorithm.
- A fully convolutional neural network for 3D registration prediction. Our architecture is able to detect global features and deformations that could not be detected with a sliding-window approach (for ex. [Yang 2016]). It also proves to be faster at testing time as only one pass of the whole image is required.
- A fast and robust end-to-end framework for myocardium segmentation leveraging the speed of the SVF-net registration with a large number of atlases.

Part - II

Low-Dimensional Representation of Cardiac Motion

Part I was focused on the segmentation of the myocardium using the image of the first frame of the sequence (end-diastole). In this part, we develop methods to efficiently compute the motion of the heart during a cardiac cycle. The main clinical

application is to analyze the cardiac deformations at a local or global level. This can provide an estimate of the cardiac function efficiency and possibly indicators of heart disease. Tracking the tissue with physical markers is not directly possible in a clinical setting because it is invasive. Therefore, noninvasive methods relying on medical images have been developed in recent years to track automatically the myocardium.

The aim of this part can be summed up with the general question:

- *Can we improve cardiac motion tracking algorithms in order to have a more accurate estimation of the main clinical indices and an efficient representation?*

In Chapter 5, we first give an overview of the traditional methods for cardiac motion tracking. Then, we propose an innovative approach to study the cardiac motion. Standard methods rely on the registration of each frame with respect to the first frame taken as a single reference. We challenge this methodology by considering multiple references in the cardiac sequence. We build a *Barycentric Subspace* using these references and use it in the registration process as an additional a-priori to improve the tracking along the sequence. Results show that this innovative approach leads to substantial improvement of the tracking accuracy at end-systole and the estimation of the ejection fraction. Then, in Chapter 6, we present the polyaffine algorithm, a reduced-order model for tracking cardiac motion. This algorithm has the advantage not only to track the motion but also to give a representation with a low-dimensional number of parameters that have physiological meaning. We introduce local basis and reduced-order affine parameters and show that they give an efficient representation of the cardiac motion with better interpretability. We also use this representation to build a reduced-order model of the cardiac motion for fast simulation and personalization.

The main contributions developed in this part are:

- The introduction of a new method for dimension reduction and low-dimensional subspace analysis: *Barycentric Subspace Analysis* [Pennec 2015] in the context of medical images.
- The methods for computing the coordinates of an image within a *Barycentric Subspace*, for choosing the reference frames building the optimal subspace, and for reconstructing an image given the coordinates and the references.
- An extension of the polyaffine parameters framework with an additional parameters reduction by keeping only the 6 most relevant parameters and the use of this parametrization of the motion to build a very fast reduced-order model of the cardiac function.

Part - III

Application to Diagnosis and Longitudinal Analysis

Previous parts were focused on the improvement of the computational methods of the heart in order to compute a robust, accurate and automatic low-dimensional representation of the cardiac motion. In this last part of the manuscript, we show how this representation can be used in several clinical problems to improve diagnosis, prognosis and therapy planning. We have seen how cardiac motion tracking can provide crucial information such as the ejection fraction and strain values to the cardiologist. But these are simple numbers which do not encompass all the relevant information of the full deformation of the myocardium during the cycle. A representation of the motion that goes beyond the classical clinical parameters, while still represented by a low number of parameters, could be used efficiently to improve the diagnosis of a disease.

Once a patient is diagnosed and its present cardiac condition evaluated, a clinician would like to have information about the possible evolution of the disease in order to chose the most adapted therapeutic option. One possible way to provide this information is to perform a longitudinal analysis of the motion. This kind of study is particularly challenging since two time-dimensions have to be coupled: the time of the cardiac motion (heartbeat) and the evolution of this motion through different time points. Furthermore, to build such a model, one needs to gather multiple cardiac motions of the same patient over a large period of time (to detect measurable differences). Waiting for years to acquire the data makes it difficult to build large enough database, especially when working with children who tend to have a large opt-out rate of clinical studies. For these reasons, there has been very few studies on longitudinal analysis of cardiac motion, even though this has been a trending topic in brain analysis for many years.

This leads to the key question:

- *Can we perform group-wise statistics on a low-dimensional representation of the cardiac motion to improve diagnosis, prognosis and therapy planning?*

This part is divided in 3 chapters focusing each on a different clinical question and applied each to a different population. In Chapter 7, we classify two populations of healthy subjects and infarcted patients based on a prior reduction of dimensionality of the motion using a reduced number of polyaffine parameters. We show that our low-dimensional representation allows for very good classification results using classical machine learning algorithms. Furthermore, we are able to quantify the importance of each of the parameters in the classification. These parameters represent clinical indices that can be understood by clinicians and provide insights into what is the main impact of an infarct on the motion. In Chapter 8, we use the Barycentric representation of the motion presented in Chapter 5 as an efficient cardiac motion signature. We apply this representation to two populations, one of healthy subjects and one of patients with Tetralogy of Fallot (ToF). We show that the signature of each population presents significant differences, corresponding in particular to a longer systolic duration for the ToF population. Finally in Chapter 9, we perform two longitudinal analyses. The first analysis concerns a population of adolescents with cardiomyopathies. The evolution of their cardiac motion has

been studied for two time points spanning one year. The second analysis was an experimental study of the evolution of the cardiac motion in response of a clinical test with several acquisitions spanning 2 hours. Healthy subjects and obese patients followed this experimental protocol and we compare their respective responses.

The main contributions of this part are:

- The use of a low-dimensional representation with polyaffine transformations to classify a population of hearts with myocardial infarction. The method favorably compares with state-of-the-art in term of classification accuracy (AUC, sensitivity and accuracy) inline with the best competitive methods.
- The group-wise analysis of the features extracted from the projection and its application in the context of the study of Tetralogy of Fallot.
- The longitudinal analysis of the cardiac motion for two medical conditions: cardiomyopathies and obesity using a low-dimensional representation based on polyaffine transformations.

Perspectives

In the last chapter of this thesis, we conclude by summarizing the key contributions of this work. We also recap the key objectives and goals of the work. Finally, we propose some perspectives for future works in order to extend the methods which are described in this thesis. Some of these suggestions are already under development in collaboration with other PhD students.

1.6 Publications

The contributions described led to 3 journal papers (1 accepted, 1 accepted subject to minor revisions and 1 in preparation) and 6 conference papers.

1.6.1 Journal Articles

- **M.-M. Rohé**, M. Sermesant, and X. Pennec. Low-Dimensional Representation of Cardiac Motion Using Barycentric Subspaces: a New Group-Wise Paradigm for Estimation, Analysis, and Reconstruction. *Accepted subject to minor revisions at Medical Image Analysis* Chapter 5 and Chapter 8.
- **M.-M. Rohé**, M. Sermesant, and X. Pennec. SVF-Net: Learning Deformable Image Registration Using Intensity and Geometric Features for the Automatic Segmentation of the Myocardium. *In preparation for submission to a journal.* Chapter 3 and Chapter 4.
- A. Suinesiaputra, **M.-M. Rohé**, et al. Statistical shape modeling of the left ventricle: myocardial infarct classification challenge. *IEEE Journal of Biomedical and Health Informatics*, 2017.

1.6.2 Selective Peer-Reviewed Conference Paper

- **M.-M. Rohé**, M. Sermesant, and X. Pennec. Barycentric Subspace Analysis: a new Symmetric Group-wise Paradigm for Cardiac Motion Tracking. *In MICCAI 2016 - the 19th International Conference on Medical Image Computing and Computer Assisted Intervention (oral presentation)*.
- **M.-M. Rohé**, N. Duchateau, M. Sermesant, and X. Pennec. Combination of Polyaffine Transformations and Supervised Learning for the Automatic Diagnosis of LV Infarct. *In Statistical Atlases and Computational Modeling of the Heart (STACOM 2015), Munich, Germany, 2015*. Chapter 7.
- **M.-M. Rohé**, R. Molléro, M. Sermesant, and X. Pennec. Highly Reduced Model of the Cardiac Function for Fast Simulation. *In IEEE - IVMSP Workshop 2016, Image, Video, and Multidimensional Signal Processing Workshop (IVMSP), 2016 IEEE 12th, Bordeaux, France*. Chapter 6.
- **M.-M. Rohé**, M. Datar, T. Heimann, M. Sermesant and X. Pennec. SVF-Net: Learning Deformable Image Registration Using Shape Matching. *In MICCAI 2017 - the 20th International Conference on Medical Image Computing and Computer Assisted Intervention*. Chapter 3.
- R. Molléro, D. Neumann, **M.-M. Rohé**, M. Datar, H. Lombaert, N. Ayache, D. Comaniciu, O. Ecabert, M. Chinali, G. Rinelli, X. Pennec, M. Sermesant, and T. Mansi. Propagation of Myocardial Fibre Architecture Uncertainty on Electromechanical Model Parameter Estimation: A Case Study. *In Functional Imaging and Modeling of the Heart, FIMH 2015, Maastricht*.
- S. Jia, C. Camaioni, **M.-M. Rohé**, P. Jaïs, X. Pennec, and M. Sermesant. Prediction of Post-Ablation Outcome in Atrial Fibrillation Using Shape Parameterization and Partial Least Squares Regression. *In Functional Imaging and Modeling of the Heart, FIMH 2017, Toronto*.

1.6.3 Awards

- A *Student Travel Award* was awarded at the conference MICCAI 2016 for the paper: Barycentric Subspace Analysis: a new Symmetric Group-wise Paradigm for Cardiac Motion Tracking.

Part I

**FROM MEDICAL IMAGES TO
3D SHAPE**

Registration in Medical Imaging

Contents

2.1	Introduction	27
2.2	Registration Algorithms	28
2.2.1	Similarity Measures	30
2.2.2	Transformation Spaces and Regularization	31
2.2.3	Optimization Methods	32
2.3	LCC Log-Domain Diffeomorphic Demons	33
2.3.1	Optimization Method: Alternate Optimization	33
2.3.2	Deformation Parametrization: Stationary Velocity Field	34
2.3.3	Similarity Metric: Local Correlation Coefficient	35
2.4	Framework of Currents for Shape Registration	36
2.4.1	The Formalism of Currents for Surface Representation	38
2.4.2	Surface Registration Using Currents and LDDMM	40

This chapter introduces the concept of registration which will be extensively used in the rest of this manuscript in different applications related to cardiac images analysis. The process of registration consists in finding correspondences between a pair (or a group) of images or meshes. Along with the segmentation of organs, it has been one of the principal challenges in medical imaging and a very important tool for analysis. We first introduce the concept, the applications in medical imaging and the mathematical formalism. Then we go deeper and present two specific different registration algorithms which we use in the rest of this work. This is not meant to be a full review of registration methods (a comprehensive review can be found for instance in [Sotiras 2013] and more recent advances in the field are discussed in [Schnabel 2016]) but an overview of the principal possible choices in order to position the methods that we use with respect to the state-of-the-art.

2.1 Introduction

Registration is a key instrument in computational anatomy and has gained an increasing importance in the past years. Many applications of registration have been developed and we give here some of the main ones. The registrations of images between different subjects (inter-patient registration) can be used to find

relevant differences in their morphology and relate them to a specific disease [Smith 2002, Cash 2004] in a *cross-sectional analysis* or to construct an average model of the anatomy and its variability [Rueckert 2001]. Another common application is the analysis of the evolution in time - *longitudinal analysis* - using multiple acquisitions of the same subject [Serag 2012, Ashburner 2013, Lorenzi 2013b, Hadj-Hamou 2016b] to study the impact of a specific disease or therapy. In *multi-atlas segmentation*, the segmentation (and more generally the labels) associated with one or multiple images are transported using registration and combined for a new patient [Lorenzo-Valdés 2002, Aljabar 2009, Iglesias 2015]. Finally, as the last application of this non-exhaustive list, we can cite the *parallel transport* of images in order to study them in a common space [Lorenzi 2013b, Pennec 2011]. The goal of registration is to find a transformation that aligns a pair of spatially dependent data. We define the formalism with the example of images [Zitova 2003], but we ask the reader to keep in mind that the same formalism can be used with little changes in the context of the registration of other data structures (landmarks, curves, surfaces, volumes...): we then refer to the registration of shapes [Van Kaick 2011]. An example of shape registration is given in this chapter with the framework of currents and, in the following chapter, a method that combines both in an innovative way will be introduced.

In general, registration involves two images but the concept can be extended to a group of images [Bhatia 2004]. For example, in the context of cardiac motion tracking, it is of great interest to use the information present in the whole sequence. Therefore, there is a need to perform the task with all the images taken together [Mahapatra 2012a] rather than to do it independently for each frame. In this introduction, we will focus on the simple case with two images: the source or moving image M , and the target or fixed image F . These two images are defined in the image domain Ω and we suppose they are related by a transformation φ that we are looking to find. In chapter 5, we define methods to extend this framework from pair-wise to group-wise registration in the context of cardiac motion tracking using Barycentric Subspace.

2.2 Registration Algorithms

The task of finding the optimal transformation mapping the two images together is traditionally cast as an optimization problem of an energy \mathcal{E} of the form:

$$\begin{aligned} \mathcal{E}(M, F, \varphi) &= \mathcal{S}(F, M \circ \varphi) + \mathcal{R}(\varphi), \\ \varphi &= \arg \min_{\varphi \in \mathcal{T}} \mathcal{E}(M, F, \varphi). \end{aligned} \quad (2.1)$$

This objective function has two distinctive terms. The first one, the similarity term \mathcal{S} evaluates how well the fixed image F and the warped moving image $M \circ \varphi$ are aligned. A standard measure in the case of images is the *Sum of Square Difference* (SSD) but we will see many more complicated choices for this metric. The second




Figure 2.1: An example of registration of a 3D cardiac short-axis image. The moving image is mapped to the fixed image with the deformation field (blue). One can notice the presence of fat on the bottom of the heart in the moving image and not in the fixed image. These differences of texture can prove challenging for a registration algorithm, and cause large deformations to be evaluated.

term, the regularization \mathcal{R} encodes the a-priori on the smoothness and properties of the transformation set by the user. It can take multiple forms as well, but usually favors small deformations over large ones.

The moving image is warped through the action of the transformation φ which acts on all the points x of the image domain Ω . One common way to characterize the transformation is through the displacement field u which is defined as φ minus the identity transformation: $\varphi(x) = x + u(x)$. In the registration, the transformation φ is optimized within a set of possible transformations \mathcal{T} . The choice of the space \mathcal{T} depends on the level of detail we want the transformation to capture. It is strongly related to the choice of the regularization. If the set of transformation has few degrees of freedoms and is parametrized by a small number of parameters θ that produce smooth deformations (for example rigid transformations), then regularization might not even be necessary. However, if the parametrization of the transformation is dense, then there is a need of a strong regularization for the problem to be well-posed.

Finally, one needs a method to solve the optimization. The choice of the optimizer is important as well, as different methods will lead different results. This is due to the inherently ill-posed nature of the registration problem. If the transformation is densely parametrized (one vector of 3 scalars per voxel), then the information in the image (one scalar per voxel) does not give enough constraints versus the number of parameters. Even in the very restricted case of rigid transformation (6 parameters in 3D), the problem of registration can be ill-posed. One can think, for example, of matching two spheres with the same radius one to another. The optimal solution would be defined up to a rotation. Also, given the non-linearity and non-convexity of the objective functions, the algorithm will always reach a local minimum. For these reasons, the choice of the optimization scheme and its initialization is of great importance to ensure to get a relevant solution.

Therefore, a registration algorithm is mainly defined by these three components: (i) the energy criteria with the similarity term to match the data, (ii) the choice of

the transformation space and a regularization term to constraint the problem, (iii) an optimization scheme to solve it. We review the main possible choices for these components and their impact of the registration process.

2.2.1 Similarity Measures

Following the classification of [Sotiras 2013], we can distinguish three different type of choice for the similarity criteria. Firstly, geometric approaches [Gil 2010, Morel 2009, Chui 2003], which look for the correspondences between landmarks of the images. Landmarks are chosen as to correspond to salient image location, in order to be easily detected by the algorithm. Once these landmarks have been extracted in both images, the solution of the registration can be found by minimizing the energy on the, generally small, number of landmarks. One advantage of this method is its relative robustness to large deformations. Indeed, as landmarks are detected on the whole image, large deformations mapping the images can be found as long as the respective landmarks are detected correctly. Among the disadvantages, one can cite the difficulty of detecting landmarks in a reliable way. Finding correspondences between landmarks is also a difficult problem and errors in the correspondences can greatly impact the final results. Furthermore, the landmarks give only a sparse set of correspondence, which does not allow to find dense transformation of the space and requires the use of interpolation between the landmarks. Therefore, the accuracy of the transformation found decreases as the distance to the landmarks increases.

Secondly, iconic methods evaluate the accuracy of the registration by looking at a criteria based on the intensity of the images. With respect to geometric approaches, they evaluate an energy criteria based on the dense image grid rather than a sparse set of landmarks. This might lead to a better quantification of the transformations but it comes at the cost of increased computational expense. This computational cost often prevents these methods to search extensively to all possible parameters. Therefore, they might not be able to find large deformations and get stuck in a local minimum of the criteria. Also, since the criteria is computed on the whole images, the registration can be influenced by part of the images with little information, instead of being driven by the main landmarks of the anatomy. The most common similarity measures is the sum of squared (SSD) differences which is a sound choice under the assumption that intensity between images differ only by a Gaussian noise process. As images are often acquired with different machines and setups, this might not be a realistic assumption so that it is important to normalize the intensity of both images. Other metrics can also, often more efficiently, deal with this problem. For example, correlation coefficient (CC) assumes an affine relationship between intensities. The relation of the intensity between both images can be more complex, for example in multi-modal registration. For these cases, approaches derived from information theory [Maes 1997, Viola 1997] use a criterion called mutual information (MI) which does not assume any particular relationship between the image intensities. It is also possible to combine multiple similarity mea-

tures, in [Cifor 2013] each voxel is described by three image features: intensity, local phase, and phase congruency. Finally, recent advances in deep learning have proposed to directly learn the similarity criteria in a machine learning approach [Simonovsky 2016, Cheng 2016].

Thirdly, hybrid methods that use both types of information in order to combine the advantages and prevent the drawbacks of each method. One way to do it is to use the geometric information as an initialization. Then, each information is taken into account in a separate and sequential way. The geometric information is used first to align both images and get the large deformations, increasing the robustness of the method, then an iconic methods follows to refine the matching [Yin 2010, Postelnicu 2009]. However, there is no guarantee that the correspondence found in the first step will be conserved in the second step. Another way to combine both methods is to use the geometric information as a constraint on the algorithm [Roche 2001, Cachier 2003, Biesdorf 2009, Lu 2010]. For example by adding a penalization on the similarity criteria or by constraining the transformation to lie within a subspace that is consistent with the result of the geometric approach.

2.2.2 Transformation Spaces and Regularization

The choice of the transformation space and the regularization is of great importance and comes down to a trade-off between the complexity of the deformation and the computational efficiency. The degrees of freedom of the transformation space encodes the richness of the description of the transformation and related parameters will have to be estimated by the algorithm. The number of parameters (or degrees of freedom) that a registration algorithm can estimate in a robust way depends closely on the data to noise ratio of the images. When millions parameters have to be estimated (for example non-parametric dense transformations), one needs to add a regularization term for the problem to be well-posed. We review quickly the most common choices for the transformation spaces, starting from the most constrained to the less constrained.

The simplest type of transformations are the rigid (or linear) transformations [Ourselin 2000, Ashburner 2007b]. They include translations and rotations. In 3D they are parametrized with a vector of 3 parameters for the translation and 3 angles for the rotations around each axis. Regularization penalization can be added to encode prior knowledge on the parameters, although the model is usually supposed to be constrained enough. Extension of rigid transformations to incorporate scaling and shearing can be done with affine transformations. These transformations are encoded by a 3×4 matrix acting on the homogeneous coordinates (12 parameters: the parameters of the rigid transformation plus 3 parameters of scaling on each axis, 3 parameters of shearing and 3 parameters of rotation). Then, one can combine affine transformations in order to reduce the constraint and model more complicated transformations. In [Feldmar 1996], the authors introduce locally affine transformations that are attached at each point of the surface. In piecewise-affine transformations [Pitiot 2006], these affine transformations are defined at a

regional level of the image. Finally, with polyaffine transformations [Arsigny 2005], the regions are now defined with weights on the image. The weight at each point of the image can be interpreted as the probability of a point to belong in a specific region. These affine-transformations have the advantage to be parametrized with a reduced number of parameters which can be interpreted as regional translation/rotation/scaling/shearing.

Many other parametrization of non-rigid (or elastic) deformations exist and instead of giving an exhaustive list, we give here properties that are of great important for biomedical imaging. The first important property of a parametrization is the diffeomorphism: the transformation is invertible and both the function and its inverse are differentiable. One important consequence of this property is that these deformations preserve the topology of the images. The transformation is always one-to-one and no crossings can appear. A main example of such diffeomorphic parametrization is the Large Deformation Diffeomorphic Metric Mapping (LDDMM) framework [Beg 2005]. In LDDMM, diffeomorphisms are parameterized using the end points of the flow of a time varying velocity fields according to the Euler-Lagrange transport equation. Another important property of the parametrization is the symmetry. This means that changing the order of the moving and fixed image should gives the inverse transformation. Additionally, one would like to have a simple way to compute the inverse of the transformation. Even though the LDDMM framework gives diffeomorphic transformation, the algorithm is not symmetric and there is no wasy way to compute the inverse. To make the algorithm symmetric multiple solutions have been proposed [Beg 2007, Avants 2008]. In [Arsigny 2006], the author proposes to consider Stationary Velocity Fields (SVF) as opposed to time-varying ones. It provides an efficient formulation of diffeomorphic registration and an easy way to compute the inverse, while still maintaining all the desirable properties of LDDMM. On the other side, not all diffeomorphisms can be represented with a SVF.

2.2.3 Optimization Methods

The choice of the optimization strategy to solve equation 2.1 can have an important impact on the result. Indeed, due to the complexity of the problem (especially for non-rigid registration), it is unrealistic to search exhaustively the whole subspace of possible transformations. One has to reduce the space on which to minimize the function which might lead to find a sub-optimal solution or local minimum. Therefore, an optimization strategy has to be defined that is able to scan a sufficiently large part of the transformation space while still being computationally tractable.

Optimization methods are generally classified in two separate categories: continuous optimization which deal with variables that take real values and discrete optimization where variables are taken from a discrete set. A popular example of a registration method using discrete values is the discrete Markov random field [Shekhovtsov 2008] and the discrete optimization is usually done using a graph-based method [Boykov 2001]. In [Heinrich 2016], the authors represent the image domain as multiple comprehensive supervoxel layers and the optimisation acts on

the coarse scale representation of supervoxels. The case where the parameters are continuous variables is more common. The most common way to solve the energy is gradient descent and its extensions. Starting from an initial guess of the transformation (usually identity), the algorithm follows the direction given by the gradient of the energy which is supposed to be differentiable [Klein 2007]. The step size at which one follows the gradient direction has to be chosen efficiently as to adapt as one gets closer to the solution. In conjugate gradient descent, the knowledge of the previous computed gradients are used as well to compute the next update. As another example of many adaptation of gradient descent, one can cite stochastic gradient descent that alleviates the computational burden of computing the exact gradient with an approximation [Klein 2009b].

2.3 LCC Log-Domain Diffeomorphic Demons

In this section, we present the *LCC Log-Domain Diffeomorphic Demons* [Lorenzi 2013a] algorithm that will be used with several applications through this manuscript and also extended to other problems such as multi-reference registration (Chapter 5). This algorithm builds on the demons forces proposed in [Thirion 1998]. Inspired from the optical flow equations [Barron 1994], non-parametric non-rigid registration is considered as a diffusion process with the demons forces that push according to local difference of intensities between both images. The methods alternates between these forces and a regularization with a Gaussian smoothing kernel.

2.3.1 Optimization Method: Alternate Optimization

[Pennec 1999] reformulates the method in order to provide a better theoretical framework of what were heuristic and intuitive ideas. It is shown that the update demons forces can be considered as an approximation of a second order gradient descent on the sum of square of intensity differences. The regularization with Gaussian smoothing is considered as a minimization problem of an energy criteria. We present here the framework of [Cachier 2003, Vercauteren 2008] where the problem is cast into the minimization of a well-posed criterion by introducing a hidden variable c in the optimization for point correspondences. The interest of this variable is to separate the optimization into two easily tractable sub-problems. Each iteration walks towards the optimum of the global energy:

$$E(M, F, c, s) = \frac{1}{\sigma_i^2} Sim(F, M \circ c) + \frac{1}{\sigma_x} dist(s, c)^2 + \frac{1}{\sigma_T^2} Reg(s),$$

σ_x accounts for a spatial uncertainty on the correspondences, Sim is a similarity metric that classically can be SSD such that $Sim(F, M \circ c) = \|F - M \circ c\|^2$ but we will see an extension to a more advanced and robust metric, $dist(s, c) = \|c - s\|$ and $Reg(s)$ can have multiple forms depending on the problem.

In the classic demons algorithm, the optimization is performed within the complete space of non-parametric transformations using additive updates of the form

Algorithm 1 Exponential $\phi = \exp(v)$ with scaling and squaring.

- 1: Input: Velocity field v .
 - 2: Choose N such that $2^N v$ is close to 0.
 - 3: Scaling: Scale velocity field $\phi = 2^N v$.
 - 4: **for** N times **do**
 - 5: Squaring: Square $\phi = \phi \circ \phi$
 - 6: Output: Diffeomorphic map $\phi = \exp(v)$.
-

$c = s + u$. The demons forces solves for the similarity energy with respect to u :

$$E_{SSD}^{sim} = \frac{1}{\sigma_s^2} Sim(F, M \circ (s + u)) + \|u\|^2.$$

For the regularization, different choice of Reg will lead to different solutions. One common choice is the Tikhonov regularizer proposed in [Cachier 2004]:

$$Reg(s) = \int_{\Omega} \sum_{k=1}^{\infty} \alpha_k \left(\sum_{i_1 + \dots + i_k = k} \|\partial_{i_1} \dots \partial_{i_k} s\|^2 / (\sigma_d^{2k} k!) \right),$$

which leads to an efficient and easy way to solve the regularization step by convolution with a Gaussian smoothing kernel: $s = G_{\sigma} * c$.

2.3.2 Deformation Parametrization: Stationary Velocity Field

One of the limitations of the deformation field parametrization is the lack of constraints to ensure that the transformations computed are one-to-one or invertible. Also this parametrization does not force any smoothness whereas it is of considerable interest to work with transformations that preserve properties such as smoothness of curves, surfaces or other features associated to anatomy. In [Vercauteren 2008] the authors propose to use Stationary Velocity Fields (SVF) [Arsigny 2006] instead of deformation fields used for the additive demons to parametrize the transformations. We present now the formalism of SVF which are derived from time-varying velocity fields from the LDDMM framework.

The large deformation diffeomorphic metric mapping setting (LDDMM) [Beg 2005] provides a framework where diffeomorphisms are parameterized using the end point of the flow ϕ of a time varying velocity fields $\phi'(t) = v_t(\phi_t)$ such that the transformation φ can be found as $\varphi(x) = x + \int_0^1 v_t(\phi_t) dt$. While time-varying velocity fields provide a sound framework to parametrize diffeomorphisms, it is very computationally intensive and difficult to implement in practice. In [Arsigny 2006] the authors propose to simplify the condition and use Stationary Velocity Fields (SVF), which provide a more efficient framework while still maintaining all the desirable properties of LDDMM. The one-parameter subgroup of a SVF v is defined as the unique solution of:

$$\frac{\partial \phi(x, t)}{\partial t} = v(\phi(x, t)),$$

with initial condition $\phi(x, t) = id$. The transformation φ is then defined as the Lie group exponential map $\exp(v) = \phi(x, 1) = \phi(x)$ and can be efficiently computed with the scaling and squaring algorithm (see Algorithm 1). The use of SVFs leads to a good compromise between theory and efficiency for computationally tractable registrations. The composition of transformations in the log-space are done using the Baker-Campbell-Hausdorff (BCH) formula [Bossa 2007]:

$$\begin{aligned} BCH(v, w) &= \log(\exp(v) \circ \exp(w)) \\ &= v + w + \frac{1}{2}[v, w] + \frac{1}{12}[w, [v, w]] - \frac{1}{12}[w, [v, w]] + \dots \end{aligned} \quad (2.2)$$

In practice, for numerical efficiency, one usually keeps only the first two terms of the decomposition which gives already a good approximation if one of the terms is small. The optimization problem can then be written in the log-space with $c = \exp(v_x)$ and $s = \exp(v)$ as:

$$E(M, F, v, v_x) = \frac{1}{\sigma_i^2} Sim(F, M \circ \exp(v_x)) + \frac{1}{\sigma_x} \|v_x - v\|^2 + \frac{1}{\sigma_T^2} Reg(v).$$

2.3.3 Similarity Metric: Local Correlation Coefficient

Finally, [Lorenzi 2013a] introduces the local correlation coefficient (LCC) in the log-demons framework replacing the classical SSD similarity measure. SSD has the benefit to be simple to implement numerically but is very sensitive to intensity biases in the images, which can be very important if the images were acquired with different machines or under different settings. In the case the bias is simply a global affine relation between the intensities of the images, it can be either pre-processed prior to the registration. Or criteria such as the (normalized) correlation criteria can account for multiplicative and additive bias [Dong 1995]. But it might not be sufficient if the relation is more complex and local.

In this case, the LCC can provide a more robust metric by estimating the local affine scaling parameters of the intensities (additive plus multiplicative) [Pennec 2003, Avants 2008] and was used in several successful registration algorithms as a good trade-off between the simple SSD and the very unconstrained Mutual Information. Considering the fixed image F , and its local mean image $\bar{F} = G_\sigma * F(x)$ defined by Gaussian smoothing G_σ with kernel size σ . The LCC is defined as:

$$\rho(F, M) = \int_{\Omega} \frac{\overline{FM}}{\sqrt{\overline{F^2 M^2}}}.$$

The LCC similarity varies between -1 (perfect negative correlation), 0 (no correlation) and 1 (perfect correlation). It measures how the intensities of the two images are correlated within a local neighborhood of size σ . In [Lorenzi 2013a], the authors show that the LCC similarity measure can be efficiently used in the Log-Demons framework with a solution of the update $\delta v = v_x - v$ of the similarity energy:

Algorithm 2 LCC Log-Demons algorithm.

- 1: Input: A fixed image F and a moving image M . (Optional: an initialization of the velocity field v_0).
 - 2: **for** n until convergence **do**
 - 3: Given the current estimation of the deformation $\varphi_n = \exp(v_n)$, compute δv using the symmetric forces in Eq. 2.3, the fixed image F , the moving image M and the current warped moving image M_n and warped fixed image F_n ,
 - 4: For fluid-like regularization, smooth the update with a Gaussian Kernel: $\delta v = G_\sigma * \delta v$,
 - 5: Compose the update with the current estimation $v_x = BCH(v_n, \delta v)$ using BCH Eq. 2.2,
 - 6: For diffusion-like regularization, smooth the update with a Gaussian Kernel: $v_x = G_\sigma * v_x$,
 - 7: $v_{n+1} = v_x$, compute the exponential of v_{n+1} and $-v_{n+1}$ using Algorithm 1, and $M_{n+1} = M \circ \exp v_{n+1}$, $F_{n+1} = F \circ \exp -v_{n+1}$,
 - 8: Output: Diffeomorphic transformation φ parametrized by a SVF: $\varphi = \exp(v)$.
-

$$\delta v = \frac{2\Lambda}{\|\Lambda\|^2 - \frac{4\sigma_i^2}{\rho^2\sigma_x^2}}, \quad (2.3)$$

with:

$$\Lambda = \left(\frac{G_\sigma * (F\nabla G^T)}{G_\sigma * (FG)} - \frac{G_\sigma * (G\nabla F^T)}{G_\sigma * (FG)} + \frac{G_\sigma * (G\nabla G^T)}{G_\sigma * (G^2)} - \frac{G_\sigma * (F\nabla F^T)}{G_\sigma * (F^2)} \right).$$

Therefore, the LCC criterion preserves the structure of the original Log-Demons algorithm. In particular, it is still symmetric (the demons forces are computed in a symmetric way) and the SVF parametrization is diffeomorphic.

Finally, it should be noted that the algorithm allows for multi-resolution in order to look at the transformation at different scales. One main benefit, is that larger transformation can be estimated avoiding the problem of being stuck in a local minima. Another benefit of multi-resolution is to speed up the registration. The higher levels of the multi-resolution scheme are faster because the sizes of the images processed are lower. Usually, 3 levels of multi-resolution are used. The LCC Log-Demons algorithm is summed up in Algorithm 2.

2.4 Framework of Currents for Shape Registration

A shape $S \in \mathcal{S}$ is a generic geometrical object that can be represented by different structures of data. The most simple one is the *Point Cloud* where the shape is represented by a set of points $(x_i)_{i \leq n}$. If the number of points is fixed and each label i supposed to represent a specific location in space, then we will usually refer to them as *Landmarks*. If we add a connectivity matrix linking pairs of points (or vertices),




Figure 2.2: Two example of shapes from segmentation of cardiac images. (Left): short axis view of a surface segmentation of the myocardium: LV in red and RV in green. (Right): 3D view of a volumetric segmentation of the myocardium

we get *Curves* and one can cite for example the representation of white matter bundles in medical imaging [Durrleman 2011]. The *Surface Meshes* are triangular meshes in 3D made of a set of vertices and a connectivity matrix linking triplets of vertices. Similarly *Volumetric Meshes* are tetrahedral meshes with connectivity matrix linking quadruplets of vertices. In the context of cardiac imaging, these two structures are used to represent segmentations of the myocardium (see Fig. 2.4). The surface mesh representation is often used in statistical shape analysis of the heart [McLeod 2013b, Frangi 2002, Bruse 2017b, Bruse 2017a] whereas the volumetric mesh representation of the myocardium is used for bio-mechanical model [Bestel 2001, Marchesseau 2013] of the heart function.

The formalism presented before in the context of medical image registration can be extended with little change to the registration of shapes, with no regards of the specific representation chosen. Let's consider two shapes, a moving M and a fixed F . Similarly to image registration, we consider the registration as the problem of finding the transformation φ minimizing the energy:

$$\begin{aligned} \mathcal{E}(M, F, \varphi) &= \mathcal{S}(F, M \circ \varphi) + \mathcal{R}(\varphi), \\ \varphi &= \arg \min_{\varphi \in \mathcal{T}} \mathcal{E}(M, F, \varphi). \end{aligned} \quad (2.4)$$

The main challenge when considering shapes is the difficulty to derive a similarity metric. In the case of shapes represented by landmarks, some simple similarity metric can be defined and the aligning transformation can be recovered as the solution of a system of equations constructed from the established correspondences [Guo 2004]. However, in the general case of a point clouds with no correspondence (and generally a different number of points), the correspondence problem itself is far from trivial, especially in the case of strong deformations. Rigid registration can be solved efficiently with the *Iterative Closest Point* (ICP) algorithm, which updates




Figure 2.3: A triangulated mesh computed from a segmentation of the myocardium of a MRI image (blue). At each face (triangle) of the surface is associated a Direct delta currents whose vector is the normal at the surface and the norm is the area of the surface.

the transformation by finding the closest point on the target shape at each iteration [Besl 1992, Chen 1992]. For non-rigid diffeomorphic transformations with an infinite number degree of freedom, one has to define a relevant metric between shapes. In this context, the framework of currents [Vaillant 2005, Glaunes 2005] provides an unifying method to process any sets of shapes whatever the chosen representation (sets of points, curves, surfaces...) with a metric that does not assume any kind of correspondence between structure. We briefly summarize the mathematical theory of the currents and then present the algorithm.

2.4.1 The Formalism of Currents for Surface Representation

We present the formalism of currents in the specific case of surfaces, and refer the reader to [Durrleman 2010] for a comprehensive review of currents in the generic case. Mathematically, a current is a continuous linear mapping $\mathcal{L}_W(w)$ from a vector space W to \mathbb{R} , i.e. it is an application that integrates vector fields. For any surface S , we can measure the flux of a vector field $w \in W$ going through this surface:

$$\mathcal{L}_W(w) = \int_{x \in S} \langle w(x) | n(x) \rangle d\sigma(x), \quad (2.5)$$

where $n(x)$ is the normal of the surface at the point x , and $d\sigma(x)$ is the infinitesimal surface element at x . The shape S is then characterised by how it integrates these vector fields. A common analogy made is the one of a 3D scanner used to acquire the geometry of an object. It digitalises the geometry by probing its surface with laser beams. The object is reconstructed through the reflection of the beams on the surface which can be thought as an equivalent of the mathematical integration of a flux given by Eq. 2.5. A surface is now embedded in the vector space of currents which means usual linear operations such as the addition or subtraction are now possible.

The space of test vector field W has not yet be defined, it will have a great impact of how the currents will represent the surfaces. In [Glaunes 2005], the authors propose to define W as a vector space generated by a Gaussian kernel $K_W(x, y) = \exp(-\|x - y\|^2/\lambda_W^2)$ which turns W into a reproducible kernel Hilbert space (RKHS). W is the dense span of basis vector fields which can be written as $w_x^\alpha(y) = K_W(x, y)\alpha$. The vector α represents the direction of the basis field and y represents its position. The kernel K_W defines an inner product in W with the formula:

$$\langle w_x^\alpha, w_y^\beta \rangle_W = \alpha^T K_W(x, y)\beta.$$

There is a close relation between the RKHS W and the space of currents W^* which is its dual: a basis vector field $w_x^\alpha(y) = K_W(x, y)\alpha$ in W is associated to a Dirac delta currents $\delta_x^\alpha = \mathcal{L}_W(w_x^\alpha)$ in W^* . Intuitively, a Dirac delta current is an infinitesimal vector α that is concentrated at the spatial position x . We can decompose the current \mathcal{T} of the surface S as an infinite sum of Direct delta currents defined infinitesimally at each point of the surface and oriented along the normal.

In practice, surfaces will be represented by a finite set of triangles (triangulated meshes). Then the representation will be a finite sum of Dirac delta currents:

$$\mathcal{T} = \sum_k \delta_{x_k}^{\alpha_k}(w),$$

where x_k are the barycentres of the mesh faces and α_k their normal (Fig 2.3). The vector field w dual of the current $\mathcal{T}(w)$ is the spatial convolution of every normal vector α_k with the kernel K_W , $w(x) = \sum_k K_W(x, x_k)\alpha_k$. The scalar product in W^* for Delta dirac currents derives directly from the scalar product in W :

$$\langle \delta_x^\alpha(w), \delta_y^\beta(w) \rangle_{W^*} = \langle K_W(x, \cdot)\alpha, K_W(y, \cdot)\beta \rangle_W = \alpha^T K_W(x, y)\beta.$$

The distance and the scalar-product between two surfaces can be calculated using their basic representation as discrete sum of Dirac delta currents. An important computational property of the distance is its asymptotic behavior when the distance between x and y is low:

$$\|\delta_x^\alpha - \delta_y^\beta\|_{W^*}^2 = \|\alpha - \beta\|^2 + 2\|x - y\|^2/\lambda_W\alpha^T\beta + O(\|x - y\|^4/\lambda_W^4).$$

In this case the difference of orientation $\|\alpha - \beta\|^2$ will mainly drives the distance. On the other side, when the points x and y are more than a few λ_W apart, there is an upper bound of the distance:

$$\|\delta_x^\alpha - \delta_y^\beta\|_{W^*}^2 \leq \|\alpha\|^2 + \|\beta\|^2.$$

It means that delta currents located far away will have little impact on the optimization of the distance, as the penalty is almost constant, so that an efficient algorithm can focus on optimizing the distance of currents that are within a neighborhood of some λ_W .

We have defined a efficient metric on the space of surfaces represented by currents. This metric depends on the choice of the size of the kernel λ_W . This choice of the width λ_W is of great importance as it defines at which scale one looks at the features of the shapes. The larger the value of λ_W , the coarser the representation of the shape with currents. Difference of shapes at scale lower than λ_W will be considered as noise and will contribute weakly to the distance. This parameter sets the level of details we want to study and a low parameter will accounts for more details of the surfaces. On the other side, algorithm trying to match two fine representations of shapes with currents might be less robust computationally.

2.4.2 Surface Registration Using Currents and LDDMM

The transformation φ which maps the moving mesh M to the fixed mesh F is parametrized using the Large Deformation Metric Mappings (LDDMM) framework [Dupuis 1998, Miller 2002, Beg 2005]. φ is computed as the end-point of the integration of a time-varying velocity field $\phi'(t) = v_t(\phi_t)$ and $\varphi(x) = x + \int_0^1 v_t(\phi_t) dt$. This parametrization ensures that the transformation is diffeomorphic which means it is an invertible function and both the function and its inverse are smooth. One can compute the inverse transformation by following the reverse path.

The velocity belongs to a RKHS V generated by the Gaussian Kernel $K_V(x, y) = \exp(-\|x - y\|^2 / \lambda_V^2)$. The vector fields v_t are therefore defined as the convolutions of sets of moments vectors β^i by the kernel K_V . The regularization in the LDDMM framework is done by penalizing the length of the deformation trajectory. The energy of a velocity field at deformation $\phi_t(y)$ is defined using a right-invariant metric $\|v_t\|_{\phi_t}^2 = \|v_t \circ \phi_t^{-1}\|_V^2$, where $\|\cdot\|_V$ is the norm of the RKHS V generating the velocity fields, which the associated scalar product:

$$\langle v, w \rangle = \sum_{i,j} \beta_i^T K_V(x_i, x_j) \beta_j.$$

The energy of the trajectory is computed as $E(\phi_t) = \int_0^1 \|v_t \circ \phi_t^{-1}\|_V^2$. Deformations minimizing this energy are geodesics and therefore are entirely defined by their initial velocity fields $v_0(x)$ and the momentums β_i^0 through the geodesic shooting property [Miller 2006, Ashburner 2011].

The flow of deformations is entirely parameterized by a set of initial positions of control points \mathbf{c}_0 and initial momenta α_0 . Finally, the registration problem is cast into the minimization an energy made up of a matching term (the norm of the previously defined inner-product between currents representation of shapes) and a regularity term on the initial momenta (which defines entirely the transformation by geodesic shooting):

$$E(M, F, \alpha_0, \mathbf{c}_0) = \frac{1}{2\sigma_k^2} \|\varphi_1(M) - F\|_W^2 + \alpha_0^T \mathbf{K}(\mathbf{c}_0, \mathbf{c}_0) \alpha_0,$$

where σ_k is a trade-off parameter between regularization and matching, \mathbf{K} is a gaussian kernel. This is now a finite dimensional optimization problem which is

optimized using a gradient descent scheme (details of the implementation can be found in [Durrleman 2014]).

SVF-Net: Learning Deformable Image Registration Using Shape Matching

Contents

3.1 Motivations	43
3.2 Reference Deformations Computation	46
3.2.1 Currents for Shape matching	47
3.2.2 Elastic Body Spline Interpolation	47
3.2.3 Stationary Velocity Fields for Diffeomorphism Parametrization	48
3.3 SVF-Net: Fully Convolutional Neural Network Architecture	49
3.4 Experiments	50
3.4.1 Training	51
3.4.2 Evaluation	51
3.5 Conclusions	53

In Chapter 2, we introduced registration in medical imaging. We showed two examples of registrations algorithms, one working with images and one working with shapes. This chapter is meant to bridge the gap between shapes (usually resulting from segmentations of images) and images. We show how we can efficiently use segmentations and shape matching in order to learn deformable image registration.

Part of this work was submitted to the MICCAI 2017 conference and is currently under review. The version presented here is an extended version that is in preparation of a journal submission.

3.1 Motivations

Non-linear registration - the process of finding dense voxel correspondence between pairs of images - is a key instrument in computational anatomy and has gained an increasing importance in the past years. Many applications of registration have been developed ranging from the analysis of different subjects (cross-sectional analysis) to find relevant differences in their morphology, the analysis of the evolution in time

(longitudinal analysis) using multiple acquisitions of the same subject, atlas-to-image registration to perform an atlas-based segmentation or the parallel transport of images in order to study them in a common space.

Traditional methods to find the optimal deformation field mapping two images rely on the optimization of a matching criteria - controlling the local correspondence of the voxel intensities - together with a regularization term - making the problem well-posed and imposing regularity on the deformation to find. These methods usually have several drawbacks preventing them from being used in clinical practice. One main problem is the high computational cost (time and memory) of these methods as they often requires many iterations and evaluations of the energy function [Ceritoglu 2013]. Another problem is the possibility of the optimization to remain stuck in a local minimum because of the non-convexity of the objective function. Finally, these methods usually require to set up multiple parameters which can considerably impact the quality of the registration. While it may not be a problem for an experienced user, this may be unpractical to non-experienced users or to automatize the process for data with different levels of noise.

In an attempt to overcome these problems, new approaches to predict registration parameters based on machine learning have been proposed. The main challenge to predict such a dense deformation field is the non-linearity of the relation between the voxel intensities of the pairs of images and the parametrization of the transformation. In this context, Convolutional Neural Networks have set new standards for several tasks in computer vision where there is a need to predict highly non-linear function. Whereas these methods have gained large popularity for medical image segmentation and classification, they are still under-represented in the context of image registration.

This under-representation may be due to the inherent difficulty to provide a ground truth mapping between pairs of images. While it is possible for a human to classify an image or to draw the contours of the segmentation, the task of finding pairwise voxel matching is very difficult if not impossible, especially when the correspondances have to be found in a 3D image. Therefore, to train learning-based methods, we have to find alternative ways to define ground truth than manual input. One way is to compute synthetic images deformed with known transformations, but it is hard for these images to account for the inherent noise and artifacts present between pairs of medical images, often leading to over-simplistic synthetic examples which would not be as challenging as real pair of images. In [Yang 2016], the prediction is trained based on ground truth provided by a registration algorithm previously run on pairs of images. Doing so effectively speeds up the registration algorithm. However, the problems seen in the algorithm used for computing the ground truth will likely be learned by the learning method if one does not add additional information.

We chose to take another approach and train on reference deformations defined from the registration of previously segmented region of interests instead of the result of a registration algorithm on the images. Therefore, our approach does not try to replicate a classical registration algorithm matching voxel intensities, but learn




Figure 3.1: Overview of the method. From images with segmented ROIs (here cardiac images with segmented myocardium), we derive ground truth transformations parametrized by SVF (section 2). In section 3, we present our SVF-Net architecture for deterministic prediction. Finally, in section 4 we evaluate our method using dice scores and distances on surfaces.

to align more global and structural information of the images. This choice is supported by the fact that a good matching of intensities is not always correlated with physiologically plausible deformations. In a classical optimization approach, it is difficult to avoid minimizing a matching criteria based on the difference of voxel intensities. However, learning-based methods give us the opportunity to learn on different type of information than intensities of the images. For example, in the context of inter-patient cardiac registration, one is mostly interested in getting a very accurate matching of the contours of both myocardiums rather than a good intensity matching so it seems natural to use this information to define our references deformations. These deformations are then used to train a fully convolutional network building our registration method.

From a database of images where the region of interests have been segmented, we consider the segmentations as surfaces and we register these surfaces to a common template giving us dense correspondences between shapes. The framework of currents [Durrleman 2014] provides an elegant mathematical method to perform this task. Once the point correspondences between pair of images is defined, we use an elastic body spline to interpolate this deformation on the whole image grid. Finally, an iterative log-approximation scheme is used in order to derive a dense Stationary Velocity Field (SVF) [Arsigny 2006] which provides an efficient framework to parametrize diffeomorphisms. These reference SVFs are then used to train a convolutional neural network directly predicting the transformation between pair of images parametrized by a SVF. We compare our method with a state of the art

registration algorithm [Lorenzi 2013a] and show that, not only the accuracy is increased, but also it is more robust and way faster. The organization of the paper is summed up in Figure 3.1.

Contributions

- A method for computing reference transformations between pair of images, using mesh segmentations which are registered in the space of currents.
- As the shapes can be defined or corrected manually, our method provides an efficient way to introduce manual input in the computation of the deformations to train a learning-based registration algorithm. This is not the case with approaches that rely on synthetic images or the result of a registration algorithm on the images. This approach comes nearer to the classical definition of a ground truth defined manually.
- A fully convolutional neural network for 3D registration prediction. Our architecture is able to detect global features and deformations that could not be detected with a sliding-window approach (for ex. [Yang 2016]). It also proves to be faster at testing time as only one pass of the whole image is required.
- A more robust and faster registration method validated on a large database of 187 segmented cardiac images.

3.2 Reference Deformations Computation

In this section, we detail our methodology to derive a reference deformation parametrized by a dense SVF $v(x)$ mapping symmetrically two shapes together. We place ourselves in the context of shapes defined by surfaces, as this is a traditional output of segmentation algorithms, but the method defined here is generic and could also be applied to other types of data structures such as point clouds, landmarks, curves and volumes. From a database of N images I_n where the region of interests have been segmented, we consider the segmentations as surfaces S_n and we register these surfaces to a common template T giving us a deformation φ_n mapping the template to each of the segmented shape. Each point of the template T can then be thought of as a landmark which is mapped into each of the surface meshes of our database. Using the registration with respect to a common template avoids the need to do all the pairwise shape registrations thereby limiting the number of shape registration to N . Then, this correspondence with the template can be used for the N^2 pairs of images and the point correspondence between pair of images gives us a displacement field defined for the set of landmarks. We extrapolate this deformation field to the whole grid using spline interpolation (sec. 3.2.2). Finally, we choose to parametrize the deformations using SVF instead of the traditional deformation fields and explain this choice in section 3.2.3. The whole process is summed up in Algorithm 3 and we detail each of these steps in the following subsections.

Algorithm 3 Computation of the SVF symmetric parametrization of the registration between two images using ground truth segmentations.

- 1: Using the framework of currents, compute the transformations $\varphi_{1,2}$ between one surface mesh template T and both surface meshes S_1 and S_2 to get $\tilde{S}_1 = \varphi_1(T) \approx S_1$ and $\tilde{S}_2 = \varphi_2(T) \approx S_2$. Each point p_i of the template is mapped to $\varphi_1(p_i)$ and $\varphi_2(p_i)$ giving us a dense point correspondence between the pair of images.
 - 2: Compute the displacement field extrapolation $u_{1,2}(x)$ on the grid of the image, such that on each point p_i of T we have $u_1(\varphi_1(p_i)) = \varphi_2(p_i) - \varphi_1(p_i)$ (resp. with u_2).
 - 3: Initialize $v^0 = (u_1 - u_2)/2$.
 - 4: **for** i while convergence **do**
 - 5: Compute the exponential fields $w_1 = \exp(v^i)$ and $w_2 = \exp(-v^i)$
 - 6: Estimate the error of the transformations $e_1 = u_1 - w_1$ and $e_2 = u_2 - w_2$.
 - 7: Compose the SVF with the errors: $v^i = BCH(v^{i-1}, (e_1 - e_2)/2)$ where BCH is the Baker Campbell-Hausdorff approximation of the composition of SVF [Lorenzi 2013a]
 - 8: **return**: the SVF v mapping symmetrically the two segmentations .
-

3.2.1 Currents for Shape matching

We describe here the methodology to derive correspondences between pair of surfaces segmented from images. In our case, we want to avoid the burden of having to define landmark correspondences and directly learn in an automatic way from the surface segmentations. To do so, the framework of currents [Vaillant 2005] has gained increased popularity in the medical imaging field and allows to avoid a strict point-wise matching and treat the problem as true surface matching without the point correspondence issue. In this work, we will use an extension of currents to non-oriented surfaces called varifolds [Durrleman 2014]. We refer the reader to Chapter 2 for a description of the LDDMM registration using currents.

3.2.2 Elastic Body Spline Interpolation

From two shapes, both registered to the common template, a point The point correspondence between pair of images gives us a displacement field defined for the set of landmarks. To interpolate it to something defined on the whole image grid an elastic body spline interpolation is used. This choice is motivated because this interpolation is driven by a physical model, making it a natural choice for regions where no landmarks are found. Given the two sets of n deformed points from the template: $\varphi_1(p_i)$ and $\varphi_2(p_i)$, which are mapping the shapes S_1 to the shape S_2 , we look for a transformation u which fulfills the interpolation conditions: $u(\varphi_1(p_i)) = \varphi_2(p_i) - \varphi_1(p_i)$, for all $i = 1, \dots, n$.

The elastic body spline is a 3- D spline that is based on a physical model (the Navier equations) of an elastic material. It provides a curved mapping between




Figure 3.2: Example of a ground truth SVF (red vector field scaled at 0.3) computed from two segmentations. We show here the moving image together with the segmentation of the myocardium (blue) and the segmentation of the myocardium of the fixed image (orange). (Left): Short-axis view. (Right): Longitudinal view.

two 3-D spaces that is consistent with the physical model given this set of pairs of corresponding landmarks [Davis 1997]. Finally, the deformation field u parametrizes a transformation that approximately maps the shapes S_1 and S_2 , and the pair of images on the landmarks.

3.2.3 Stationary Velocity Fields for Diffeomorphism Parametrization

At this point we have a displacement field u which parametrizes the transformation φ such that $\varphi(x) = x + u(x)$. One of the limitations of the parametrization with displacement fields is the lack of constraints to ensure that the transformations computed are invertible and smooth whereas it is of considerable interest to work with transformations that preserve properties such as smoothness of curves, surfaces or other features associated to anatomy. Therefore, it is crucial to work with a parametrization that creates diffeomorphic transformations. In this context, the large deformation diffeomorphic metric mapping setting (LDDMM) [Beg 2005] provides an elegant framework where diffeomorphisms are parameterized using the end point of the flow ϕ of a time varying velocity fields $\phi'(t) = v_t(\phi_t)$ such that $\varphi(x) = x + \int_0^1 v_t(\phi_t) dt$.

While time-varying velocity fields provide a sound framework to parametrize diffeomorphisms, it is very computationally intensive and not so easy to implement in practice. In order to provide an efficient formulation of diffeomorphic registration, while still maintaining all the desirable properties of LDDMM, [Arsigny 2006] proposes to consider Stationary Velocity Fields (SVF) as opposed to time-varying ones. The one-parameter subgroup of a SVF v is defined as the unique solution of:

$$\frac{\partial \phi(x, t)}{\partial t} = v(\phi(x, t)),$$

with initial condition $\phi(x, 0) = id$. The transformation φ is then defined as the




Figure 3.3: Fully convolutional neural networks for 3D registration: The inputs are the fixed and the moving 3D images. The output is a dense SVF symmetrically mapping the two images defined on the initial image grid. The first layers are 2D to account for the non-isotropy of the input images (short-axis cardiac MRI acquisition usually have slice Z spacing 4 times bigger than the in plane X/Y spacing). The number of channels is doubled after each max-pooling layer.

Lie group exponential map $\exp(v) = \phi(x, 1) = \phi(x)$. The use of SVFs leads to a good compromise between theory and efficiency for computationally tractable registrations.

3.3 SVF-Net: Fully Convolutional Neural Network Architecture

Convolutional Neural Networks (CNN) provide a very efficient way to learn highly non-linear functions but these methods have yet to gain widespread dissemination in image registration. Traditional methods tackle the problem in a patch-based approach [Yang 2016], which are easy and fast to train but are looking only locally at each patch and therefore miss global information about the registration. For image segmentation, fully convolutional networks [Long 2015] have been developed in order to process the whole image in a stream, therefore having the advantage to also look at global features, instead of looking only locally at each patch. It also has the benefit to be faster in test time as there is only one prediction to perform for the whole image instead of predicting each patch individually in a sliding-window approach. In this work, we propose to adapt the fully convolutional architecture to the task of registration prediction by training on the dense ground truth SVF previously computed.

Figures 3.3 illustrates our proposed network architecture. Similar to the standard U-Net architecture [Ronneberger 2015], it has a contracting part where the network detects the features and an expanding path building the SVF parametrization of the




Figure 3.4: Boxplot of the similarity score between the ROIs. (Red): proposed method, (Blue): Log-Demons with LCC metric. The ROIs we look at are: Left Ventricle Blood Pool (LVBP), Left Ventricle Myocardium (LVM), Right Ventricle Blood Pool (RVBP) and Right Ventricle Myocardium (RVM).

registration. The input of the network is both images (moving and fixed) stacked together. In our application, we study cardiac MRI short-axis images, which are acquired with a non-isotropic resolution in the Z axis (slice spacing ranges from 6 to 10mm whereas in-plane spacing ranges from 1.5 to 2.5mm). To account from this discrepancy, the first two layers are 2D layers, then our features map is isotropic in all directions and we apply 3D layers for the 3rd and 4th layers. The features from the contracting path are directly concatenated to the deconvoluted features from the expanding paths. The number of features is double in the first convolution after the max-pooling in the contracting path, whereas we symmetrically halve the number of features in the expanding path. Finally, in the last layer, a simple 2D convolution builds the 3 layers corresponding to the SVF parametrization in the X, Y, Z axis at the same grid as the initial input images. Our architecture has a total of approx. 10 millions parameters.

3.4 Experiments

We test our method on the problem of inter-patients registration of 3D MRI cardiac images. Our dataset consists of $N = 187$ short-axis acquisitions of end-diastolic cardiac images. These data were acquired in multiple clinical centers (University College London Hospitals, Ospedale Pediatrico Bambino Gesù Roma and Deutsches Herzzentrum Berlin) and with different scanners causing a large inter-center vari-

ability which further increases the difficulty of the registration. For each of these images, the myocardium was segmented based on a data-driven machine learning approach [Georgescu 2005], from which surfaces corresponding to four regions of interest were extracted. As a preprocessing step, and to study all the data in a common space, we first rigidly align the images and resample them to have a consistent image size of $64 \times 64 \times 16$ through all the dataset.

3.4.1 Training

We divide the our dataset into a training set (80% = 150 images) and a testing set (20% = 37 images). On the training set we compute the reference SVFs for all the pair-wise registrations based on the framework we described. The computation of the surface matching with the framework of currents took 6 hours on a single core CPU for each of the 150 training images (a cluster of CPU was used). Note that it would not have been technically possible to perform the matching with currents for all the pairs (150^2 pairs) of images due to the computational cost. This further supports our choice not to rely directly on the LDDMM parametrization of the currents registration but to use SVFs instead. Once the surface are mapped, the process to compute the SVF parametrization took 4 minutes for each of the pair of images. Because our method already gives us a large database of ground truth data, we only use small translations in the X and Y axis for data augmentation. To train the network, we use the ADAM solver [Kingma 2014], an algorithm for first-order gradient-based optimization of stochastic objective functions, based on adaptive estimates of lower-order moments. We implement the network using Tensorflow¹ and we train it on a NVIDIA TitanX GPU with 100,000 iterations which took approximately 20 hours.

3.4.2 Evaluation

We compare the results with state-of-the-art registration algorithm *LCC Log-Demons* [Lorenzi 2013a] for which we use a set of parameters optimized in a trial and error approach on a subset of the training set. To evaluate the accuracy of the registration, we compute the registration of all pair of images in the test set for a total of $37^2 = 1369$ registrations. On average, one registration with the *LCC Log-Demons* algorithm took approximately 4 minutes with a Intel Core i7-4600U CPU 2.10GHz whereas our SVF-Net algorithm took 6 seconds on the same CPU and less than 30 ms on a NVidia TitanX GPU, increasing the speed by $\times 40 / \times 8000$ with the CPU/GPU implementation.

We apply the registration to the 4 regions of interests defined using our segmented meshes. These regions are: Left Ventricle Blood Pool (LVBP), Left Ventricle Myocardium (LVM), Right Ventricle Blood Pool (RVBP) and Right Ventricle Myocardium (RVM). Then we look at three standard measures of similarity between

¹www.tensorflow.org

the deformed region of interests i of the moving images M_i (either seen as a surface or as a volume) and the corresponding region of interest of the fixed image F_i . The measures we compute three different similarity metrics: Dice coefficient, the Hausdorff distance, the Local Correlation Coefficient (LCC) metric (which is the metric optimized with the Log-Demons algorithm) and one metric measuring the smoothness of the deformations.

We computed the metric for the two different methods. Our proposed approach performed better with respect to the Dice coefficient and the Hausdorff distance in approx 75% of the cases. The detailed boxplots results can be seen in Figure 3.4. We see that our method performs consistently better than the other tested methods for the dice index and the hausdorff distance and the difference was found to be statistically significant (paired t-test). In particular, the difference in accuracy is larger for the RV than for the LV. It is not surprising since the texture for the RV is usually less consistent between different patients, therefore a traditional registration method can have difficulty to match voxel intensities without a shape prior.

As for the LCC metric, which measures voxel intensity matching, although on average it is better for the LCC Log-Demons algorithm (which optimizes this metric), there is many outliers for which our method performs better as well for these cases, probably because the optimization algorithm gets stuck in a local minimum of the function. Finally, we compare the smoothness of the deformations. The difference of shapes seen in the images can be important, therefore Jacobians tend to be large even for regular deformations. We are more interested in evaluating how variable the Jacobian values are inside each region and we show the variance of the Log-Jacobians normalized by its mean value in Figure 3.4 (bottom right). Statistically significant differences (p-value of Welch’s test) can only be seen for the RVM (our proposed method has lower variance) and the LVBP (our method has higher variance). Overall, both methods output deformations with similar regularity.

Finally, it is also important to note two other important qualities of our method with respect to state-of-the-art optimization beyond accuracy improvement. Firstly, this method does not require the setup of any parameter (such as fidelity/smoothness trade-off, size of kernels, number of iterations...) at test time, making it possible to use for a non-experienced user. The parameters are set at training time by an experienced user. Secondly, it is way more robust and never gives completely incoherent results (in our experiments) whereas optimization approaches can get stuck in a local minima (especially at the RV location). Figure 3.5 shows different cases and the matching given by both methods. First row sees a typical case where optimization fails because of the large difference in texture and shape seen in both images, it gets stuck in a local minima and does not manage to retrieve the fixed image shape. Second row, we see that the optimization mistakes the blood at the left of the right ventricle blood pool as being part of the right ventricle, making the matching of both right-ventricles completely wrong, whereas our method does not have this problem. Finally, in the last row, we show one of the case where the optimization performs the best with respect to our method. One can see that, although our matching is less precise, it is still acceptable.




Figure 3.5: Three examples of the results of our registration method versus optimization approach. (Left column): the moving image with the myocardium segmentation in cyan. (2nd column): the fixed image with the segmentation in orange. (3rd column): our proposed registration with the deformed myocardium segmentation of the moving image in cyan and the target segmentation in orange. (Right column): the optimization comparison (*Log-Demons LCC*).

3.5 Conclusions

In this chapter, we propose a novel methodology to build ground truth deformation from pair of segmented images and to train a predictive algorithm. The method relies on a convolutional network whose architecture consists of contracting layers to detect relevant features and a symmetric expanding path that matches them together and outputs the transformation parametrization. Whereas convolutional networks have seen a widespread expansion and have been already applied to many medical imaging problems such as segmentation and classification, its application to build registration algorithm has so far faced the challenge of defining ground truth data on which to train the algorithm. Indeed, it is an easy task for a human to provide ground truth classification of an image, and tools have been developed to manually define contours of a region of interest to train segmentation, but it is nearly-impossible to manually give an accurate pair-wise voxel matching between two images to train a registration method.

Here, we present a novel training strategy which uses reference deformations

defined with segmented regions of interest. Using the framework of currents, the surface mesh of these regions of interest are matched to a common template giving us the point correspondence between pairs of images. Based on this correspondence, a parametrization of the deformation by stationary velocity fields is then interpolated on the whole image grid and used to train the network. Compared to other methods, our method has the benefit not to rely on either synthetic images (which often builds pair of images that fail to account for the inter-patients texture variability and the artifacts seen in a MRI acquisition) or on the previous computation from an optimization method (training an algorithm on such ground truth will likely carry all the drawbacks seen in the optimization). We directly use segmentation of ROIs together with the framework of currents to build point-correspondence between shapes, which are then interpolated to get a SVF on the whole image grid.

We show that our method significantly outperforms a state-of-the-art optimization method in term of registration accuracy as estimated with similarity metrics on 4 ROIs of the heart. We also highlighted multiple advantages of our method: (i) it is way faster with a speed-up of $\times 40/\times 8000$ with the CPU/GPU implementation, (ii) it does not require the choice of any parameter at test time making it easy to use for a non-experienced user, (iii) it does not get stuck in local minima in the case of differences in textures especially at the RV location and (iv) it is more robust and never gives completely inconsistent results. These qualities are probably more important than just raw accuracy improvement as they represent the main problems currently holding back registration methods to be used for large scale database.

One possible perspective that is opened by our method is the multi-atlas mesh segmentation with a very large number of templates. Whereas traditional methods have to a-priori choose a set of probable templates due to the computational cost, the rapidity and robustness of our method gives us the possibility to use many more atlases in order to compute a whole probabilistic distribution of segmented shapes. In the next chapter we build a complete pipeline to segment the myocardium in an automatic and robust way using a multi-atlas segmentation approach with the fast registration algorithm defined here.

An Automatic Multi-Atlas Myocardium Segmentation

Contents

4.1	Chapter Overview	56
4.2	Motivations	57
4.3	Segmentation in Clinical Practice	57
4.4	Segmentation Methods In Medical Imaging	59
4.5	Multi-Atlas Segmentation: Overview of the Method	60
4.6	Rigid Alignment by Landmarks Detection	61
4.6.1	Method for Rigid Alignment	62
4.6.2	CNNs and Heatmap for Landmarks Detection	63
4.6.3	Training and Results	63
4.6.4	Alignment of the Images and Cropping	65
4.7	Registration of Images and Fusion of Segmentation	66
4.7.1	Points Propagation	66
4.7.2	Points Combination	66
4.7.3	Supervised Learning of Local Distance	67
4.7.4	Computation of Weights Based on Local Distance	68
4.8	Conclusions	70

This chapter concludes part I of this thesis with a segmentation pipeline derived from the previously defined registration algorithm. In Chapter 2, we introduced registration of shapes and images with the example of two algorithms. Then in Chapter 3, we use the information of the shape registration in order to learn image registration. This allows the improvement of the registration quality and robustness while being very fast and efficient. Finally, in this chapter, we leverage the speed of this registration algorithm to build a robust and accurate mesh segmentation of the myocardium using a multi-atlas approach. It can be seen as the opposite task of what has been done in the previous chapter: we were using segmentation to learn registration and now we are using registration to compute segmentation.




Figure 4.1: Example of two segmentations of the myocardium. (Left): segmentation of the myocardium with a tetrahedral mesh. On top of the delineation of the contours with the shape, each point and cell of the mesh contains multiple labels about its localization. Here we show the classification RV/LV in red/blue for each point. (Right): mask segmentation. The mask is an image where each voxel is given a value usually between 0 and 1. Here we show a mask with values corresponding to 4 classes: RV/LV blood pool (green/orange), RV/LV myocardium (purple/blue).

4.1 Chapter Overview

In this chapter, we start by giving an overview of segmentation in medical imaging. We first present the use in clinical practice and the different possible representation of a segmented shape in section 4.3. In section 4.4, we go on by presenting the main classes of methods and do a quick overview of the state-of-the-art of the segmentation methods. Then, we present our multi-atlas approach for the automatic segmentation of the myocardium. It is made of two consecutive steps. First, in section 4.6, we detect the location and the orientation of the heart in the image. We perform this task thanks to a CNN trained to detect two landmarks. One landmark gives the position of the heart and the other one the orientation. Using these landmarks, the target image is rigidly aligned with a database of images with ground-truth segmentation. This is a mandatory step before applying diffeomorphic registration. In section 4.5, we define a multi-segmentation method to segment the aligned images using the diffeomorphic registration method presented in Chapter 3. We show that we can further improve the accuracy by fusing the different registered atlas with specific weights which are functions of 3 metrics evaluating the registration accuracy. The contribution of each metric on the weight of a point is defined using a supervised learning approach on a training set.

The main contributions presented in this chapter are:

- A fast and robust end-to-end framework for myocardium segmentation leveraging the speed of the SVF-net registration with a large number of atlases.

- A CNN to detect landmarks and use them to rigidly align an image with respect to a database.
- A supervised learned method to define weight in order to fuse the segmentation results of multiple registered atlas.

4.2 Motivations

Most of the methods of cardiac function analysis rely on a correct segmentation of the myocardium and we need such a segmentation on multiple occasions in this manuscript. Bio-mechanical models of the heart motion which are used in Chapter 6 are based on a segmented 3D tetrahedral mesh of the LV and RV ventricles. Many cardiac motion tracking algorithms use either the mask of the segmentation or even a more complex labeling of the regions of the heart. Polyaffine cardiac motion tracking, presented in Chapter 6, requires a segmentation of the heart with the AHA regions labels. Finally, the statistical shape analysis of the myocardium to classify and detect infarct in Chapter 7 requires the surface shapes of the LV at end-diastole and end-systole. More generally, segmentation of the myocardium and blood pool are also often used in clinical practice to derive clinical indexes.

4.3 Segmentation in Clinical Practice

Both ventricles play a fundamental role in the body function with the circulation of oxygenated blood to the body. The most commonly used clinical indices to evaluate a cardiac disease are based on geometrical measurements of regions of the hearts: the blood pool volume, the wall thickness of the myocardium, or the myocardial mass. In particular, a major clinical measure to assess the efficiency of the heart as a pumping machine is the stroke volume (the volume of blood pumped by the left ventricle at each beat). This volume is very often expressed as the ejection fraction which is a ratio of the stroke volume to the end-diastolic volume. The ejection fraction is an important clinical factor to estimate how important a reduction of the efficiency of the heart function is (see Fig. 4.2). To compute this value one often relies on the manual segmentation of the LV blood pool at end-diastole and end-systole. Then the values computed are compared to reference values found in the literature [Kilner 2010, Fratz 2013, Kramer 2013].

Manual segmentation of the contours of the myocardium and the blood pool is still the most used method in clinical practice with respect to semi-automatic or automatic methods. However, it has several drawbacks. This task is very time-consuming, especially if one wants to segment each frame of a cardiac sequence, whereas the clinical work-flow is usually very time-constrained and clinicians do not have the necessary time to perform these segmentations. It also requires a minimum training for the clinicians to perform accurately the delineations of the myocardium. Even with trained experts, the inter-rater variability of manual segmentations can be important [Suinesiaputra 2015] which will impact the measures




Figure 4.2: The importance of the assessment of the ejection fraction in cardiac medical imaging. Normal ejection fraction ranges between 46% and 70%. An ejection fraction below 35% will be considered as a quite severe reduction of the cardiac function. *Images from <http://www.hrsonline.org/>.*

derived from the segmentations. Finally, there is a lack of common guidelines for the segmentation, and different experts can choose to include or exclude papillary muscles and trabeculations [Han 2008, Chuang 2011], or to choose a different basal slice [Van der Geest 1997, Paknezhad 2016] as the top slice to segment. This choice can have an important impact on the estimation of the ventricular volume and undermine the comparison with computed using different methods.

Therefore it is of great importance to have an automatic, accurate and robust method for the segmentation of the myocardium. The accuracy of the following analysis will be strongly dependent on the quality of the initial segmentation. In this chapter, we use the previously defined registration method and the segmented templates to get the segmentation of a new, yet-unseen cardiac image. Segmentation of a structure can be represented in different ways. The most common representation is a binary image of the same size than the original image and with voxel values of 1 for foreground and 0 for background. Segmentations can also be represented as shapes (surfaces or volumes) delineating the contours of the regions of interest. This representation is usually more interesting, as one can add different labels at each point of the mesh giving additional information. An example of both representation can be seen in Fig. 4.1. Although methods have been developed to compute tetrahedral meshes representing a shape from a binary mask [Yerry 1984, Courchesne 2007, Fang 2009], it is more convenient and more efficient to directly compute a tetrahedral mesh. Furthermore, segmentation methods based on the deformation of an atlas with registration, such as the one we will use, can easily work with either binary mask or tetrahedral mesh. Finally, a binary mask can easily be extracted from a segmentation represented with a tetrahedral mesh. These considerations support the choice of working with tetrahedral meshes rather than binary masks.

4.4 Segmentation Methods In Medical Imaging

Segmentation is a major topic in medical imaging and it would be too long to give an exhaustive review of all the techniques even restricting the topic to cardiac segmentation. We give here a short review of the main class of algorithms and invite the reader to refer to the survey of [Petitjean 2011, Suinesiaputra 2014] for a thorough review of the techniques in the context of short-axis cardiac images and [Dolz 2014] for brain structures.

Image-based methods. Most image based methods sequentially process each of the contours of the endocardium and epicardium using priori information about their shapes. For endocardial contour, it can be done efficiently by thresholding the blood pool voxels intensities which is quite higher than the myocardium [Katouzian 2006]. Other methods use dynamic programming to search for the optimal path in a cost matrix that assigns low cost to object frontiers [Üzümcü 2006]. Using the endocardial contour, the epicardium is then found in a second step, often using a spatial model with a-priori on the myocardial thickness to improve robustness in the case the contrast of the wall is not sufficiently defined [Cousty 2010, Jolly 2009].

Voxel-wise semantic methods. Instead of finding contours and delineation of the surfaces, the problem is posed as the classification of each voxel (myocardium versus background, or myocardium blood pool and background). Machine learning methods with classification forest [Shotton 2008, Lempitsky 2009, Margeta 2011] have been developed with low-level local features and high-level context-rich features. Smoothness constraints are often imposed based on semantic information as a post-processing step [Mahapatra 2014].

Deformable models methods. Deformable models [Delingette 1992, Montagnat 2005] are surfaces defined within the image domain that can deform under the influence of internal forces, defined within the surface itself, and external forces, computed from the image data. The internal forces ensure the smoothness of the model whereas the external forces drives the matching of the model to the data. With respect to voxel-wise segmentation, these methods have the benefit to directly represent the segmentation with a surface, thereby making it possible to achieve subvoxel accuracy, a highly desirable property for medical imaging segmentation. Reviews of deformable models methods can be found in [McInerney 1996, Xu 2000].

Atlas-based methods. These methods use either one template, or multiple references for which segmentation have already been computed (usually manually). This image or sets of images are usually called an atlas, and a new-image can be segmented by mapping its coordinate to that of the atlas. First applied to brain segmentation [Collins 1995] it has also been used extensively for cardiac segmentation [Lorenzo-Valdés 2002, Lorenzo-Valdés 2004, Isgum 2009]. In a first step, the

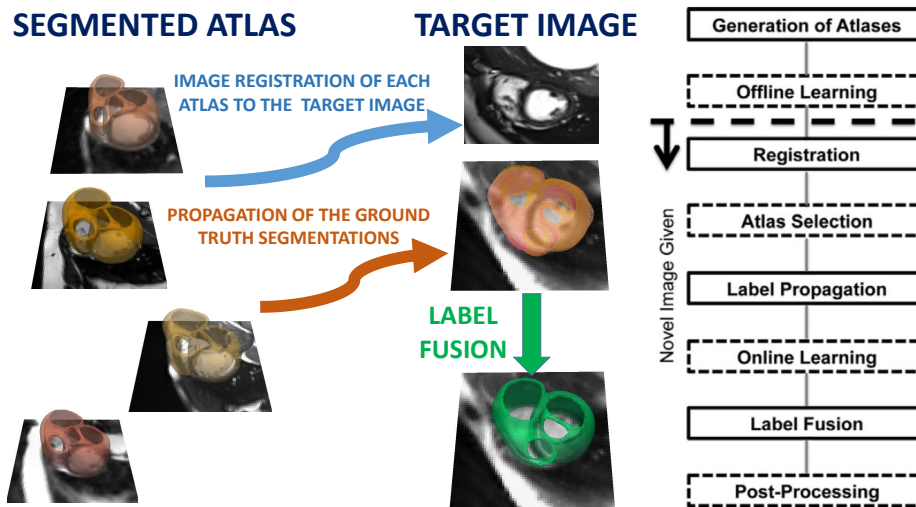


Figure 4.3: Schematic representation of the Multi-atlas segmentation pipeline. A registration method is used to register each atlas to the target image. Then the deformations are applied to the ground truth segmentations to transport them in the coordinates of the target image space. Finally, a method to merge all the labels is used. Right schema from [Iglesias 2015] shows additional optional steps.

template is registered to the target image and then the segmentation of the template is propagated using this deformation. If multiple templates are used, a strategy to combine the registration has to be defined and we will present these methods in more details in a next section of this chapter.

Deep learning methods. More recently, deep learning methods based on convolutional neural networks [Krizhevsky 2012] have been adapted to the problem of image segmentation [Carneiro 2012, Avendi 2016]. The classification is done in a voxel-wise manner where a fixed-size patch around each voxel is used to output the label for the voxel. Patches extracted from manually segmented images are used to train the network which is then applied to each voxel in a sliding-window approach. Other architecture, such as the U-architecture proposed in [Ronneberger 2015] and extended to 3D images [Çiçek 2016, Fritscher 2016], process the image in a single stream by using upsampling layers. This kind of architecture is more difficult to train, but the computational time for inference is largely decreased and it also has the advantage to process higher level features that are not limited by the size of the patch.

4.5 Multi-Atlas Segmentation: Overview of the Method

We give in this section a complete end-to-end overview of our method for myocardium segmentation. We chose to perform segmentation using the framework of multi-atlas segmentation (MAS) (see Fig. 4.3), an extension of atlas-based

methods with multiple templates. First introduced by [Rohlfing 2004, Klein 2005, Heckemann 2006], the use of multiple atlas extends the range of anatomy that can be predicted with such model. If the database of atlases is large enough, one can expect that at least one or multiple atlases will have a shape close to the one of the target geometry. Also errors due to mis-labeling or failed registration are averaged in the final segmentation making the process more robust. The benefit of having a more important prior knowledge of the shapes (and therefore to better capture anatomical variation) comes at a high computational cost since each of the atlas image has to be registered to the target image.

MAS algorithm can be decomposed into several components (see Fig. 4.3 - right). Due to the computational burden of performing the registration with respect to multiple images, methods that use several atlases for segmentation tend to select a-priori the atlases to be used. Only the atlases that are supposed to be the closest to the target image are used in order to reduce the number of registration done. These atlases are identified and registered to the target image [Rohlfing 2003a]. While this selection can efficiently decrease the computational complexity, a lot of information is lost by not using the whole database of atlases. Furthermore, it is often difficult to decide if an atlas is close to the target image without performing the registration (which can give a lot of information such as similarity between the registered images and distance of the deformation).

The registration algorithm presented in the previous chapter gives us a very fast and robust method with which we can easily perform many registrations in a short period of time. Therefore, we chose to leverage this speed by not selecting any atlas a-priori and we use the whole database of ground truths to predict the segmentation. The method can be decomposed in three consecutive steps. First, in section 4.6, we extract a ROI of the target image that is aligned with our database. This is necessary because the diffeomorphic registration algorithm we use, as most of diffeomorphic registration methods, requires that the pairs of images are rigidly aligned. Then, the extracted ROI of the target image is registered with each of the atlas. Finally, we develop a weighting scheme to combine the estimations coming from each of the ground truth segmentation in section 4.7.

4.6 Rigid Alignment by Landmarks Detection

Rigid alignment of images is required before running the diffeomorphic registration. For the images of the atlases, we can easily define a common pose and alignment using the information from the segmented ground truth. For example, we can compute the barycenter of the LV and the RV defined by the ground truth. Then, using the center of the LV, we define a common ROI for all ground truth images by cropping with a fixed size window around this landmark. The center of the RV is used to compute the direction of the LV to the RV defining orientation of the heart in the X/Y plan. A $2D$ rotation on this plan is applied to all images of the atlases (see Figure 4.5) so that the RV and LV are aligned. We describe in more details the




Figure 4.4: (Left): Image from a slice of a 3D acquisition. (Right): Same image after pre-processing (cropping the ROI and 2D rotation around the Z axis to align with the pose of the atlases). Most of the background has been removed from the image and only the important information remains. The pose of the LV/RV ventricles and the heart position is aligned with the atlas making the registration step easier. To do so, a landmark corresponding to the LV and one corresponding to the RV are detected.

choice of the transformation using in section 4.6.4. Using this method, we align and crop our database of ground truth images and segmentations in a common space.

With a new image, for which we don't have a ground truth segmentation to easily compute these landmarks, we need to define a method to align this image to the reference space of the database of ground truth images. We first describe some possible methods for performing this alignment and we give reasons that support the choice of our method which rely on the detection of two landmarks. Then we present the CNN and the architecture used to detect these landmarks and compute the corresponding 3D transformation that will be applied to the image.

4.6.1 Method for Rigid Alignment

To perform this task, multiple approaches can be taken. A possible approach would be to perform rigid-registration of the target image to one (or multiple) images of the atlas. To do so, a large range of algorithms have been developed [Ourselin 2000, Ino 2006, Ashburner 2007b]. While this method is possible, it might be computationally costly to rigidly register the target image to each of the atlas. One possible way to reduce this burden would be to perform the rigid-registration with respect to only one atlas supposed to be pre-aligned with respect to all the other references. In this case, the alignment of the target image rely completely on one registration which could fail in the case both images present large differences in

texture.

We chose to take another approach and align the position and the pose of the heart by the automatic detection of two landmarks of the images: the center of the RV and the center of the LV (see Fig. 4.4). From the position of these two landmarks, one can define a rigid transformation similarly to what has been done with the ground truth images. Our method has the advantage not to rely on any specific atlas and therefore does not introduce a bias with respect to a reference. Moreover, these landmarks can be easily defined on the atlas images using the ground truth segmentation meshes. Therefore the set of atlas gives us a set of images and landmarks positions on which to train algorithms to detect these landmarks with machine learning methods.

4.6.2 CNNs and Heatmap for Landmarks Detection

To detect landmarks of the heart, several methods have been developed. [Pavani 2010, Shi 2011] propose to use Haar-like features which are assigned optimal weights so as to maximize their ability to discriminate objects. In [Jolly 2006], the authors uses techniques such as maximum discrimination, thresholding and connected component analysis in a global localization step to detect the left ventricle. Supervised machine learning techniques can also be used such as in [Mahapatra 2012b], where random forests classifiers are trained on low level features (intensity, texture, shape asymmetry and context information) derived from the neighborhood of annotated landmarks on training images. More recently, methods based on reinforcement learning have been developed [Ghesu 2016] with effective hierarchical feature extraction achieved through deep learning.

Inspired by recent works [Pfister 2015, Bulat 2016] using heatmaps regression for landmark detection, we chose to take a similar approach to detect both landmarks (LV and RV centers) using CNNs. In particular, the work of [Payer 2016] investigates the idea of directly estimating multiple landmark locations from 3D image using a single fully-convolutional CNN, trained in an end-to-end manner to regress heatmaps for landmarks instead of absolute landmark coordinates. This approach has multiple advantages over the ones previously proposed. It is a learning-based method so that we can efficiently leverage our large database of more than 180 images with ground truth landmarks derived from segmentations. Also, the prediction of heatmaps is an easier task for a CNNs than the prediction of absolute coordinates of landmarks, as the localization of the responses in the successive layers can be directly used to predict the heatmap. Especially, a U-Net type architecture [Ronneberger 2015] already used for registration can be easily adapted to the task of heatmap prediction.

4.6.3 Training and Results

To train the algorithm, we see the same database of 187 segmented images which was presented in previous chapter. We divide them into a training set of 150 images and




Figure 4.5: Heatmaps for both landmarks (RV center: left image, LV center: right image).

a testing sets of 37 images. For all the images of the training sets, both landmarks are computed using the ground truth segmentation as the barycenter of all the points labeled as RV/LV. Then for each landmark p , heatmaps are defined on the image grid as:

$$H_p(\mathbf{x}) = \exp\left(-\frac{\|\mathbf{x} - p\|^2}{\sigma^2}\right),$$

where σ is the decaying factor of the heatmap (set to 20 in our case: multiple values were tested with little impact on the accuracy of the method). Examples of such heatmap for both the RV center and LV center landmarks are shown in Fig. 4.5. An architecture similar to the one presented in Fig. 3.3 is used to learn and predict the maps. The input of the network is the complete image. It has two 2D-layers and two 3D layers. The features from the contracting path are directly concatenated to the deconvoluted features from the expanding paths. In the last layer, 2D convolution builds the 2 layers corresponding to both heatmaps at the same grid as the initial input images. For data augmentations, random translation on the X and Y axis are used, as well as translation of the Z axis corresponding to multiple of the gap between slices. The network is trained using Tensorflow on a NVIDIA TitanX GPU with 10,000 iterations which took approximately 10 hours.

At test time, the landmark position p is inferred by computing the point that minimizes the least-square distance to the predicted map H_{pred} :

$$p = \arg \min_{\mathbf{x}} \sum \|H_{pred}(\mathbf{x}) - H_p(\mathbf{x})\|^2.$$

The result of the absolute error for the 37 test images can be seen in Fig. 4.6. The median error along the X and Y axis is below 3 mm for each landmarks. It is within the range of the voxel spacing (between 2 and 3 mm) for these images. There is few outliers and, even for these cases, the error stays below 10mm which




Figure 4.6: Absolute error on the test set for the detection of both landmarks along each axis. The error of the detection of the landmarks is comparable to the voxel spacing (on average $2.5mm \times 2.5mm \times 8mm$ for these images).

is still largely acceptable. For the Z axis detection, the error is a little bit higher, due to the higher voxel spacing on this direction compared to the others. The error has a median below 5 mm for both landmarks which is to be compared to the slice gap (between 8 and 10mm). The results were not compared with other state-of-the-art methods, as the goal was not necessarily to get the most accurate landmark detection, but rather to get an efficient and robust method to pre-align the images before the registration process.

4.6.4 Alignment of the Images and Cropping

We do not perform complete rigid 3D alignment with 6 degrees of freedom. Indeed, we restrict ourselves to only rotation around the Z -axis and we do not consider rotations along the X and Y axis. In MRI images, the distance between slices is usually relatively large (6 – 10mm). Therefore, we want to avoid any pre-processing that would require interpolations of values in between slices (such as would rotations around X or Y axis). Similarly, to avoid any interpolation, we restrict the Z axis translation to multiple of the slice distance. Finally, we are looking for 4 degrees of freedom (rotation θ around the Z axis, two unrestricted translation on the X, Y plane and one *slice alignment* on the Z axis which can be considered as just finding a common reference slice).

As a last step before the segmentation pipeline, we crop the target image. Extraction of a cropped region around the heart is a necessary pre-processing task to decrease the complexity of the registration process by providing an image whose background has been removed. It is even more important since the registration method we use rely on 3D convolutional networks which require a fixed size image as input and whose complexity does not scale well with the size of the image. The images were cropped using a fixed size region based on the detected landmarks so that the myocardiums of all images are completely included in the ROI. The result

of the pre-processing step (landmark detection, rigid alignment and cropping) can be seen in Fig. 4.4.

4.7 Registration of Images and Fusion of Segmentation

4.7.1 Points Propagation

We consider a database of M images I_j , $j = 1, \dots, M$, or atlases for which we have ground truth segmentations. These segmentations are defined by meshes S_j with points p_k^j . Note that with respect to traditional multi-atlas segmentation methods, our method does not work with segmentation represented with a binary mask of the image but with a tetrahedral mesh. We suppose that these meshes are defined with the same number of points which are in correspondence. A possible way to get such a representation from general meshes can be done by performing the shape registration with a common template as we have done in the previous chapter. We perform the registration of each of these images with respect to the target image I . The resulting deformation field is applied to the segmentation of the atlas S_j . We note \hat{S}_j the estimation of the segmentation of the image I using the deformation of the segmentation of the atlas j and \hat{p}_k^j the estimation of the position of the point k of the segmentation of the image I using the atlas j . Once each image of the atlas is registered to the target image, one has to define a method to combine the estimations defined by each deformed segmented mesh.

4.7.2 Points Combination

The estimations \hat{S}_j with points \hat{p}_k^j of the segmentation of the image I using atlas $j = 1, \dots, M$ can be fused together to the estimation \bar{S} with points \bar{p}_k :

$$\bar{p}_k = \sum_{j=1}^M \omega_k^j \hat{p}_k^j,$$

where ω_k^j is a weighting function representing our confidence of the estimation of \hat{p}_k^j of atlas j and is constrained by $\sum_j \omega_k^j = 1$. The choice of this function ω is often referred as the label fusion (the step of combining propagated atlas labels) and is one of the core components of MAS [Artaechevarria 2009, Aljabar 2009]. A natural and widely used choice for this function, without any a-priori knowledge on the accuracy of the registration computed, is the average over all atlases: $\omega_k^j = 1/M$. All the atlases are given the same weight and this can be considered as the equivalent of majority voting [Heckemann 2006, Klein 2005, Rohlfing 2004] in the binary mask case. Another common choice is the best atlas selection, where only one atlas S_q , which is supposed to be the closest one with respect to some metric, is used to compute the propagated labels.

Choosing only one atlas, even if this is supposed to be the best one, might not be the optimal choice if one has a very large database. To incorporate a-priori




Figure 4.7: Visualization of the δ values for 4 ground truth meshes at one slice. Segmentations that are a little bit less accurate (such as the 2 bottom ones) tend to have larger values for the δ coefficients.

knowledge about the accuracy of the segmentation while still using all the atlases, [Artaechevarria 2008] introduce global weight to each atlas proportional to the normalized mutual information between the registered atlas image and novel image intensities. While global weights are able to account for global dissimilarities between registered images, they cannot model the spatially-varying nature of registration accuracy. Therefore, more advanced methods use local and semi-local weighting schemes [Isgum 2009, Iglesias 2009] based on similarity metrics on the intensity of the registered images. Finally, local and global metrics can be combined such as in [Wolz 2013] which use three different terms, reflecting global, organ-level and local similarities. In this work, we chose to define a local metric and adapt this framework to the case where segmentations are defined with points instead of binary masks. This local metric is learned with a supervised method which we describe in the next section.

4.7.3 Supervised Learning of Local Distance

Definition of the Local Distance. In this work, we define the weights based on an estimation of the error at each point. In this section we describe how we learn this local distance. This local distance will be used in the next section to define the weight of each point of each atlas in the final segmentation. The error at each point is estimated locally as a linear function of 3 different metrics related to registration accuracy.

- The first one $A_k^j = LCC(\hat{p}_k^j)$ is the *Local Correlation Coefficient* (LCC) [Lewis 1995, Cachier 2000, Lorenzi 2013a]. This metric locally estimates the similarity between the voxel intensities of the images. This metric is preferred to the standard SSD because we are dealing with MRI acquisitions from different patients and scanners whose intensities can largely differ.
- The second metric $B_k^j = \|p_k^j - \hat{p}_k^j\|^2$ is the square norm of the displacement of the transformation at point p_k^j . The rationale is that we have more confidence in small transformations (corresponding to similar images) than in large transformations (corresponding to pair of hearts whose shapes differ largely).
- Finally, the last metric $C_k^j = \|\hat{p}_k^j - \text{mean}_{i \neq j} \hat{p}_k^i\|^2$ corresponds to the agreement of the estimation of an atlas with respect to the rest of the database.
- On top of these parameters corresponding to the confidence we have on the registration, we define for a point of each ground truth atlas a parameter δ_k^j which measures the intrinsic confidence on the position of the point k of atlas j .

Therefore local distance $d(\hat{p}_k^j)$ of the point k of the atlas j is defined as:

$$d(\hat{p}_k^j) = \delta_k^j + \alpha A_k^j + \beta B_k^j + \gamma C_k^j.$$

Supervised Learning of the Coefficients. The coefficients $(\alpha, \beta, \gamma, \delta_k^j)$ of this distance are learned on a training set as to minimize the square difference of the distance versus the ground truth error. We consider a training set of N images I_i for which we apply the registration with respect to each atlas and compare the estimated point $\hat{p}_k^{i,j}$ with the ground truth atlas $q_k^{i,j}$

$$(\alpha, \beta, \gamma, \delta_k^j) = \underset{i=1}{\text{argmin}} \sum \left(\|q_k^{i,j} - \hat{p}_k^{i,j}\|^2 - d(\hat{p}_k^{j,i}) \right)^2.$$

The parameters are estimated using our database with $M = 150$ data taken as atlases j and 37 data used for the training i . Optimal values obtained for the parameters are $\alpha = -187.34, \beta = 0.0790, \gamma = 1.39$. This is consistent with the intuition of the impact of each feature: a high value for the LCC coefficient is correlated with a good registration accuracy whereas the two other coefficients are negatively correlated. Figure 4.7 shows the values for δ_k^j for 4 different meshes. Points that are located in the right ventricle have higher values of δ_k as the registration tends to be less accurate because myocardium of the right ventricle is less clearly defined.

4.7.4 Computation of Weights Based on Local Distance

We define the local weight of each point of each atlas as a function of the local distance with a kernel σ_{metric} . Furthermore, we also smooth the weights spatially with a kernel $\sigma_{spatial}$. Finally, the weights are normalized in order to sum to 1 :




Figure 4.8: (Left): Error on the testing set with different values of σ_{metric} with $\sigma_{spatial} = 10$. With a very low value of σ_{metric} , only one template is taken into account at each point (the best template). With very large values, all the templates are equally weighted. (Right): Error on the testing set with different values of $\sigma_{spatial}$ with $\sigma_{metric} = 4$. Very low values correspond to completely local and non-smooth, very large values correspond to global metric which means smoothing on the whole mesh. Result for multi-atlas fused with average (all the templates are equally weighted) and best (only the best template is taken into account) are shown.

$$\tilde{\omega}_k^j = G_{\sigma_{spatial}} \star \exp\left(-d(\hat{p}_k^j)/\sigma_{metric}^2\right),$$

$$\omega_k^j = \frac{\tilde{\omega}_k^j}{\sum_j \tilde{\omega}_k^j}.$$

The kernel σ_{metric} corresponds to the confidence on our estimation of the local distance $d(\hat{p}_k^j)$. Large values for the kernel corresponds to small confidence. At the limit when σ_{metric} becomes large enough, all the weights become equal and we get the simple method of averaging the labels and the local distance does not have any impact. The spatial kernel $\sigma_{spatial}$ is to ensure spatial consistency of the resulting segmentation. Large values of $\sigma_{spatial}$ will make the weights more global whereas small values will make them more local. At the limit when $\sigma_{spatial}$ becomes large enough, each ground truth template will have all the points with exactly the same value for the weight.

Figure 4.9 shows the results on the error on the test set depending on different value for σ_{metric} and $\sigma_{spatial}$. Results show that this method to fuse the labels improve the accuracy of the segmentation with respect to the simple average of the points and to the best template. Furthermore, we achieve the best accuracy with a spatial kernel around 10, with the accuracy decreasing afterwards. It shows that evaluating locally the error still increase the accuracy with respect to a completely global metric.

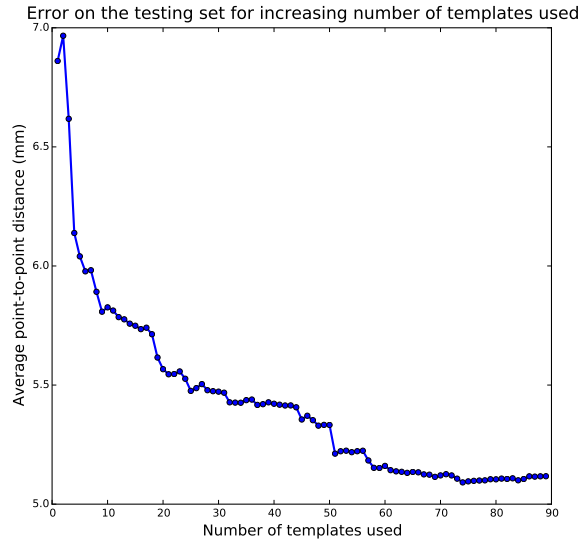


Figure 4.9: Error on the testing set with respect to the number of templates used.

Finally, Fig. 4.9 shows the accuracy of the segmentation on the test sets with respect of the number of templates on the training set used. Results show major improvement of the accuracy of the segmentation as the number of templates increase. While standard registration algorithms would not allow the use of such large database because of the computational cost of the registration, the speed of our registration method makes it possible in our case. While classical registration algorithms need around 1 – 10 minutes to perform one registration, our method takes less than 1 second to perform the same registration. Therefore, we can leverage large database of ground truth in this multi-atlas segmentation method, increasing efficiently the accuracy of the segmentation.

4.8 Conclusions

Most of the work done in this thesis relies on a correct segmentation of the myocardium whether it is in the form of a binary mask (for cardiac motion tracking), in the form of a tetrahedral mesh for cardiac bio-mechanical simulations (Chapter 6) or in the form of a labeling of the AHA regions (polyaffine framework in Chapter 7 and Chapter 9). Therefore, developing a segmentation pipeline was crucial to perform the analysis done in the rest of this manuscript. With respect to most state-of-the-art method, our method computes directly a tetrahedral mesh representation for the segmentation. It is an important quality of our method because of the flexibility of this representation which can be used directly for shape statistics and bio-mechanical simulation.

To combine the different segmentation into the final result, we define local weights that are a-priori learned on a training sample. These weights are based on an estimation of the confidence of the evaluation of a specific point by each atlas.

We show that fusing the points with weights defined with this local estimation of the error efficiently increase the accuracy of the resulting segmentation. Furthermore, results show that adding templates in the estimation of the segmentation in the multi-atlas method increases the accuracy.

The method presented here has several important qualities. It is completely automatic and does not even require the location of the heart as user input, thanks to the landmarks detection network. With respect to traditional multi-atlas segmentation algorithm, the speed of the registration method defined in Chapter 3 allows us to use a large database of atlases. We show that increasing the database efficiently improves the accuracy of the final segmentation. This supports the necessity of having a very fast and robust registration algorithm that can perform many segmentation with a reasonable computational cost. In order to fully validate the accuracy of our method, comparison with other state-of-the-art methods needs to be done using metric such as dice scores and hausdorff distances.

Part II

**LOW-DIMENSIONAL
REPRESENTATION OF
CARDIAC MOTION**

Barycentric Subspace for Cardiac Motion Tracking

Contents

5.1 Chapter Overview	75
5.2 Background: Cardiac Motion Tracking	76
5.2.1 Cardiac Motion Tracking Algorithms	77
5.2.2 Evaluation of Tracking Methods	79
5.3 Motivations: Low-Dimensional Representation of the Cardiac Motion	79
5.4 Methodology: Barycentric Subspaces	82
5.4.1 Definition of the Subspace	82
5.4.2 Projection of an Image to the Subspace	83
5.4.3 Computation of the optimal References of a sequence of Images	84
5.4.4 Computation of an Image within the Subspace	84
5.5 Using Barycentric Subspaces as a prior on the Registration	85
5.5.1 Barycentric Log-Demons Algorithm	88
5.5.2 Evaluation using a Synthetic Sequence	88
5.5.3 Towards Symmetric Transitive Registration	89
5.6 Conclusion	90
5.A Projection of an image to the Barycentric Subspace	91
5.B Frame-To-Frame Barycentric Registration Formulation	92

Part I of this thesis focused on the segmentation of the myocardium from the end-diastole image using multi-atlas registration. In this part, from the segmentation defined at the first frame, we follow the motion of the myocardium through all the frames of a cardiac sequence to study the cardiac motion and define its low-dimensional representation. This representation will be used in Part III of this manuscript in which we show applications to diagnosis and longitudinal analysis.

5.1 Chapter Overview

In this chapter, we first give an overview of cardiac motion tracking. We present the challenges of adapting a registration method to this specific task. Then, we propose

an innovative methodology to challenge the traditional approach used in cardiac motion analysis. Instead of relying on the registration of each frame with respect to the end-diastolic frame chosen as the reference, we develop a new group-wise paradigm which relies on *Barycentric Subspaces* computed using multiple references. We first introduce the theoretical framework of barycentric subspaces and the methods to project an image and to compute its coefficients. Then we show how this subspace can be used efficiently as a prior on the cardiac motion in the registration process.

The work presented in this chapter is extracted from a journal article [Rohé X] submitted to Medical Image Analysis and accepted subject to minor revisions. This journal article is an extension of the work published at the MICCAI 2016 conference [Rohé 2016b].

The main contributions with respect to state-of-the-art are:

- The introduction of a new method for dimension reduction and low-dimensional subspace analysis: the *Barycentric Subspace* [Pennec 2015] in the context of medical images.
- The methods for computing the coordinates of an image within a *Barycentric Subspace*, for choosing the reference frames building the optimal subspace, and for reconstructing an image given the coordinates and the references.
- The use of this subspace to build an efficient prior for a cardiac motion registration algorithm, reducing the error for the estimation of the large deformations.

5.2 Background: Cardiac Motion Tracking

Many pathologies of the heart affect its motion during the cardiac cycle and therefore it is crucial for clinicians to have methods to understand and analyze the different patterns of motions seen in a population [Konstam 2011]. The motion dynamics observed in pathological cases can be compared to that of normal subjects in order to better understand the impact of a given disease. Efficient classification and quantification of the cardiac motion of a patient can help clinicians to have additional insights in order to help in diagnosis, therapy planning, and to determine the prognosis for a given patient [Bijnens 2007, Young 2006]. For example it can be used to extract relevant clinical indices such as the ejection fraction or strain values at different locations of the heart [Rohé 2015], to compare two different populations [McLeod 2015a] based on the pattern of their cardiac motions, and to perform longitudinal analysis during the development of a disease or following a therapy.

Cardiac Images Tracking the myocardial tissue with physical markers is invasive and can not be used in clinical practice. Therefore, non-invasive methods have been developed using medical imaging. To capture and analyze this 3D motion over a cardiac cycle, clinicians rely on acquisitions of a sequence (or frames) of images representing a cardiac cycle with Magnetic Resonance, Echocardiography or Computed Tomography. These acquisitions are becoming more and more widely

used in clinical practice as a means to assess the efficiency of the cardiac motion. A visual judgment from an expert can give a first idea about the specificities of the cardiac motion of a given patient. However, clinicians need a more precise quantitative evaluation because the differences seen in the images can be small and hard to detect but still have important physiological consequences. It could be possible for clinician to manually track a given set of points for all frames and derive clinical indices from this motion. But this would be very time-consuming, subject to bias between different rates, and difficult since the correspondences have to be found in 3D images. Therefore, it is almost never performed in a clinical workflow. A convenient method to track cardiac tissue is to use a modality derived from MRI called tagged resonance imaging (tagged MRI) [Park 1996, Huang 1999, Young 1999, Garot 2000]. The intersections of the tagged planes can be used as control points to infer the motion of the myocardium. But this modality is not used routinely in clinical practice. In this work, we will restrict ourselves to the more complicated problem of tracking traditional MRI without the information from the tagged lines. However, the methods presented can be extended with little changes to the other modalities: tagged MRI, Echo or CT.

5.2.1 Cardiac Motion Tracking Algorithms

With tracking algorithms, a quantitative measure of the motion is estimated as a result of registration. Non-rigid diffeomorphic transformations are usually preferred to rigid transformations which are too simple to encode the complex dynamic of the myocardium. In general, the deformation from each frame with respect to a reference is computed to describe how each image differ from the reference image. Therefore, all the registration methods described in Chapter 2 can be extended naturally to cardiac motion tracking. In this section, we will present the specific challenges of registration in the context of cardiac motion tracking and give some example of how classical registration algorithms are adapted this specific task. We refer the reader to the review articles [Makela 2002] and more recently [Wang 2012] for a quite exhaustive list of state-of-the-art methods for cardiac motion tracking.

Incompressibility Constraint One of the major challenge of cardiac motion tracking over classical registration is how to incorporate physiological constraints into the algorithm. The myocardium has been shown to be close to incompressible [TSUIKI 1980, Wenk 2012]. Therefore, algorithms which guarantee exactly or nearly incompressible deformations have been proposed to account for this constraint. In [Bistoquet 2008a, Bistoquet 2008b], a model using a matrix-valued radial basis function to represent divergence free displacement fields, which is a first order approximation of incompressibility, is used. [Zhu 2010] employ two surfaces, each driven by the image-derived information, which a evolving together in a coupled way, keeping the volume of the myocardium constant to ensure incompressibility. The log-demons algorithms was also adapted to account for the incompressibility of the myocardium. The iLogDemons algorithm [Mansi 2011, McLeod 2011] incorporates

this constraint by changing the way the regularization is performed and replacing the traditional Gaussian smoothing by an efficient elastic-like regulariser based on isotropic differential quadratic forms of vector fields. The registration energy functional is finally minimised under the divergence-free constraint. In [McLeod 2012], polyaffine parameters are estimated based on the Demons optical flow approach adding a penalization for incompressibility.

Choice of Parametrization Another challenge is the dimensionality of the representation of the motion. Log-demons algorithm (and other methods such as [Tautz 2011, Curiale 2016]) uses transformations represented by a dense fields parametrized at the voxel level. When these transformations have to be estimated for all the frames of the sequence, the number of parameters representing the motion can become very large and encode lots of redundant information. Therefore, representations using a lower number of parameters have been developed. One of the most popular representation are transformations parameterized using B-spline [Schnabel 2001, Ledesma-Carbayo 2005, Shi 2012, Rueckert 2015]. In [Rueckert 1999] a B-spline based method was introduced in the context of breast MR sequences and was later extended to incorporate a incompressibility constraint in [Rohlfing 2003b].

Polyaffine transformations [Commowick 2008, Arsigny 2009] which are extensions of Polyrigid transformations [Arsigny 2003, Arsigny 2005] are another way to define transformations with a reduced number of parameters. In this framework, rigid or affine transformations are computed at the regional level and fused together in a smooth manner to obtain a diffeomorphic global transformation. Therefore, the global transformation is only represented by a number of parameters equal to the 12 times the number of regions. In Chapter 6 of this work, we will present these transformations in more details.

Group-Wise Registration and Temporal Consistency Finally, the tracking of the cardiac motion should not be considered as the registration of each frame individually but rather as a unique problem of finding the most probable motion taking into account the information of all the group of frames. Therefore, methods to do the registration in a way that ensure temporal consistency have been developed. In [De Craene 2010, De Craene 2011, De Craene 2012], a temporal diffeomorphic framework based on B-spline was proposed where the velocities rather than displacement were represented by a series of free form deformations (FFD). In [Huang 1999] tag line features are extracted from the sequence and used as input of a four-dimensional (4-D) time-varying B-spline model for 4-D tracking while [Wang 2015] proposes a method based on two spatial phases of the signal combined to get an analytical estimator for 2D local displacements. In [Schaerer 2010], the authors use a dynamic model to enforce periodicity and temporal smoothness constraints, while in [Clarysse 2000], the spatio-temporal myocardial displacement field is modeled by a cosine series model fitted to the entire tagged dataset.

The work that will be presented in this chapter takes an innovative approach to tackle this problem of temporal consistency using group-wise registration. Instead of performing the registration with respect to a single reference, a subspace made of multiple reference representing optimally the cardiac sequence is built to serve as a low-dimensional representation of the sequence. We introduce this notion in the following section.

5.2.2 Evaluation of Tracking Methods

Finally, we give a little word on how to evaluate the accuracy of these methods as this is always a burning issue. Both real image sequences (where ground truth motion is defined with landmarks annotated by experts) and synthetic image sequences can be used. Real images provide more realistic and difficult data because all the artifacts and noise inherent to the acquisition are present whereas synthetic sequences do not always reproduce the same level of noise. However, the annotation of the ground-truth motion by experts is a very difficult task and there is usually a low agreement as to what the motion is. This is why we will prefer to use synthetic sequences to evaluate the accuracy of our tracking methods.

An important resources to evaluate the performance are challenges where multiple algorithms are tested on the same data to compare the qualities and the weaknesses of each method. In part of the STACOM workshop at the MICCAI conference, two challenges related to cardiac motion tracking have been organized. The first one, held at STACOM 2011 [Tobon-Gomez 2013], concerned MRI and Echo sequences. The second, held at STACOM 2012 [De Craene 2013], was exclusively using Echo images with a focus on strain assessment. Each of these challenges was open to any team wanting to test, validate and compare their algorithms. We refer the reader to these challenges for a comparative view of the state-of-the-art methods for cardiac motion tracking.

5.3 Motivations: Low-Dimensional Representation of the Cardiac Motion

Usually, the analysis of the motion of the heart is done by performing a statistical study on the deformations computed from time-sequences of medical images. To do so, one needs to cope with the the non-linearity of the space of deformations of medical images and to cast traditional linear statistics in the Riemaniann space of deformations. One elegant framework to study motion is the one defined of Joshi et al. [Joshi 2004]. This framework works in the space of images \mathcal{M} , and we note I a particular point of the manifold: in our case an image. In a nutshell, images are mapped together by deformations which are geodesics paths from one image to the other. Then statistics can be performed on each point of the manifold by analyzing the tangent vector of the geodesic at this point [Rao 2003, Rougon 2004]. In the context of cardiac motion analysis, the different frames of the motion are compared




Figure 5.1: (Left): representation of the classical methodology. A mean point (image in green) is computed and the statistical analysis for each data is done with respect to this point. The reference is not a point of the data. (Right): our proposed multi-reference approach. Each data is directly looking at a set of references chosen among the data, the subspace is not built based on a central point.

and deformation fields matching the voxel intensities of the images are computed - the registration step - in order to have an estimation of the motion through the cardiac cycle. Once these deformations are computed, the statistical analysis relies on finding a good low-dimensional representation of these deformations. In a Euclidean space, one would simply have to compute the *Principal Component Analysis* [Hammer 2009] to get subspaces maximizing the unexplained variance. In the space of deformations on medical images, we need to consider extensions of the PCA to manifolds.

There are different ways to extend the concept of principal affine spaces from a Euclidian space to something defined on manifolds. The simplest generalization is tangent PCA, where a covariance matrix is build on the tangeant space of the Karcher or Fréchet mean [Qiu 2012, Vaillant 2004, Sweet 2010]. In the context of the study of a cardiac cycle, this method would require the definition of a mean point on which the deformation to the rest of the sequence are computed. While this is possible, the mean is often a poor descriptor of the whole cycle due to its circular pattern (see Figure 8.4: top-left). Because of that, the study is often simply done based on the first frame which corresponds to the end-diastole (ED). This frame is gated with the ECG and corresponds to the start of the propagation of the electrical wave, therefore it might seem to be a natural choice to use as a reference to study the whole cardiac motion. But this choice is empirical and introduces biases to the whole study. Moreover, one single reference is often not enough to study the whole motion especially when deformations are large [Tobon-Gomez 2013]. Other methods such as Principal Geodesic Analysis (PGA) [Fletcher 2004] and extension [Zhang 2016] define subspaces which are spanned by the geodesics going through a point. The tangent vector is then restricted to belong to a linear space of the tangent space. But these methods also rely on only one single reference and therefore are

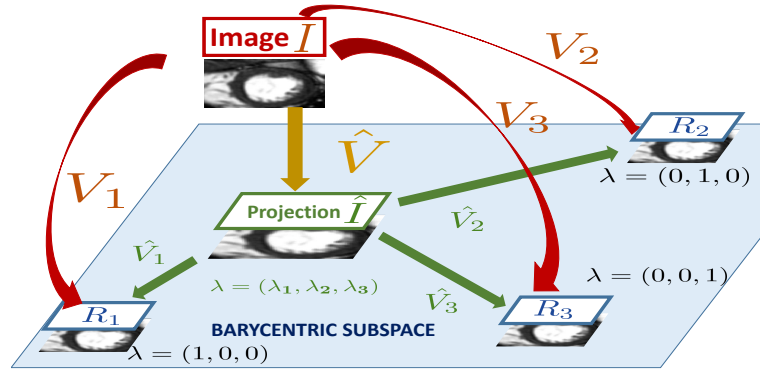


Figure 5.2: Barycentric subspace of dimension 2 built from 3 references images (R_1, R_2, R_3). \hat{I} is the projection of the image I within the barycentric subspace such that $\|\hat{v}\|^2$ is minimum under the conditions $\sum_j \lambda_j \hat{v}_j = 0$ and $\hat{v} + \hat{v}_j = v_j$.

subject to the same limitations when studying the whole cardiac cycle.

For these reasons, there is a need for a new, multi-reference framework to study the cardiac motion. In this paper, we use a more general type of family of subspaces on manifolds called Barycentric Subspaces which was first introduced in [Pennecc 2015]. The point of view taken to construct these subspaces is different from the one traditionally seen in statistical analyses. Instead of building a mean value and subspaces based on the data to study a new point with respect to the mean value, each point is directly analyzed with respect to multiple references as is schematically represented in Figure 5.1. Therefore, the analysis is not performed with respect to a single template: the subspace is built based on multiple reference images chosen among the frames of the sequence. This alleviates the problem of relying on a central value which might not be a good descriptor of all the data. This also gives a more consistent framework to study data in the case where the underlying distribution is either multimodal or simply not sufficiently centered.

We first introduce the methodological background that builds on the theory of [Pennecc 2015] and we apply it to medical images. We define the method to compute the optimal reference images as well as the barycentric coordinates of a projection inside the Barycentric Subspace. Then, we introduce the method to compute an image within the subspace given coordinates and references. We show that this multiple references approach leads to substantial improvement over the traditional single reference methodology. We improve the estimation of the motion through the registration by using our barycentric subspace as a better prior for the cardiac motion tracking, reducing the registration error for the large deformation between end-diastole and end-systole by 30%. Finally, we reconstruct the sequence of images from our low-dimensional representation and show that our method has better results both quantitatively and qualitatively than traditional single reference methods based on tangent PCA.

Algorithm 4 Computation of the projection of an image I to the subspace spanned by the references R_j and the associated barycentric coefficients λ_j^* .

- 1: Given an image I , and a set of references (R_1, \dots, R_{k+1}) compute the registration v_j of the image I to each reference.
 - 2: Derive the closed-form solution $\lambda^* = S^{-1}\mathbf{1} / \sum_i (S^{-1}\mathbf{1})_i$ (see Appendix 5.A).
 - 3: Return: the barycentric coefficients λ^* , the projection vector $\hat{v} = \sum_i \lambda_i^* v_i$, and the projection of the image $\hat{I} = I \circ \exp \hat{v}$.
-

5.4 Methodology: Barycentric Subspaces

In this section, we detail the methods and algorithms to introduce Barycentric Subspaces in the context of medical images. Barycentric Subspaces were first presented for generic Riemannian Manifolds in [Penec 2015] and then adapted in the context of medical imaging in [Rohé 2016b]. We remind here the main steps and notations defined in the previous works to adapt the theoretical framework from Riemannian Manifolds to the context of computational anatomy (image deformation analysis). We follow the generic framework of Joshi et al. [Joshi 2004]: we work in the space of images \mathcal{M} and we use the notation I for a particular point of the manifold, which in our case corresponds to a specific frame within a 3D+t sequence of images of cardiac motion during a cycle. Two images I_1 and I_2 are mapped one onto the other by deformations: the geodesic which is the optimal path from one image to another. Geodesics are represented by the initial velocity field $\overrightarrow{I_1 I_2}$ of the geodesic path. In practice, the geodesic is the result of the registration of the two images which gives us an inexact matching that approximates the tangent vector of the geodesic shooting one image to another. In the following, we will place ourselves in stationary velocity fields (SVF) framework [Vercauteren 2008] which gives a simple and yet effective way to parametrize smooth deformations along geodesics using one-parameter sub-group. We use $v_{i,j}$ as the notation to represent the stationary velocity field parametrization of the deformation mapping image I_i to I_j and we suppose that this SVF is inverse consistent: the inverse mapping of I_j to I_i can be obtained by taking the opposite: $v_{j,i} = -v_{i,j}$.

5.4.1 Definition of the Subspace

A *Barycentric Subspace* of dimension k is defined with respect to a set of $(k+1)$ reference images $(R_j)_{j=1, \dots, k+1}$. While traditional subspaces are defined explicitly, Barycentric Subspaces are defined implicitly as the set of points \hat{I} for which there exists $(k+1)$ *Barycentric coefficients* λ_j which fulfill the condition: $\sum_{j=1}^{k+1} \lambda_j \overrightarrow{\hat{I} R_j} = 0$. Using the notation with SVFs, we write the condition $\sum_{j=1}^{k+1} \lambda_j v_j = 0$, where v_j is the SVF mapping the image \hat{I} to the reference R_j . These notations are schematically represented in Figure 5.2. Since images of the Barycentric Subspace are defined implicitly, we need to introduce specific methods to find the projection of an image within the subspace as well as to compute an image based on its coordinates.

Algorithm 5 Derivation of the optimal references R_j of a set of N images I_n .

- 1: Given N images I_n , compute the cross registration $v_{i,j}$ of all pair of images.
 - 2: **for** all $k + 1$ combinations of references (R_1, \dots, R_{k+1}) within the set of images **do**
 - 3: For $n = 1, \dots, N$, compute the projection vector $\hat{v}(I_n)$ within the barycentric subspace defined by (R_1, \dots, R_{k+1}) with Algorithm 4.
 - 4: Sum the norm of the projection vector: $\mathcal{E}(R_1, \dots, R_{k+1}) = \sum_{n=1, \dots, N} \|\hat{v}(I_n)\|^2$.
 - 5: **return:** the set of references (R_1, \dots, R_{k+1}) realizing the minimum of $\mathcal{E}(R_1, \dots, R_{k+1})$.
-

5.4.2 Projection of an Image to the Subspace

We denote $\hat{I}_{(R_j)_{j=1, \dots, k+1}}$ (or simply \hat{I}) the projection of an image I on the Barycentric Subspace spanned by the reference images $(R_j)_{j=1, \dots, k+1}$. This projection is associated with coefficients $\lambda(\hat{I}) = \lambda_{j=1, \dots, k+1}$ representing the coordinates of \hat{I} (and by extension the low-dimensional representation of I within the Subspace). The projection \hat{I} of I is defined as the closest point to I that belongs to the barycentric subspace. The distance to the subspace is encoded by the norm of the SVF \hat{v} which parametrizes the deformation of I to \hat{I} realizing the minimum distance as shown in Figure 5.2. As seen previously, the constraint that \hat{I} belongs to the barycentric subspace can be written as $\sum_j \lambda_j \hat{v}_j = 0$. Using the Baker-Campbell-Hausdorff (BCH) [Vercauteren 2008] formula, we get a first order approximation of $v_i = \hat{v} + \hat{v}_i$. The problem can now be written as:

$$\min_{\hat{v}} \|\hat{v}\|^2, \quad \text{subject to } \sum_i \lambda_i (v_i - \hat{v}) = 0, \quad \sum_i \lambda_i \neq 0.$$

Since the weights λ_j are defined up to a global scale factor, we can add the additional condition that they should sum to one: $\sum_j \lambda_j = 1$. This way, the coefficients are normalized which make them easier to analyze and they can be seen as Barycentric coefficients in the traditional Euclidean meaning. We also add the constraint $\lambda_i \leq 1$: it forces the projection to lie within a border defined by the references. This forces the references to be the extremal points of the subspace. The condition on \hat{v} and λ becomes:

$$\min_{\hat{v}} \|\hat{v}\|^2, \quad \text{subject to } \hat{v} = \sum_i \lambda_i v_i, \quad \sum_i \lambda_i = 1, \quad \lambda_i \leq 1.$$

A closed-form solution λ^* of this optimization problem can be found by solving the Lagrangian [Bertsekas 2014] (we leave the details at the Appendix 5.A). Finally, the projection vector \hat{v} is simply equal to the weighted sum of the SVF from the registration: $\hat{v} = \sum_i \lambda_i^* v_i$. The computation of the projection vector \hat{v} , the projection of the image \hat{I} and the coefficients λ_j are summed up in Algorithm 4.

Algorithm 6 Computation of an image point from a set of coordinates and optimal references.

- 1: Given $k + 1$ references R_j and $k + 1$ barycentric coefficients λ_j , set $I_0 = R_p$ where p is the index of the largest λ_j .
 - 2: **for** n until convergence **do**
 - 3: Compute the registration of the image I_n with respect to each current reference to get the SVFs v_i .
 - 4: Project the SVFs on the barycentric subspaces to get $\hat{v} = \sum_i \lambda_i v_i$.
 - 5: Warp the references in the direction of the projection: $\hat{R}_i = R_i \circ \exp(\hat{v} - v_i)$.
 - 6: Update the intensity by computing the λ -weighted average: $I_n = \sum_i \lambda_i \hat{R}_i$.
 - 7: **return**: the image I .
-

5.4.3 Computation of the optimal References of a sequence of Images

At this point, we have not defined a methodology to choose the references R_j used to build the subspace among all possible points of the space. Using the fact that the norm of the projection \hat{v} to the subspace encodes its distance to each image, we propose an optimization approach by choosing the references R_j minimizing the average distance to the space:

$$(R_1, \dots, R_{k+1}) = \arg \min \mathcal{E}(R_1, \dots, R_{k+1}) = \arg \min \sum_j \|\hat{v}_j\|^2,$$

where \hat{v}_j is the projection of the point j to the barycentric subspace defined by (R_1, \dots, R_{k+1}) . We sum up the process to find the optimal references in Algorithm 5.

5.4.4 Computation of an Image within the Subspace

A method to synthetically compute an image given a set of coordinates is also needed, in order to extrapolate data within the subspace or to reconstruct a sequence from its low-dimensional representation. Given a set of coordinates λ_j in a Barycentric Subspace defined by $k + 1$ references R_j , we want to compute the image I which fulfills the condition $\sum_i \lambda_i v_i = 0$. This condition alone could lead to multiple solutions: we could start from any of the reference, deform it, and find a different image. Therefore, in order to get a single consistent solution we compute the λ -weighted average of the intensity of the warped images. It also has the benefit to enforce a smooth change of texture as the coefficients change. The algorithm is described in Algorithm 6.




Figure 5.3: (Left): Example of the norm of the projection vector to the Subspace for 1 patient over the cardiac sequence with $1D$, $2D$ and $3D$ –Barycentric Subspace. The projection is null for each of reference frames. (Right): Mean norm over all the frame of the sequences and averaged over all the patients for subspaces of increasing number of references/dimensionality.




Figure 5.4: Schematic representation of the proposed approach, references used for each approach are shown in green. Left: traditional approach where only the ED image is used as a reference leading to large displacements to be evaluated between the ED and ES frame. Right: proposed approached where multiple images are used as references, all the images of the sequence are close to at least one reference.

5.5 Using Barycentric Subspaces as a prior on the Registration

The key instrument of cardiac motion analysis is non-linear registration: it allows to track the motion of the myocardium and compute deformation fields representing the motion during a cycle. These deformation fields can then be used to perform statistical analysis and compare different motions. A great variety of registration algorithms [Klein 2009a, Ashburner 2007a] have been developed in medical imaging, some of which have been adapted and improved to the specific task of cardiac motion tracking [Chandrashekar 2004, Tautz 2011, De Craene 2011]. These algorithms can be roughly classified in two categories. On one side, algorithms which look for smooth deformations by reducing the degrees of freedom and parametrize the deformation using a reduces number of variables, such as for example [Arsigny 2009]. On the

Algorithm 7 Use of the Barycentric Subspace as a prior in the computation of the registration of a frame with respect to $k + 1$ references.

- 1: Registration of the image I with respect to $k + 1$ references R_i .
 - 2: Initialization of the algorithm with the standard registration of the image with respect to each reference to get v_i^0 .
 - 3: Compute the projection \hat{v}^0 , the coefficients λ_i and the barycentric velocity fields \hat{v}_i^0 such that $v_i^0 = \hat{v}_i^0 + \hat{v}^0$ using Algorithm 4.
 - 4: Compute the warped references $\hat{R}_i^0 = R_i \circ \exp \hat{v}^0$,
 - 5: **for** j until convergence **do**
 - 6: Compute the update field \hat{u}_i^j from each warped reference \hat{R}_i^{j-1} to the current image I with the demons forces.
 - 7: Project the update field \hat{u}_i^j to the set of current barycentric velocities \hat{v}_k^{j-1} and updates the barycentric velocity: $\hat{v}_i^j = \hat{v}_i^{j-1} + \sum_k c_k \hat{v}_k^{j-1}$.
 - 8: Compute the warped references $\hat{R}_i^j = R_i \circ \exp \hat{v}_i^j$,
 - 9: Compute \hat{v} mapping each warped reference \hat{R}_i^j to the current image and compose this SVF with the barycentric velocity to get an estimation of the full deformation: $v_i^j = \hat{v}_i^j + \hat{v}$.
 - 10: Extract the barycentric coefficients such that $\sum_i \lambda_i \hat{v}_i^j = 0$
 - 11: Return: the estimated barycentric coefficients λ_i and the $k + 1$ deformations v_i^j mapping symmetrically each reference to the image I .
-

other side, algorithms whose deformations are parametrized by dense vector fields, for example [Beg 2005]. This latter category provides a more generic and less restrictive framework capable of representing any type of deformation of the myocardium at the cost of complexity. In this framework, the direct optimization similarity criterion Sim (measuring the resemblance of two images by comparing voxel intensities) on the whole space of non-parametric transformations leads to an ill-posed problem due to the number of degrees of freedom. To overcome this problem and impose some spatial regularity [Simpson 2015] on the solution, most of the registration methods add a regularization term *Reg* corresponding to the a priori knowledge one has on the transformation to find. Using the parametrization of the deformations by a SVF v as we introduced in Section 5.4:

$$E(v) = Sim(F, M \circ \exp(v)) + Reg(v),$$

where F is the fixed image, M the moving image. The regularization term can take multiple forms, but most of the time registration algorithms consider slowly varying deformation as our prior knowledge of the transformation, thereby forcing the transformation to be as smooth and as small as possible. This methodology is efficient to find small deformations, for example the one mapping the ED to nearby frames, but this kind of regularization often leads to an underestimation of the large deformations as the one happening between the ED and ES frame (see Fig. 5.4 for a schematic representation).

To correct this bias, one possible solution is to perform the registration in




Figure 5.5: (Left): Average point-to-point error on the meshes over the whole cycle between the ground-truth and the deformed meshes compared for the two methods (our proposed approach in dotted lines, and the standard one in plain lines). The registration is performed with respect to the three references for both cases. (Right): volume curves induced by the registration and comparison with the ground truth volume. Our proposed approach (red - dotted) performs a better approximation of the ground truth volume curve (green).

a group-wise manner where a group of images are simultaneously considered and an additional criteria is set up to ensure temporal-consistency [Balci 2007, Perperidis 2005]. We could use multiple images as references and perform registration so that the reference is close to the frame analyzed (see Fig. 5.4). But there is now another problem to be dealt with: the registration of each frame being done with respect to a different reference, there is no a common framework and space to analyze the cardiac motion as a whole. In order to have something comparable for all the references, we proceed differently: instead of performing the registration with respect to one image - the closest reference - we build a subspace containing these references and use it as a prior on the registration process.

To do so, we propose to use the barycentric template defined by 3 reference frames of the sequence as the additional prior on the transformations. Instead of applying the regularization energy $Reg(v)$ to the full deformation v , we compute the projection on the Barycentric Subspace \hat{v} and we apply the regularization only to the reduced portion of the full deformation not explained by the projection corresponding to the difference between v and \hat{v} . This SVF encodes the distance of the image to the subspace and is smaller than the full deformation (see Figure 5.2). By relaxing the regularization this way, we let the registration move freely within the a 2-D barycentric template representing the cardiac motion.

5.5.1 Barycentric Log-Demons Algorithm

The methodology defined is quite generic, we detail in this section one way to implement it in practice in the case of the *LCC Log-Domain Diffeomorphic Demons* algorithm [Lorenzi 2013a]. This algorithm proceeds in multiple iterations of two successive steps: the first step optimizes the matching criteria by computing the so-called *demons forces*, then, in the second step, the estimated velocity field is smoothed by applying a gaussian filter. To allow the registration to freely move in the barycentric subspace defined by a set of images - instead of being constrained by the regularization on the whole velocity field (the "standard" method) - we proceed in two steps: first we evaluate the barycentric subspace structure and then we iterate on this subspace. To do so, we perform one standard registration of the current image I with respect to each of the references R_i to get the velocity fields v_i . Then, we decompose the velocity field v_i as the sum of the barycentric velocity \hat{v}_i warping the reference R_i to the projection \hat{I} inside the barycentric subspace and the residual velocity field \hat{v} of the projection (see Fig. 5.2). Finally, we iterate by projecting the update *demons forces* on the barycentric velocity until convergence. The barycentric template is therefore used as a prior on the cardiac motion and we perform the regularization only with respect to the projecting field. The methodology is described in Algorithm 7.

5.5.2 Evaluation using a Synthetic Sequence

We evaluate the method using one synthetic time series of $T = 30$ cardiac image frames computed using the method described in [Prakosa 2013]. The use of a synthetic sequence has the important advantage to provide a dense point correspondence field following the motion of the myocardium during the cardiac cycle which can be used to evaluate the accuracy of the tracking. Another option could be to use point correspondence manually defined by experts, but they tend to be inconsistent and not reliable [Tobon-Gomez 2013]. First, we compute the optimal references using the methodology described in Algorithm 5, giving us the three reference frames spanning the barycentric subspace: #1 is frame 1, #2 is frame 11 and #3 is frame 21. Then we register each frame i of the sequence using the method described above to get the deformations from each of the three references to the current images using both the *standard* method and our approach using Barycentric Subspaces as a prior. We deform each of the 3 ground truth meshes corresponding to the reference frames (1,11 and 21) with the deformation from the reference frame to the current frame. We compare our approach with the standard approach where the registration between one of the reference and the current frame is done directly. In Figure 5.5 (left), we show the point-to-point registration error of the deformed mesh using the 3 different deformations (one with respect to each references). Substantial reduction of the error (of about 30%) can be seen for the largest deformations (between end-systole and the first reference for the blue curve corresponding to the frame 1 chosen as reference). This comes at the cost of additional error for the small deformations

evaluated at the frame near the respective references. In Figure 5.5 (right), we show the estimation of the volume curve (which is one of the most important cardiac feature used in clinical practice). Our better estimation of the large deformation leads to a substantial improvement of the volume curve estimation. In particular the estimation of the ejection fraction goes from 32% with the standard method to 38%, closer to the ground truth (43%), reducing the estimation error by half.

We have two different methods to compute the coefficients λ , either integrated directly in the registration process (Algorithm 7), or after the registration using Algorithm 5. We apply both of these methods to the previous population of Healthy/ToF patients. The set of curves computed with either methods are very similar with a mean standard deviation of 1.34×10^{-3} between the coefficients. In a context where we just want to extract features of the cardiac motion and not get an improvement on the registration, this shows that the method defined in section 5.4.2 already gives a very robust estimate whereas being computationally more tractable than the method described in this section.

5.5.3 Towards Symmetric Transitive Registration

Traditionally in cardiac motion tracking, two different method for computing the motion deformations can be used. The first method, which we have seen in the previous section, estimates the motion by computing the deformation from each frame to a common reference. The second method computes the deformation mapping each successive frame and then derive the full motion by composing these deformations one by one. The problem encountered by most registration algorithms in this context is the lack of *transitivity* [Škrinjar 2008]: the deformation given by the registration between two images is different when it is done directly or by the composition of the result of the registration with an intermediate image (in our setting this condition can be written as $v_i^j = BCH(v_i^k, v_k^j) \simeq v_i^k + v_k^j$: this is an approximation using the BCH decomposition at the first order as was done before in this paper). This is due to the accumulation of the registration errors at each step of the registration and can lead to large errors at the end of the cycle.

In this last section, we use the barycentric subspace representation to derive a method to get approximately transitive registration (at the first-order of the BCH approximation), an important property of the registration ensuring robustness. This method is schematically represented in Figure 5.6 in the case of a Barycentric Subspace with 2 references. Using Barycentric Subspaces as a basis for the registration at each step, we define the symmetric registration using the following formula 5.1 which is schematically represented in Fig. 5.6 (left):

$$W_s^t = \hat{v}^s - \hat{v}^t + 1/2(\sum_i \lambda_i^t \hat{v}_i^s - \sum_i \lambda_i^s \hat{v}_i^t). \quad (5.1)$$

The equations leading to this formula are detailed in 5.B, and we will only give here a simple interpretation of the meaning of the formula. The first two SVFs on the left represent the residual projections from the barycentric subspace to the two




Figure 5.6: (Left): Schematic representation of the symmetric multi-references barycentric registration in the case of a 1-D barycentric subspace spanned by 2 references. I_t and I_s are two frames of the sequence and \hat{I}_t, \hat{I}_s corresponds to the respective projection to the barycentric subspace. (Right): comparison of the error between the standard registration (blue plain), the barycentric method presented in section 3.1 (blue dotted) and the symmetric-barycentric extension presented in section 3.2 (red dotted).

frames. This encodes what cannot be represented in the subspace. The sum on the right is a symmetric estimation of the SVF \hat{W}_s^t (the vector representing the deformation from the projection of one image to the other one) within the barycentric subspace. This estimation is done by going through each reference image forward and backward and weighting the barycentric velocity by the coefficients on the subspace. We apply it to frame-to-frame registration and show a significant improvement of the accuracy of the registration with respect to a non-transitive method. In figure 5.6 (right), the error for frame-to-frame registration (starting for frame 0) of our method compared with the standard one can be seen. Even though our method has higher error for the first frames of the sequence, the transitivity property ensures that there is not an accumulation of the error for the last frames of the sequence as opposite to the traditional method, meaning that we have substantially less error for frames wt the end of the cycle.

5.6 Conclusion

In this chapter, we have challenged the traditional framework for studying the cardiac motion based on the ED frame chosen as the single reference. Instead, we have proposed a multi-references methodology introducing a new type of subspaces called Barycentric Subspaces. Intuitively, this multi-reference approach is more adapted to the cardiac motion circular pattern and overcomes the drawback of having either to chose one specific template such as the ED frame (whose choice might introduce bias) or to compute a mean image (which might be a poor descriptor of the whole distribution). The practical methods and algorithms to use these subspaces in the

context of the manifold of medical images have been derived: how to compute the projection and the coefficients of a specific image and how to compute an image from the coefficients. We have shown how one can use these subspaces as an additional prior within a registration algorithm. By relaxing the regularization constraint and let the moving image move freely in a $2-D$ subspace we have improved the accuracy of the registration for the large deformations between end-diastole and end-systole. Finally, we have used the Barycentric Subspace to derive a method to get approximately transitive registration (at the first-order of the BCH approximation) for frame-to-frame registration.

5.A Projection of an image to the Barycentric Subspace

Starting from the optimization problem defined in 5.4.2:

$$\min_{\hat{v}} \|\hat{v}\|^2, \quad \text{subject to } \sum_i \lambda_i (v_i - \hat{v}) = 0, \quad \sum_i \lambda_i \neq 0.$$

Without loss of generality, we can add the additional condition that the λ should sum to one: $\sum_i \lambda_i = 1$. We also add the constraint $\lambda_i \leq 1$. keeping in mind that what follows can be easily extended to the unconstrained case. Then, the optimization becomes:

$$\min_{\hat{v}} \|\hat{v}\|^2, \quad \text{subject to } \hat{v} = \sum_i \lambda_i v_i, \quad \sum_i \lambda_i = 1, \quad \lambda_i \leq 1,$$

whose Lagrangian is:

$$\mathcal{L}(\boldsymbol{\lambda}, \mu, \boldsymbol{\kappa}) = \left\| \sum_i \lambda_i v_i \right\|^2 + \mu(1 - \sum_i \lambda_i) + \boldsymbol{\kappa}(1 - \boldsymbol{\lambda}),$$

where $\boldsymbol{\kappa}$ and $\boldsymbol{\lambda}$ are vectors, μ is scalar. The solution can be found by solving the set of equations [Bertsekas 2014]:

$$\begin{cases} \forall j : \langle \sum_i \lambda_i v_i | v_j \rangle &= (\mu + \kappa_j)/2 \\ \forall j : \kappa_j(1 - \lambda_j) &= 0 \\ \forall j : \kappa_j &\geq 0 \\ \sum_i \lambda_i &= 1. \end{cases}$$

If all κ_i are equal to zero, which is equivalent to the inequality constraint not being filled, then we simply have to solve : $\langle \sum_i \lambda_i v_i | v_j \rangle = \mu$ for all j . Denoting S the matrix of the scalar product $S_{i,j} = \langle v_i | v_j \rangle$, this is equivalent to:

$$S\boldsymbol{\lambda} = \mu\mathbf{1}.$$

Finally adding the condition $\sum_i \lambda_i = 1$ gives us the optimal solution $\boldsymbol{\lambda}^*$:

$$\boldsymbol{\lambda}^* = S^{-1}\mathbf{1} / \sum_i (S^{-1}\mathbf{1})_i.$$

If some κ_i are not null, then $\lambda_i = 1$ for these indices. We simply have to solve the lower-dimensional problem removing the satisfied inequality constrains.

5.B Frame-To-Frame Barycentric Registration Formulation

Given two images I_t and I_s , images of a cardiac sequence at frame number t and s , we want to derive the formula for the SVF W_t^s mapping one image to the other using their projection on to the barycentric subspace. We use the notation schematically shown in Figure 5.6: W_t^s will map the two images together, \hat{W}_t^s will map the projections \hat{I}_s and \hat{I}_t , \hat{v}^s will be the projection of one frame s to the subspace, v_i^s will be the SVF mapping the image to the reference i and \hat{v}_i^s will be the SVF mapping the projected image to the reference.

Using a BCH approximation at the first order, we have the following equality with is true with respect to each reference i :

$$\hat{W}_s^t = BCH(\hat{v}_i^s, -\hat{v}_i^t) \approx \hat{v}_i^s - \hat{v}_i^t.$$

Taking the λ -weighted sum and using the fact that $\sum_i \lambda_i^s \hat{v}_i^s = 0$ and $\sum_i \lambda_i^t \hat{v}_i^t = 0$ give us the two following equalities:

$$\hat{W}_s^t = \sum_i \lambda_i^t \hat{v}_i^s = - \sum_i \lambda_i^s \hat{v}_i^t.$$

Finally, we take the average of these two equalities and add the projection vector to get our frame-to-frame formulation of the registration using the subspace:

$$W_s^t = \hat{v}^s - \hat{v}^t + 1/2(\sum_i \lambda_i^t \hat{v}_i^s - \sum_i \lambda_i^s \hat{v}_i^t).$$

Highly Reduced Model of the Cardiac Function for Fast Simulation and Personalization

Contents

6.1	Chapter Overview	93
6.2	Motivations	94
6.3	Reduced-Polyaffine Projection for Compact Cardiac Motion Representation	95
6.4	Biophysical Model of the Heart Simulation Database	96
6.5	Parameters Mapping through PLS Regression	97
6.6	Applications	99
6.6.1	Direct Highly Reduced Cardiac Function Model	99
6.6.2	Personalization of Model Parameters	100
6.7	Conclusion	100

In this chapter, we introduce the polyaffine parameters representation of the cardiac motion. This representation will be used again in Chapter 7 and Chapter 9. Here, we present a drastic dimension reduction method to link the biophysical parameters of an electromechanical of the heart with a compact representation of cardiac motion with the polyaffine transformations. The work presented here was published at the IVMSM conference [Rohé 2016a].

6.1 Chapter Overview

Our approach relies on a projection of the displacement fields along the whole cardiac motion to the space of reduced-polyaffine transformations. Using these transformations, not only we describe the motion using a very small number of parameters but we show that each of these parameters has a physiological meaning. Moreover, using a PLS regression on a learning set made of a large number of simulations, we are able to find which of the input parameters of the model most impact the motion and what are the main relations mapping the polyaffine representation to the parameters of the model. We illustrate the potential of this method both for

building a direct and very fast model characterized by a highly reduced number of parameters and for personalization of such a model from motion.

The main contributions of this chapter with respect to state-of-the-art are:

- The extension of the polyaffine reduction to keep only the 6 most relevant parameters.
- A surrogate highly reduced model for fast forward computation.
- The illustration of the method on the personalization of a cardiac model.

6.2 Motivations

Modelling of the heart had an increasing interest in the recent years as it provides a way to better assess the cardiac function and predict its evolution. In order to estimate subject-specific model parameters, cardiac motion is often used, as it can be extracted from medical images. However, given the complex dynamics of cardiac motion, its analysis and its link with the underlying physical parameters of the cardiac tissue are difficult to achieve. This complexity of the model mostly results from: the number of input parameters, the complexity of the output motion which lies in an infinite dimension space and the non-linear relationship between the input and output. In this article, we introduce a new method to reduce a cardiac motion model to very few parameters by reducing the dimension of the output motion and estimating the main modes of variation linking biophysical parameters and cardiac motion.

A cardiac biophysical model can be described as a function f which maps a geometry S (discretely represented by a segmented mesh which is a finite set of points in \mathbb{R}^3) and a set of input parameters $\Theta = (\theta_i)_{i=1,\dots,M}$ to the simulation of the motion of this geometry through the cardiac cycle $\mathcal{M} = (S_t)_{0 < t < 1}$:

$$\mathcal{M} = (S_t)_{0 < t < 1} = f(S_0, \theta_1, \dots, \theta_i, \dots, \theta_M). \quad (6.1)$$

In order to reduce such model, we first need to be able to express the cardiac motion $\mathcal{M} = (S_t)_{0 < t < 1}$ by a reduced number of parameters $M = (m_j)_{j=1,\dots,N}$. To do so, we use a polyaffine projection [Mcleod 2013c] which we further develop by expressing the parameters on a basis adapted to the heart geometry and his motion. In this new frame we only project on the 6 (instead of 12 for standard affine) most relevant parameters. We show that, not only the parameters of this reduced-polyaffine projection gives a very good approximation of the whole cardiac motion, but also that each of these parameters is physiologically interpretable. We then learn the relation between both the reduced-polyaffine parameters M and the input parameters Θ of the model with a Partial Least Square (PLS) regression. This relation will be the basis of our two applications. First, we build a direct forward model using only a highly reduced number of parameters: we show for 100 simulations with different input parameters Θ that we are able to reconstruct

a motion directly from the parameters M given by the first modes of the PLS regression that is very close to the full model computation. Then we tackle the inverse problem of personalization: given a motion M , we can find parameters Θ for which the simulation through the full model gives a motion approximating the original one.

Such approach is a hybrid method between recently developed hyper-reduced models, looking for a lower dimensional representation of the spatial state variables [Ryckelynck 2009] and the meta-models, looking for correlations between input and output of models [Tøndel 2012]. One important difference is that we explicitly define the first dimension reduction in the spatial domain in order to have meaningful parameters, related to regional deformations.

6.3 Reduced-Polyaffine Projection for Compact Cardiac Motion Representation

In this section, we propose a method to project a given cardiac motion \mathcal{M} to a subspace of finite dimension in which it will be represented by a set of parameters $M = (m_j)_{j=1, \dots, N}$:

$$\pi : \mathcal{M} \rightarrow \tilde{\mathcal{M}} \sim M = (m_1, \dots, m_j, \dots, m_N).$$


We suppose we have a cardiac motion represented by T displacements fields $(\vec{D}_t)_{t=1, \dots, T}$ mapping each point of the initial mesh S_0 to the corresponding point of the mesh S_t at frame t . Instead of looking at displacements fields, we choose to represent the cardiac motion by stationary velocity fields (SVF) such that $\vec{v}_t = \log \vec{D}_t$. As described in [Arsigny 2006], working with SVF allows us to perform vectorial statistics on diffeomorphisms, while preserving the invertibility constraint, contrary to the Euclidian statistics on displacement fields. Since the space of the SVF is dense, we need to reduce the dimension by projection onto a lower dimensional space. To do so, we will use the polyaffine projection [Mcleod 2013c].

By defining K regions and associated smooth weights, we describe locally affine deformations using few parameters while still being globally invertible. The polyaffine transformation is the weighted sum of these locally affine transformations:

$$v_{poly}(x) = \sum_{k=1}^K \omega_k(x) M_k \tilde{x}.$$

We use the standard American Heart Association (AHA) 17 regions for the left ventricle. We also define 10 additional regions dividing the right ventricle in a similar way for a total of $K = 27$ regions. The weights ω_i are normalized Gaussian function around the barycentre of each region [Mcleod 2013c]. The parameters of the optimal projection of a Stationary Velocity Fields v onto the space of polyaffine transformations has an analytical solution $m = \hat{m} = \Sigma^{-1}b$ [Seiler 2012] which minimizes in the least-square sense:

$$C(m_1, \dots, m_N) = \int_{\Omega} \|\sum_i v_{poly}(x) - v(x)\|^2 dx \simeq \frac{1}{2}(m - \hat{m})^T \Sigma (m - \hat{m}) - \frac{1}{2} \hat{m} \Sigma \hat{m}.$$



In order to get interpretable parameters for each region we chose to express them in a local coordinate system adapted to the geometry of the heart. If we call $\mathcal{R} = (O, e_1, e_2, e_3)$ the original Cartesian coordinate system, we define the local coordinate of the region k as $\mathcal{R}'_i = (O_k, e_1^k, e_2^k, e_3^k)$ where O_k is the barycenter of the region (the red point in the enclosed figure), e_1 the radial vector (green vector), e_2 the longitudinal vector (purple vector) and e_3 the circumferential vector (blue vector).

We can express the polyaffine parameters $M = (R, T)$, where R is the 3×3 matrix of the rotational parameters and T is the translation, in this new frame. We propose a method to reduce further the model by keeping only the 3 parameters of the translation and the 3 parameters of the strains (diagonal). We first introduce the projection matrix \mathcal{Q} which is a $12N \times 6N$ matrix giving the relation between the $6N$ translation and diagonal parameters expressed in the local coordinates m_L and the $12N$ parameters expressed in the original coordinates m such that $\mathcal{Q}m_L = m$. When expressing m this way, we constrain it to be within the subspace spanned by \mathcal{Q} which is exactly the subspace of the polyaffine transformation whose non-diagonal and non-translation parameters are equal to zeros in the local coordinates. The least-square minimization can now be rewritten as:

$$C(M_1, \dots, M_N) \simeq \frac{1}{2}(\mathcal{Q}m_L - \hat{m})^T \Sigma (\mathcal{Q}m_L - \hat{m}) - \frac{1}{2} \hat{m} \Sigma \hat{m}$$

$$\frac{\partial C}{\partial m_L} = \mathcal{Q}^T \Sigma (\mathcal{Q}m_L - \hat{m}) = 0 \implies m = \mathcal{Q}m_L = \mathcal{Q}(\mathcal{Q}^T \Sigma \mathcal{Q})^{-1} \mathcal{Q}^T \Sigma \hat{m} \quad (6.2)$$

We show the results of a polyaffine projection for a whole cardiac motion in Fig. 6.1. The projection has a mean absolute error below $1mm$, and we explain more than 80% of the original motion (in \mathcal{L}^2 norm of the velocity field with respect to the whole displacement). We can also see that the 3 diagonal parameters gives a good account of the 3 strains. We show both the original motion and the polyaffine projection as a supplementary material.

6.4 Biophysical Model of the Heart Simulation Database

We use a cardiac mechanical model based on the Bestel-Clement-Sorine (BCS) modeling [Bestel 2001]. This model describes the heart as a Mooney Rivlin passive material, and model the stress along the cardiac fibres according to microscopic scale phenomena. Particularly, this model is compatible with the laws of thermodynamics. It also integrates a circulation model representing the 4 phases of the cardiac model, where the aortic pressure is modelled by a 4-parameter Windkessel model. Finally, the electrophysiological pattern of activity is simulated using




Figure 6.1: Polyaffine projection for a cardiac motion. (Left) figure: error of the projection versus absolute displacement. (Right) figures: diagonal parameters over time of the polyaffine matrix for each of the region of the left ventricle from which we can recognize (resp. from left to right) the radial, longitudinal and circumferential strains.

an Eikonal model, describing the depolarization front propagation from the endocardial surface to the whole myocardium. Fourteen parameters are used by the model: $(\sigma_0, k_{rs}, k_{atp}, k_0, \alpha, \mu, E_s)$ active parameters, (K, c_1, c_2) passive parameters and (R_p, τ, Z_c, L) for the valve model. We simulate $S = 500$ cardiac cycles whose input parameters are drawn randomly according to a uniform distribution within a range of parameters chosen as to obtain a physiologically realistic behaviour (following the sensitivity analysis of the model done in [Marchesseau 2012]). In total, each simulation has therefore $M = 14$ specific parameters $\theta_{i=1, \dots, M}$ of the model. On the other side, for each of these motions we calculate the $N = 6 \times K \times T$ reduced-polyaffine parameters $m_j = (m_{l,k,t})_{l=1, \dots, 6, k=1, \dots, K, t=1, \dots, T}$.

6.5 Parameters Mapping through PLS Regression

In this section, we are looking to link the polyaffine parameters M with the input parameters of the cardiac model Θ . PLS regression finds the multidimensional direction in the X (the predictors variables) space that explains the maximum multidimensional variance direction in the Y space (the dependant variables) [Rosipal 2006]. It combines both features from the PCA (the projection of Y and X into subspaces of high variance) and standard linear regression (by the search of linear relations between the modes of X and Y). With X the parameters of the polyaffine and Y the parameters of the model, PLS returns the modes in the space of the polyaffines transformation which have maximum variance and maximum covariance with the parameters of the model Y . It is more robust than standard regression when the space of X and Y is highly-dimensional because of the embedded dimension reduction when looking for linear relations.

With our two sets of parameters $\Theta = \theta_i$ and $M = m_j$, we can either try to predict the parameters M from the parameters Θ and use this relation to build a highly reduced direct model as we show in section 6.6.1, or we can estimate the relation between the parameters Θ of the model from the polyaffine parameters M of the motion and apply this relation to the problem of personalization as we present in section 6.6.2. We therefore compute the PLS regression both with $X = \Theta$ the




Figure 6.2: (Top) row: PLS regression with $X = \theta$ the predictor variable and $Y = p$ the dependant variable. (Bottom) row: $X = M$ and $Y = \theta$. (Left) figure: variance explained by the PLS regression with respect to the number of modes. (Right) figure: VIP of each of the parameters.

predictor variable and $Y = M$ the dependant variable and with $X = M$ and $Y = \Theta$.

In Fig. 6.2, we show the variance explained by the PLS regression with respect to the number of modes used. We see that, with only 5 modes, we can explain more than 90% of the polyaffine parameters variance using only 30% of the variance of the cardiac model parameters. This shows that we can explain most of the impact of the parameters of the model by only few of them and that 70% of the variance of the parameters of the model impacts the motion represented by the polyaffine parameters by only 10%. On the other side, we do not predict as well the parameters $Y = \Theta$, which could explain problems of identifiability, but the parameters we predict are highly correlated to actual change in the motion looking at the explained variance of $X = M$ with only few modes.

We also calculate the Variable Importance in the Projection (VIP) [Tran 2014] for both of our regression with 5 modes. VIP compares the variance explained by the modes of the PLS regression for each of the variable in X . Variables with high impact on the dependant variables will be well estimated by the modes of X and have high VIP factor. The parameters shown in red are the one that were considered as the most important for personalization in [Marchesseau 2012]. Our statistical analysis confirms these findings by quantifying the importance of each of




Figure 6.3: (Left) figure: mean absolute displacement of the motion and mean absolute error of the direct model. (Right) and (Middle) figures: volume curve defined by the direct model versus the full model for a case of simulation with high ejection fraction and one with low ejection fraction.

these parameters with the VIP. Finally, the analysis of the VIP for the other way shows that radial strain and contraction are the most important motion features to explain the parameters of the model, which is also logical.

6.6 Applications

6.6.1 Direct Highly Reduced Cardiac Function Model

In this section, we build a direct cardiac model using the PLS modes of X and Y and the linear relation found previously. Our reduced model can be expressed as a linear function g that approximates the function f from equation 6.1 and that can be expressed with only L parameters, the PLS modes, which linearly map the M inputs of the model Θ to the N polyaffine parameters. This reduced model has two sources of error compared to the full model. First, the subspace of polyaffine transformations already gives an approximation of the full motion as seen in Section 7.3, so that even the projection of the motion on this subspace is an approximation. Secondly, the polyaffine parameters are estimated from the PLS modes and are only approximating the best polyaffine parameters given by the projection.

We perform 100 new simulations and show in Fig. 6.3 how well our highly reduced model approximates the motion given by the full model. The mean error of the points of the mesh along motion stays below $2mm$ which means that we explain more than 75% of the complete motion. The impact on the volume curve of using an increasing number of modes on top of just the mean parameters is shown for two simulations: one with high ejection fraction and one with low ejection fraction. The first mode already approximates quite well the volume curve, giving a physiological meaning of this mode, whereas adding more modes to the model improves the approximation. We show the simulations for these two cases as supplementary materials.




Figure 6.4: (Left) figure: percentage of explanation of each of the model parameters by the calibration for 1 mode (blue bars), 5 modes (red bars) and 50 modes (green bars). (Right) and (Middle) figures: volume curve defined by the direct model versus the full model for a case of simulation with high ejection fraction and one with low ejection fraction.

6.6.2 Personalization of Model Parameters

We now look at the inverse problem: we want to see if we can recover the model parameters (or, at least, parameters that give a close motion) from a given motion and its polyaffine parameters representation. We ran 100 simulations with different parameters on which we calculate the polyaffine projection of the motion. Then, using the linear relation from the PLS regression coefficients, we find the estimated personalised parameters and compare the motion with these parameters to the original one. We show the results in Fig. 6.4. Only the contraction parameter σ_0 is well approximated with the first mode, which confirms that this mode is linked to the contraction of the heart and the volume curve as seen in Sec. 6.6.1. Even though we get some errors in the estimation of the parameters of the model, the first modes of the PLS give model parameters that generate a motion whose volume curve is very close to the original one.

6.7 Conclusion

In this chapter, we proposed an innovative methodology to reduce a whole biophysical cardiac model using polyaffine projection and PLS regression. While the initial model is very complex (as a non-linear application from a large number of parameters to the whole space of displacement fields), we are able to explain it with good accuracy with only a couple of modes. Moreover, we have shown the potential of this method for personalization and for direct forward computation of the motion. Extension of this method to other inputs of a cardiac function model such as fibers orientations or electrophysiology might be possible and give further insight into how each input component impacts the model. Using the inverse relation found by the PLS between the modes of the polyaffine parameters and the input parameters, we can use the methodology presented to efficiently calibrate the model based on the polyaffine transformations.

Part III

APPLICATION TO DIAGNOSIS
AND LONGITUDINAL
ANALYSIS

Polyaffine Transformations for the Automatic Diagnostic of LV Infarct

Contents

7.1	Motivations	104
7.2	Database of Asymptomatic Subjects and Patients with Myocardial Infarction	105
7.3	Extraction of features of interest through shape and motion dimensionality reduction	106
7.3.1	Polyaffine projection	106
7.3.2	Thickness parameters	108
7.4	Dimensionality reduction of the parameters and classification	110
7.4.1	Learnt Dimensionality Reduction	110
7.4.2	Cross-Validation on Training Set: Classifier Selection	112
7.5	Results and Validation on Testing Set	113
7.6	Conclusion	114

Part I and II, we show how to process medical images of cardiac motion in order to get a low-dimensional representation. Finally, in this part, we apply this low-dimensional representation to the clinical problems of diagnostic and longitudinal analysis. This chapter tackles the problem of diagnosis in the context of Left-Ventricle infarct using a low-dimensional representation of the heart with polyaffine transformations.

The work presented in this chapter was published in [Rohé 2015]. It was done in the context of the *LV statistical shape modelling challenge: myocardial infarction*, a challenge to model the statistical shape of the left ventricle to detect myocardial infarction. 11 research centers participated at this challenge which took place during the Statistical Atlases and Computational Modeling of the Heart (STACOM) workshop at MICCAI 2015 (Munich). A journal paper [Suinesiaputra 2017] collates the challenge results together with the baseline prediction model, and discusses the main advantages and disadvantages of the different approaches. We first present the motivations and the data of the challenge, then our method and the results on the training set and finally how our method compare with the other competitive methods with the final results on the test set.

7.1 Motivations

Because of normal aging or medical events, the shape and function of the heart is modified in order to maintain sufficient cardiac output. This process is known as remodeling. It can either be positive (adaptive)[Kallaras 2001, Muhl 2008] as for example the gradual improvement of the heart function in response to its increasing solicitation when practicing regularly sport. On the other side, maladaptive remodeling is often the effect of a disease or an adverse cardiac event. Usually, this remodeling is the answer of the body following a difficulty to provide sufficient cardiac output. It is often at first positive and beneficial before cardiac function deteriorates into adverse remodeling.

One of the main pathologies affecting cardiac motion and output is myocardial infarction (MI). Ventricular remodeling due to MI is extensively studied in the literature [McKay 1986, Yousef 2000, Gaudron 1993]. At the early stage of infarction, wall stress increases due to infarct expansion, which forces the left ventricle (LV) to dilate to maintain the supply of blood to the circulatory system. Myocardial infarction occurs when blood flow to the heart muscle is lowered and the myocardial cells in the territory start dying. The local contractility is reduced and can lead, if prolonged, to severe remodelling of the heart to maintain physiological constraints [Konstam 2011]. The function of the heart is then impaired [Bijnens 2007], and is no longer able to pump as efficiently as it used to, which might cause complications. Acute complications may include heart failure if the damaged heart is no longer able to pump blood adequately around the body.

Therefore, a quantitative understanding of this pathology and how the heart function changes with an infarct is highly desired. Several methods for computer-aided diagnosis of infarction have already been developed using echocardiographic images of the heart coupled with pattern recognition algorithms [Sudarshan 2013] although none of the features used are explicitly related to physiological characteristics of cardiac function. Hence, automated characterization and quantification of adverse remodeling would be a valuable tool for clinicians to quantify the progression of heart disease or to estimate the benefit of a medical treatment.

Due to the complexity and high-dimensionality of these data, we try to quantify both shape and motion using a limited number of parameters, which we combine and use to compare patients and learn the main modes characterizing both populations. The features of interest characterizing the shape of the patients consist of the regional thickness at both end-diastole and end-systole. We also use features representing the deformation along the cycle. Our approach relies on statistics on the motion of the heart between end-diastole and end-systole. We project the motion on the subspace of polyaffine transformations [Arsigny 2009]. With these transformations, we can express a deformation with a limited number of parameters [McLeod 2015a]. We develop further the methodology by reducing the transformations to keep only the most relevant parameters.

Then, we test classical machine learning algorithms on our set of combined shape/motion parameters and compare the performance of each algorithm using




Figure 7.1: Example of the supplied data points: the epicardial surface is shown in red and the endocardium in black/green. The contraction seen between end-diastole (ED) and end-systole (ES) is the main factor we evaluate to detect infarction.

cross-validation techniques. Validating the method with 10-fold cross-validation, we get results of 95% correct labeling on yet-unseen cases of the training set. In addition, our method notably highlights the relative importance of the different features for the classification of this population. Finally, our method is applied on a testing set of 200 patients together with 10 other competitive methods. We show state-of-the-art results with respect to classification accuracy metrics. Notably, our method manages to achieve a 100% sensitivity, meaning it did not miss the classification of any patient of the infarcted patient.

7.2 Database of Asymptomatic Subjects and Patients with Myocardial Infarction

In this challenge, the goal is to classify between control subjects and patients with infarct in an automatic way, based on the STACOM statistical shape modeling labeled data consisting of a segmentation of the myocardium (both epi and endo) wall at end-diastole and at end-systole.

MESA and DETERMINE database. The database is made of 200 asymptomatic volunteers from the MESA study (Multi-Ethnic Study of Atherosclerosis) [Bild 2002]. These subjects did not present any clinical symptoms of cardiovascular disease at recruitment. The other half is made of 200 patients with clinical evidence of myocardial infarction from the DETERMINE database [Kadish 2009] (Defibrillators to Reduce Risk by Magnetic Resonance Imaging Evaluation). Both

these datasets are available on request from the Cardiac Atlas Project database [Fonseca 2011]. The 400 subjects were randomly split into 200 cases for training and 200 cases for validation. These two categories of subjects may differ both in the shape of the heart and in the deformation along the cycle. Indeed, after an infarct the damaged region will tend to shrink and the deformation along the cycle will be lower. Similar studies have already been done with the same dataset such as [Zhang 2015], which focuses on the shape differences between both population whereas we use both shape and motion features.

Segmentation method. A set of 3D points in Cartesian coordinates from the surface (both endocardium and epicardium) of LV shapes were provided for all the patients at End-Systole (ES) and End-Diastole (ED) (see Fig. 7.1). These surfaces have point-to-point correspondences between all the patients and between ED and ES. To compute these segmentations, a custom-made software package (CIM version 6.0, University of Auckland, New Zealand) was used to fit a finite element LV model [Nielsen 1991] onto cine MR images. The fitting was performed interactively and manually using guide-point modeling technique [Li 2010]. The models were registered rigidly to the anatomical landmarks of each heart, therefore the points of the model are considered to be in correspondence for each patient. Since images were acquired using a different imaging protocol between the healthy group (Gradient Recalled Echo or GRE) than the infarcted group (Steady-State Free Precession or SSFP), the segmentation were corrected for protocol bias using the method described in [Medrano-Gracia 2013].

7.3 Extraction of features of interest through shape and motion dimensionality reduction

In this section, we introduce the first dimensionality reduction that is applied to the studied data (made of one segmentation at end-diastole and one at end-systole, each comprised of 1089 points both for the endocardium and the epicardium). It consists in a non-learning approach to project the data of these segmentations to a limited number of regional parameters representing motion and shape.

7.3.1 Polyaffine projection

Due to point to point correspondence of the meshes and prior registration, we already have an estimate of the displacement field ϕ mapping each point at end-diastole to the corresponding point at end-systole. Instead of looking at displacements fields, we choose to represent the cardiac motion by the stationary velocity fields (SVF) \mathbf{v} such that $\mathbf{v} = \log \phi$. Working with SVF allows to perform vectorial statistics on diffeomorphisms, while preserving the invertibility constraint, contrary to the Euclidian statistics on displacement fields.

In [Arsigny 2009], the authors introduce the space of Log-Euclidean Polyaffine

Transformations (LEPT). By defining K regions and smooth weights $\omega_k(x)$, these transformations have the properties to describe locally affine deformations using few parameters while still being invertible. The polyaffine transformation is the weighted sum of these locally-affine transformations \mathbf{M}_k :


$$\mathbf{v}_{poly}(x) = \sum_{k=1}^K \omega_k(x) \mathbf{M}_k \tilde{x}.$$

In the case of cardiac motion, we have a standardized regional decomposition into the standard American Heart Association (AHA) 17 regions for the left ventricle. We define the weights ω_k as normalized Gaussian functions around the barycenter \bar{x}_k of each region such that:

$$\tilde{\omega}_k(x) = \exp\left(\frac{\kappa}{2}(x - \bar{x}_k)^T \phi_k^{-1}(x - \bar{x}_k)\right), \quad \omega_k(x) = \frac{\tilde{\omega}_k(x)}{\sum_{j=1}^N \tilde{\omega}_j(x)}.$$

If we gather the parameters of the polyaffine transformation into a large vector m such that $\mathbf{m} = \text{vect}(\mathbf{M}_1, \dots, \mathbf{M}_K)$. The parameters of the optimal projection of a Stationary Velocity Fields \mathbf{v} onto the space of polyaffine transformations has an analytical solution [McLeod 2015a] $\mathbf{m} = \hat{\mathbf{m}} = \Sigma^{-1}b$, which minimizes in the least-squares sense:

$$C(\mathbf{M}_1, \dots, \mathbf{M}_K) = \int_{\Omega} \|v_{poly}(x) - v(x)\|^2 dx \simeq \frac{1}{2}(\mathbf{m} - \hat{\mathbf{m}})^T \Sigma (\mathbf{m} - \hat{\mathbf{m}}) - \frac{1}{2} \hat{\mathbf{m}} \Sigma \hat{\mathbf{m}}.$$



In order to get interpretable parameters for each region, we choose to express them in a local coordinate system adapted to the geometry of the heart. If we call $\mathcal{R} = (\mathbf{O}, \mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$ the original Cartesian coordinate system, we define the local coordinate of the region k as $\mathcal{R}'_i = (\mathbf{O}_k, \mathbf{e}_1^k, \mathbf{e}_2^k, \mathbf{e}_3^k)$ where \mathbf{O}_k is the barycenter of the region (the red point in the enclosed figure), \mathbf{e}_1 the radial vector (green vector), \mathbf{e}_2 the longitudinal vector (purple vector) and \mathbf{e}_3 the circumferential vector (blue vector). We can express the polyaffine parameters $\mathbf{M} = (\mathbf{R}, \mathbf{T})$, where \mathbf{R} is the 3×3 matrix of the rotational parameters and \mathbf{T} is the translation, in this new frame through the equations:

$$\begin{aligned} \mathbf{R}'_k &= \mathbf{P}_k^{-1} \mathbf{R}_k \mathbf{P}_k \\ \mathbf{T}'_k &= \mathbf{P}_k^{-1} (\mathbf{R}_k \mathbf{O}_k + \mathbf{T}_k), \end{aligned}$$

where \mathbf{P}_k is the transfer matrix from the base $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$ to the base $(\mathbf{e}_1^k, \mathbf{e}_2^k, \mathbf{e}_3^k)$. Then, the new expression of the parameters in this local coordinates system:

$$\mathbf{M}_k = \begin{bmatrix} s_r & a_{1,2} & a_{1,3} & t_r \\ a_{2,1} & s_l & a_{2,3} & t_l \\ a_{3,1} & a_{3,2} & s_c & t_c \end{bmatrix},$$

can be related to physiological deformation. The 3 translation parameters correspond to the motion along the 3 local axes (radial, longitudinal, and circumferential) whereas the diagonal coefficients correspond to the strain along these directions.

We propose a method to further reduce the model by keeping only the 3 parameters of the motion and the 3 parameters of the strain. This defines a polyaffine projection that, when expressed in the local basis previously defined, has only these parameters not equal to zero. We first introduce the projection matrix \mathbf{Q} which is a $12K \times 6K$ matrix giving the relation between the $6K$ translation and diagonal parameters expressed in the local coordinates \mathbf{m}_L and the $12K$ parameters expressed in the original coordinates \mathbf{m} , such that $\mathbf{Q}\mathbf{m}_L = \mathbf{m}$. When expressing \mathbf{m} this way, we constrain it to be within the subspace spanned by \mathbf{Q} . This subspace corresponds exactly to the polyaffine transformation whose non-diagonal and non-translation parameters are equal to zero in the local coordinates. The least-square minimization can now be rewritten as:

$$C(\mathbf{m}) \simeq \frac{1}{2}(\mathbf{Q}\mathbf{m}_L - \hat{\mathbf{m}})^T \Sigma (\mathbf{Q}\mathbf{m}_L - \hat{\mathbf{m}}) - \frac{1}{2} \hat{\mathbf{m}} \Sigma \hat{\mathbf{m}}$$

$$\frac{\partial C}{\partial \mathbf{m}_L} = \mathbf{Q}^T \Sigma (\mathbf{Q}\mathbf{m}_L - \hat{\mathbf{m}}) = 0 \implies \mathbf{m} = \mathbf{Q}\mathbf{m}_L = \mathbf{Q}(\mathbf{Q}^T \Sigma \mathbf{Q})^{-1} \mathbf{Q}^T \Sigma \hat{\mathbf{m}}$$

For each of the 200 training data we compute the LEPT projection of the deformation field. We are able to parametrize the 3D displacement fields (made of 6534 parameters: 3 parameters for each of the 2178 points of the mesh) by only $6K = 102$ polyaffine parameters. Despite this large reduction of dimensionality, these parameters explain on average more than 70% of the original displacement. Box-plots of each of the 6 parameters are shown in Figure 7.2, where the most discriminant parameters (p value < 0.001) are highlighted in bold. The radial displacement as well as the strain are significantly lower (in absolute value) for the infarcted subjects, which is consistent with what would be clinically expected. Similar differences can be seen for the longitudinal parameters. On the other side, the circumferential motion is less significant, mostly due to the fact that it is very hard to track it accurately with clinical images and therefore not reflected in the provided meshes.

7.3.2 Thickness parameters

On top of the polyaffine parameters that characterize the deformation of the heart during a cardiac cycle, we also introduce parameters representing the overall shape of the heart. We choose to study the thickness of the wall within each of the AHA zones at ED and ES. These parameters correspond to the initial and final stages of the




Figure 7.2: Parameters of the polyaffine projection both for infarcted patients (red) and control subjects (blue). (Top row): radial parameters for both the diagonal parameters - representing strain - and the translation parameters - representing motion. (Middle row): longitudinal parameters. (Bottom row): circumferential parameters. In bold the most significant parameters (p -value < 0.001).

transformation from ED to ES, and therefore complement the above-described parameters. We define the thickness as the local distance between endocardial points and their corresponding epicardial locations. These values are also averaged per




Figure 7.3: Box-plots of the thickness of the myocardium wall per AHA region. In blue, the control set and in red the patients with infarcts. (Left): end-diastole. (Right): end-systole. In bold the most significant parameters (p -value < 0.001).

AHA zone, and summarized in Fig. (7.3). Significant differences are observed in the thickness of the myocardium wall at end-systole in most of the regions, especially near the apex, for the diseased patients with respect to the control group. On the other side, thickness ED diastole is less discriminant between both groups. Other parameters related to shape were considered (such as the height of the heart at ED/ES and the diameter at the base at ED/ES) but no significant differences between both populations were seen and therefore we do not use them for classification.

7.4 Dimensionality reduction of the parameters and classification

In this section, we use both polyaffine and thickness parameters previously introduced in order to classify between healthy and infarcted subjects. We use the machine-learning toolbox Scikit-Learn [Pedregosa 2011] to test a collection of standard state-of-the-art algorithms on our dataset and compare their performance in predicting yet unseen data. The features that serve to feed the tested learned algorithms were considered in four different ways: either polyaffine or thickness parameters separately (sets # 1 and # 2), or concatenated without normalization (# 3) or with normalization so that they have a mean of 0 and a variance of 1.

7.4.1 Learnt Dimensionality Reduction

Complementary to the a-priori reduction of dimensionality imposed by the polyaffine model and the use of 17 AHA regions, we also evaluated the influence of a second dimensionality reduction of the data both with a Principal Component Analysis (PCA) and a Principal Least Square (PLS) decomposition [Wold 2001, Rosipal 2006] prior to the tested algorithm.

Algorithm 8 Partial Least Square Regression with PLS1 Algorithm

- 1: Input: X , an $n \times m$ matrix of predictors, Y an $n \times p$ matrix of responses, $l \leq n-1$
the number of modes we want to compute
 - 2: Center the variables: $X_c^0 = X - \bar{X}, Y_c^0 = Y - \bar{Y}$
 - 3: (Optional) Normalize the variables: $X_c^0 = X_c^0 / \text{std}(X_c^0)$ and $Y_c^0 = Y_c^0 / \text{std}(Y_c^0)$
 - 4: $w_0 = X_c^{0T} Y_c^0 / \|X_c^{0T} Y_c^0\|$
 - 5: $t_0 = X_c^0 w_0$
 - 6: **for** $k = 0$ to $l - 1$ **do**
 - 7: $t_k = t_k / \|t_k\|^2$
 - 8: $p_k = X_c^{kT} t_k$
 - 9: $q_k = Y_c^T t_k$
 - 10: $X_c^{k+1} = X_c^k - t_k \|t_k\|^2 p_k^T$
 - 11: $w_{k+1} = X_c^{T^{k+1}} Y_c$
 - 12: $t_{k+1} = X_c^{k+1} w_{k+1}$
 - 13: Output: the matrices P, Q, W, T whose columns are the $(p_k, q_k, w_k, t_k)_{k \leq l-1}$.
-

PCA and PLS for dimensionality reduction. PCA is designed to spread the data according to the main modes of variability and is known to be a useful dimension reduction pre-processing to prevent over-fitting and improve the performance of some machine-learning algorithm.

PLS looks at modes of the input variables that correlate the most with an output variable (in our case the pathology label 0 or 1). Therefore, in contrast with PCA, the modes also correlate with our classification. The method can be considered as the optimal estimation of two latent variables u and v such that:

$$\max_{|u|=|v|=1} \text{cov}(Xu, Yv) = \max_{|u|=|v|=1} \text{var}(Xu) \text{corr}(Xu, Yv)^2 \text{var}(Yv).$$

Using the centered variable $X_c = X - \bar{X}$ and $Y_c = Y - \bar{Y}$, the general underlying model of multivariate PLS can be written as:

$$\begin{aligned} X_c &= TP^T + E \\ Y_c &= UQ^T + F, \end{aligned}$$

where X is an $n \times m$ matrix of predictors (n number of samples and m number of features), Y is an $n \times p$ matrix of responses, T and U are $n \times l$ (l number of modes) matrices of projections on X and Y corresponding to the scores on both the input and output space, Q and P are $m \times l$ orthogonal loadings matrices and finally E and F are errors terms. With respect to PCA, one can see that the main difference comes from the fact that the dimension reduction is not only done for the predictors but for both the predictors and the responses. Several algorithms exist to compute the PLS modes, we present the *PLS1* algorithm [Gil 1998, Trygg 2002] which is a widely used recursive algorithm (Algorithm 8).




Figure 7.4: Difference between PCA and PLS for dimensionality reduction and classification (first 2 modes of the decomposition plotted in 2D space. Both decomposition make 2 clusters appears but the PLS decomposition does it more efficiently.

Fig. 7.4 shows the first 2 modes of the decomposition and one can see the difference between both methods. PLS is a supervised method and therefore can more efficiently finds modes that separates both the healthy and diseased subjects. PCA, even though it is unsupervised, still finds modes that correlates quite highly with the clinical diagnosis.

Finally, the coefficients of the modes of the PLS reduction can give further insight into the main contributors of the infarct classification. Fig. 7.5 summarizes the loadings of the first mode of the PLS with respect to each parameter. Notably, this can be used to assess which of the parameters is the most important for the classification. The radial parameters are the most prevalent, whereas both the circumferential parameters and the thickness at ED provide very little contribution to the first mode and therefore to the classification.

7.4.2 Cross-Validation on Training Set: Classifier Selection

All algorithms were tested with 10-fold cross validation on the dataset made of 200 patients. Fig. 7.6 summarizes the results of the different algorithms. Combining both sets of parameters improves the performance of most of the algorithms showing that these sets give different kind of information about the data. We also see that PLS regression, by preprocessing the data and orienting the modes of the input variables upto the best correlation with the pathology labels, improves the performance of all machine learning algorithms especially for Decision Tree and Nearest Neighbors. With more than 95% of correct labeling, SVM-SVC algorithm used on the PLS reduction with 5 modes is the method that performs the best. It is interesting to see that increasing the number of PLS modes further does not improve the classification. Our interpretation is that the subsequent modes of the PLS are not correlated to the classification and can therefore induce over-fitting of the data. We also tested the method with different cross-validation such as leave-one-out, 2-fold

		STRAIN			TRANSLATION			THICKNESS		SUM OVER ALL PARAMETERS	
		RADIAL	LONG	CIRC	RADIAL	LONG	CIRC	END DIASTOLE	END SYSTOLE		
AHA REGIONS	BASE	1	0.9%	0.3%	0.0%	2.1%	0.9%	0.5%	0.0%	0.2%	5%
		2	0.4%	0.9%	0.0%	2.3%	0.9%	0.1%	0.2%	0.2%	5%
		3	0.0%	0.6%	0.0%	2.0%	1.2%	0.0%	0.3%	0.2%	4%
		4	0.0%	0.0%	0.2%	1.5%	1.3%	0.0%	0.0%	0.0%	3%
		5	0.1%	0.1%	0.3%	1.2%	1.2%	0.2%	0.0%	0.2%	3%
		6	0.5%	0.0%	0.0%	1.5%	1.0%	0.5%	0.0%	0.2%	4%
	MIDDLE	7	1.9%	1.4%	0.2%	2.3%	0.8%	0.5%	0.0%	0.8%	8%
		8	1.3%	1.6%	0.2%	2.6%	0.9%	0.0%	0.0%	0.8%	7%
		9	0.9%	1.6%	0.1%	2.4%	1.1%	0.1%	0.0%	0.6%	7%
		10	0.8%	1.3%	0.0%	1.9%	1.2%	0.0%	0.0%	0.7%	6%
		11	0.9%	0.5%	0.0%	1.5%	1.2%	0.0%	0.0%	0.7%	5%
		12	1.5%	0.6%	0.0%	1.7%	1.0%	0.4%	0.0%	0.6%	6%
	APICAL	13	1.9%	1.3%	0.4%	2.5%	0.8%	0.1%	0.3%	1.6%	9%
		14	1.6%	1.2%	0.2%	2.2%	1.0%	0.1%	0.2%	1.2%	8%
		15	1.4%	0.7%	0.0%	1.6%	1.1%	0.0%	0.1%	1.1%	6%
		16	1.7%	0.7%	0.1%	1.8%	0.9%	0.4%	0.3%	1.1%	7%
	APEX	17	1.2%	1.1%	0.4%	1.5%	1.0%	0.0%	0.6%	1.1%	7%
SUM OVER ALL REGIONS		17.0%	13.8%	2.2%	32.7%	17.8%	3.0%	2.1%	11.4%		

Figure 7.5: Loadings of the first PLS mode showing the contribution of each of the parameters and each of the AHA zone. In green the most important parameters and in red the less important.

or 5-fold in order to see the robustness of the method with respect to the size of the training set and got similar performance.

7.5 Results and Validation on Testing Set

A prediction model (which will be referenced as the *baseline*) was introduced to provide a benchmark to assess the clinical benefit of the participating algorithms. A binary multiple logistic regression was applied to the 200 cases from the training set. The features selected to build the model were EDV, ESV, LVM and EF. They were chosen because of their wide clinical use and they are known to be significantly impacted by myocardial infarction. The baseline model was given by (see [Suinesiaputra 2017] for more details):

$$\ln\left(\frac{p(\mathbf{X})}{1-p(\mathbf{X})}\right) = \beta_0 + \beta_1 X_{\text{EDV}} + \beta_2 X_{\text{ESV}} + \beta_3 X_{\text{LVM}} + \beta_4 X_{\text{EF}} \quad (7.1)$$

where the intercept $\beta_0 = 12.35$ and the contributions of each variable were $\beta_1 = 0.11$, $\beta_2 = -0.09$, $\beta_3 = -0.03$ and $\beta_4 = -0.31$. The largest effect was given by LVM ($P < 0.001$), followed by EDV ($P < 0.05$).

Table 7.1 compares classification performances between the participating methods after the optimal cut-off value, except for one method (L2GF) where the cut-off value was defined by the participant. All methods achieved very good classification results with accuracy ranges from 0.83 to 0.98. Our method gets an accuracy of

		Decision Tree	Random Forest	Logistic Regression	Nearest Neighbors	SVM Linear	SVM svc
	PolyAffine (PA)	79%	88%	85%	81%	89%	86%
	Thickness (TH)	76%	84%	86%	89%	85%	86%
	PA + TH	79%	91%	89%	90%	90%	92%
	Norm. PA + TH	77%	88%	92%	85%	90%	94%
PCA	2 Modes	84%	87%	90%	89%	87%	89%
	5 Modes	84%	87%	90%	87%	88%	93%
	10 Modes	84%	89%	90%	87%	88%	93%
	All modes	84%	80%	93%	87%	92%	96%
PLS	2 Modes	91%	93%	94%	94%	93%	94%
	5 Modes	91%	96%	96%	97%	95%	97%
	10 Modes	91%	97%	96%	94%	95%	96%
	All modes	91%	94%	94%	94%	93%	94%

Figure 7.6: Cross-validation results (10-fold) of the classification with respect to different state-of-the-art machine learning algorithms and different sets of input data. Combination of algorithms and parameters that have the best performance are shown in green whereas the worst are shown in red.

96% on the test set which is in line with the cross-validated results on the training set. With respect to the other competitive methods, only two methods obtains a better accuracy (with the best at 98%). It is worth to note that our method achieves a 100% sensitivity meaning that we classified correctly all of the patients with myocardial infarctions. All the other methods classified at least 2 false negatives. In a clinical context, it is more problematic fail to diagnose a disease (False Negative) rather than to be too conservative and diagnose a problem when there is none (False Positive). Our method also performs significantly better than the baseline model ($p < 0.05$). It shows that simple clinical indexes are not enough to detect myocardial infarction and there is significant added value in looking at more complex remodeling as we do with our method.

Figure 7.7 (left) visually compares the participating methods with the Baseline model (shown as a black curve). The ROC curves of 6 methods are either lower or in line with the baseline curve. The curves of the other 5 methods (including ours) were generally above the Baseline curve, indicating some benefits of shape information for predicting MI. At least seven methods and also the Baseline model failed to predict two cases, which were both MI shapes. These two difficult cases are shown in Fig. 7.7 (right). Both cases have cardiac volume and function within the normal range, but detailed geometrical shape visualization shows reduced contraction in local area (pointed by green arrows). Four methods correctly identified one of the case but only two methods (including ours) classified both cases correctly.

7.6 Conclusion

In this chapter, we evaluated the contribution of prior reduction of dimensionality to the classification of high-dimensional motion data. One of the assets of our work is an innovative methodology to project a motion on a reduced number of polyaffine parameters. We apply the methodology to classify a population and detect an

Table 7.1: Classification performance results. The positive values were MI shapes, while the negative values were AV shapes. NA denotes not available value. [†] denotes a method that the AUC is statistically higher than the Baseline at $p < 0.05$. Our proposed method is labeled as PT.

	AUC	cutoff	spec	sens	acc	FP	FN
GMPT [†]	0.994	0.76	0.96	0.95	0.95	4	5
SSM-PLS [†]	0.996	0.56	0.99	0.97	0.98	1	3
IC-ShapeMotion [†]	0.989	0.42	0.96	0.98	0.97	4	2
MS	0.901	0.53	0.83	0.83	0.83	17	17
ASMSVM	0.977	0.85	0.99	0.92	0.95	1	8
FM	0.931	0.58	0.90	0.84	0.87	10	16
JCCA	0.939	0.12	0.93	0.91	0.92	7	9
RF	0.977	0.55	0.93	0.91	0.92	7	9
PT [†]	0.991	0.29	0.92	1.00	0.96	8	0
HeAT-RDF	0.976	0.72	1.00	0.86	0.93	0	14
L2GF	NA	NA	0.89	0.97	0.93	11	3
Baseline	0.970	0.74	0.99	0.87	0.93	1	13

infarct based on the segmentations at end-systole and end-diastole. Following the first dimensionality reduction given by the polyaffine parameters, we use traditional statistical reductions on our sets of parameters with PCA and PLS. Using 10-fold cross validation, we show that the resulting parameters have good predictive power with more than 95 % correct classification on 200 infarcted/control cases. We are also able to quantify the importance of each of the parameters in the classification. Notably, this provides insights into what is the main impact of an infarct both in terms of motion and shape.

Our method performs state-of-the-art results with respect to classification accuracy (AUC, sensitivity and accuracy) inline with the best competitive methods of the challenge. In particular, it classified correctly all the patients with myocardial infarction with no false negative result. With respect to the other method, a major quality of our method is the interpretability of the parameters in term of local strain and displacement on the AHA regions. These regions are used clinically which means our classification can be justified using quantitative parameters that are known to clinicians. This is important in term of clinical application, as clinicians will be reluctant to rely on methods without being able to understand the result.

A possible clinical application of our method would be to score patient against our PLS modes. This would give us an index quantifying the severity of the infarct with respect to the reference population seen. With 0 meaning average healthy and 1 average infarcted. Severity of infarct could be quantified by their score, with values higher than 1 meaning severe infarct. This information could be used in clinical




Figure 7.7: (Left): ROC curves of the participating methods. Our proposed method based on Polyaffine Transformations (PT) in green achieves accuracy among the best methods. (Right): The two most difficult cases which are both infarcted but with large ejection fraction. The green arrows point to less contractile area of the myocardium, indicating infarction. Our method gave correct classification for both cases.

practice to adjust the level of treatment with respect to the severity of the infarct. Finally, a possible extension of the method presented here is the detection of the localization of an infarct. Because the parameters we extract are regional parameters, one could also try to learn which AHA region is infarcted, provided that this information is available in the training data. Our regional polyaffine decomposition provides an efficient framework to study locally the motion and detect local abnormalities of the deformation.

Cardiac Motion Signature Using Barycentric Subspaces

Contents

8.1 Chapter Overview	117
8.2 Cardiac Motion Signature from Low-Dimensional Representation	118
8.2.1 Data	118
8.2.2 Methods	118
8.2.3 Optimal References Frames and Barycentric Curves	119
8.2.4 Group-Wise Analysis of Differences	120
8.3 Reconstruction of Cardiac Sequences	121
8.3.1 Qualitative results	122
8.3.2 Quantitative results	123
8.4 Conclusion	124

Similarly to what has been done in the previous chapter, we tackle here the problem of diagnosis using a low-dimensional representation of the cardiac motion. We build on the Barycentric Subspace framework defined in Chapter 5 to define a signature of the cardiac motion. The signature is made up of both the optimal references frames and the coefficients of the projection on the Barycentric Subspace. This low-dimensional representation is applied to the diagnosis of Tetralogy of Fallot (ToF).

8.1 Chapter Overview

We show that this multiple references approach leads to substantial improvement over the traditional single reference methodology in two different problems related to cardiac motion. We build a low-dimensional representation of the cardiac motion signature which actually separates two different populations of healthy subjects and patients with Tetralogy of Fallot, showing that our method extracts features allowing the efficient analysis of the cardiac motion. Then, we reconstruct the sequence of images from our low-dimensional representation and show that our method has better results both quantitatively and qualitatively than traditional single reference methods based on tangent PCA.

The main contributions of the work presented here are:

- The group-wise analysis of the features extracted from the projection and its application in the context of the study of Tetralogy of Fallot.
- The reconstruction of a sequence using three reference frames and the coefficients, opening the way to extensions for synthetic sequences computation.

The work presented in this chapter is a journal extension of the work published in the MICCAI conference [Rohé 2016b] by going deeper into the analysis of the barycentric subspace.

8.2 Cardiac Motion Signature from Low-Dimensional Representation

In this section, we compute the low-dimension representation of a set of cardiac sequences within Barycentric Subspaces. We show that the features extracted from the representation define relevant cardiac motion signature capturing the main phases of a cardiac cycle. In the last part of this section, we compare these features between two populations and we show that they lead to an efficient discrimination.

8.2.1 Data

We applied the previously defined methodology to compare the cardiac motion signature using two different populations. The acquisition consists of a cine sequence in the short-axis view of steady-state free precession magnetic resonance images of two different groups of subjects:

- The first group consists of 15 adult control subjects (12, males, 3 females, mean age \pm SD = 28 ± 5) from the STACOM 2011 cardiac motion tracking challenge dataset [Tobon-Gomez 2013] which is openly available. The temporal resolution is 30 frames with a variable spacial resolution ranging from $1.15mm \times 1.15mm$ to $1.25mm \times 1.25mm$ (in-plane) and slice thickness of $8mm$.
- The second group is made of 10 Tetralogy of Fallot (ToF) patients (5 males, 5 females, mean age \pm SD = 21 ± 7) [McLeod 2015b]. TOF is a congenital heart defect that is present at birth. These patients all had a full repair early in infancy, resulting in the destruction of the pulmonary valves. The temporal resolution ranges from 15 to 19 frames with a spacial resolution of $1.21mm \times 1.21mm$ to $1.36mm \times 1.36mm$ (in-plane) and slice thickness of $8mm$.

8.2.2 Methods

The stationary velocity fields $v_{i,j}$ parametrizing the deformations between each frames of the sequences were computed using registrations with the open-source




Figure 8.1: Choice of the optimal references within Barycentric Subspaces of increasing dimensionality. (Left): 1D-Subspace with 2 reference images. (Right): 2D-Subspace with 3 reference images.

algorithm *LCC-Log Demons* [Lorenzi 2013a]. In order to improve the registration accuracy, all the images were resampled as a preprocessing step to have an isotropic resolution in the X, Y, Z spatial directions. We apply the methodology described in section 5.4.2 and project each of the T frames of the cardiac motion to a Barycentric Subspace of dimension k spanned by $k + 1$ references. We build the optimal Barycentric Subspace by choosing the references realizing the minimum of the projection energy as described in Algorithm 5. Figure 5.3 (left) shows the error of the approximation of the subspace with 2, 3 and 4 references over one sequence for one patient. We see a lower error for frames around the references (for which the projection and the error is null) and a larger error for frames further away from the references. Figure 5.3 (right) shows the error averaged over all the frames and how it varies when the dimension of the subspace is increased. As with traditional dimension reduction methods, the error decreases rapidly for the first dimensions which explain most of the variability of the data. For the remaining of this section, we will focus more specifically on the 1D/2D subspaces as they give a good trade-off between the complexity of the subspace and the accuracy of the low-dimensional representation of the motion.

8.2.3 Optimal References Frames and Barycentric Curves

The frame chosen as optimal references are shown in Figure 8.1 for the Barycentric Subspaces with 2 reference (left) and 3 references (right). In order to describe them, we are going to number them as $(\#1, \#2)$ and $(\#1, \#2, \#3)$ although we ask the reader to keep in mind that there is no specific order in the definition of the references. For the 2 references (1-D) case (left figure) we see that reference $\#1$ (blue) is chosen as either one of the first frame of the sequence, or one of the last. This value is very consistent with the end-diastolic frame, which is one of the two extreme point of the cycle where the heart is fully relaxed. The reference $\#2$ (red) is always a frame close to a third of the time of the full cardiac cycle, which corresponds closely to the end-systolic time, where the heart is fully contracted.




Figure 8.2: Curves of the lambda coefficients within Barycentric Subspaces of increasing dimensionality. Both the mean (plain lines) over the complete population of 25 patients and the $+\sigma$ and $-\sigma$ curves (dotted lines) are shown. (Left): 1D-Subspace with 2 reference images. (Center): 2D-Subspace with 3 reference images. (Right): 3D-Subspace with 4 reference images.

Looking now at the subspace built with 3 references (right), we see that reference #2 (red) corresponding to the ES is still quite consistent, with little changes compared to the 1D case (for all the patients the 2nd reference is either the same for both the 1D and 2D case or differs by just one frame, a result consistent with the variability seen by other methods for the detection of ES [Kong 2016]). This confirms that the ES frame is well captured by one reference. For the two other references (#1 in blue and #3 in green), one is at the beginning and can be recognized as one of the ED frame whereas the other one is at the end of the sequence and can be related to the frame corresponding to the diastasis even though it is less consistently defined.

The mean barycentric coefficients together with the variation at $+/- 1SD$ are shown in Figure 8.2 for the 1D (left) and 2D (right) cases. For the 1D case, the pattern of the coefficients closely relates to the volume curves, with the λ_1 peaking on average at a frame close to the ED and λ_2 peaking at the ES frame. When adding a 3rd reference and curve (right), the second curve corresponding to the ES reference only marginally changes whereas the coefficient which corresponded to the ED frame is now divided between the 1st and 3rd coefficients at the beginning and end of the cycle.

8.2.4 Group-Wise Analysis of Differences

Finally, we compare the parameters of each population (healthy vs. ToF) in the case of a 2-D Barycentric Subspace. In the left plot of Figure 8.3, one can see the difference between the optimal reference frames chosen for each of the population. For the first reference (corresponding to the end-diastole), the frame chosen for each population is quite consistent, with the ToF patients just having a slightly higher value. For the 2nd reference, corresponding to the end-systolic frame, significant




Figure 8.3: (Left): comparison of the optimal references estimated for the healthy controls (blue) and the ToF patients (ref). Significant differences can be seen for the 2nd corresponding to a larger systolic duration. The curves represent two typical patterns of the volume curve for each population. (Right): 2-D plot representing the projection of the two population on the first 2 modes of the PLS. Any linear classifier would be able to separate both populations as two clear clusters appear.

differences can be seen between the two populations: the time value of this reference for the ToF population is way higher than the one for the healthy controls. This is a sign that the ToF patients shows higher systolic contraction time, a fact that is confirmed by clinical experience. Reference #3 shows also differences between the two population, with higher intra-population variability.

To investigate further the differences between the two populations, we perform a *Partial Least Square* (PLS) [Rosipal 2006] decomposition with the classification (Healthy vs. ToF) as the response variable. PLS decomposes the data into multiple modes which both explains the variability and the correlation with the response. In our case, the features used were both the 3 optimal reference frames (normalized by the total number of frames) and the 30 barycentric coefficients corresponding to the curves in the 2-D space. For normalization, we scale the barycentric coefficients so that their importance in the decomposition is similar to the choice of the optimal frames. The projection of the data on a 2-D space can be seen in Figure 8.3-right. Two clear clusters can be seen, showing that the PLS manages to discriminates the two populations based on the features from the barycentric subspace projection. All these results show that this low-dimensional signature of the motion is encoding relevant features of the cardiac motion.

8.3 Reconstruction of Cardiac Sequences

In this section, we evaluate the reconstruction of the sequence using our low-dimensional representation. From one sequence of 30 cardiac images we compute the Barycentric coefficients and reference frames using the previously defined methodol-




Figure 8.4: Reconstructed sequences using the barycentric representation compared with other methods. (Top row): initial true cardiac sequence. (Second row): reconstruction using PCA decomposition with the ED frame as reference. (Third row): reconstruction using PCA with the mean image as reference. (Bottom row): proposed approach using Barycentric compact representation of the motion. Each column corresponds to one frame of the sequence which are from left to right: 0 (ED), 5, 10 (ES), 15, and 25.

ogy. Then we reconstruct the sequence simply using these coefficients (3 coefficients for each frame) which represent the position of each frame in the 2-D Barycentric subspace - and the 3 reference images. We compare our method with two single reference approaches. The first one, we start from the ED image and we perform the PCA covariance analysis on the tangent vectors of the deformations. We build a 2-D subspace by taking the first two modes of the PCA decomposition. The cardiac motion is reconstructed by applying this low-dimensional representation of the deformation to the ED image. The second one, the same methodology is applied but with the analysis done at the mean image computed using all the images of the sequence. To sum up, both these methods are encoded by the scores of the 2 modes, 1 scalar image for the reference (the ED frame for the first method or the mean for the second) and 2 vector images representing the coefficients of the modes.

8.3.1 Qualitative results

A qualitative comparison of the different methods can be seen in Figure 8.4. PCA with the ED image as reference (2nd row) performs a good reconstruction for the early frame of the sequence but fails to reconstruct well the large deformation of




Figure 8.5: (Top): comparison of the 2-D representation of single reference PCA versus multi-reference Barycentric. (Bottom): Error of the reconstruction of the sequence for the different methods.

the ES (3rd column) compared to the initial sequence (1st row): the fact that there is only one reference image limits its ability to update the change of appearance of the images during the cardiac cycle. PCA at the mean image (3rd row) performs better to recover these frames, but the overall appearance of the image is blurred because the initial texture is not defined using a single frame but an average of images. Finally, our proposed approach using multiple references (4th row) has a similar appearance to the initial sequence for all the frames of the cycle.

8.3.2 Quantitative results

A quantitative comparison of the performance of the three methods is shown in Figure 8.5. Top figures show a comparison of the 2-D PCA (with ED image as reference) and Barycentric low-dimensional representation. The chosen reference frames for the Barycentric methods are reference #1 is frame 0, reference #2 is frame 10 and reference #3 is frame 17 and represent the different phases of the cycle. For the PCA method, there is only one reference which only accounts correctly for the part of the motion close to this single reference. We can also see that the PCA curve never gets close to the mean point (coordinates 0,0) which is therefore a poor descriptor of the motion. Finally, the mean voxel intensity error with respect to the initial sequence, Figure 8.5 (bottom), shows a better performance of the multi-reference Barycentric approach than for the traditional PCA (reduction of the error

on average of approx 30%).

8.4 Conclusion

In this chapter, we have used the parameters of the Barycentric Subspace framework defined in Chapter 5 to build a Cardiac Motion Signature encoding relevant features of the motion. We have presented two possible applications where this framework show clear advantages over the traditional method. The first application is the computation of a low-dimensional representation of the cardiac motion from the projection on to the subspace. This representation captures relevant differences in the cardiac motion and lead to an efficient discrimination between two populations of healthy and diseased patients. Then, we have compared the reconstruction of the sequence using our multi-reference low-dimensional representation with traditional representation using one reference: our reconstruction performs better both qualitatively and quantitatively. This confirms that our low-dimensional representation encodes important characteristics of the motion.

Longitudinal Analysis of the Cardiac Motion

Contents

9.1	Chapter Overview	125
9.2	Motivations	126
9.3	Cardiac Motion Features Extraction	127
9.4	Cardiomyopathies	129
9.4.1	Data	129
9.4.2	Mean Motion Model	130
9.4.3	Longitudinal Motion Analysis of CMP patients	131
9.5	Obesity	132
9.5.1	Data: Obesity Patients (CVD)	132
9.5.2	Mean Motion Model Parametrized by BMI	133
9.5.3	Longitudinal Motion Analysis of CVD patients	134
9.6	Conclusion	136

In Chapter 7 and Chapter 8 we analyzed the cardiac motion of a population at a specific time in order to build group-wise statistics and automatic diagnosis. In this chapter, we tackle an additional challenge by looking at the longitudinal evolution of the motion for both patients with cardiomyopathies and subjects with different levels of obesity. We analyze the cardiac motion acquired at multiple time points in order to understand its evolution using a low-dimensional representation.

9.1 Chapter Overview

We first present the motivations behind the longitudinal analysis with state-of-the-art examples of such studies applied to brain analysis. The challenges to apply such analysis to the heart motion are discussed with the problem of data acquisition and motion representation. Then, we sum up the methodology for extracting the polyaffine parameters which will serve as features describing the cardiac motion. We present methods to align spatially and temporally these parameters so that the features can be compared for different patients. Finally, two different studies are performed, one focusing on each medical condition: cardiomyopathies and obesity.

In each of these studies, we compute a mean model using cross-sectional analysis before performing a longitudinal analysis using a low-dimensional representation of the motion. We show preliminary results whose significance still need to be confirmed using statistical validation on larger dataset.

9.2 Motivations

Cardiovascular disease (CVD) is a worldwide issue, being one of the leading cause of death. In particular, in Chapter 1, we have described two medical conditions: cardiomyopathies and obesity. In paediatric cardiovascular disease, predicting the reaction of patients to treatments, which treatments are the best to use, and when to treat can be difficult to define due to small patient numbers and limited outcome data. Cardiomyopathy is a group of diseases that affect the heart muscle. Those affected are at an increased risk of sudden cardiac death. Obesity is a medical condition which is one of the primary risk factors for cardiovascular disease. This risk factor originates early in life and increases depending on the level of the condition (classification of Obesity can be found in Chapter 1).

Understanding the evolution and the impact of these medical conditions on the cardiac motion is therefore very important in a clinical perspective. It can help clinicians to detect which patients are at risk of heart failure and possibly require heart transplant or mechanical support. The goal is to give a predictive view of what will be the cardiac motion in a future time-point under the assumption that the evolution stays the same. It can also provide further insight into the understanding of a particular disease and its impact on the cardiac motion over time. Longitudinal studies have been a trending topic in the brain analysis. One can cite as example the studies of [Giedd 1999, Evans 2006, Lebel 2011, Hadj-Hamou 2016a] which look at the brain development during childhood or adolescence. Another hot topic of longitudinal studies for the brain is the analysis of the evolution of Alzheimer disease over time [Coben 1985, Lorenzi 2010, Miller 2016, Hadj-Hamou 2016b]. While there are many longitudinal studies focusing on the brain, this kind of study has yet to be applied routinely for cardiac motion analysis. In [Mcleod 2013a], a preliminary approach for a $5D$ cardiac model is proposed as a first step towards the longitudinal analysis of cardiac motion. Defining the longitudinal evolution of something that is already temporal makes the topic complex. Furthermore, to build such a model, one needs to gather multiple cardiac motions of the same patient over a large period of time (to detect measurable differences). Waiting for years to acquire the data makes it difficult to build large enough database, especially when working with children who tend to have a large opt-out rate of clinical studies.

To study the evolution of cardiac motion and get insight into the $3D$ dynamics of the myocardium, one has to track it quantitatively. Numerous methods for tracking cardiac tissue were proposed and already presented in this manuscript. Among these methods we chose to use the polyaffine framework presented in Chapter 6. With respect to classical non-rigid registration algorithm parametrized with dense

displacements field, this framework has the advantage to parametrize the cardiac motion with a relatively low number of parameters. For comparative population-wise studies, it is very important to have a parametrization with a small number of parameters if one wants to get robust group-wise analysis. The Polyaffine model provides such a parametrization by combining regional transformation in a smooth, global manner and has previously been applied successfully in Chapter 7 for the classification of infarct patients but also in many other clinical studies [Zhang 2007, Hansen 2012, McLeod 2012, McLeod 2013c, McLeod 2015a]. Another advantage of this parametrization is that the regional decomposition with the AHA regions is understood clinically and is common to all patients. This makes the comparison between different subjects possible without relying on spatial normalization such as parallel transport [Qiu 2008], which would have been necessary if we were working with displacement fields.

9.3 Cardiac Motion Features Extraction

Polyaffine representation. The polyaffine framework was described in Chapter 6. We only give a very quick summary here. We define $K = 25$ regions using the standard American Heart Association (AHA) 17 regions for the left ventricle and 8 additional regions dividing the right ventricle in a similar way. An affine transformation is defined on each region and these regions are associated with smooth weights. These weights are used to fuse together the locally affine transformations into a global one in a smooth way. The polyaffine transformation is the weighted sum of these locally affine transformations with the weights defined at each point:

$$v_{poly}(x) = \sum_{k=1}^K \omega_k(x) M_k \tilde{x},$$

and the displacement field is defined as the exponential of the SVF: $u(x) = \exp(v_{poly})$.

Tracking motion from MRI sequences. Depending on the source of the data, two different methods were used to track the motion of the myocardium from MRI sequences. For the CVD and CMP patients, a pipeline combining segmentation and tracking is used to compute a mesh of the myocardium at the first frame and the tracking of the mesh along the sequence using the methods described in [Georgescu 2005, Ionasec 2010, Wang 2013]. The process involves three steps: landmark detection, segmentation and tracking. After each step an expert check the results and refine, if needed. Then, a representation with polyaffine parameters that best described the motion is computed. For the healthy subjects (HTY), the segmentation was done using the multi-atlas algorithm described in Chapter 4 and the tracking with the log-demons polyaffine algorithm [McLeod 2012]. In both case, we obtain polyaffine parameters. These parameters serve as features representing the cardiac motion.




Figure 9.1: Volume curves for 20 CMP patients before (left) and after (right) temporal realignment. A mean motion computed using the left representation would underestimate the mean ejection fraction as the peak contraction does not happen at the same frames for all the patients.

Temporal resampling. As seen in the description of the populations, there is a wide variability in the number of frames of each acquisition of a MRI sequence. Normally, one acquisition corresponds to one cardiac cycle from the R peak to the next R peak. It is divided into a given number of frames (typically from 10 to 40). The temporal discretization corresponding to the number of frames differs from one subject to the other and we need to perform temporal resampling. In order to study all the patients together, the frames of the sequence are resampled to have a common number of 30 frames for each acquisition by using linear interpolation.

Temporal alignment. Cine MR images sequences are gated with electrocardiogram (ECG) signals. Frames are supposed to be acquired at a given time after the peak of the R wave which initiates the sequence. Nevertheless, some frames are sometimes missing to get the whole cardiac cycle. Also the first frame does not always corresponds to the end-diastolic time. Finally, the ECG signal may sometimes be weak so that a precise detection of the R peak is not possible. For these reasons we perform a temporal alignment as a pre-processing step in order to have comparable parameters. To do so, we use the LV volume curves and we detect the ES frame (which is well defined as the frame of minimum volume). All the parameters are aligned using a constant temporal shift. Examples of the volume curves before and after the temporal alignment are shown in Figure 9.1.

Parameter normalization with scale. One of the challenge of the populations we study is the large range of the ages as we study children and adolescents (from 3 months to more than 20 years). This causes an important variability in the size of the hearts as some patients have not finished their growth while others are already adults (see the histogram of volumes in Fig. 9.2). In order to make the polyaffine

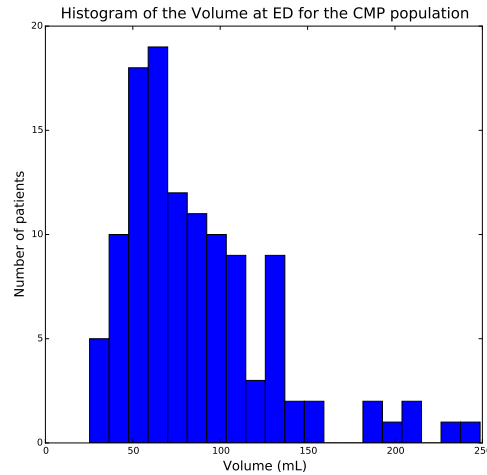


Figure 9.2: Histogram of the left-ventricular ED volumes for the patients of the CMP populations. Large variability can be seen, therefore we need to normalize the parameters in order to study all the patients together.

parameters comparable between different patients of this population, one has to normalize them in order to remove the scale effect. The 9 parameters from the linear part of the affine matrix are not independent on the scale and they do not have any unit. However, the 3 translation parameters are dependent on the scale as they represent displacements in mm . A standard choice is to use the Body Surface Area (BSA) index. But, we have a direct measure of the size of the heart with the volume of the LV extracted from the segmentation. Therefore, we chose to normalize the translation parameters with respect to the cubic root of the volume of the LV at ED giving us a set of parameters without unit that are comparable between patients of different size. Finally for each patient i , the set of (3×4) affine matrices $M_{i,t,k}$ for $t \leq T = 30$ and $k \leq K = 25$ is flattened to a feature vector p_i of size $F = T \times K \times 12$ which serves as reduced-order representation of the cardiac motion.

9.4 Cardiomyopathies

9.4.1 Data

This section focuses on the analysis of cardiomyopathies in children and adolescents. We use two different populations. The first population is a data-set of healthy controls (HTY), described below, in order to establish normal motion patterns to compare the different population with. The other data-set was acquired in the context of the MD-Paedigree project and consists of patients suffering from different forms of cardiomyopathies (CMP). This medical condition was described in Chapter 1 and we describe here the images acquired and used in this study.




Figure 9.3: (Left): normalized volumes curves of the mean models applied to a template geometry. (Right): average strain values on the left-ventricle for the mean model. Differences between the mean model for the CMP patients (red) and the healthy subjects (blue) can be seen.

Healthy Subjects (HTY). This population consists of 15 controls adults subjects (12, males, 3 females, mean age \pm SD = 28 ± 5) from the STACOM 2011 cardiac motion tracking challenge dataset [Tobon-Gomez 2013] which is openly available. The temporal resolution is 30 frames with a variable spacial resolution ranging from $1.15mm \times 1.15mm$ to $1.25mm \times 1.25mm$ (in-plane) and slice thickness of $8mm$.

Cardiomyopathies Patients (CMP). The second population is made of 84 children and adolescent with cardiomyopathies with ages ranging from 3 months to 19 years with an average of 8.15 and standard deviation of 6.28. These patients were enrolled in 3 different clinical centers (OPBG at Rome, DHZB at Berlin and UCL at London). Each of these patient had a MRI acquisition at baseline and 8 of them had a follow-up acquisition after 1 year as well. The procedure for the MRI acquisition and the scanners were different at each center. Therefore, there is a huge variability in the temporal resolution (from 10 frames to 50 and 25.7 on average), in the spatial resolution (X, Y spacing ranging from $0.85mm$ to $3.125mm$ and Z spacing from $5mm$ to $10mm$) and in the size of the images (from $128 \times 128 \times 7$ to $432 \times 432 \times 17$).

9.4.2 Mean Motion Model

After having pre-processed the parameters, all the polyaffine representations of the motion are aligned spatially and temporally. We can now compute the mean motion model. For the CMP and HTY populations a simple arithmetic mean is computed by taking the average of all the feature vectors p_i for i ranging over the population of N subjects:




Figure 9.4: (Left): myocardium at ES for both the healthy mesh (blue) and the CMP mesh (red). Contraction is more important for the healthy motion. (Right): radial strain maps at ES for each of the two mean model. The healthy one (left) shows higher radial strain values than the CMP one (right).

$$\bar{p} = \frac{1}{N} \sum_{i=1}^N p_i.$$

Given the estimated mean parameters, the motion can be simulated on a new geometry. Because the polyaffine parameters we computed are defined using local basis on each region, we do not have to perform any spatial alignment even though each subject does not lie on the same coordinate space. We can apply the set of polyaffine parameters to a template geometry defined by the segmentation at end-diastole S_0 .

Figure 9.3 shows classical clinical metrics computed on the mean model for each population. For the volume curve (left), one can see a larger ejection fraction for the mean model computed on the healthy population. This is in accordance with clinical knowledge and CMP patients tend to have a heart function less efficient than healthy patients. Strain values (right) on each of the three directions show also significant differences between the two populations. Especially for the radial strain curve, one can see that higher values for the healthy motion at ES (37% versus 26%). Finally, Figure 9.4 shows the meshes at ES for each of the two mean model and the local differences of the radial strain.

9.4.3 Longitudinal Motion Analysis of CMP patients

In order to study the longitudinal evolution of the motion despite the limited number of data we have (8), we choose to use a low-dimensional space computed on the whole dataset of CMP patients. This low-dimensional space represents the variability of the motion with respect to the pathology. To build this space, we use Partial Least Square regression that was already introduced in Chapter 6. We apply the PLS on the HTY and CMP populations with the dependent variable chosen as 0 for healthy and 1 for CMP. Therefore we get modes corresponding to the main variations of the motion of the CMP population with respect to the healthy population. Figure 9.5 left shows a representation in $2D$ of the two populations in this low-dimensional




Figure 9.5: (Left): 2D plot of the scores on the first 2 PLS modes for both the HTY (blue) and CMP (red) populations. (Right): Longitudinal evolution for 8 CMP patients on the 2D-space made of the two first modes of the PLS.

space. One can see, that this representation already performs a good separation of the two populations.

Then, we use this low-dimensional representation to analyze the evolution of the 8 patients for which we have a follow-up MRI acquisition. Figure 9.5 right shows the plot of the baseline and follow-up acquisition in this space with the corresponding evolution. An evolution toward the upper-right corner (corresponding to the CMP space) can be interpreted as a worsening of the condition of the patient whereas an evolution toward the bottom-left corner (corresponding to the HTY space) corresponds to an improvement of the cardiac motion.

9.5 Obesity

9.5.1 Data: Obesity Patients (CVD)

In this section, we study a population with various level of obesity which was part of the clinical protocol of the MD-Paedigree project. The population is made of 92 patients with ages varying from 2 years to 20 years with an average of 16.2 and standard deviation 4.4 years. Similarly to the Cardiomyopathies patients, these patients were enrolled in the 3 clinical centers: OPBG, DHZB and UCL. Most of these patients present some level of obesity, although some controls subjects were included in the enrollment of patients at the UCL clinical center. The BMI (index measuring the level of obesity of a patient as described in Chapter 1) of the population ranges from 15 to 50 with an average of 28.7 and standard deviation of 7.12. We recall that a patient is classified as pre-obese if its BMI lies between 25 and 30 and is considered as obese if it is greater than 30. It should be noted that this population contains some control subjects which are not overweight (BMI




Figure 9.6: (Left): normalized volumes curves of the mean models applied to a template geometry. (Right): average strain values on the left-ventricle for the mean model. Differences between the mean model for $BMI = 22.5$ corresponding to a normal range (blue) and for $BMI = 45$ corresponding to obese class III (red) can be seen.

≤ 25). Therefore, we don't need to add any control population in this study.

9.5.2 Mean Motion Model Parametrized by BMI

For the CVD population, we compute a mean motion model parametrized by the BMI. Indeed, as previously described, obesity is a medical condition which can be evaluated with the BMI as a clinical factor. Therefore, rather than computing a mean motion for the whole population, we compute a mean motion dependent on this factor:

$$\bar{p}_{BMI} = \alpha + \beta \times BMI.$$

Using ordinary least squares, we compute the parameters of the mean motion model dependent on the BMI. Then, we apply it to a common geometry for two levels of the factor. A level which corresponds to an average BMI for a healthy subject ($BMI = 22.5$) and a level which corresponds to an obesity of class III ($BMI = 45$) in order to compare the differences between both populations.

Figure 9.6 (left) shows the volume curves for the two levels of BMI. One can see very few differences, showing that BMI is not impacting significantly the volume curve and the ejection fraction. On Figure 9.6 (right) we compare the strain values over time on the three principal directions. One can see differences of the strain for the radial and longitudinal directions. The radial and longitudinal strains are higher for the healthy population than for the obese one. The significance of this finding should be confirmed with statistical testing.




Figure 9.7: Curves computed with the longitudinal generative model for 30 patients. Curves are colored with the BMI value of the patient (blue are low-BMI/healthy subjects and red are high BMI/obese patients). The curves model an exponential decay of the impact of the experiment on the cardiac motion of a patient.

9.5.3 Longitudinal Motion Analysis of CVD patients

Clinical experiment protocol. A subset of the CVD population made of 77 patients (containing healthy, overweight and obese patients) took part in a clinical study whose aim was to assess the cardiovascular response after the ingestion of a high-energy (1635 kcal), high-fat (142g) meal. MRI short-axis were acquired before the ingestion of the meal and after 20, 40, 60, 90, 120 and 240 minutes, for a total of up to $T_2 = 7$ longitudinal acquisitions. The clinical protocol of the study is described in detail in [Hauser 2016]. We extract the representations of the motion with polyaffine parameters for each of these timepoints when available. For a patient i at timepoint $0 \leq t \leq T_2$, we use the notation $p_{i,t}$ for the flatten vector of polyaffine parameters (which is our vector of features characterizing each motion) of size $F = T_1 \times K \times 12$ and F is the number of features. We emphasize the fact that two different temporal dimensions are defined: the number of longitudinal acquisitions during the experiment $T_2 = 7$ and the number of frames on which each cycle is discretized $T_1 = 30$.

We define a generative model for the longitudinal evolution of the motion. The underlying assumption about the impact of the experiment on the motion is that it is important on the beginning of the experiment after the ingestion of the meal and then decays as time goes by. In [Jiao 2014], a pharmacokinetic model is proposed and incorporated into the registration process in order to allow for a groupwise registration of the temporal time frames. We take a similar approach to model the temporal dependency of our feature vectors using the following model with an exponential decay:




Figure 9.8: (Left): 2D scatter-plot of the α and t_{max} parameters for all the patients. The colormap corresponds to the BMI factor of each patient (from dark blue $BMI = 16$ to dark red $BMI = 50$). (Right): 2D plot of the scores of the first two modes of the pls on the feature vectors q_i with respect to the BMI. These PLS scores separates the population with respect to the value of the BMI (high scores corresponding to high BMI).

$$\hat{p}_{i,t} = p_{i,0} + \frac{t}{t_{max}} \exp\left(-\frac{t}{t_{max}}\right) q_i,$$

q_i is a vector of size F which encodes the impact of the experiment on the motion of the patient. The parameter t_{max} accounts the temporal dependency. It corresponds to a reference time for the decaying impact of the experiment on the motion to get back at the rest state. This parameter and the vector of features q_i are estimated for each patient using ordinary least squares to minimize the distance of the model with respect to the measurements $p_{i,t}$. We define $\alpha = \|q_i\|$ the amplitude of the change on the feature vectors. Example of the curves $f(t) = \alpha t/t_{max} \exp(-t/t_{max})$ are given in Figure 9.7 for α and t_{max} parameters from a set of patients.

We show the results of the optimization of the parameters α and t_{max} for our population in Figure 9.8 (left) colored by the value of the BMI. The patients for which the experience has the greatest impact on the motion are plotted on the upper-right part of the graph. This corresponds to patients for which both the amplitude α and relaxation time t_{max} are high. There does not seem to be an important correlation with the BMI, as patients with high BMI can have either small reaction (bottom-left) or high reaction (upper-right). It would be interesting to test the correlation with other factors (e.g. genetic) as they may reveal a different digestive process at work.

Then we look at the variation of the feature vector factor q_i with respect to the level of BMI. To do so, we perform a PLS decomposition whose dependant variable is the BMI and the predictor variables are the feature vector q_i . The scores of the two first modes of the PLS are shown in a 2D plot in Figure 9.8 (right). We see

that this $2D$ subspace separates the data with respect to the level of BMI, with the high-BMI patients on the upper-right corner of the plot and the low-BMI controls on the bottom-left.

Finally, we use the generative model we defined to compute synthetically the evolution for two extreme points. We note q_{m1} and q_{m2} the two first modes of the PLS for the vectors q and similarly t_{m1} and t_{m2} the two first modes for the PLS decomposition of the parameters t_{max} . Two sets of feature vectors corresponding to a extreme healthy and extreme obese longitudinal evolution are computed using these parameters weighting by 3σ . Thus, the parameters for the healthy ($-$) and obese ($+$) simulations are:

$$q_{\pm} = q_{mean} \pm 3q_{m1} \pm 3q_{m2}$$

$$t_{\pm} = t_{mean} \pm 3t_{m1} \pm 3t_{m2}$$

$$\hat{p}_{\pm,t} = p_{0,mean} + \frac{t}{t_{\pm}} \exp\left(-\frac{t}{t_{\pm}}\right) q_{\pm}.$$

These polyaffine parameters are applied to a common geometry for each of the 7 timepoints of the study. In Figure 9.9, we show the volume for 3 of these longitudinal timepoints ($T = 0, T = 90, T = 240$) for each of the two archetypal points (left corresponding to the healthy longitudinal motion, and right to the obese longitudinal motion). One can see significant difference in the evolution of the cardiac motion. For the healthy subject, the ejection fraction increases during the experiment at $T = 90$ whereas there is a decrease of the ejection fraction for the obese patient. At the end of the experiment $T = 240$ both curves have almost returned to the original level.

9.6 Conclusion

Throughout this thesis, a major theme was population-based modeling of different phenomena impacting cardiac motion. One of the main challenge in this kind of study is the difficulty to analyze and model a cardiac motion when it is parametrized with a large number of parameters. This is what motivated the development of reduced-order models, in particular the polyaffine representation of the cardiac motion. We show in this chapter how this representation can be used to perform group-wise analysis of the parameters and to construct a mean motion model by averaging spatio-temporally aligned transformation parameters.

Naturally, this problem becomes even more difficult to tackle when an additional dimension, the longitudinal evolution, is added to the analysis. We face the challenge to model coupled temporal phenomena: the temporality of the cardiac cycle and the longitudinal dimension. To overcome this problem, we proposed in this chapter two different methods. For the CMP population, due to the small size of the population for which we have access to longitudinal data, we chose to study it using a low-dimensional subspace built on the two first modes of a PLS defined on a larger



Figure 9.9: Volume curves for the longitudinal ($T = 0, T = 90$ and $T = 240$) generative model at -3σ on the healthy direction (left) and at $+3\sigma$ on the obese direction (right). At $T = 90$ the volume curves show a more important contraction for the healthy population and an opposite effect (less contraction) for the obese population. At $T = 240$ both longitudinal models output a curve close to the original one due to the exponential decay.

database. For the CVD population, a general longitudinal model is defined with a reduced number of parameters. Then the parameters are analyzed using PLS regression on the BMI, the standard index to evaluate the obese medical condition.

The work presented here shows promising preliminary results that need to be confirmed by further analysis. For the study on the CMP patients, these results still need to be validated on a larger database than the current population of 8 patients with longitudinal evolution. This is a major challenge as the acquisition of longitudinal data in clinical practice is very complicated due to the important proportion of patients who opt-out of clinical studies, especially in children and adolescents. For the CVD population, further analysis could be done to study the impact of additional etiology parameters than just the BMI.

Conclusions and Perspectives

Contents

10.1 Summary of the Main Contributions	139
10.2 Perspectives and Future Applications	141
10.3 Virtual Patient in the Age of Artificial Intelligence: the Future of Medicine	144

Cardiac motion and shape analysis is a very challenging area of research that can provide crucial information in the clinical workflow for diagnosis, classification and therapy planning. The study of cardiac motion is usually divided in 3 consecutive steps: the extraction of the myocardium geometry with segmentation, the tracking of this shape through all the frames with cardiac motion tracking algorithms and finally the application to the clinical problem using a low-dimensional representation of the motion. This corresponds to the 3 parts that were presented in this manuscript.

In the following sections we summarize the main contributions of this thesis. This manuscript does not pretend to provide final answers to the problem of cardiac motion analysis but introduces some innovative ideas which constitute small steps toward the automatic model-based analysis of the cardiac motion. There is still a lot of research work possible to improve these methods for which we also propose some general research perspectives. These perspectives either aim at translating the methods to other organs and modalities or at developing new improvements of the method presented here. Finally, we conclude this manuscript with some words on what could be the future of computational medicine and the challenges that lie ahead.

10.1 Summary of the Main Contributions

SVF-Net: Fast and Robust Registration Method Standard methods using deep learning in the context of registration define ground truth transformations using a classical registration algorithm on the pair of images. In Chapter 3, we propose an innovative methodology to build reference deformations from pairs of segmented images that are matched using a shape registration algorithm. These reference transformations are used to train a predictive registration algorithm based on convolutional neural networks. With respect to the standard method using only images, using segmentations opens the possibility for a user to correct and validate the results of the segmentation (which was done for our ground truth segmentations).

We show that our method significantly outperforms a state-of-the-art optimization method in term of registration accuracy. Maybe more important than just accuracy improvement, our method is way faster and more robust to outliers. These qualities opens the way for the use of registration in larger database, which was not possible before due to the computational cost and lack of robustness of classical registration algorithms.

Automatic Myocardium Segmentation In Chapter 4, we propose a framework for the automatic segmentation of the myocardium. Our pipeline is made of two successive modules. First, a convolutional neural network detects the RV and LV locations. This result is used to define a ROI around the myocardium and a standardized orientation of the heart. The image is cropped around this ROI which is used for the second module of the pipeline. We leverage the speed of the registration algorithm defined in Chapter 3 to perform multi-atlas registration with a large number of templates. We combine the different segmentation into the final result by defining local weights. Results show that we benefit the large database of ground truth is improving the accuracy. Such a large database can be used thanks to the speed of our registration algorithm.

Barycentric Subspace for Cardiac Motion Analysis In Chapter 5 and Chapter 8, we have challenged with a new symmetric group-wise paradigm the traditional framework for studying the cardiac motion based on the ED frame chosen as the single reference. Our approach relies on using subspaces as the reference for registration instead of choosing a specific arbitrary single image which can introduce bias. These subspaces represent the cardiac motion by meaningful parameters showing different clear patterns between two populations. Using this subspace as a prior within the registration, we achieve a better evaluation of the deformation between ED and ES frames and in particular we improve the estimation of the ejection fraction. Finally, the methodology can also be used to perform symmetric transitive registration, for better tracking along the sequence. This is the first application of Barycentric Subspaces in the context of medical imaging.

Reduced Model of the Cardiac Function for Fast Simulation and Personalization In Chapter 6, we proposed an innovative methodology to reduce a whole biophysical cardiac model using polyaffine projection and PLS regression. This method can be used for personalization and for direct forward computation of the motion given by the model. Using the inverse relation found by the PLS between the modes of the polyaffine parameters and the input parameters, we can use the proposed methodology to efficiently personalize a model based on polyaffine transformations. These transformations encompass a lot of information about the motion whereas standard methods for personalization usually only use volume curve or regional volumes. With the direct forward computation, our contribution opens the way to very fast simulation (less than 2 minutes) and the use of these simula-

tions in clinical setting. This is hardly possible with the current model regarding the time needed (≥ 2 h) to perform one simulation.

Reduced-Order Polyaffine Transformations for the Automatic Diagnostic of LV Infarct In Chapter 7, we presented an innovative methodology to project a motion on a reduced number of polyaffine parameters. We apply the methodology to classify a population and detect an infarct based on the segmentations at end-systole and end-diastole. Results show that our method has a very good predictive power with more than 95 % correct classification on 200 infarcted/control cases. Our method performs state-of-the-art results with respect to classification accuracy (AUC, sensitivity and accuracy) inline with the best competitive methods of the challenge. With respect to the competitive methods of the challenge, a major quality of our method is the interpretability of the parameters in term of local strain and displacement on the AHA regions. We are also able to quantify the importance of each of the parameters in the classification. Notably, this provides insights into what is the main impact of an infarct both in terms of motion and shape.

10.2 Perspectives and Future Applications

One of the important focus was to build methods that are robust with respect to the data and reproducible with a new database and population. Thus, the contributions we made in this work can also be applied to other organs or other heart conditions. Similarly, although the proposed methods were based on the analysis of MRI images of the heart, they could be applied to other modalities with small changes. The choice of MRI images was convenient because of the large availability of this modality in the MD-Paedigree project which delivered most of the data. But there are reasons pushing for an extension of these methods to other modalities. In clinical practice, the Echo acquisition is way more common than MRI and there are many cases when MRI images are not available. CT images have a better resolution and could permit an analysis with better precision than MRI images.

Therefore, we propose perspectives of extension of the methods developed in this manuscript to other organs or modalities. Some of these perspectives are already being developed with promising preliminary results.

SVF-Net and Multi-Atlas Segmentation: Application to Left Atrial Wall Segmentation Atrial fibrillation (AF) is the most common type of cardiac arrhythmia, characterized by uncoordinated electrical activation and disorganized contraction of the atria. In the case the patient does not answer to drug treatments, atrial ablation might be an effective treatment for AF. In this case, an accurate segmentation of the left atrial endocardium and epicardium [Jia 2016] can give further insight to the clinician. The method applied here to MRI images for the segmentation of the myocardium could also be used for CT images and the segmentation of left atrial. Preliminary work has already been done in this direction with the




Figure 10.1: Preliminary result: a segmentation computed on a CT image of the atrial with a multi-atlas method based on SVF-Net.

results shown in Fig. 10.1. More generally, it shows that the methods presented in this work can be translated and adapted to other imaging modalities (MRI / Echo / CT) or to other organs.

Barycentric Subspace Registration for Multi-Atlas Brain Segmentation

The Barycentric Subspace framework presented in this manuscript and applied in the context of cardiac motion could also provide a relevant framework for multi-atlas segmentation. One possible extension would be to use them to realize a probabilistic segmentation of brain images [Aljabar 2009]. With multiple templates, corresponding for example to different brain pathologies and represent different shape, one can compute the barycentric coefficient of a new image. These coefficients can be used as weights on each template. Large weights correspond to templates that are close to the target image and therefore correspond to templates on which we have a good confidence on the registration. The labels of the ground truth segmentation can be transported with the registration and fused with the weights corresponding to the barycentric coefficients.

Barycentric Subspaces for Geometric Dimension Reduction in Brain Computer Interface Signals

Electroencephalography (EEG)-based Brain Computer Interfaces (BCI) are interfaces (see Fig. 10.2) where the acquisition is carried out by an Electroencephalograph (EEG) and the application-specific information is extracted from the resulting signal [Clerc 2016]. EEG data have an extremely low Signal-to-Noise Ratio, while being subject to an important amount of cross-session and cross-subject variability. Hence, state-of-the-art classifiers generalize poorly,

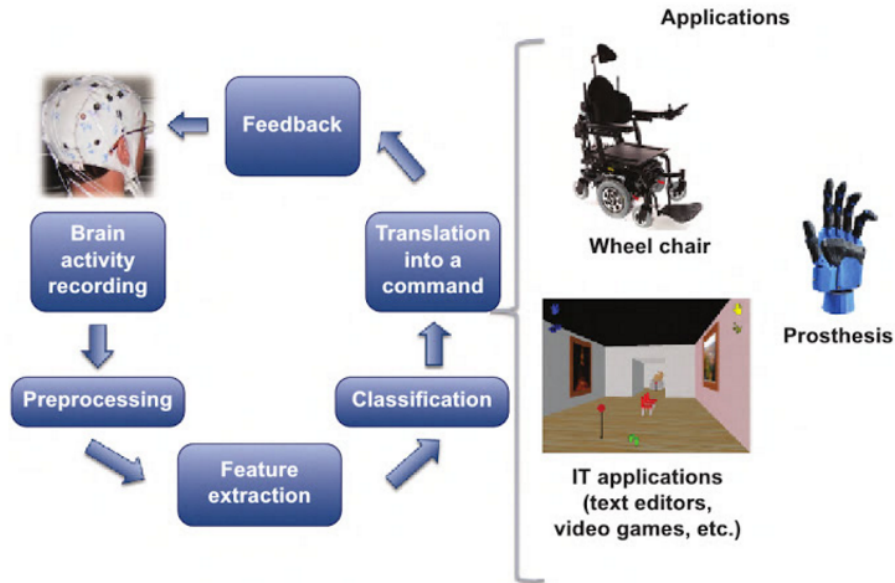


Figure 10.2: Architecture of a BCI working in real time with examples of applications. Taken from [Clerc 2016]

and calibration needs to precede each and every usage session.

Recent researches have shown that the classification task is more robust to cross-session / cross-subject variability when working the framework of Riemannian geometry [Barachant 2012]. Under the assumption that each mental task follows a specific distribution, the feature vector associated to a signal segment consists of the coefficients of a large covariance matrix. The resulting feature space is the high dimensional Riemannian manifold of Symmetric Positive Definite matrices, where an affine-invariant metric is well defined. While Riemannian distance-based classifiers show promising results, the high dimensionality of the manifold is still an issue [Gayraud 2016].

Principal Component Analysis (PCA) is the ubiquitous tool to obtain low dimensional representation of the data in linear spaces. Methods have been developed to generalize PCA to Riemannian manifolds. One can for example perform the analysis of the covariance matrix of the data in the tangent space at the Fréchet mean (Tangent PCA). However, the covariance matrix of the EEG signal has a large variability with respect to the curvature. This makes traditional methods to generalize PCA for non-linear subspace approximation in non-linear manifolds such as tangent PCA not adapted. In this context the Barycentric Subspaces presented in Chapter 5 could provide an efficient framework to overcome this problem and perform dimension reduction on the EEG signal.

10.3 Virtual Patient in the Age of Artificial Intelligence: the Future of Medicine

We conclude this work with a long-term view on what could be the future of computational medicine and how the work presented here could perform a small step in this direction. Machine learning has seen a spectacular development over the last 20 years and it has become crucial for solving complex problems in most sciences. The same methods are opening up many new possibilities in medicine. By now, there is little doubt that artificial intelligence and machine learning is going to transform medicine and the way a patient will be treated from the diagnosis of the disease to the choice of the most suitable treatment. It's essential to remember, however, that one needs to acquire large amount of data to build relevant statistical model. Once acquired, data can be analyzed and interpreted.

There has been an on-going effort to collect and regroup cardiac data in order to form large-scale database that can be used for research to build and validate models. An important initiatives in the creating of such database is the Cardiac Atlas Project [Fonseca 2011] containing the two database that were used in Chapter 7 for the automatic diagnosis of infarct: the MESA database of asymptomatic hearts [Bild 2002] and the DETERMINE database of infarcted patients [Kadish 2009]. More recently, the UK biobank [Petersen 2013] is a large-scale database made of MRI images from more than 100,000 british participants and the Cardiovascular Magnetic Resonance (EuroCMR) registry [Bruder 2013] is a collaboration of 57 centers in 15 countries to gather imaging data.

Creating such large database comes with multiple challenges such as storage, analysis, and information privacy. However, with the increasing interest of the scientific, industrial and medical community towards Big Data, there will be more and more initiatives of this type aiming at gathering very large amounts of data. Statistical and machine learning methods such as the one presented in this thesis will therefore gain relevance and it will be possible to model even more complex effects. In this thesis, we put a particular emphasize into the automation and the speed of the methods used. We believe they will provide efficient, automatic and fast tools to analyze and process the large databases that will become available to researchers and clinicians in the future.

Clinical medicine has always required clinicians to handle huge amounts of data, from macro-level physiology and behavior to laboratory and imaging studies. In particular, imaging data prove to be very complex to analyze as the difference between two cardiac sequence, for example, can be very small. In this context, we believe machine learning will become a pivotal tool for clinicians seeking to truly understand their patients and to take decision using as much information as possible in order to choose the optimal therapeutical option for the patient. As patients conditions and medical technologies become more complex, the role of machine learning will grow, and clinical medicine will be challenged to grow with it. The combination of medical imaging, statistical analysis and biophysical modeling will provide clinicians

with powerful tools for a personalised computer-aided medicine.

Bibliography

- [Aljabar 2009] Paul Aljabar, Rolf A Heckemann, Alexander Hammers, Joseph V Hajnal and Daniel Rueckert. *Multi-atlas based segmentation of brain images: atlas selection and its effect on accuracy*. Neuroimage, vol. 46, no. 3, pages 726–738, 2009. (Cited on pages [28](#), [66](#) and [142](#).)
- [Anderson 2004] Robert H Anderson, Reza Razavi and Andrew M Taylor. *Cardiac anatomy revisited*. Journal of Anatomy, vol. 205, no. 3, pages 159–177, 2004. (Cited on page [3](#).)
- [Arsigny 2003] Vincent Arsigny, Xavier Pennec and Nicholas Ayache. *Polyrigid and polyaffine transformations: A new class of diffeomorphisms for locally rigid or affine registration*. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 829–837. Springer, 2003. (Cited on page [78](#).)
- [Arsigny 2005] Vincent Arsigny, Xavier Pennec and Nicholas Ayache. *Polyrigid and polyaffine transformations: a novel geometrical tool to deal with non-rigid deformations—application to the registration of histological slices*. Medical image analysis, vol. 9, no. 6, pages 507–523, 2005. (Cited on pages [32](#) and [78](#).)
- [Arsigny 2006] Vincent Arsigny, Olivier Commowick, Xavier Pennec and Nicholas Ayache. *A log-euclidean framework for statistics on diffeomorphisms*. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 924–931. Springer, 2006. (Cited on pages [32](#), [34](#), [45](#), [48](#) and [95](#).)
- [Arsigny 2009] Vincent Arsigny, Olivier Commowick, Nicholas Ayache and Xavier Pennec. *A fast and log-euclidean polyaffine framework for locally linear registration*. Journal of Mathematical Imaging and Vision, vol. 33, no. 2, pages 222–238, 2009. (Cited on pages [78](#), [85](#), [104](#) and [106](#).)
- [Artaechevarria 2008] Xabier Artaechevarria, Arrate Muñoz-Barrutia and Carlos Ortiz-de Solorzano. *Efficient classifier generation and weighted voting for atlas-based segmentation: Two small steps faster and closer to the combination oracle*. In Medical Imaging, pages 69141W–69141W. International Society for Optics and Photonics, 2008. (Cited on page [67](#).)
- [Artaechevarria 2009] Xabier Artaechevarria, Arrate Munoz-Barrutia and Carlos Ortiz-de Solórzano. *Combination strategies in multi-atlas image segmentation: Application to brain MR data*. IEEE transactions on medical imaging, vol. 28, no. 8, pages 1266–1277, 2009. (Cited on page [66](#).)
- [Ashburner 2007a] John Ashburner. *A fast diffeomorphic image registration algorithm*. Neuroimage, vol. 38, no. 1, pages 95–113, 2007. (Cited on page [85](#).)

- [Ashburner 2007b] John Ashburner and Karl J Friston. *Rigid body registration. Statistical parametric mapping: The analysis of functional brain images*, pages 49–62, 2007. (Cited on pages 31 and 62.)
- [Ashburner 2011] John Ashburner and Karl J Friston. *Diffeomorphic registration using geodesic shooting and Gauss–Newton optimisation*. *NeuroImage*, vol. 55, no. 3, pages 954–967, 2011. (Cited on page 40.)
- [Ashburner 2013] John Ashburner and Gerard R Ridgway. *Symmetric diffeomorphic modeling of longitudinal structural MRI*. *Frontiers in neuroscience*, vol. 6, page 197, 2013. (Cited on page 28.)
- [Avants 2008] Brian B Avants, Charles L Epstein, Murray Grossman and James C Gee. *Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain*. *Medical image analysis*, vol. 12, no. 1, pages 26–41, 2008. (Cited on pages 32 and 35.)
- [Avendi 2016] MR Avendi, Arash Kheradvar and Hamid Jafarkhani. *A combined deep-learning and deformable-model approach to fully automatic segmentation of the left ventricle in cardiac MRI*. *Medical image analysis*, vol. 30, pages 108–119, 2016. (Cited on page 60.)
- [Ayache 2006] N Ayache, JP Boissel, S Brunak, G Clapworthy, G Lonsdale, J Fingberg, A Frangi, G Deco, P Hunter, P Nielsen *et al.* *Towards virtual physiological human: Multilevel modelling and simulation of the human anatomy and physiology*. *Virtual Physiological Human: White paper, EC-DG INFSO and DG JRC*, vol. 156, pages 213–220, 2006. (Cited on page 7.)
- [Balci 2007] Serdar K Balci, Polina Golland and W Wells. *Non-rigid groupwise registration using B-spline deformation model*. *Open source and open data for MICCAI*, pages 105–121, 2007. (Cited on page 87.)
- [Barachant 2012] Alexandre Barachant, Stéphane Bonnet, Marco Congedo and Christian Jutten. *Multiclass brain–computer interface classification by Riemannian geometry*. *IEEE Transactions on Biomedical Engineering*, vol. 59, no. 4, pages 920–928, 2012. (Cited on page 143.)
- [Barron 1994] John L Barron, David J Fleet and Steven S Beauchemin. *Performance of optical flow techniques*. *International journal of computer vision*, vol. 12, no. 1, pages 43–77, 1994. (Cited on page 33.)
- [Beg 2005] M Faisal Beg, Michael I Miller, Alain Trouvé and Laurent Younes. *Computing large deformation metric mappings via geodesic flows of diffeomorphisms*. *International journal of computer vision*, vol. 61, no. 2, pages 139–157, 2005. (Cited on pages 32, 34, 40, 48 and 86.)

- [Beg 2007] Mirza Faisal Beg and Ali Khan. *Symmetric data attachment terms for large deformation image registration*. IEEE transactions on medical imaging, vol. 26, no. 9, pages 1179–1189, 2007. (Cited on page 32.)
- [Belohlavek 1993] Marek Belohlavek, DAVID A FOLEY, THOMAS C GERBER, THOMAS M KINTER, James F Greenleaf and JAMES B SEWARD. *Three- and four-dimensional cardiovascular ultrasound imaging: a new era for echocardiography*. In Mayo Clinic Proceedings, volume 68, pages 221–240. Elsevier, 1993. (Cited on page 13.)
- [Berrington de Gonzalez 2010] Amy Berrington de Gonzalez, Patricia Hartge, James R Cerhan, Alan J Flint, Lindsay Hannan, Robert J MacInnis, Steven C Moore, Geoffrey S Tobias, Hoda Anton-Culver, Laura Beane Freeman *et al.* *Body-mass index and mortality among 1.46 million white adults*. N Engl J Med, vol. 2010, no. 363, pages 2211–2219, 2010. (Cited on page 12.)
- [Bertsekas 2014] Dimitri P Bertsekas. *Constrained optimization and lagrange multiplier methods*. Academic press, 2014. (Cited on pages 83 and 91.)
- [Besl 1992] Paul J Besl and Neil D McKay. *Method for registration of 3-D shapes*. In Robotics-DL tentative, pages 586–606. International Society for Optics and Photonics, 1992. (Cited on page 38.)
- [Bestel 2001] Julie Bestel, Frédérique Clément and Michel Sorine. *A biomechanical model of muscle contraction*. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 1159–1161. Springer, 2001. (Cited on pages 37 and 96.)
- [Bhatia 2004] Kanwal K Bhatia, Joseph V Hajnal, Basant K Puri, A David Edwards and Daniel Rueckert. *Consistent groupwise non-rigid registration for atlas construction*. In Biomedical Imaging: Nano to Macro, 2004. IEEE International Symposium on, pages 908–911. IEEE, 2004. (Cited on page 28.)
- [Biesdorf 2009] Andreas Biesdorf, Stefan Wörz, Hans-Jürgen Kaiser, Christoph Stippich and Karl Rohr. *Hybrid spline-based multimodal registration using local measures for joint entropy and mutual information*. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 607–615. Springer, 2009. (Cited on page 31.)
- [Bijnens 2007] Bart Bijnens, P Claus, F Weidemann, J Strotmann and GR. Sutherland. *Investigating cardiac function using motion and deformation analysis in the setting of coronary artery disease*. Circulation, 2007. (Cited on pages 76 and 104.)
- [Bild 2002] Diane E Bild, David A Bluemke, Gregory L Burke, Robert Detrano, Ana V Diez Roux, Aaron R Folsom, Philip Greenland, David R Jacobs Jr, Richard Kronmal, Kiang Liu *et al.* *Multi-ethnic study of atherosclerosis: objectives*

- and design*. American journal of epidemiology, vol. 156, no. 9, pages 871–881, 2002. (Cited on pages 105 and 144.)
- [Bistoquet 2008a] Arnaud Bistoquet. *Cardiac motion recovery from magnetic resonance images using incompressible deformable models*. PhD thesis, Georgia Institute of Technology, 2008. (Cited on page 77.)
- [Bistoquet 2008b] Arnaud Bistoquet, John Oshinski and Oskar Škrinjar. *Myocardial deformation recovery from cine MRI using a nearly incompressible biventricular model*. Medical image analysis, vol. 12, no. 1, pages 69–85, 2008. (Cited on page 77.)
- [Bogaert 2005] Jan Bogaert, Steven Dymarkowski and Andrew M Taylor. *Clinical cardiac mri*. Taylor & Francis US, 2005. (Cited on page 15.)
- [Boron 2012] Walter F Boron and Emile L Boulpaep. *Medical physiology*, 2e updated edition: with student consult online access. Elsevier Health Sciences, 2012. (Cited on page 4.)
- [Bossa 2007] Matias Bossa, Monica Hernandez and Salvador Olmos. *Contributions to 3D diffeomorphic atlas estimation: application to brain images*. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 667–674. Springer, 2007. (Cited on page 35.)
- [Boykov 2001] Yuri Boykov, Olga Veksler and Ramin Zabih. *Fast approximate energy minimization via graph cuts*. IEEE Transactions on pattern analysis and machine intelligence, vol. 23, no. 11, pages 1222–1239, 2001. (Cited on page 32.)
- [Bruder 2013] Oliver Bruder, Anja Wagner, Massimo Lombardi, Jürg Schwitter, Albert van Rossum, Günter Pilz, Detlev Nothnagel, Henning Steen, Steffen Petersen, Eike Nagelet *al*. *European Cardiovascular Magnetic Resonance (EuroCMR) registry—multi national results from 57 centers in 15 countries*. Journal of Cardiovascular Magnetic Resonance, vol. 15, no. 1, page 9, 2013. (Cited on page 144.)
- [Bruse 2017a] Jan L Bruse, Giuliano Giusti, Catriona Baker, Elena Cervi, Tain-Yen Hsia, Andrew M Taylor and Silvia Schievano. *Statistical Shape Modeling for Cavopulmonary Assist Device Development: Variability of Vascular Graft Geometry and Implications for Hemodynamics*. Journal of Medical Devices, vol. 11, no. 2, page 021011, 2017. (Cited on page 37.)
- [Bruse 2017b] Jan L Bruse, Abbas Khushnood, Kristin McLeod, Giovanni Biglino, Maxime Sermesant, Xavier Pennec, Andrew M Taylor, Tain-Yen Hsia, Silvia Schievano, Modeling of Congenital Hearts Alliance Collaborative Group *et al*. *How successful is successful? Aortic arch shape after successful aortic coarctation repair correlates with left ventricular function*. The Journal of thoracic

- and cardiovascular surgery, vol. 153, no. 2, pages 418–427, 2017. (Cited on page 37.)
- [Bulat 2016] Adrian Bulat and Georgios Tzimiropoulos. *Convolutional aggregation of local evidence for large pose face alignment*. 2016. (Cited on page 63.)
- [Cachier 2000] Pascal Cachier and Xavier Pennec. *3D non-rigid registration by gradient descent on a gaussian-windowed similarity measure using convolutions*. In *Mathematical Methods in Biomedical Image Analysis*, 2000. Proceedings. IEEE Workshop on, pages 182–189. IEEE, 2000. (Cited on page 68.)
- [Cachier 2003] Pascal Cachier, Eric Bardinet, Didier Dormont, Xavier Pennec and Nicholas Ayache. *Iconic feature based nonrigid registration: the PASHA algorithm*. *Computer vision and image understanding*, vol. 89, no. 2, pages 272–298, 2003. (Cited on pages 31 and 33.)
- [Cachier 2004] Pascal Cachier and Nicholas Ayache. *Isotropic energies, filters and splines for vector field regularization*. *Journal of Mathematical Imaging and Vision*, vol. 20, no. 3, pages 251–265, 2004. (Cited on page 34.)
- [Carneiro 2012] Gustavo Carneiro, Jacinto C Nascimento and António Freitas. *The segmentation of the left ventricle of the heart from ultrasound data using deep learning architectures and derivative-based search methods*. *IEEE Transactions on Image Processing*, vol. 21, no. 3, pages 968–982, 2012. (Cited on page 60.)
- [Cash 2004] Thomas F Cash, Jennifer A Morrow, Joshua I Hrabosky and April A Perry. *How has body image changed? A cross-sectional investigation of college women and men from 1983 to 2001*. *Journal of consulting and clinical psychology*, vol. 72, no. 6, page 1081, 2004. (Cited on page 28.)
- [Ceritoglu 2013] Ceritoglu and et al. *Computational analysis of LDDMM for brain mapping*. *Frontiers in neuroscience*, 2013. (Cited on page 44.)
- [Chandrashekhara 2004] Raghavendra Chandrashekhara, Raad H Mohiaddin and Daniel Rueckert. *Analysis of 3-D myocardial motion in tagged MR images using nonrigid image registration*. *IEEE Transactions on Medical Imaging*, vol. 23, no. 10, pages 1245–1250, 2004. (Cited on page 85.)
- [Chen 1992] Yang Chen and Gérard Medioni. *Object modelling by registration of multiple range images*. *Image and vision computing*, vol. 10, no. 3, pages 145–155, 1992. (Cited on page 38.)
- [Cheng 2016] Xi Cheng, Li Zhang and Yefeng Zheng. *Deep similarity learning for multimodal medical images*. *Computer Methods in Biomechanics and Biomedical Engineering: Imaging & Visualization*, pages 1–5, 2016. (Cited on page 31.)

- [Chuang 2011] Michael L Chuang, Philimon Gona, Gilion LTF Hautvast, Carol J Salton, Susan J Blease, Susan B Yeon, Marcel Breeuwer, Christopher J O'Donnell and Warren J Manning. *Impact of left ventricular trabeculations and papillary muscles on measures of cavity volume and ejection fraction*. Journal of Cardiovascular Magnetic Resonance, vol. 13, no. S1, page P36, 2011. (Cited on page 58.)
- [Chui 2003] Haili Chui and Anand Rangarajan. *A new point matching algorithm for non-rigid registration*. Computer Vision and Image Understanding, vol. 89, no. 2, pages 114–141, 2003. (Cited on page 30.)
- [Çiçek 2016] Özgün Çiçek, Ahmed Abdulkadir, Soeren S Lienkamp, Thomas Brox and Olaf Ronneberger. *3d u-net: learning dense volumetric segmentation from sparse annotation*. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 424–432. Springer, 2016. (Cited on page 60.)
- [Cifor 2013] Amalia Cifor, Laurent Risser, Daniel Chung, Ewan M Anderson and Julia A Schnabel. *Hybrid feature-based diffeomorphic registration for tumor tracking in 2-d liver ultrasound images*. IEEE transactions on medical imaging, vol. 32, no. 9, pages 1647–1656, 2013. (Cited on page 31.)
- [Clarysse 2000] Patrick Clarysse, C Basset, Leila Khouas, Pierre Croisille, Denis Friboulet, Christophe Odet and Isabelle E Magnin. *Two-dimensional spatial and temporal displacement and deformation field fitting from cardiac magnetic resonance tagging*. Medical Image Analysis, vol. 4, no. 3, pages 253–268, 2000. (Cited on page 78.)
- [Clerc 2016] Maureen Clerc, Laurent Bougrain and Fabien Lotte. Anatomical evaluation of colin 27 against a database of labeled brain scans. John Wiley & Sons, 2016. (Cited on pages 142 and 143.)
- [Coben 1985] Lawrence A Coben, Warren Danziger and Martha Storandt. *A longitudinal EEG study of mild senile dementia of Alzheimer type: changes at 1 year and at 2.5 years*. Electroencephalography and clinical neurophysiology, vol. 61, no. 2, pages 101–112, 1985. (Cited on page 126.)
- [Collins 1995] D Louis Collins, Colin J Holmes, Terrence M Peters and Alan C Evans. *Automatic 3-D model-based neuroanatomical segmentation*. Human brain mapping, vol. 3, no. 3, pages 190–208, 1995. (Cited on page 59.)
- [Commowick 2008] Olivier Commowick, Vincent Arsigny, Aurélie Isambert, Jimena Costa, Frédéric Dhermain, François Bidault, P-Y Bondiau, Nicholas Ayache and Grégoire Malandain. *An efficient locally affine framework for the smooth registration of anatomical structures*. Medical Image Analysis, vol. 12, no. 4, pages 427–441, 2008. (Cited on page 78.)

- [Courchesne 2007] O Courchesne, François Guibault, Julien Dompierre and Farida Cheriet. *Adaptive mesh generation of MRI images for 3D reconstruction of human trunk*. In International Conference Image Analysis and Recognition, pages 1040–1051. Springer, 2007. (Cited on page 58.)
- [Cousty 2010] Jean Cousty, Laurent Najman, Michel Couprie, Stéphanie Clément-Guinaudeau, Thomas Goissen and Jérôme Garot. *Segmentation of 4D cardiac MRI: Automated method based on spatio-temporal watershed cuts*. Image and Vision Computing, vol. 28, no. 8, pages 1229–1243, 2010. (Cited on page 59.)
- [Curiale 2016] Ariel H Curiale, Gonzalo Vegas-Sánchez-Ferrero and Santiago Aja-Fernández. *Influence of ultrasound speckle tracking strategies for motion and strain estimation*. Medical image analysis, vol. 32, pages 184–200, 2016. (Cited on page 78.)
- [Davis 1997] Malcolm H Davis, Alireza Khotanzad, Duane P Flamig and Steven E Harms. *A physics-based coordinate transformation for 3-D image matching*. IEEE transactions on medical imaging, 1997. (Cited on page 48.)
- [De Craene 2010] Mathieu De Craene, Gemma Piella, Nicolas Duchateau, Etel Silva, Adelina Doltra, Hang Gao, Jan DâĂžhooge, Oscar Camara, Josep Brugada, Marta Sitges *et al.* *Temporal diffeomorphic free-form deformation for strain quantification in 3D-US images*. Medical Image Computing and Computer-Assisted Intervention–MICCAI 2010, pages 1–8, 2010. (Cited on page 78.)
- [De Craene 2011] Mathieu De Craene, Catalina Tobon-Gomez, Constantine Butakoff, Nicolas Duchateau, Gemma Piella, Kawal S Rhode and Alejandro F Frangi. *Temporal diffeomorphic free form deformation (TDFFD) applied to motion and deformation quantification of tagged MRI sequences*. In International Workshop on Statistical Atlases and Computational Models of the Heart, pages 68–77. Springer, 2011. (Cited on pages 78 and 85.)
- [De Craene 2012] Mathieu De Craene, Gemma Piella, Oscar Camara, Nicolas Duchateau, Etelvino Silva, Adelina Doltra, Jan DâĂžhooge, Josep Brugada, Marta Sitges and Alejandro F Frangi. *Temporal diffeomorphic free-form deformation: Application to motion and strain estimation from 3D echocardiography*. Medical Image Analysis, vol. 16, no. 2, pages 427–450, 2012. (Cited on page 78.)
- [De Craene 2013] Mathieu De Craene, Stéphanie Marchesseau, Brecht Heyde, Hang Gao, M Alessandrini, Olivier Bernard, Gemma Piella, AR Porras, Lennart Tautz, Anja Hennemuth *et al.* *3D strain assessment in ultrasound (straus): A synthetic comparison of five tracking methodologies*. IEEE transactions on medical imaging, vol. 32, no. 9, pages 1632–1646, 2013. (Cited on page 79.)

- [Delingette 1992] Hervé Delingette, Martial Hebert and Katsushi Ikeuchi. *Shape representation and image segmentation using deformable surfaces*. Image and vision computing, vol. 10, no. 3, pages 132–144, 1992. (Cited on page 59.)
- [Dolz 2014] J Dolz, L Massoptier and Maximilien Vermandel. *Segmentation algorithms of subcortical brain structures on MRI: a review*. Journal of Neuroimage, pages 200–212, 2014. (Cited on page 59.)
- [Dong 1995] Lei Dong and Arthur L Boyer. *An image correlation procedure for digitally reconstructed radiographs and electronic portal images*. International Journal of Radiation Oncology* Biology* Physics, vol. 33, no. 5, pages 1053–1060, 1995. (Cited on page 35.)
- [Donnelly 2001] Lane F Donnelly, Kathleen H Emery, Alan S Brody, Tal Laor, Victoria M Gylys-Morin, Christopher G Anton, Stephen R Thomas and Donald P Frush. *Minimizing radiation dose for pediatric body applications of single-detector helical CT: strategies at a large children’s hospital*. American Journal of Roentgenology, vol. 176, no. 2, pages 303–306, 2001. (Cited on page 17.)
- [Dupuis 1998] Paul Dupuis, Ulf Grenander and Michael I Miller. *Variational problems on flows of diffeomorphisms for image matching*. Quarterly of applied mathematics, pages 587–600, 1998. (Cited on page 40.)
- [Durrleman 2010] Stanley Durrleman. *Statistical models of currents for measuring the variability of anatomical curves, surfaces and their evolution*. PhD thesis, Université Nice Sophia Antipolis, 2010. (Cited on page 38.)
- [Durrleman 2011] Stanley Durrleman, Pierre Fillard, Xavier Pennec, Alain Trouvé and Nicholas Ayache. *Registration, atlas estimation and variability analysis of white matter fiber bundles modeled as currents*. NeuroImage, vol. 55, no. 3, pages 1073–1090, 2011. (Cited on page 37.)
- [Durrleman 2014] Stanley Durrleman, Marcel Prastawa, Nicolas Charon, Julie R Korenberg, Sarang Joshi, Guido Gerig and Alain Trouvé. *Morphometry of anatomical shape complexes with dense deformations and sparse parameters*. NeuroImage, 2014. (Cited on pages 41, 45 and 47.)
- [Ecabert 2008] Olivier Ecabert, Jochen Peters, Hauke Schramm, Cristian Lorenz, Jens von Berg, Matthew J Walker, Mani Vembar, Mark E Olszewski, Krishna Subramanyan, Guy Laviet *al*. *Automatic model-based segmentation of the heart in CT images*. IEEE transactions on medical imaging, vol. 27, no. 9, pages 1189–1201, 2008. (Cited on page 2.)
- [Elliott 2007] Perry Elliott, Bert Andersson, Eloisa Arbustini, Zofia Bilinska, Franco Cecchi, Philippe Charron, Olivier Dubourg, Uwe Köhl, Bernhard Maisch, William J McKenna *et al*. *Classification of the cardiomyopathies: a position*

- statement from the European Society Of Cardiology Working Group on Myocardial and Pericardial Diseases.* European heart journal, 2007. (Cited on pages 9 and 10.)
- [Elliott 2012] Perry Elliott, Bert Andersson, Eloisa Arbustini, Zofia Bilinska, Franco Cecchi, Philippe Charron, Olivier Dubourg, Uwe Kühl, Bernhard Maisch, William J McKenna *et al.* *Cardiomyopathy Classification: Ongoing Debate in the Genomics Era.* Biochemistry Research International, 2012. (Cited on pages 9 and 10.)
- [Evans 2006] Alan C Evans, Brain Development Cooperative Group *et al.* *The NIH MRI study of normal brain development.* Neuroimage, vol. 30, no. 1, pages 184–202, 2006. (Cited on page 126.)
- [Fang 2009] Qianqian Fang and David A Boas. *Tetrahedral mesh generation from volumetric binary and grayscale images.* In Biomedical Imaging: From Nano to Macro, 2009. ISBI'09. IEEE International Symposium on, pages 1142–1145. IEEE, 2009. (Cited on page 58.)
- [Feldmar 1996] Jacques Feldmar and Nicholas Ayache. *Rigid, affine and locally affine registration of free-form surfaces.* International journal of computer vision, vol. 18, no. 2, pages 99–119, 1996. (Cited on page 31.)
- [Ferre 1999] Maurice R Ferre, Peter D Jakab and James S Tieman. *Automatic registration system for use with position tracking and imaging system for use in medical applications,* February 23 1999. US Patent 5,873,822. (Cited on page 2.)
- [Fletcher 2004] P Thomas Fletcher, Conglin Lu, Stephen M Pizer and Sarang Joshi. *Principal geodesic analysis for the study of nonlinear statistics of shape.* IEEE transactions on medical imaging, 2004. (Cited on page 80.)
- [Fonseca 2011] Carissa G Fonseca, Michael Backhaus, David A Bluemke, Randall D Britten, Jae Do Chung, Brett R Cowan, Ivo D Dinov, J Paul Finn, Peter J Hunter, Alan H Kadish *et al.* *The Cardiac Atlas Project—An imaging database for computational modeling and statistical atlases of the heart.* Bioinformatics, vol. 27, no. 16, pages 2288–2295, 2011. (Cited on pages 106 and 144.)
- [Frangi 2002] Alejandro F Frangi, Daniel Rueckert, Julia A Schnabel and Wiro J Niessen. *Automatic construction of multiple-object three-dimensional statistical shape models: Application to cardiac modeling.* IEEE transactions on medical imaging, vol. 21, no. 9, pages 1151–1166, 2002. (Cited on page 37.)
- [Fratz 2013] Sohrab Fratz, Taylor Chung, Gerald F Greil, Margaret M Samyn, Andrew M Taylor, Emanuela R Valsangiacomo Buechel, Shi-Joon Yoo and Andrew J Powell. *Guidelines and protocols for cardiovascular magnetic resonance in children and adults with congenital heart disease: SCMR expert*

- consensus group on congenital heart disease*. Journal of Cardiovascular Magnetic Resonance, vol. 15, no. 1, page 51, 2013. (Cited on page 57.)
- [Fritscher 2016] Karl Fritscher, Patrik Raudaschl, Paolo Zaffino, Maria Francesca Spadea, Gregory C Sharp and Rainer Schubert. *Deep Neural Networks for Fast Segmentation of 3D Medical Images*. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 158–165. Springer, 2016. (Cited on page 60.)
- [Garot 2000] Jérôme Garot, David A Bluemke, Nael F Osman, Carlos E Rochitte, Elliot R McVeigh, Elias A Zerhouni, Jerry L Prince and Joao AC Lima. *Fast determination of regional myocardial strain fields from tagged cardiac images using harmonic phase MRI*. Circulation, vol. 101, no. 9, pages 981–988, 2000. (Cited on page 77.)
- [Gaudron 1993] Peter Gaudron, Christoph Eilles, I Kugler and G Ertl. *Progressive left ventricular dysfunction and remodeling after myocardial infarction. Potential mechanisms and early predictors*. Circulation, vol. 87, no. 3, pages 755–763, 1993. (Cited on page 104.)
- [Gayraud 2016] Nathalie Gayraud, Nathanael Foy and Maureen Clerc. *A Separability Marker Based on High-Dimensional Statistics for Classification Confidence Assessment*. In IEEE International Conference on Systems, Man, and Cybernetics October 9-12, 2016. (Cited on page 143.)
- [Georgescu 2005] Bogdan Georgescu, Xiang Sean Zhou, Dorin Comaniciu and Alok Gupta. *Database-guided segmentation of anatomical structures with complex appearance*. In Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on, volume 2, pages 429–436. IEEE, 2005. (Cited on pages 51 and 127.)
- [Ghesu 2016] Florin C Ghesu, Bogdan Georgescu, Tommaso Mansi, Dominik Neumann, Joachim Hornegger and Dorin Comaniciu. *An artificial agent for anatomical landmark detection in medical images*. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 229–237. Springer, 2016. (Cited on page 63.)
- [Giedd 1999] Jay N Giedd, Jonathan Blumenthal, Neal O Jeffries, F Xavier Castellanos, Hong Liu, Alex Zijdenbos, Tomáš Paus, Alan C Evans and Judith L Rapoport. *Brain development during childhood and adolescence: a longitudinal MRI study*. Nature neuroscience, vol. 2, no. 10, pages 861–863, 1999. (Cited on page 126.)
- [Gil 1998] Juan A Gil and Rosario Romera. *On robust partial least squares (PLS) methods*. Journal of Chemometrics, vol. 12, no. 6, pages 365–378, 1998. (Cited on page 111.)

- [Gil 2010] Arturo Gil, Oscar Martinez Mozos, Monica Ballesta and Oscar Reinoso. *A comparative evaluation of interest point detectors and local descriptors for visual SLAM*. Machine Vision and Applications, vol. 21, no. 6, pages 905–920, 2010. (Cited on page 30.)
- [Glaunes 2005] Joan Glaunes. *Transport par difféomorphismes de points, de mesures et de courants pour la comparaison de formes et l’anatomie numérique*. These de sciences, Université Paris, vol. 13, 2005. (Cited on pages 38 and 39.)
- [Goland 2008] Sorel Goland, Lawrence SC Czer, Daniel Luthringer and Robert J Siegel. *A case of arrhythmogenic right ventricular cardiomyopathy*. Canadian Journal of Cardiology, vol. 24, no. 1, pages 61–62, 2008. (Cited on page 14.)
- [Goshtasby 2005] A Ardeshir Goshtasby. *2-d and 3-d image registration: for medical, remote sensing, and industrial applications*. John Wiley & Sons, 2005. (Cited on page 2.)
- [Guo 2004] Hongyu Guo, Anand Rangarajan, S Joshi and Laurent Younes. *Non-rigid registration of shapes via diffeomorphic point matching*. In Biomedical Imaging: Nano to Macro, 2004. IEEE International Symposium on, pages 924–927. IEEE, 2004. (Cited on page 37.)
- [Hadj-Hamou 2016a] Mehdi Hadj-Hamou. *Beyond Volumetry in Longitudinal Deformation-Based Morphometry: Application to Sexual Dimorphism during Adolescence*. PhD thesis, Université Nice Côte d’Azur, 2016. (Cited on page 126.)
- [Hadj-Hamou 2016b] Mehdi Hadj-Hamou, Marco Lorenzi, Nicholas Ayache and Xavier Pennec. *Longitudinal analysis of image time series with diffeomorphic deformations: a computational framework based on stationary velocity fields*. Frontiers in Neuroscience, vol. 10, 2016. (Cited on pages 28 and 126.)
- [Hammer 2009] Øyvind Hammer, DAT Harper and PD Ryan. *PAST-PALaeontological STatistics, ver. 1.89*. University of Oslo, Oslo, pages 1–31, 2009. (Cited on page 80.)
- [Han 2008] Yuchi Han, Eric Olson, Martin S Maron, Warren J Manning and Susan B Yeon. *2075 Papillary muscles and trabeculations significantly impact ventricular volume, ejection fraction, and regurgitation assessment by cardiovascular magnetic resonance in patients with hypertrophic cardiomyopathy*. In Journal of Cardiovascular Magnetic Resonance, volume 10, page A344. Springer, 2008. (Cited on page 58.)
- [Hansen 2012] Michael S Hansen, Signe S Thorup and Simon K Warfield. *Polyaffine parametrization of image registration based on geodesic flows*. In Mathematical Methods in Biomedical Image Analysis (MMBIA), 2012 IEEE Workshop on, pages 289–295. IEEE, 2012. (Cited on page 127.)

- [Hauser 2016] Jakob A Hauser, Vivek Muthurangu, Jennifer A Steeden, Andrew M Taylor and Alexander Jones. *Comprehensive assessment of the global and regional vascular responses to food ingestion in humans using novel rapid MRI*. American Journal of Physiology-Regulatory, Integrative and Comparative Physiology, vol. 310, no. 6, pages R541–R545, 2016. (Cited on page 134.)
- [Heckemann 2006] Rolf A Heckemann, Joseph V Hajnal, Paul Aljabar, Daniel Rueckert and Alexander Hammers. *Automatic anatomical brain MRI segmentation combining label propagation and decision fusion*. NeuroImage, vol. 33, no. 1, pages 115–126, 2006. (Cited on pages 61 and 66.)
- [Heinrich 2016] Mattias P Heinrich, Ivor JA Simpson, BartŁomiej W Papież, Michael Brady and Julia A Schnabel. *Deformable image registration by combining uncertainty estimates from supervoxel belief propagation*. Medical image analysis, vol. 27, pages 57–71, 2016. (Cited on page 32.)
- [Heller 2002] Gary V Heller, Manuel D Cerqueira, Neil J Weissman, Vasken Dilisizian, Alice K Jacobs, Sanjiv Kaul, Warren K Laskey, Dudley J Pennell, John A Rumberger, Thomas Ryan *et al.* *Standardized myocardial segmentation and nomenclature for tomographic imaging of the heart: a statement for healthcare professionals from the Cardiac Imaging Committee of the Council on Clinical Cardiology of the American Heart Association*. Journal of Nuclear Cardiology, vol. 9, no. 2, pages 240–245, 2002. (Cited on page 2.)
- [Huang 1999] Jiantao Huang, Dana Abendschein, Victor G Davila-Roman and Amir A Amini. *Spatio-temporal tracking of myocardial deformations with a 4-D B-spline model from tagged MRI*. IEEE Transactions on Medical Imaging, vol. 18, no. 10, pages 957–972, 1999. (Cited on pages 77 and 78.)
- [Hunter 2003] Peter J Hunter and Thomas K Borg. *Integration from proteins to organs: the Physiome Project*. Nature reviews Molecular cell biology, vol. 4, no. 3, pages 237–243, 2003. (Cited on page 7.)
- [Hunter 2010] Peter Hunter, Peter V. Coveney, Bernard de Bono, Vanessa Diaz, John Fenner, Alejandro F. Frangi, Peter Harris, Rod Hose, Peter Kohl, Pat Lawford, Keith McCormack, Miriam Mendes, Stig Omholt, Alfio Quarteroni, John Skår, Jesper Tegner, S. Randall Thomas, Ioannis Tollis, Ioannis Tsamardinos, Johannes H. G. M. van Beek and Marco Viceconti. *A vision and strategy for the virtual physiological human in 2010 and beyond*. Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences, vol. 368, no. 1920, pages 2595–2614, 2010. (Cited on page 7.)
- [Iglesias 2009] Juan Eugenio Iglesias and Nico Karssemeijer. *Robust initial detection of landmarks in film-screen mammograms using multiple FFDM atlases*. IEEE transactions on medical imaging, vol. 28, no. 11, pages 1815–1824, 2009. (Cited on page 67.)

- [Iglesias 2015] Juan Eugenio Iglesias and Mert R Sabuncu. *Multi-atlas segmentation of biomedical images: a survey*. Medical image analysis, vol. 24, no. 1, pages 205–219, 2015. (Cited on pages 28 and 60.)
- [Ino 2006] Fumihiko Ino, Jun Gomita, Yasuhiro Kawasaki and Kenichi Hagihara. *A GPGPU approach for accelerating 2-D/3-D rigid registration of medical images*. In International Symposium on Parallel and Distributed Processing and Applications, pages 939–950. Springer, 2006. (Cited on page 62.)
- [Ionasec 2010] Razvan Ioan Ionasec, Ingmar Voigt, Bogdan Georgescu, Yang Wang, Helene Houle, Fernando Vega-Higuera, Nassir Navab and Dorin Comaniciu. *Patient-specific modeling and quantification of the aortic and mitral valves from 4-D cardiac CT and TEE*. IEEE transactions on medical imaging, vol. 29, no. 9, pages 1636–1651, 2010. (Cited on page 127.)
- [Isgum 2009] Ivana Isgum, Marius Staring, Annemarieke Rutten, Mathias Prokop, Max A Viergever and Bram Van Ginneken. *Multi-atlas-based segmentation with local decision fusion—Application to cardiac and aortic segmentation in CT scans*. IEEE transactions on medical imaging, vol. 28, no. 7, pages 1000–1010, 2009. (Cited on pages 59 and 67.)
- [Jia 2016] Shuman Jia, Loïc Cadour, Hubert Cochet and Maxime Sermesant. *STACOM-SLAWT Challenge: Left Atrial Wall Segmentation and Thickness Measurement Using Region Growing and Marker-Controlled Geodesic Active Contour*. In Statistical Atlases and Computational Modeling of the Heart (STACOM 2016), 2016. (Cited on page 141.)
- [Jiao 2014] Jieqing Jiao, Graham E Searle, Andri C Tziortzi, Cristian A Salinas, Roger N Gunn and Julia A Schnabel. *Spatio-temporal pharmacokinetic model based registration of 4D PET neuroimaging data*. Neuroimage, vol. 84, pages 225–235, 2014. (Cited on page 134.)
- [Jolly 2006] Marie-Pierre Jolly. *Automatic segmentation of the left ventricle in cardiac MR and CT images*. International Journal of Computer Vision, vol. 70, no. 2, pages 151–163, 2006. (Cited on page 63.)
- [Jolly 2009] M Jolly. *Fully automatic left ventricle segmentation in cardiac cine MR images using registration and minimum surfaces*. The MIDAS Journal-Cardiac MR Left Ventricle Segmentation Challenge, vol. 4, 2009. (Cited on page 59.)
- [Joshi 2004] Sarang Joshi, Brad Davis, Matthieu Jomier and Guido Gerig. *Unbiased diffeomorphic atlas construction for computational anatomy*. NeuroImage, 2004. (Cited on pages 79 and 82.)
- [Kachelrieß 2000] Marc Kachelrieß, Stefan Ulzheimer and Willi A Kalender. *ECG-correlated image reconstruction from subsecond multi-slice spiral CT scans of*

- the heart*. Medical physics, vol. 27, no. 8, pages 1881–1902, 2000. (Cited on page 17.)
- [Kadish 2009] Alan H Kadish, David Bello, J Finn, Robert O Bonow, Andi Schaechter, Haris Subacius, Christine Albert, James P Daubert, Carissa G Fonseca and Jeffrey J Goldberger. *Rationale and design for the defibrillators to reduce risk by magnetic resonance imaging evaluation (DETERMINE) trial*. Journal of cardiovascular electrophysiology, vol. 20, no. 9, pages 982–987, 2009. (Cited on pages 105 and 144.)
- [Kallaras 2001] Konstantinos Kallaras, Elizabeth A Sparks, Dara P Schuster, Kwame Osei, Charles F Wooley and Harisios Boudoulas. *Cardiovascular Effects of Aging Interrelationships of Aortic, Left Ventricular, and Left Atrial Function*. Herz, vol. 26, no. 2, pages 129–139, 2001. (Cited on page 104.)
- [Katouzian 2006] Amin Katouzian, Ashwin Prakash and Elisa Konofagou. *A new automated technique for left-and right-ventricular segmentation in magnetic resonance imaging*. In Engineering in Medicine and Biology Society, 2006. EMBS'06. 28th Annual International Conference of the IEEE, pages 3074–3077. IEEE, 2006. (Cited on page 59.)
- [Kilner 2010] Philip J Kilner, Tal Geva, Harald Kaemmerer, Pedro T Trindade, Juerg Schwitler and Gary D Webb. *Recommendations for cardiovascular magnetic resonance in adults with congenital heart disease from the respective working groups of the European Society of Cardiology*. European heart journal, page ehp586, 2010. (Cited on page 57.)
- [Kingma 2014] Diederik P. Kingma and Jimmy Ba. *Adam: A Method for Stochastic Optimization*. CoRR, 2014. (Cited on page 51.)
- [Klein 2005] Arno Klein, Brett Mensh, Satrajit Ghosh, Jason Tourville and Joy Hirsch. *Mindboggle: automated brain labeling with multiple atlases*. BMC medical imaging, vol. 5, no. 1, page 7, 2005. (Cited on pages 61 and 66.)
- [Klein 2007] Stefan Klein, Marius Staring and Josien PW Pluim. *Evaluation of optimization methods for nonrigid medical image registration using mutual information and B-splines*. IEEE transactions on image processing, vol. 16, no. 12, pages 2879–2890, 2007. (Cited on page 33.)
- [Klein 2009a] Arno Klein and et al. *Evaluation of 14 nonlinear deformation algorithms applied to human brain MRI registration*. Neuroimage, 2009. (Cited on page 85.)
- [Klein 2009b] Stefan Klein, Josien PW Pluim, Marius Staring and Max A Viergever. *Adaptive stochastic gradient descent optimisation for image registration*. International journal of computer vision, vol. 81, no. 3, page 227, 2009. (Cited on page 33.)

- [Kong 2016] Bin Kong, Yiqiang Zhan, Min Shin, Thomas Denny and Shaoting Zhang. *Recognizing End-Diastole and End-Systole Frames via Deep Temporal Regression Network*. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 264–272. Springer, 2016. (Cited on page 120.)
- [Konstam 2011] Marvin A. Konstam, Daniel G. Kramer, Ayan R. Patel, Martin S. Maron and James E. Udelson. *Left Ventricular Remodeling in Heart Failure. Current Concepts in Clinical Significance and Assessment*. JACC: Cardiovascular Imaging, 2011. (Cited on pages 76 and 104.)
- [Kramer 2013] Christopher M Kramer, Jörg Barkhausen, Scott D Flamm, Raymond J Kim and Eike Nagel. *Standardized cardiovascular magnetic resonance (CMR) protocols 2013 update*. Journal of Cardiovascular Magnetic Resonance, vol. 15, no. 1, page 91, 2013. (Cited on page 57.)
- [Krizhevsky 2012] Alex Krizhevsky, Ilya Sutskever and Geoffrey E Hinton. *Imagenet classification with deep convolutional neural networks*. In Advances in neural information processing systems, pages 1097–1105, 2012. (Cited on page 60.)
- [Lauterbur 1973] Paul C Lauterbur. *Image formation by induced local interactions: examples employing nuclear magnetic resonance*. 1973. (Cited on page 14.)
- [Lebel 2011] Catherine Lebel and Christian Beaulieu. *Longitudinal development of human brain wiring continues from childhood into adulthood*. Journal of Neuroscience, vol. 31, no. 30, pages 10937–10947, 2011. (Cited on page 126.)
- [Ledesma-Carbayo 2005] María J Ledesma-Carbayo, Jan Kybic, Manuel Desco, Andrés Santos, M Suhling, Patrick Hunziker and Michael Unser. *Spatio-temporal nonrigid registration for ultrasound cardiac motion estimation*. IEEE transactions on medical imaging, vol. 24, no. 9, pages 1113–1126, 2005. (Cited on page 78.)
- [Lempitsky 2009] Victor Lempitsky, Michael Verhoek, J Alison Noble and Andrew Blake. *Random forest classification for automatic delineation of myocardium in real-time 3D echocardiography*. In International Conference on Functional Imaging and Modeling of the Heart, pages 447–456. Springer, 2009. (Cited on page 59.)
- [Lewis 1995] JP Lewis. *Fast normalized cross-correlation*. In Vision interface, volume 10, pages 120–123, 1995. (Cited on page 68.)
- [Li 2010] Bo Li, Yingmin Liu, Christopher J Occleshaw, Brett R Cowan and Alistair A Young. *In-line automated tracking for ventricular function with magnetic resonance imaging*. JACC: Cardiovascular Imaging, vol. 3, no. 8, pages 860–866, 2010. (Cited on page 106.)

- [Liang 1999] Zhi-Pei Liang and PC Lauterbur. *Principles of magnetic resonance imaging: a signal processing perspective. 2000.* New York: IEEE, 1999. (Cited on page 15.)
- [Long 2015] Jonathan Long, Evan Shelhamer and Trevor Darrell. *Fully convolutional networks for semantic segmentation.* In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2015. (Cited on page 49.)
- [Lorenzi 2010] Marco Lorenzi, Nicholas Ayache, Giovanni Frisoni, Xavier Pennec *et al.* *4D registration of serial brain's MR images: a robust measure of changes applied to Alzheimer's disease.* In Spatio Temporal Image Analysis Workshop (STIA), MICCAI. Citeseer, 2010. (Cited on page 126.)
- [Lorenzi 2013a] Marco Lorenzi, Nicholas Ayache, Giovanni B Frisoni and Xavier Pennec. *LCC-Demons: a robust and accurate symmetric diffeomorphic registration algorithm.* NeuroImage, vol. 81, pages 470–483, 2013. (Cited on pages 33, 35, 46, 47, 51, 68, 88 and 119.)
- [Lorenzi 2013b] Marco Lorenzi and Xavier Pennec. *Geodesics, parallel transport & one-parameter subgroups for diffeomorphic image registration.* International journal of computer vision, vol. 105, no. 2, pages 111–127, 2013. (Cited on page 28.)
- [Lorenzo-Valdés 2002] Maria Lorenzo-Valdés, Gerardo I Sanchez-Ortiz, Raad Mohiaddin and Daniel Rueckert. *Atlas-based segmentation and tracking of 3D cardiac MR images using non-rigid registration.* In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 642–650. Springer, 2002. (Cited on pages 28 and 59.)
- [Lorenzo-Valdés 2004] Maria Lorenzo-Valdés, Gerardo I Sanchez-Ortiz, Andrew G Elkington, Raad H Mohiaddin and Daniel Rueckert. *Segmentation of 4D cardiac MR images using a probabilistic atlas and the EM algorithm.* Medical Image Analysis, vol. 8, no. 3, pages 255–265, 2004. (Cited on page 59.)
- [Lu 2010] Huanxiang Lu, Mauricio Reyes, Amira Šiñerifović, A Šerifović, S Weber, Y Sakurai, H Yamagata and Ph C Cattin. *Multi-modal diffeomorphic demons registration based on point-wise mutual information.* In Biomedical Imaging: From Nano to Macro, 2010 IEEE International Symposium on, pages 372–375. IEEE, 2010. (Cited on page 31.)
- [Lung 2016] National Heart Lung and Blood Institute. *What Is Cardiomyopathy?*, 2016. (Cited on page 10.)
- [Maes 1997] Frederik Maes, Andre Collignon, Dirk Vandermeulen, Guy Marchal and Paul Suetens. *Multimodality image registration by maximization of mutual information.* IEEE transactions on medical imaging, vol. 16, no. 2, pages 187–198, 1997. (Cited on page 30.)

- [Mahapatra 2012a] Dwarikanath Mahapatra. *Groupwise registration of dynamic cardiac perfusion images using temporal dynamics and segmentation information*. In SPIE Medical Imaging, pages 83141X–83141X. International Society for Optics and Photonics, 2012. (Cited on page 28.)
- [Mahapatra 2012b] Dwarikanath Mahapatra. *Landmark detection in cardiac mri using learned local image statistics*. In International Workshop on Statistical Atlases and Computational Models of the Heart, pages 115–124. Springer, 2012. (Cited on page 63.)
- [Mahapatra 2014] Dwarikanath Mahapatra. *Automatic cardiac segmentation using semantic information from random forests*. Journal of digital imaging, vol. 27, no. 6, pages 794–804, 2014. (Cited on page 59.)
- [Makela 2002] Timo Makela, Patrick Clarysse, Outi Sipila, Nicoleta Pauna, Quoc Cuong Pham, Toivo Katila and Isabelle E Magnin. *A review of cardiac image registration methods*. IEEE Transactions on medical imaging, vol. 21, no. 9, pages 1011–1021, 2002. (Cited on page 77.)
- [Mansfield 1977] Peter Mansfield. *Multi-planar image formation using NMR spin echoes*. Journal of Physics C: Solid State Physics, vol. 10, no. 3, page L55, 1977. (Cited on page 14.)
- [Mansi 2011] Tommaso Mansi, Xavier Pennec, Maxime Sermesant, Hervé Delingette and Nicholas Ayache. *iLogDemons: A demons-based registration algorithm for tracking incompressible elastic biological tissues*. International journal of computer vision, vol. 92, no. 1, pages 92–111, 2011. (Cited on page 77.)
- [Marchesseau 2012] Stéphanie Marchesseau, Hervé Delingette, Maxime Sermesant and Nicholas Ayache. *Fast Parameter Calibration of a Cardiac Electromechanical Model from Medical Images based on the Unscented Transform*. Biomechanics and Modeling in Mechanobiology, 2012. (Cited on pages 97 and 98.)
- [Marchesseau 2013] Stéphanie Marchesseau, Hervé Delingette, Maxime Sermesant, R Cabrera-Lozoya, Catalina Tobon-Gomez, Philippe Moireau, RM Figueras i Ventura, Karim Lekadir, Alfredo Hernandez, Mireille Garreau *et al.* *Personalization of a cardiac electromechanical model using reduced order unscented Kalman filtering from regional volumes*. Medical image analysis, vol. 17, no. 7, pages 816–829, 2013. (Cited on page 37.)
- [Marcus 1982] Frank I Marcus, Guy H Fontaine, Gerard Guiraudon, Robert Frank, Jean L Laurenceau, Christine Malergue and Yves Grosogeat. *Right ventricular dysplasia: a report of 24 adult cases*. Circulation, vol. 65, no. 2, pages 384–398, 1982. (Cited on pages 9 and 10.)

- [Margeta 2011] Ján Margeta, Ezequiel Geremia, Antonio Criminisi and Nicholas Ayache. *Layered spatio-temporal forests for left ventricle segmentation from 4D cardiac MRI data*. In International Workshop on Statistical Atlases and Computational Models of the Heart, pages 109–119. Springer, 2011. (Cited on page 59.)
- [McCollough 1999] Cynthia H McCollough and Frank E Zink. *Performance evaluation of a multi-slice CT system*. Medical physics, vol. 26, no. 11, pages 2223–2230, 1999. (Cited on page 17.)
- [McInerney 1996] Tim McInerney and Demetri Terzopoulos. *Deformable models in medical image analysis: a survey*. Medical image analysis, vol. 1, no. 2, pages 91–108, 1996. (Cited on page 59.)
- [McKay 1986] RAYMOND G McKay, MARC A Pfeffer, RICHARD C Pasternak, JOHN E Markis, PATRICIA C Come, Shoichiro Nakao, James D Alderman, JAMES J Ferguson, Robert D Safian and William Grossman. *Left ventricular remodeling after myocardial infarction: a corollary to infarct expansion*. Circulation, vol. 74, no. 4, pages 693–702, 1986. (Cited on page 104.)
- [McLeod 2011] Kristin McLeod, Adityo Prakosa, Tommaso Mansi, Maxime Sermesant and Xavier Pennec. *An incompressible log-domain demons algorithm for tracking heart tissue*. In International Workshop on Statistical Atlases and Computational Models of the Heart, pages 55–67. Springer, 2011. (Cited on page 77.)
- [McLeod 2012] Kristin McLeod, Christof Seiler, Maxime Sermesant and Xavier Pennec. *A near-incompressible poly-affine motion model for cardiac function analysis*. In International Workshop on Statistical Atlases and Computational Models of the Heart, pages 288–297. Springer, 2012. (Cited on pages 78 and 127.)
- [McLeod 2013a] Kristin McLeod. *Reduced-order statistical models of cardiac growth, motion and blood flow: application to the tetralogy of Fallot heart*. PhD thesis, Université Nice Sophia Antipolis, 2013. (Cited on page 126.)
- [McLeod 2013b] Kristin McLeod, Tommaso Mansi, Maxime Sermesant, Giacomo Pongiglione and Xavier Pennec. *Statistical shape analysis of surfaces in medical images applied to the Tetralogy of Fallot heart*. In Modeling in Computational Biology and Biomedicine, pages 165–191. Springer, 2013. (Cited on page 37.)
- [McLeod 2013c] Kristin McLeod, Christof Seiler, Nicolas Toussaint, Maxime Sermesant and Xavier Pennec. *Regional analysis of left ventricle function using a cardiac-specific polyaffine motion model*. In International Conference on Functional Imaging and Modeling of the Heart, pages 483–490. Springer, 2013. (Cited on pages 94, 95 and 127.)

- [McLeod 2015a] Kristin McLeod, Maxime Sermesant, Philipp Beerbaum and Xavier Pennec. *Spatio-temporal tensor decomposition of a polyaffine motion model for a better analysis of pathological left ventricular dynamics*. IEEE transactions on medical imaging, vol. 34, no. 7, pages 1562–1575, 2015. (Cited on pages 76, 104, 107 and 127.)
- [McLeod 2015b] Kristin McLeod, Maxime Sermesant, Philipp Beerbaum and Xavier Pennec. *Spatio-Temporal Tensor Decomposition of a Polyaffine Motion Model for a Better Analysis of Pathological Left Ventricular Dynamics*. IEEE Transactions on Medical Imaging, vol. 34, no. 7, pages 1562–1675, July 2015. (Cited on page 118.)
- [Medrano-Gracia 2013] Pau Medrano-Gracia, Brett R Cowan, David A Bluemke, J Paul Finn, Alan H Kadish, Daniel C Lee, Joao AC Lima, Avan Suinesiaputra and Alistair A Young. *Atlas-based analysis of cardiac shape and function: correction of regional shape bias due to imaging protocol for population studies*. Journal of Cardiovascular Magnetic Resonance, vol. 15, no. 1, page 80, 2013. (Cited on page 106.)
- [Mihl 2008] C Muhl, WRM Dassen and H Kuipers. *Cardiac remodelling: concentric versus eccentric hypertrophy in strength and endurance athletes*. Netherlands Heart Journal, vol. 16, no. 4, pages 129–133, 2008. (Cited on page 104.)
- [Miller 2002] Michael I Miller, Alain Trouvé and Laurent Younes. *On the metrics and Euler-Lagrange equations of computational anatomy*. Annual review of biomedical engineering, vol. 4, no. 1, pages 375–405, 2002. (Cited on page 40.)
- [Miller 2006] Michael I Miller, Alain Trouvé and Laurent Younes. *Geodesic shooting for computational anatomy*. Journal of mathematical imaging and vision, vol. 24, no. 2, pages 209–228, 2006. (Cited on page 40.)
- [Miller 2016] Michael I Miller. *Reducing Variability in Anatomical Definitions Over Time Using Longitudinal Diffeomorphic Mapping*. In Spectral and Shape Analysis in Medical Imaging: First International Workshop, SeSAMI 2016, Held in Conjunction with MICCAI 2016, Athens, Greece, October 21, 2016, Revised Selected Papers, volume 10126, page 51. Springer, 2016. (Cited on page 126.)
- [Montagnat 2005] Johan Montagnat and Hervé Delingette. *4D deformable models with temporal constraints: application to 4D cardiac image segmentation*. Medical Image Analysis, vol. 9, no. 1, pages 87–100, 2005. (Cited on page 59.)
- [Morel 2009] Jean-Michel Morel and Guoshen Yu. *ASIFT: A new framework for fully affine invariant image comparison*. SIAM Journal on Imaging Sciences, vol. 2, no. 2, pages 438–469, 2009. (Cited on page 30.)

- [Morin 2003] Richard L Morin, Thomas C Gerber and Cynthia H McCollough. *Radiation dose in computed tomography of the heart*. *Circulation*, vol. 107, no. 6, pages 917–922, 2003. (Cited on page 17.)
- [Nacif 2012] Marcelo Souto Nacif, Anna Zavodni, Nadine Kawel, Eui-Young Choi, João AC Lima and David A Bluemke. *Cardiac magnetic resonance imaging and its electrocardiographs (ECG): tips and tricks*. *The international journal of cardiovascular imaging*, vol. 28, no. 6, pages 1465–1475, 2012. (Cited on page 15.)
- [Nielsen 1991] PM Nielsen, IJ Le Grice, BH Smaill and PJ Hunter. *Mathematical model of geometry and fibrous structure of the heart*. *American Journal of Physiology-Heart and Circulatory Physiology*, vol. 260, no. 4, pages H1365–H1378, 1991. (Cited on page 106.)
- [Norouzi 2014] Alireza Norouzi, Mohd Shafry Mohd Rahim, Ayman Altameem, Tanzila Saba, Abdolvahab Ehsani Rad, Amjad Rehman and Mueen Uddin. *Medical image segmentation methods, algorithms, and applications*. *IETE Technical Review*, vol. 31, no. 3, pages 199–213, 2014. (Cited on page 2.)
- [Odeblad 1955] Erik Odeblad and Gunnar Lindström. *Some preliminary observations on the proton magnetic resonance in biologic samples*. *Acta Radiologica*, no. 6, pages 469–476, 1955. (Cited on page 14.)
- [Organization 1987] World Health Organization *et al.* *Measuring obesity—Classification and description of anthropometric data. Report on a WHO consultation of the epidemiology of obesity. Warsaw 21-23 October 1987. Copenhagen: WHO, 1989*. Nutrition Unit document, EUR/ICP/NUT, vol. 123, 1987. (Cited on page 11.)
- [Ourselin 2000] Sébastien Ourselin, Alexis Roche, Sylvain Prima and Nicholas Ayache. *Block matching: A general framework to improve robustness of rigid registration of medical images*. In *International Conference on Medical Image Computing And Computer-Assisted Intervention*, pages 557–566. Springer, 2000. (Cited on pages 31 and 62.)
- [Paknezhad 2016] Mahsa Paknezhad, Stephanie Marchesseau and Michael S Brown. *Automatic basal slice detection for cardiac analysis*. *Journal of Medical Imaging*, vol. 3, no. 3, pages 034004–034004, 2016. (Cited on page 58.)
- [Pan 2004] Tinsu Pan, Ting-Yim Lee, Eike Rietzel and George TY Chen. *4D-CT imaging of a volume influenced by respiratory motion on multi-slice CT*. *Medical physics*, vol. 31, no. 2, pages 333–340, 2004. (Cited on page 17.)
- [Park 1996] Jinah Park, Dimitri Metaxas, Alistair A Young and Leon Axel. *Deformable models with parameter functions for cardiac motion analysis from tagged MRI data*. *IEEE Transactions on Medical Imaging*, vol. 15, no. 3, pages 278–289, 1996. (Cited on page 77.)

- [Paul 2004] Asit K Paul and Hani A Nabi. *Gated myocardial perfusion SPECT: basic principles, technical aspects, and clinical applications*. Journal of nuclear medicine technology, vol. 32, no. 4, pages 179–187, 2004. (Cited on pages 15 and 16.)
- [Pavani 2010] Sri-Kaushik Pavani, David Delgado and Alejandro F Frangi. *Haar-like features with optimally weighted rectangles for rapid object detection*. Pattern Recognition, vol. 43, no. 1, pages 160–172, 2010. (Cited on page 63.)
- [Payer 2016] Christian Payer, Darko Štern, Horst Bischof and Martin Urschler. *Regression Heatmaps for Multiple Landmark Localization Using CNNs*. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 230–238. Springer, 2016. (Cited on page 63.)
- [Pearce 2012] Mark S Pearce, Jane A Salotti, Mark P Little, Kieran McHugh, Choonsik Lee, Kwang Pyo Kim, Nicola L Howe, Cecile M Ronckers, Preetha Rajaraman, Alan W Craft *et al.* *Radiation exposure from CT scans in childhood and subsequent risk of leukaemia and brain tumours: a retrospective cohort study*. The Lancet, vol. 380, no. 9840, pages 499–505, 2012. (Cited on page 17.)
- [Pedregosa 2011] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot and E. Duchesnay. *Scikit-learn: Machine Learning in Python*. Journal of Machine Learning Research, vol. 12, pages 2825–2830, 2011. (Cited on page 110.)
- [Penneç 1999] Xavier Penneç, Pascal Cachier and Nicholas Ayache. *Understanding the “demon’s algorithm”: 3D non-rigid registration by gradient descent*. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 597–605. Springer, 1999. (Cited on page 33.)
- [Penneç 2003] Xavier Penneç, Pascal Cachier and Nicholas Ayache. *Tracking brain deformations in time sequences of 3D US images*. Pattern Recognition Letters, vol. 24, no. 4, pages 801–813, 2003. (Cited on page 35.)
- [Penneç 2011] Xavier Penneç, Marco Lorenziet *al.* *Which parallel transport for the statistical analysis of longitudinal deformations*. In Colloque GRETSI, volume 11, page 2. Citeseer, 2011. (Cited on page 28.)
- [Penneç 2015] Xavier Penneç. *Barycentric Subspaces and Affine Spans in Manifolds*. In Geometric Science of Information GSI’2015, Second International Conference, Lecture Notes in Computer Science, Palaiseau, France, 2015. (Cited on pages 21, 76, 81 and 82.)
- [Pennell 2004] Dudley J Pennell, Udo P Sechtem, Charles B Higgins, Warren J Manning, Gerald M Pohost, Frank E Rademakers, Albert C van Rossum,

- Leslee J Shaw and E Kent Yucel. *Clinical indications for cardiovascular magnetic resonance (CMR): Consensus Panel report*. Journal of Cardiovascular Magnetic Resonance, vol. 6, no. 4, pages 727–765, 2004. (Cited on page 15.)
- [Perperidis 2005] Dimitrios Perperidis, Raad H Mohiaddin and Daniel Rueckert. *Spatio-temporal free-form registration of cardiac MR image sequences*. Medical image analysis, vol. 9, no. 5, pages 441–456, 2005. (Cited on page 87.)
- [Petersen 2013] Steffen E Petersen, Paul M Matthews, Fabian Bamberg, David A Bluemke, Jane M Francis, Matthias G Friedrich, Paul Leeson, Eike Nagel, Sven Plein, Frank E Rademakers *et al.* *Imaging in population science: cardiovascular magnetic resonance in 100,000 participants of UK Biobank—rationale, challenges and approaches*. Journal of Cardiovascular Magnetic Resonance, vol. 15, no. 1, page 46, 2013. (Cited on page 144.)
- [Petitjean 2011] Caroline Petitjean and Jean-Nicolas Dacher. *A review of segmentation methods in short axis cardiac MR images*. Medical image analysis, vol. 15, no. 2, pages 169–184, 2011. (Cited on page 59.)
- [Pfister 2015] Tomas Pfister, James Charles and Andrew Zisserman. *Flowing convnets for human pose estimation in videos*. In Proceedings of the IEEE International Conference on Computer Vision, pages 1913–1921, 2015. (Cited on page 63.)
- [Pitiot 2006] Alain Pitiot, Eric Bardinet, Paul M Thompson and Grégoire Mallat. *Piecewise affine registration of biological images for volume reconstruction*. Medical image analysis, vol. 10, no. 3, pages 465–483, 2006. (Cited on page 31.)
- [Poh 2008] Kian Keong Poh, Robert A Levine, Jorge Solis, Liang Shen, Mary Flaherty, Yue-Jian Kang, J Luis Guerrero and Judy Hung. *Assessing aortic valve area in aortic stenosis by continuity equation: a novel approach using real-time three-dimensional echocardiography*. European heart journal, vol. 29, no. 20, pages 2526–2535, 2008. (Cited on page 14.)
- [Postelnicu 2009] Gheorghe Postelnicu, Lilla Zollei and Bruce Fischl. *Combined volumetric and surface registration*. IEEE transactions on medical imaging, vol. 28, no. 4, pages 508–522, 2009. (Cited on page 31.)
- [Prakosa 2013] A. Prakosa and *et al.* *Generation of Synthetic but Visually Realistic Time Series of Cardiac Images Combining a Biophysical Model and Clinical Images*. IEEE Transactions on Medical Imaging, 2013. (Cited on page 88.)
- [Qiu 2008] Anqi Qiu, Laurent Younes, Michael I Miller and John G Csernansky. *Parallel transport in diffeomorphisms distinguishes the time-dependent pattern of hippocampal surface deformation due to healthy aging and the demen-*

- tia of the Alzheimer's type*. NeuroImage, vol. 40, no. 1, pages 68–76, 2008. (Cited on page 127.)
- [Qiu 2012] Anqi Qiu, Laurent Younes and Michael I Miller. *Principal component based diffeomorphic surface mapping*. IEEE transactions on medical imaging, vol. 31, no. 2, pages 302–311, 2012. (Cited on page 80.)
- [Rao 2003] Anil Rao, Gerardo I Sanchez-Ortiz, Raghavendra Chandrashekara, Maria Lorenzo-Valdés, Raad Mohiaddin and Daniel Rueckert. *Construction of a cardiac motion atlas from MR using non-rigid registration*. In International Workshop on Functional Imaging and Modeling of the Heart, pages 141–150. Springer, 2003. (Cited on page 79.)
- [Razavi 2003] Reza Razavi, Derek LG Hill, Stephen F Keevil, Marc E Miquel, Vivek Muthurangu, Sanjeet Hegde, Kawal Rhode, Michael Barnett, Joop van Vaals, David J Hawkes *et al.* *Cardiac catheterisation guided by MRI in children and adults with congenital heart disease*. The Lancet, vol. 362, no. 9399, pages 1877–1882, 2003. (Cited on page 14.)
- [Robb 1942] Jane Sands Robb and Robert Cumming Robb. *The normal heart: anatomy and physiology of the structural units*. American Heart Journal, vol. 23, no. 4, pages 455–467, 1942. (Cited on page 3.)
- [Roche 2001] Alexis Roche, Xavier Pennec, Grégoire Malandain and Nicholas Ayache. *Rigid registration of 3-D ultrasound with MR images: a new approach combining intensity and gradient information*. IEEE transactions on medical imaging, vol. 20, no. 10, pages 1038–1049, 2001. (Cited on page 31.)
- [Rohé 2015] Marc-Michel Rohé, Nicolas Duchateau, Maxime Sermesant and Xavier Pennec. *Combination of Polyaffine Transformations and Supervised Learning for the Automatic Diagnosis of LV Infarct*. In International Workshop on Statistical Atlases and Computational Models of the Heart, pages 190–198. Springer, 2015. (Cited on pages 76 and 103.)
- [Rohé 2016a] Marc-Michel Rohé, Roch Molléro, Maxime Sermesant and Xavier Pennec. *Highly reduced model of the cardiac function for fast simulation*. In Image, Video, and Multidimensional Signal Processing Workshop (IVMSP), 2016 IEEE 12th, pages 1–5. IEEE, 2016. (Cited on page 93.)
- [Rohé 2016b] Marc-Michel Rohé, Maxime Sermesant and Xavier Pennec. *Barycentric Subspace Analysis: a new Symmetric Group-wise Paradigm for Cardiac Motion Tracking*. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 300–307. Springer, 2016. (Cited on pages 76, 82 and 118.)
- [Rohé X] Marc-Michel Rohé, Maxime Sermesant and Xavier Pennec. *Low-Dimensional Representation of Cardiac Motion Using Barycentric Subspaces:*

- a New Group-Wise Paradigm for Estimation, Analysis, and Reconstruction*. Medical image analysis, accepted subject to minor revisions, X. (Cited on page 76.)
- [Rohlfing 2003a] Torsten Rohlfing, Robert Brandt, Randolph Menzel and Calvin R Maurer Jr. *Segmentation of three-dimensional images using non-rigid registration: Methods and validation with application to confocal microscopy images of bee brains*. In Medical Imaging 2003, pages 363–374. International Society for Optics and Photonics, 2003. (Cited on page 61.)
- [Rohlfing 2003b] Torsten Rohlfing, Calvin R Maurer, David A Bluemke and Michael A Jacobs. *Volume-preserving nonrigid registration of MR breast images using free-form deformation with an incompressibility constraint*. IEEE transactions on medical imaging, vol. 22, no. 6, pages 730–741, 2003. (Cited on page 78.)
- [Rohlfing 2004] Torsten Rohlfing, Robert Brandt, Randolph Menzel and Calvin R Maurer. *Evaluation of atlas selection strategies for atlas-based image segmentation with application to confocal microscopy images of bee brains*. NeuroImage, vol. 21, no. 4, pages 1428–1442, 2004. (Cited on pages 61 and 66.)
- [Ronneberger 2015] Olaf Ronneberger, Philipp Fischer and Thomas Brox. *U-net: Convolutional networks for biomedical image segmentation*. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 234–241. Springer, 2015. (Cited on pages 49, 60 and 63.)
- [Ropers 2003] Dieter Ropers, Ulrich Baum, Karsten Pohle, Katharina Anders, Stefan Ulzheimer, Bernd Ohnesorge, Christian Schlundt, Werner Bautz, Werner G Daniel and Stephan Achenbach. *Detection of coronary artery stenoses with thin-slice multi-detector row spiral computed tomography and multiplanar reconstruction*. Circulation, vol. 107, no. 5, pages 664–666, 2003. (Cited on page 17.)
- [Rosipal 2006] Roman Rosipal and Nicole Krämer. *Overview and recent advances in partial least squares*. In Subspace, latent structure and feature selection. 2006. (Cited on pages 97, 110 and 121.)
- [Rougon 2004] Nicolas F Rougon, Caroline Petitjean and Françoise J Preteux. *Building and using a statistical 3D motion atlas for analyzing myocardial contraction in MRI*. In Medical Imaging 2004, pages 253–264. International Society for Optics and Photonics, 2004. (Cited on page 79.)
- [Rueckert 1999] Daniel Rueckert, Luke I Sonoda, Carmel Hayes, Derek LG Hill, Martin O Leach and David J Hawkes. *Nonrigid registration using free-form deformations: application to breast MR images*. IEEE transactions on medical imaging, vol. 18, no. 8, pages 712–721, 1999. (Cited on page 78.)

- [Rueckert 2001] Daniel Rueckert, A Frangi and J Schnabel. *Automatic construction of 3D statistical deformation models using non-rigid registration*. In Medical Image Computing and Computer-Assisted Intervention–MICCAI 2001, pages 77–84. Springer, 2001. (Cited on page 28.)
- [Rueckert 2015] D Rueckert and P Aljabar. *Non-rigid registration using free-form deformations*. In Handbook of Biomedical Imaging, pages 277–294. Springer, 2015. (Cited on page 78.)
- [Ryckelynck 2009] D. Ryckelynck. *Hyper-reduction of mechanical models involving internal variables*. International Journal for Numerical Methods in Engineering, vol. 77, no. 1, pages 75–89, 2009. (Cited on page 95.)
- [Schaerer 2010] Joël Schaerer, Christopher Casta, Jérôme Pousin and Patrick Clarysse. *A dynamic elastic model for segmentation and tracking of the heart in MR image sequences*. Medical Image Analysis, vol. 14, no. 6, pages 738–749, 2010. (Cited on page 78.)
- [Schiller 1989] Nelson B Schiller, Pravin M Shah, Michael Crawford, Anthony De-Maria, Richard Devereux, Harvey Feigenbaum, Howard Gutgesell, Nathaniel Reichek, David Sahn, Ingela Schnittger *et al.* *Recommendations for quantitation of the left ventricle by two-dimensional echocardiography*. Journal of the American Society of Echocardiography, vol. 2, no. 5, pages 358–367, 1989. (Cited on page 13.)
- [Schnabel 2001] Julia Schnabel, Daniel Rueckert, Marcel Quist, Jane Blackall, Andy Castellano-Smith, Thomas Hartkens, Graeme Penney, Walter Hall, Haiying Liu, Charles Truwit *et al.* *A generic framework for non-rigid registration based on non-uniform multi-level free-form deformations*. In Medical image computing and computer-assisted intervention–MICCAI 2001, pages 573–581. Springer, 2001. (Cited on page 78.)
- [Schnabel 2016] Julia A Schnabel, Mattias P Heinrich, Bartłomiej W Papież and J Michael Brady. *Advances and challenges in deformable image registration: From image fusion to complex motion modelling*, 2016. (Cited on page 27.)
- [Schoenhagen 2005] Paul Schoenhagen, Arthur E Stillman, Sandra S Halliburton and Richard D White. *CT of the heart: principles, advances, clinical uses*. Cleve Clin J Med, vol. 72, no. 2, pages 127–138, 2005. (Cited on page 17.)
- [Schulze 2014] D Schulze, M Heiland, H Thurmann and G Adam. *Radiation exposure during midfacial imaging using 4- and 16-slice computed tomography, cone beam computed tomography systems and conventional radiography*. Dentomaxillofacial Radiology, 2014. (Cited on page 17.)
- [Seiler 2012] Christof Seiler, Xavier Pennec and Mauricio Reyes. *Capturing the Multiscale Anatomical Shape Variability with Polyaffine Transformation Trees*.

- Medical Image Analysis, vol. 16, no. 7, pages 1371–1384, 2012. (Cited on page 95.)
- [Serag 2012] Ahmed Serag, Paul Aljabar, Serena Counsell, James Boardman, Jo V Hajnal and Daniel Rueckert. *LISA: Longitudinal image registration via spatio-temporal atlases*. In Biomedical Imaging (ISBI), 2012 9th IEEE International Symposium on, pages 334–337. IEEE, 2012. (Cited on page 28.)
- [Shekhovtsov 2008] Alexander Shekhovtsov, Ivan Kovtun and Václav Hlaváč. *Efficient MRF deformation model for non-rigid image matching*. Computer Vision and Image Understanding, vol. 112, no. 1, pages 91–99, 2008. (Cited on page 32.)
- [Shi 2011] Wenzhe Shi, Xiahai Zhuang, Haiyan Wang, Simon Duckett, Declan Oregan, Philip Edwards, Sebastien Ourselin and Daniel Rueckert. *Automatic segmentation of different pathologies from cardiac cine MRI using registration and multiple component EM estimation*. In International Conference on Functional Imaging and Modeling of the Heart, pages 163–170. Springer, 2011. (Cited on page 63.)
- [Shi 2012] Wenzhe Shi, Xiahai Zhuang, Haiyan Wang, Simon Duckett, Duy VN Luong, Catalina Tobon-Gomez, KaiPin Tung, Philip J Edwards, Kawal S Rhode, Reza S Razavi *et al.* *A comprehensive cardiac motion estimation framework using both untagged and 3-D tagged MR images based on nonrigid registration*. IEEE transactions on medical imaging, vol. 31, no. 6, pages 1263–1275, 2012. (Cited on page 78.)
- [Shotton 2008] Jamie Shotton, Matthew Johnson and Roberto Cipolla. *Semantic texton forests for image categorization and segmentation*. In Computer vision and pattern recognition, 2008. CVPR 2008. IEEE Conference on, pages 1–8. IEEE, 2008. (Cited on page 59.)
- [Simonovsky 2016] Martin Simonovsky, Benjamín Gutiérrez-Becker, Diana Mateus, Nassir Navab and Nikos Komodakis. *A Deep Metric for Multimodal Registration*. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 10–18. Springer, 2016. (Cited on page 31.)
- [Simpson 2015] Ivor JA Simpson, Manuel Jorge Cardoso, Marc Modat, David M Cash, Mark W Woolrich, Jesper LR Andersson, Julia A Schnabel, Sébastien Ourselin, Alzheimer’s Disease Neuroimaging Initiative *et al.* *Probabilistic non-linear registration with spatially adaptive regularisation*. Medical image analysis, vol. 26, no. 1, pages 203–216, 2015. (Cited on page 86.)
- [Škrinjar 2008] Oskar Škrinjar, Arnaud Bistoquet and Hemant Tagare. *Symmetric and transitive registration of image sequences*. Journal of Biomedical Imaging, 2008. (Cited on page 89.)

- [Smith 2002] Stephen M Smith, Yongyue Zhang, Mark Jenkinson, Jacqueline Chen, PM Matthews, Antonio Federico and Nicola De Stefano. *Accurate, robust, and automated longitudinal and cross-sectional brain change analysis*. *Neuroimage*, vol. 17, no. 1, pages 479–489, 2002. (Cited on page 28.)
- [Sotiras 2013] Aristeidis Sotiras, Christos Davatzikos and Nikos Paragios. *Deformable medical image registration: A survey*. *IEEE transactions on medical imaging*, vol. 32, no. 7, pages 1153–1190, 2013. (Cited on pages 27 and 30.)
- [Sudarshan 2013] Vidya Sudarshan, U. R. Acharya, Eddie Yin-Kwee Ng, Chou Siaw Meng, Ru San Tan and Dhanjoo N. Ghista. *Automated Identification of Infarcted Myocardium Tissue Characterisation Using Ultrasound Images: A Review*. *IEEE Reviews in Biomedical Engineering*, 2013. (Cited on page 104.)
- [Suinesiaputra 2014] Avan Suinesiaputra, Brett R Cowan, Ahmed O Al-Agamy, Mustafa A Elattar, Nicholas Ayache, Ahmed S Fahmy, Ayman M Khalifa, Pau Medrano-Gracia, Marie-Pierre Jolly, Alan H Kadish *et al.* *A collaborative resource to build consensus for automated left ventricular segmentation of cardiac MR images*. *Medical image analysis*, vol. 18, no. 1, pages 50–62, 2014. (Cited on page 59.)
- [Suinesiaputra 2015] Avan Suinesiaputra, David A Bluemke, Brett R Cowan, Matthias G Friedrich, Christopher M Kramer, Raymond Kwong, Sven Plein, Jeanette Schulz-Menger, Jos JM Westenberg, Alistair A Young *et al.* *Quantification of LV function and mass by cardiovascular magnetic resonance: multi-center variability and consensus contours*. *Journal of Cardiovascular Magnetic Resonance*, vol. 17, no. 1, page 63, 2015. (Cited on page 57.)
- [Suinesiaputra 2017] Avan Suinesiaputra, Pierre Ablin, Xenia Alba, Martino Alessandrini, Jack Allen, Wenjia Bai, Serkan Cimen, Peter Claes, Brett R Cowan, Jan D’hooge *et al.* *Statistical shape modeling of the left ventricle: myocardial infarct classification challenge*. *IEEE Journal of Biomedical and Health Informatics*, 2017. (Cited on pages 103 and 113.)
- [Sweet 2010] Andrew Sweet and Xavier Pennec. *Log-domain diffeomorphic registration of diffusion tensor images*. In *International Workshop on Biomedical Image Registration*, pages 198–209. Springer, 2010. (Cited on page 80.)
- [Tautz 2011] Lennart Tautz, Anja Hennemuth and Heinz-Otto Peitgen. *Motion analysis with quadrature filter based registration of tagged MRI sequences*. In *International Workshop on Statistical Atlases and Computational Models of the Heart*, pages 78–87. Springer, 2011. (Cited on pages 78 and 85.)
- [Thirion 1998] J-P Thirion. *Image matching as a diffusion process: an analogy with Maxwell’s demons*. *Medical image analysis*, vol. 2, no. 3, pages 243–260, 1998. (Cited on page 33.)

- [Tobon-Gomez 2013] Catalina Tobon-Gomez, Mathieu De Craene, Kristin Mcleod, Lennart Tautz, Wenzhe Shi, Anja Hennemuth, Adityo Prakosa, Hengui Wang, Gerry Carr-White, Stam Kapetanakis *et al.* *Benchmarking framework for myocardial tracking and deformation algorithms: An open access database*. *Medical image analysis*, vol. 17, no. 6, pages 632–648, 2013. (Cited on pages 2, 79, 80, 88, 118 and 130.)
- [Tøndel 2012] Kristin Tøndel, Ulf G Indahl, Arne B Gjuvsland, Stig W Omholt and Harald Martens. *Multi-way metamodelling facilitates insight into the complex input-output maps of nonlinear dynamic models*. *BMC Systems Biology*, vol. 6, no. 1, page 88, 2012. (Cited on page 95.)
- [Topol 2007] Eric J Topol and Robert M Califf. *Textbook of cardiovascular medicine*. Lippincott Williams & Wilkins, 2007. (Cited on page 17.)
- [Townsend 2016] Nick Townsend, Lauren Wilson, Prachi Bhatnagar, Kremlin Wickramasinghe, Mike Rayner and Melanie Nichols. *Cardiovascular disease in Europe: epidemiological update 2016*. *European Heart Journal*, 2016. (Cited on page 3.)
- [Tran 2014] Thanh N. Tran, Nelson Lee Afanador, Lutgarde M.C. Buydens and Lionel Blanchet. *Interpretation of variable importance in Partial Least Squares with Significance Multivariate Correlation (sMC)*. *Chemometrics and Intelligent Laboratory Systems*, vol. 138, pages 153–160, 2014. (Cited on page 98.)
- [Trygg 2002] Johan Trygg and Svante Wold. *Orthogonal projections to latent structures (O-PLS)*. *Journal of chemometrics*, vol. 16, no. 3, pages 119–128, 2002. (Cited on page 111.)
- [TSUIKI 1980] KAI TSUIKI and Erik L Ritman. *Direct evidence that left ventricular myocardium is incompressible throughout systole and diastole*. *The Tohoku journal of experimental medicine*, vol. 132, no. 1, pages 119–120, 1980. (Cited on page 77.)
- [Üzümcü 2006] Mehmet Üzümcü, Rob J van der Geest, Cory Swingen, Johan HC Reiber and Boudewijn PF Lelieveldt. *Time continuous tracking and segmentation of cardiovascular magnetic resonance images using multidimensional dynamic programming*. *Investigative radiology*, vol. 41, no. 1, pages 52–62, 2006. (Cited on page 59.)
- [Vaillant 2004] Marc Vaillant, Michael I Miller, Laurent Younes and Alain Trouvé. *Statistics on diffeomorphisms via tangent space representations*. *NeuroImage*, vol. 23, pages S161–S169, 2004. (Cited on page 80.)
- [Vaillant 2005] Marc Vaillant and Joan Glaunès. *Surface matching via currents*. In *Biennial International Conference on Information Processing in Medical Imaging*, pages 381–392. Springer, 2005. (Cited on pages 38 and 47.)

- [Van der Geest 1997] Rob J Van der Geest, Vincent GM Buller, Eric Jansen, Hildo J Lamb, Leo HB Baur, Ernst E van der Wall, Albert de Roos and Johan HC Reiber. *Comparison between manual and semiautomated analysis of left ventricular volume parameters from short-axis MR images*. Journal of computer assisted tomography, vol. 21, no. 5, pages 756–765, 1997. (Cited on page 58.)
- [Van Kaick 2011] Oliver Van Kaick, Hao Zhang, Ghassan Hamarneh and Daniel Cohen-Or. *A survey on shape correspondence*. In Computer Graphics Forum, volume 30, pages 1681–1707. Wiley Online Library, 2011. (Cited on page 28.)
- [Vercauteren 2008] Tom Vercauteren, Xavier Pennec, Aymeric Perchant and Nicholas Ayache. *Symmetric log-domain diffeomorphic registration: A demons-based approach*. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 754–761. Springer, 2008. (Cited on pages 33, 34, 82 and 83.)
- [Viola 1997] Paul Viola and William M Wells III. *Alignment by maximization of mutual information*. International journal of computer vision, vol. 24, no. 2, pages 137–154, 1997. (Cited on page 30.)
- [Wang 2012] Hui Wang and Amir A Amini. *Cardiac motion and deformation recovery from MRI: a review*. IEEE Transactions on Medical Imaging, vol. 31, no. 2, pages 487–503, 2012. (Cited on page 77.)
- [Wang 2013] Yang Wang, Bogdan Georgescu, Terrence Chen, Wen Wu, Peng Wang, Xiaoguang Lu, Razvan Ionasec, Yefeng Zheng and Dorin Comaniciu. *Learning-based detection and tracking in medical imaging: a probabilistic approach*. In Deformation Models, pages 209–235. Springer, 2013. (Cited on page 127.)
- [Wang 2015] Liang Wang, Adrian Basarab, Patrick R Girard, Pierre Croisille, Patrick Clarysse and Philippe Delachartre. *Analytic signal phase-based myocardial motion estimation in tagged MRI sequences by a bilinear model and motion compensation*. Medical image analysis, vol. 24, no. 1, pages 149–162, 2015. (Cited on page 78.)
- [Wenk 2012] Jonathan F Wenk, Doron Klepach, Lik Chuan Lee, Zhihong Zhang, Liang Ge, Elaine E Tseng, Alastair Martin, Sebastian Kozerke, Joseph H Gorman, Robert C Gorman *et al.* *First evidence of depressed contractility in the border zone of a human myocardial infarction*. The Annals of thoracic surgery, vol. 93, no. 4, pages 1188–1193, 2012. (Cited on page 77.)
- [Wold 2001] Svante Wold, Michael Sjöström and Lennart Eriksson. *PLS-regression: a basic tool of chemometrics*. Chemometrics and intelligent laboratory systems, vol. 58, no. 2, pages 109–130, 2001. (Cited on page 110.)

- [Wolz 2013] Robin Wolz, Chengwen Chu, Kazunari Misawa, Michitaka Fujiwara, Kensaku Mori and Daniel Rueckert. *Automated abdominal multi-organ segmentation with subject-specific atlas generation*. IEEE transactions on medical imaging, vol. 32, no. 9, pages 1723–1730, 2013. (Cited on page 67.)
- [Xu 2000] Chenyang Xu, Dzung L Pham and Jerry L Prince. *Image segmentation using deformable models*. Handbook of medical imaging, vol. 2, pages 129–174, 2000. (Cited on page 59.)
- [Yang 2016] Xiao Yang, Roland Kwitt and Marc Niethammer. *Fast Predictive Image Registration*. In International Workshop on Large-Scale Annotation of Biomedical Data and Expert Label Synthesis, pages 48–57. Springer, 2016. (Cited on pages 20, 44, 46 and 49.)
- [Yerry 1984] Mark A Yerry and Mark S Shephard. *Automatic three-dimensional mesh generation by the modified-octree technique*. International Journal for Numerical Methods in Engineering, vol. 20, no. 11, pages 1965–1990, 1984. (Cited on page 58.)
- [Yin 2010] Youbing Yin, Eric A Hoffman, Kai Ding, Joseph M Reinhardt and Ching-Long Lin. *A cubic B-spline-based hybrid registration of lung CT images for a dynamic airway geometric model with large deformation*. Physics in medicine and biology, vol. 56, no. 1, page 203, 2010. (Cited on page 31.)
- [Young 1999] Alistair A Young. *Model tags: direct three-dimensional tracking of heart wall motion from tagged magnetic resonance images*. Medical image analysis, vol. 3, no. 4, pages 361–372, 1999. (Cited on page 77.)
- [Young 2006] Alistair A Young. *Assessment of cardiac performance with magnetic resonance imaging*. Current Cardiology Reviews, vol. 2, no. 4, pages 271–282, 2006. (Cited on page 76.)
- [Yousef 2000] Zaheer Yousef, Simon Robert Redwood and Michael Stephen Marber. *Postinfarction left ventricular remodeling: a pathophysiological and therapeutic review*. Cardiovascular drugs and therapy, vol. 14, no. 3, pages 243–252, 2000. (Cited on page 104.)
- [Zerhouni 1988] Elias A Zerhouni, David M Parish, Walter J Rogers, Andrew Yang and Edward P Shapiro. *Human heart: tagging with MR imaging—a method for noninvasive assessment of myocardial motion*. Radiology, vol. 169, no. 1, pages 59–63, 1988. (Cited on page 2.)
- [Zhang 2004] Dao-Qiang Zhang and Song-Can Chen. *A novel kernelized fuzzy c-means algorithm with application in medical image segmentation*. Artificial intelligence in medicine, vol. 32, no. 1, pages 37–50, 2004. (Cited on page 2.)
- [Zhang 2007] Weiwei Zhang, J Noble and J Brady. *Spatio-temporal registration of real time 3D ultrasound to cardiovascular MR sequences*. Medical Image

- Computing and Computer-Assisted Intervention–MICCAI 2007, pages 343–350, 2007. (Cited on page 127.)
- [Zhang 2015] Xingyu Zhang, Bharath Ambale-Venkatesh, DavidA. Bluemke, BrettR. Cowan, J.Paul Finn, WilliamG. Hundley, AlanH. Kadish, DanielC. Lee, JoaoA.C. Lima, Avan Suinesiaputra, AlistairA. Young and Pau Medrano-Gracia. *Orthogonal Shape Modes Describing Clinical Indices of Remodeling*. In Functional Imaging and Modeling of the Heart. 2015. (Cited on page 106.)
- [Zhang 2016] Miaomiao Zhang, William M Wells III and Polina Golland. *Low-Dimensional Statistics of Anatomical Variability via Compact Representation of Image Deformations*. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 166–173. Springer, 2016. (Cited on page 80.)
- [Zheng 2008] Yefeng Zheng, Adrian Barbu, Bogdan Georgescu, Michael Scheuering and Dorin Comaniciu. *Four-chamber heart modeling and automatic segmentation for 3-D cardiac CT volumes using marginal space learning and steerable features*. IEEE transactions on medical imaging, vol. 27, no. 11, pages 1668–1681, 2008. (Cited on page 2.)
- [Zhu 2010] Yun Zhu, Xenophon Papademetris, Albert J Sinusas and James S Duncan. *A coupled deformable model for tracking myocardial borders from real-time echocardiography using an incompressibility constraint*. Medical Image Analysis, vol. 14, no. 3, pages 429–448, 2010. (Cited on page 77.)
- [Zitova 2003] Barbara Zitova and Jan Flusser. *Image registration methods: a survey*. Image and vision computing, vol. 21, no. 11, pages 977–1000, 2003. (Cited on page 28.)