

École doctorale n°84:
Sciences et technologies de l'information et de la communication

Doctorat européen ParisTech

T H È S E

pour obtenir le grade de docteur délivré par

l'École nationale supérieure des mines de Paris
Spécialité "Informatique temps-réel, robotique et automatique"

présentée et soutenue publiquement par

Tommaso MANSI

10 Septembre 2010

Modèles physiologiques et statistiques du cœur guidés par imagerie médicale

– Application à la tétralogie de Fallot –

~ ~ ~

Image-Based Physiological and Statistical Models of the Heart

– Application to Tetralogy of Fallot –

Directeur de thèse: **Nicholas AYACHE**

Co-encadrement de la thèse: **Xavier PENNEC, Maxime SERMESANT, Hervé DELINGETTE**

Thèse préparée à l'INRIA Sophia-Antipolis, projet Asclepios.

Jury

Dorin COMANICIU, Siemens Corporate Research
Peter HUNTER, Auckland Bioengineering Institute
Daniel RUECKERT, Imperial College London
Younes BOUDJEMLINE, AP-HP, Necker-Enfants Malades
Giacomo PONGIGLIONE, Ospedale Pediatrico Bambino Gesù
Nicholas AYACHE, INRIA Sophia Antipolis
Xavier PENNEC, INRIA Sophia Antipolis
Maxime SERMESANT, INRIA Sophia Antipolis
Hervé DELINGETTE, INRIA Sophia Antipolis
Andrew TAYLOR, Great Ormond Street Hospital

Rapporteur
Rapporteur
Rapporteur
Examineur
Examineur
Directeur
Co-directeur
Co-directeur
Invité
Invité

**T
H
È
S
E**

MINES ParisTech

Centre de Mathématiques Appliquées

Rue Claude Daunesse B.P. 207, 06904 Sophia Antipolis Cedex, France

Image-Based Physiological and Statistical Models of the Heart Application to Tetralogy of Fallot

Abstract: This thesis presents image-based methods for diagnosis, prognosis and therapy planning of patients with repaired tetralogy of Fallot, a severe congenital heart disease. This was achieved by combining advanced medical image processing with both statistical and physiological modelling.

First, we proposed a demons-based image registration algorithm to estimate the 3D myocardium strain from routine medical images, which is challenging due to the aperture problem (only the apparent motion is visible). The algorithm relies on elasticity and incompressibility constraints, rigorously implemented thanks to a new justification of the demons regularisation. Experiments on synthetic data and cine MRI of patients demonstrated that the proposed constraints improve the accuracy of the estimated cardiac displacements.

Second, a generative model of the pathological right ventricle growth was proposed using a well-posed statistical framework for shape analysis, based on the “currents” shape representation. Principal component analysis was used to identify shape features that are relevant to the disease. Partial least squares and canonical correlation analysis were employed to design a generative model of heart growth. Applied on the right ventricles of 32 patients with tetralogy of Fallot, the method identified the right ventricle dilation, basal bulging and apical dilation reported in the clinical literature. The model showed that these features appear progressively as the child grows.

Finally, we introduced an electromechanical model of the heart for personalised simulations of valve replacement in repaired tetralogy of Fallot patients. The electromechanical model simulates the main features of the cardiac function observed in these patients. Once personalised from clinical data, the model was used to predict the effects of pulmonary valve replacement on the postoperative cardiac function. Tested on two patients, the model managed to qualitatively reproduce their cardiac function. As expected, valve replacement predicted a significant improvement of the right ventricle but also of the left ventricle, suggesting a tight inter-ventricular relationship.

The combination of medical imaging, statistical analysis and biophysical modelling provides a powerful framework for a more personalised computer-aided medicine

Keywords: Medical Imaging; Image Registration; Myocardium Strain; Statistical Shape Analysis; Cardiac Remodelling; Electromechanical Model; Personalised Therapy Simulation; Tetralogy of Fallot



Modèles physiologiques et statistiques du cœur guidés par imagerie médicale. Application à la tétralogie de Fallot

Résumé: Les travaux de cette thèse sont consacrés à la quantification de cardiopathies, la prédiction de leur évolution et la planification de thérapies, avec pour application principale la tétralogie de Fallot, une malformation congénitale grave du cœur. L'idée sous-jacente est d'utiliser des modèles informatiques sophistiqués combinant traitement d'images, statistique et physiologie, pour assister la gestion clinique de ces patients.

Dans un premier temps, nous proposons une nouvelle méthode de recalage d'images plus précise pour estimer la déformation cardiaque à partir d'images médicales anatomiques, où seulement le mouvement apparent du cœur est visible. L'algorithme proposé s'appuie sur la méthode dite des démons, que l'on contraint de manière rigoureuse à être élastique et incompressible grâce à une nouvelle justification de l'étape de régularisation de l'algorithme. Les expériences réalisées sur des images synthétiques et réelles ont démontré que l'ajout de ces contraintes améliore de manière significative la précision des déformations estimées.

Nous étudions ensuite la croissance du cœur par une approche statistique basée sur les "courants". Une analyse en composantes principales permet d'identifier des altérations morphologiques dues à la pathologie. L'utilisation conjointe de la méthode PLS (moindres carrés partiels) et de l'analyse des corrélations canoniques permet de créer un modèle statistique moyen de croissance du cœur. L'analyse du ventricule droit de 32 patients avec tétralogie de Fallot a révélé une dilatation du ventricule, une déformation de sa base et un élargissement de son apex, caractéristiques que l'on retrouve dans la littérature clinique. Le modèle de croissance montre qu'elles apparaissent progressivement au fil du temps.

Enfin, nous adaptons un modèle électromécanique du cœur pour simuler la fonction cardiaque chez des patients et tester diverses stratégies de pose de valves pulmonaires. Le modèle électromécanique simule les caractéristiques principales de la fonction cardiaque. Une fois personnalisé, le modèle est utilisé pour prédire les effets postopératoires de la pose de valves chez le patient. Le modèle a été ainsi capable de reproduire, de manière qualitative, la fonction cardiaque de deux patients. Comme attendu, la fonction simulée du ventricule droit est améliorée après pose de valves, ainsi que celle du ventricule gauche, suggérant une relation étroite entre les deux ventricules du cœur.

Les méthodes de traitement d'images médicales, d'analyses statistiques et de modèles de la physiologie du cœur forment un cadre puissant pour le développement d'une médecine plus personnalisée et assistée par ordinateur.

Mots clés: Imagerie médicale; recalage d'images; déformation du myocarde; analyse statistique de formes; remodelage du cœur; modèle électromécanique; simulation personnalisée de thérapies; tétralogie de Fallot

"Essentially, all models are wrong, but some are useful", G. Box

Acknowledgements

First of all, I would like to warmly thank my advisor, Nicholas Ayache; my co-advisors, Xavier Pennec and Maxime Sermesant; and Hervé Delingette. They have guided me all along my Ph.D. period, always available to answer my questions and to help me in moments of doubt during my research work. I have learned a lot from their deep scientific knowledge, from their great faculty of presenting scientific results, from their vision of the future. I also thank them for sending me all around the world, giving me the opportunity to present my work and to build fruitful external collaborations. Thank you for making me the young scientist I am today.

I am extremely grateful to the reviewers, Dr. Dorin Comaniciu, Prof. Peter Hunter and Prof. Daniel Rueckert for having spent their precious time to read and correct this long manuscript, and for being part of my jury. I warmly thank them for their sharp and constructive comments about my work and for their encouraging compliments. I am also thankful to Prof. Giacomo Pongiglione, who accepted to be a member of my jury and came at my defense despite his clinical commitments. I thank him for his support all along the Health-e-Child project and his wise comments about the clinical implications of my research work. Dear committee, thank you, it has been a great honour for me to have you in my jury.

I owe a great amount of gratitude to all the people I had the chance to closely work with: Jean-Marc Peyrat and Marco Lorenzi, with who I had the chance to have devil discussions about demons; Stanley Durrleman, for showing me, and convincing me, that currents are powerful; Florence Billet, for our discussions about models, Ph.D. and research in general; Nicolas Toussaint, for his precious help about coding; Barbara André and Erik Pernod, for their help with SOFA; Michael Lynch and Ingmar Voigt, for being always available and kind despite my endless requests of heart segmentations; Mihaela Pop, for involving me in her exciting “gips” experiments; and Boris Bernhardt, for introducing me to the statistics.

I am also very thankful to all the people who went along with me during these years. In particular, I would like to thank Ender Konukoglu, Kristin Mcleod, Olivier Clatz, Francois Chung, Pierre Fillard, Tom Vercauteren, Olivier Colliot, Jatin Relan, Ken Wong, Florence Dru, Vincent Garcia, Aurélie Canale, Daniel Barbeau, Adityo Prakosa, Grégoire Malandain, Phani Chinchapatnam, Radomir Chabiniok, Pascal Girard, and all the Asclepios team for the warm welcome and the good time we spent together. I am also extremely grateful to Isabelle Strobant, for being always here when needed, INRIA’s Bureau des Missions, for their kind help when organising the most improbable flights around Europe, and Vincent Gomis and Franck Yampolski for helping me in setting up the videoconferencing system for the defense.

I should not forget to thank all the members of the Health-e-Child project. This Ph.D. and my life after it are also the results of the work we have done together. I thank Joerg Freund, Martin Huber, David Manset and Jérôme Revillard for the good time we spent together during the numerous Health-e-Child meetings, from the setting up of successful demos to the realisation of stunning presentations to the after-work activities everywhere in Europe. My heartfelt thanks to the Health-

e-Child cardiologists, Younes Boudjemline, Julie Blanc, Bertrand Stos, Andrew M. Taylor, Gianluca Trocchio, Giacomo Pongiglione and Benedetta Leonardi, for the numerous discussions about the clinical motivations of this research work. And of course, I should not forget to thank Edwin Fletcher Morley, Harry Dimitropoulos, Tamás Hauer, Yannick Legré, Mikk Tamme, Alexey Tsymbal, Alessandro Verri, and all the Health-e-Childers!

Finally, I would like to warmly thank Neda and Andrea Bernasconi, from the Neuroimaging of Epilepsy Laboratory, Montreal Neurological Institute. Six years ago, they transmitted me the medical imaging “virus”. Four years ago, they convinced me to do a Ph.D. Two years ago, they pushed me to follow my ideas. Everyday, they support me, guide me along my research work and personal life. And above all, they made me realise that behind any algorithms, there is a patient who is waiting for a treatment; that successful methods are those that can help patients, help clinicians, open new clinical strategies, and raise thousands of new questions.

At last but not least, I thank my family for their continuous influence and support in my life, otherwise I would not be here today. They have always guided me in the different, often difficult, choices I had to make. And I thank Marta, for being with me, supporting me, understanding me, despite the tiring, difficult, long-distance relationship. Thank you for making me the man I am today.

Thank you!

Contents

I	INTRODUCTION	1
1	Introduction	3
1.1	Clinical Context	3
1.2	Problems Investigated	5
1.3	Organisation of the Thesis and Main Contributions	6
2	From a Normal Heart to Tetralogy of Fallot	11
2.1	Structure and Function of the Normal Heart	11
2.2	Imaging the Heart	19
2.3	Tetralogy of Fallot: Physiopathology and Clinical Challenges	28
3	Semi-Automatic Delineation of the Heart	33
3.1	Motivation	33
3.2	A Pipeline for the Dynamic Segmentation of the Myocardium	35
3.3	Results on Tetralogy of Fallot	51
3.4	Discussion	55
II	DIAGNOSIS: QUANTIFICATION OF CARDIAC DEFORMATION ON ANATOMICAL IMAGES	59
4	Efficient Elastic Incompressible Image Registration	61
4.1	Motivation	62
4.2	Background: Log-Domain Diffeomorphic Demons	65
4.3	From Gaussian Regularisation to Elastic-Like Regularisation	67
4.4	Incompressible LogDemons	70
4.5	Experiments and Results	73
4.6	Discussion	80
4.7	Appendices	83
5	Estimation of 3D Myocardium Strain from Cine MRI	89
5.1	Motivation	89
5.2	Cardiac Tracking Algorithm and Strain Calculation	91
5.3	Experiments and Results	93
5.4	Discussion	102
III	PROGNOSIS: IDENTIFICATION AND MODELLING OF PATHO- LOGICAL HEART GROWTH	105
6	Statistical Modelling of the Heart Shape and its Growth	107

6.1	Motivation	108
6.2	Unbiased Template of Shapes Using Currents	109
6.3	Identification of Pathological Heart Shape Patterns	118
6.4	Estimation of a Generative Model of Heart Growth	125
7	Identification and Modelling of RV Growth in ToF	133
7.1	Clinical Motivation	134
7.2	Patient Selection and Data Preparation	134
7.3	Exploring RV Shape Patterns Related to Regurgitations	137
7.4	Generative Statistical Model of the RV Growth in ToF	141
7.5	Discussion	150
IV	THERAPY: PERSONALISED SIMULATION OF CARDIAC THERAPIES	153
8	Personalised Model of Cardiac Electromechanics	155
8.1	Motivation	156
8.2	Cardiac Anatomy Model	157
8.3	Cardiac Electrophysiology Model	166
8.4	Cardiac Biomechanics Model	169
8.5	Cardiac Hemodynamics Model	178
8.6	Discussion	187
9	<i>In-Silico</i> Pulmonary Valve Replacement in ToF	193
9.1	Motivation	193
9.2	Real-Time Simulation of Cardiac Volume Reduction	194
9.3	Personalised Simulation of Pulmonary Valve Replacement	197
9.4	Discussion	203
V	CONCLUSIONS	211
10	Conclusions and Perspectives	213
10.1	Summary of the Contributions	213
10.2	Perspectives	219
VI	Appendices	229
A	On the Image Noise Parameter of the LogDemons	231
A.1	Motivation	231
A.2	Methods	231
A.3	Preliminary Experiments and Discussions	234
A.4	Conclusions	235

Part I

INTRODUCTION

Introduction

Contents

1.1 Clinical Context	3
1.2 Problems Investigated	5
1.3 Organisation of the Thesis and Main Contributions	6

1.1 Clinical Context

Diagnosis, prognosis, therapy planning. Three steps of the clinical practice that are seriously challenged by chronic diseases. Congenital heart defects are such pathologies. They involve severe anatomical and functional abnormalities of the heart that often require surgical repair early in infancy. In some cases however, that repair yields sequels that, although initially asymptomatic, can become life-threatening decades after the intervention. The very large time-scale of cardiac failure greatly challenges the management of these patients. In particular, the diversity in child growth results in a wide variety of pathology evolutions. Repeated and long-term follow-ups are therefore necessary. A question then emerges in light of these difficulties: can computational models of the heart assist the cardiologists in the management of these patients?

The last decades have seen tremendous advances in the processing of medical images, the analysis of organ shapes and the physiological modelling of the cardiac function. Medical image analysis is providing more and more efficient and non-invasive algorithms to quantify the cardiac shape and function from medical images. In parallel, recent well-posed mathematical theories enable one to perform statistics on 3D shapes. Finally, intense research is aiming to develop detailed models of cardiac function to understand its mechanisms.

It is easy to imagine how such tools can improve the clinical workflow. On the one hand, they can help in interpreting clinical data, extracting quantitative and meaningful parameters otherwise impossible to reach. Diagnosis would be improved. On the other hand, clinical data would enable to personalise these models to the patient physiology. It would then be possible to predict the evolution of the pathology and test therapies. Nonetheless, translating these tools in the clinics is remarkably challenging due to the sparse clinical data available for a patient, especially in paediatrics. Three strategies can be distinguished. For instance, one can use models of cardiac physiology to enhance the accuracy of purely image-based algorithms.

Conversely, detailed biomechanical models of the heart can be simplified such that they can be personalised from clinical data. Finally, when no direct models can be designed due to the complexity of the studied biological phenomena, statistical approaches can be used to explore their features.

A large part of this thesis was performed within the Health-e-Child¹ project, a four-year European project partially funded by the European Community (6th Framework Program). The project was coordinated by Siemens AG, Erlangen (Germany), and involved eleven technical partners (including INRIA) and four clinical partners, namely Ospedale Pediatrico Bambino Gesù (Rome, Italy), Hôpitaux de Paris, Necker-Enfants Malades (Paris, France); Great Ormond Street Hospital (London, U.K.) and Ospedale Pediatrico Gaslini (Genoa, Italy). Health-e-Child goal was to develop a biomedical Grid platform that connects European paediatric centres and provides them with technologies that integrate the heterogeneous clinical data for data sharing, processing and knowledge discovery about paediatric diseases [Freund et al., 2006] (Figure 1.1). Health-e-Child technical developments were motivated by three paediatric diseases: brain tumours, juvenile idiopathic arthritis and congenital heart diseases. During the four years of the project, a large multi-centric database of patients was created. Nonetheless, *only standard clinical data were gathered*. In cardiology for instance, the data acquired were mainly clinical features, echocardiography evaluations, standard short-axis cine Magnetic Resonance Images (MRI) and ECG. These restrictions motivated our work.




Figure 1.1: Diagram of the Health-e-Child grid infrastructure for connecting paediatrics centres. The grid provides clinicians with integrative tools for data sharing, processing and knowledge discovery (*Image from* [Freund et al., 2006]).

¹IST-2004-027749, <http://www.health-e-child.org/>

Among the investigated congenital heart diseases, repaired tetralogy of Fallot (ToF) focused most of the work. ToF is a severe congenital heart defect of the right ventricle (RV) that requires surgical repair early in infancy. Blood leakage through the pulmonary valve subsequent to that initial repair makes the ventricle dilate in the long term, which may yield life-threatening functional abnormalities decades after repair. Nowadays, it is acknowledged that replacing the valves significantly increases life expectancy in these patients.

A major difficulty in ToF management is to determine the best timing for valve replacement. To date, no clear-cut clinical criteria are available. Metrics that quantify the intrinsic cardiac deformation could provide a more comprehensive view of the myocardium integrity. Unfortunately, this task is challenging in paediatrics as only standard anatomical images of the beating heart are usually available. *How can we estimate cardiac function from standard anatomical images?* In fact, one would want to predict if the heart will deteriorate in the near future. Yet, how the heart evolves in ToF is far from being understood. *How can we quantify heart growth in ToF?* Finally, several strategies are available for valve replacement and their effects can vary from one patient to another. *Can we predict the effects of pulmonary valve replacement therapies on a patient?* All these questions motivated the developments of this thesis, providing a unique and stimulating clinical context to the methodological contributions.

1.2 Problems Investigated

The clinical context of this thesis, centred on tetralogy of Fallot (but not only), led us to consider three specific questions on how advanced computational models of the heart could improve the clinical workflow, namely diagnosis, prognosis and therapy planning.

Q1. Diagnosis

– Quantification of the Cardiac Deformations from Anatomical Images –

Assessing the cardiac function in ToF can provide the cardiologist with crucial information to determine the best timing for re-intervention. However, in paediatrics, only standard clinical images are usually acquired. In Health-e-Child for instance, only cine MRI (cMRI) were available for processing. The question we tackle is therefore: *Can we estimate the myocardial deformation from these images?* A standard approach to estimate deformations from dynamic medical images is to use non-linear image registration to track the moving structures. However, the lack of myocardial texture in cMRI and the coarse slice thickness seriously hinder the registration algorithm. The idea is thus to constrain the registration with physiological priors about the heart constitution: provided deformations are elastic and incompressible. From a methodological point of view the question thus becomes:

- *How can we efficiently register two images such that estimated deformations are elastic and incompressible?*

Q2. Prognosis**– Identification and Prediction of Pathological Features of the Heart Shape –**

Evaluating the present cardiac condition is primordial, but predicting its evolution is decisive. This is all the more important in ToF as predicting whether the heart will fail in the near future will considerably support the decision for immediate valve replacement. Yet, cardiac remodelling is a very complex phenomenon. The intricacy and the large time scale of the phenomenon hampers the development of direct biophysical models, which would require extremely detailed methods to simulate the biological mechanisms that are involved. Recently, new mathematical methods for modelling populations of shapes are laying the foundations of well-posed statistical analyses of 3D shapes. As a result, it becomes possible to statistically relate shapes to external variables. Instead of directly modelling the pathological mechanisms, one can statistically explore the common factors between clinical indices that quantify the severity of a pathology and observed shapes. The second question we ask is therefore:

- *Can we statistically relate the 3D shape of the heart to clinical features that quantify the severity of a pathology in order to identify pathological shape patterns and model their evolution over time?*

Q3. Therapy Planning**– Personalised Simulation of Cardiac Therapies –**

Finally, at the time of treatment, several therapeutic options are usually available to the clinician. In ToF, two pulmonary valve replacement strategies are competing. However, what are their effect on the cardiac function of a specific patient is difficult to predict. *Can models help in choosing the optimal therapy for a specific patient?* The recent advances in biomechanical modelling of the heart promise *in-silico* simulations of cardiac function. The idea is to adjust a virtual heart to the physiology of a patient to test therapies, thus helping the cardiologist in deciding the optimal strategy. As we seek personalised simulations, the model must be controlled by parameters that can be estimated from clinical data. The last question we tackle is thus:

- *Can we simulate the effects of a therapy on the cardiac function of a patient by using electromechanical models of the heart?*

1.3 Organisation of the Thesis and Main Contributions

This thesis is organised around the three questions listed in the previous section. We first start with forewords about the heart and tetralogy of Fallot. We also introduce a multi-purpose segmentation pipeline that is required for the methodological contributions of this thesis. We then detail the methods to tackle the three questions addressed in this manuscript. It is worth stressing that although our contributions

are motivated by tetralogy of Fallot, they are general enough to be employed for other cardiac pathologies. For that reason, the methodological contributions are purposely separated from their applications on tetralogy of Fallot. The detailed description of the material covered by each chapter is given below.

Chapter 2 introduces the heart, its anatomy and how it works. After a short presentation of the different anatomical components that constitute the heart, we briefly describe the different phases of the cardiac cycle and the underlying electrical and biomechanical mechanisms that govern it. This chapter also outlines the different imaging modalities that are available to describe the heart anatomy, constitution and function. Finally, we present the tetralogy of Fallot, its physiopathology and the clinical challenges related to it.

Chapter 3 presents an integrated pipeline for the segmentation of the beating heart in medical images. Although cardiac segmentation is not the main topic of this thesis, this stage is crucial for our analyses and modelling. In this chapter **we propose a modular solution to the cardiac segmentation problem that is suitable for any imaging modality and can cope with the extreme anatomical variability observed in ToF patients**. In simple words, the cardiac borders are first interactively delineated using a 3D surface modeller based on variational implicit surfaces. The resulting surfaces are then automatically propagated throughout the cardiac sequence using non-linear image registration and simplex deformable models. The modularity of the pipeline enables to integrate results from other methods, like those computed by other Health-e-Child partners.

Part I – Diagnosis

Quantification of the Cardiac Deformation on Anatomical Images

Chapter 4 details our approach to the problem of estimating cardiac deformation from dynamic anatomical images. We start from the log-domain demons, an efficient non-linear image registration algorithm [Vercauteren et al., 2008] that computes diffeomorphic deformations parameterised by stationary velocity fields. First, **we propose a mathematical justification to the *ad-hoc* Gaussian regularisation of demons algorithm**. The proposed formulation is based on Tikhonov regularisation and enables us to integrate more complex regularisation schemes. With our framework, **we show how elasticity and incompressibility can be integrated rigorously in the algorithm**. We replace the Gaussian smoothing by an efficient elastic-like vector filter regularisation. We compute incompressible deformation fields by minimising the registration energy in the space of divergence-free velocity fields. The proposed algorithm, called *iLogDemons*, is compared with the original log-domain demons on synthetic images deformed by known incompressible transformations. Results demonstrate that *iLogDemons* do improve the accuracy of the estimated transformations.

Chapter 5 illustrates how iLogDemons can be used to assess the cardiac deformation from the cMRI of real patients. As the heart is a near-incompressible elastic tissue, **we investigate whether the proposed incompressibility and elasticity constraints can improve the estimation of the cardiac deformations.** With a recursive tracking algorithm, we estimate the 3D cardiac deformation in two adult patients with heart failure and one child with repaired ToF. Computed displacements and strains are compared with tagged MRI and 2D-strain echocardiography. The results confirm that iLogDemons provides more reliable results than the original log-domain demons thanks to the elastic and incompressibility model.

Part II – Prognosis

Identification and Modelling of Pathological Heart Growth

Chapter 6 proposes a statistical strategy to identify and predict pathological patterns of heart shape. The idea is to relate the heart shapes observed in a population of patients to clinical features in order to identify the morphological patterns that are relevant to the disease and model the heart growth. This approach is made possible thanks to the recent advances in computational anatomy and in particular to the method based on “currents” proposed by [Durrleman et al., 2009a], which enables statistics on shapes. **To identify pathological shape patterns,** we propose to represent the heart shape by a low-dimensional shape vector computed using principal component analysis (PCA). **We then relate the shape vectors to clinical features to identify the modes that are relevant to the disease. To model heart growth,** we instead propose to perform a cross-sectional analysis between the shapes and the patient body surface area or age. The novelty is to **use the Partial Least Squares (PLS) space decomposition to reduce the model dimensionality and to increase statistical significance.** Contrary to PCA, PLS automatically extracts the modes that are relevant to two sets of variables. Canonical correlation analysis (CCA) then provides the generative model of heart growth.

Chapter 7 reports the results of the statistical analyses of the right ventricle in a population of 32 patients with repaired ToF. **We first investigated the link between right ventricle shape and pulmonary regurgitation, an important pathological feature of these patients.** The statistical analyses with respect to three different metrics of pulmonary regurgitations identified a basal and apical bulging of the ventricle with an elongation of the outflow tract and the apparition of an aneurysm. These observations are consistent with reports found in the clinical literature. **We then estimated a generative model of right ventricle growth computed using the proposed PLS-CCA method on patient body surface area.** The model showed plausible evolution patterns, with progressive ventricle dilation, basal bulging and outflow tract aneurysm at later stages of child growth. The model was also tested on seven new patients. It managed to predict their body surface area with acceptable prediction errors. **Finally, we compared the added-**

value of PLS over PCA to compute predictive models. PLS provided a more compact and more explicative space decomposition.

Part III – Therapy

Personalised Simulation of Cardiac Therapies

In Chapter 8 we propose an electromechanical model of the beating heart for *in-silico* therapy simulations. The goal is to personalise the model from medical data to simulate the cardiac function of ToF patients and to predict the outcomes of pulmonary valve replacement. This task is challenged by the limited clinical data. Consequently, the level of detail of the models must be chosen accordingly. In this work, we rely on the electromechanical model developed by [Sermesant et al., 2006b], which is compatible with the clinical data. **Because of the complexity of ToF, we have to improve that model to simulate the specificities of that disease.** The proposed model is built as a modular workflow, from anatomy to electrophysiology, biomechanics and hemodynamics. We integrate a 3-element Windkessel model and develop a simple regurgitation constraint to simulate the impaired hemodynamics. We also employ a new boundary condition, based on contact forces, which fixes the heart in the 3D space, yielding a more realistic cardiac motion. For each element of the model, we suggest alternative models and personalisation strategies for patient-specific simulations.

Chapter 9 reports the personalised simulations of pulmonary valve replacement therapies in two patients with repaired ToF. **First, we personalise the electromechanical model of the heart using the available clinical data.** The simulation is compared to the cMRI of the patient. **Then, we use the personalised electromechanical model to simulate two pulmonary valve replacement strategies and compare their outcomes.** In the first strategy, we only replace the pulmonary valves: this is achieved by stopping the regurgitation in the model. In the second strategy, we replace the valves and reconstruct the right ventricle by resecting myocardium lesions using SOFA, a real-time soft-tissue intervention platform. Finally, we re-run the simulation on the postoperative anatomy. Despite the simplifications of the model, it qualitatively managed to recover the cardiac function of these patients. The simulations of valve replacement suggest that right ventricle reconstruction improves cardiac function. Interestingly, the experiments also indicate a tight left-right ventricle relationship: right ventricle reconstruction also improves left ventricular function.

Chapter 10 concludes this thesis by summarising the proposed contributions. Some perspectives for future work are listed.

From a Normal Heart to Tetralogy of Fallot

Contents

2.1 Structure and Function of the Normal Heart	11
2.1.1 The Cardiovascular System	11
2.1.2 Cardiac Anatomy	13
2.1.3 Cardiac Function	14
2.1.4 The Cardiac Cycle	17
2.2 Imaging the Heart	19
2.2.1 Echocardiography	20
2.2.2 Cardiac Magnetic Resonance Images	23
2.2.3 Cardiac Computed Tomography	27
2.3 Tetralogy of Fallot: Physiopathology and Clinical Challenges	28

The heart is a fantastic biological machine that pumps the blood through the human body. Its function is thus vital. On average, the heart beats 70 times and ejects 5 L of blood per minute. With a simple calculation, one can estimate about 3 billion heart beats in an average life! Despite these tremendous numbers, the heart never fails. Breakdown of cardiac function is dreadfully life-threatening. Understanding, evaluating and predicting cardiac function is therefore of great importance.


2.1 Structure and Function of the Normal Heart

2.1.1 The Cardiovascular System

The cardiovascular system is a closed loop that transports the blood towards the organ to provide them with oxygen and nutrients collected in the lungs and the digestive system. The human cardiovascular system is divided into two main parts: the *pulmonary circulation* and the *systemic circulation* (Figure 2.1). The pulmonary circulation is a closed loop that ensures the oxygenation of the blood. The oxygen-depleted blood quits the heart from the *pulmonary artery*, it goes into the lungs where it is oxygenated and then comes back to the heart through the *pulmonary veins*. The systemic circulation is much larger than the pulmonary circulation. It transports the oxygenated blood away from the heart towards the other organs. The

blood first flows through the *aorta*, then goes into other great arteries, then into the arterioles and finally into the capillaries where it releases the nutrients and the oxygen. The blood then comes back to the heart through the venules, veins and finally inferior and anterior *venae cavae*.

Consistently with the two circulatory systems, the heart is constituted of two main parts, the *left* and the *right heart*, which are themselves made up of two cavities, the atria and the ventricles (Figure 2.3). The oxygen-depleted blood first arrives into the *right atrium* through the inferior and superior venae cavae. It is then transferred into the *right ventricle* and then ejected towards the lungs through the pulmonary artery. The blood then comes back through the pulmonary veins to the *left atrium*, before entering into the *left ventricle* and pumped out through the aorta towards the other organs. Normally, the entire system is sealed, the blood cannot leave the network.



Mosby, Inc. and derived items copyright © 2003, Mosby, Inc. All Rights Reserved.

Figure 2.1: Drawing of the cardiovascular system. *In blue* Blood poor in oxygen. *In red* Blood rich in oxygen (See text for details).

Finally, each ventricle can be segmented into three segments. The blood enters first in the *inflow tract*, or *inlet*, just below the atrioventricular valves. Below the inflow tract is the *apical* segment, the larger part of the ventricle. Finally, the *outflow tract*, or *outlet*, is just below the arterial valves, from where the blood leaves the ventricle. In the right ventricle the outflow tract as a typical funnel-shape connecting the right ventricle to the pulmonary artery.




Figure 2.2: Ventricle subdivisions. The blood enters the ventricle through the inlet, or inflow tract. It then fills the apical region and is pumped through the outlet, or outflow tract.

2.1.2 Cardiac Anatomy

The heart is a powerful muscular organ whose shape and function optimise the pump function while minimising the muscular work. The heart is mainly made up of a thick *striated* muscle, called *myocardium*. The outer myocardium is often referred to as the *free wall* as it forms the exterior shape of the heart. Internally, the left and the right cavities are separated by the *septum*, which is part of the left heart. The myocardium is protected by two thin layers of endothelial tissues: the *endocardium* inside the cavities and the *epicardium* outside. The heart is further protected by the *pericardium*, a thin fibrous sac that does not contract. The pericardium isolates the heart from the other organs. A thin fluid layer between the pericardium and the epicardium reduces frictions to ensure smooth cardiac motions.

The thickness of the myocardium varies over the heart according to the speed at which the blood must be ejected. The left and right atria for instance have thin walls ($\approx 2\text{ mm}$) as they do not need to contract significantly to push the blood forward the ventricles. The free wall of the right ventricle is slightly thicker ($\approx 5\text{ mm}$), to pump the blood towards the lungs, whereas the left ventricle is $\approx 15\text{ mm}$ -thick as it must pump the blood over the systemic system. Because of these variations, the contrast of the myocardium in medical images greatly varies from one chamber to another. The left ventricle is often well-defined in MR images but the right ventricular free wall is hard to detect, which challenges its quantitative assessment.

The heart chambers are separated by one-way valves that control the direction of the blood flow during the cardiac cycle. Constituted of connective tissues, the valves open and close according to the gradient of pressure between the cavities and arteries. The atrioventricular valves, namely the *mitral* valves for the left heart and the *tricuspid* valves for the right heart, ensure that the blood flows from the atria to the ventricles and not conversely. These valves open as soon as the pressures in the atria are higher than the pressures in the ventricles. When the pressures in the ventricles




Figure 2.3: Structure and function of the normal heart. (Image modified from Mariana Ruiz Villareal, Public Domain).

are higher than in the atria, the valves close and remain sealed thanks to the *papillary muscles* and the *chordae tendineae*. Similarly, the *aortic* and the *pulmonary artery* valves control the blood flow towards the arteries. They open as soon as the pressures in the ventricles are higher than in the arteries. When the valves do not close perfectly, the blood can flow backwards, which is called *regurgitations*. Chronic regurgitations can lead to severe cardiac dilation and functional impairment, as in repaired tetralogy of Fallot (see Section 2.3).

2.1.3 Cardiac Function

Cardiac contraction and relaxation are controlled automatically by the *autonomous sympathetic system*. The heart beat is the result of two tightly calibrated biological phenomena. First, an electrical wave propagates over the myocardium to activate its contraction. Then, the activated cardiac cells, called *myofibres*, *muscle fibres* or cardiac *myocytes*, contract and then relax, resulting in the global cardiac motion. Although the heart is constituted by a succession of myocytes separated by intercalated discs, it is a functional *syncytium* as the electrical wave and the cardiac contraction happen globally over the myocardium to achieve synchronous motion. That is why local abnormalities can quickly deviate into severe global functional malfunctions.

2.1.3.1 Cardiac Electrophysiology

The cardiac activity is controlled by the periodic propagation of an electrical wave that depolarises the myocytes to make them contract [Keener and Sneyd, 1998]. The electrical wave is triggered by the autonomous sympathetic system. It starts in the atria, at the *sinoatrial nodes*, the physiological pacemaker (Figure 2.4). As

the electrical wave propagates towards the *atrioventricular nodes*, it depolarises the atrial myocytes, which contract and pump the blood to the ventricles. The electrical impulse is then stopped for a few milliseconds to synchronise the heart. This “pause” is fundamental as it enables the atria to fully contract and completely pump the blood into the ventricles. Finally, the electrical wave is propagated at very high speed ($\approx 2000\text{ mm/s}$) downwards the apex (the tip of the heart) through the left and right *bundles branches*. It reaches the *Purkinje fibres* and propagates throughout the entire myocardium, from endocardium to the epicardium, at a lower speed ($\approx 500\text{ m/s}$) [Murgatroyd and Krahn, 2002].

The electrical wave propagates from cell to cell by modifying the ion concentration on both sides of the cell membrane [Keener and Sneyd, 1998] (Figure 2.5). When a cardiac cell is stimulated (by the sympathetic system or an excited neighbouring cell), sodium ions (Na^+) flood inside the cell through specific ion channels. To counter-balance the increased membrane voltage, potassium ions (K^+) are released but at the same time calcium channels opens and calcium ions (Ca^{2+}) enter into the cell. This release of calcium keeps the cell depolarised and triggers the contraction of the myocyte through a calcium-induced calcium release mechanism (CICR) [Endo, 1977]. Once depolarised, the cell cannot be excited anymore for a small period, called *refractory period*, during which it stays contracted. Then, it repolarises by re-equilibrating the ion concentrations, and relaxes, to come back to its rest state. The transient polarisation/depolarisation of the cell is called *action potential* and its duration is commonly abbreviated APD (action potential duration). The synchrony of the cardiac motion is ensured by a tight calibration of the APD over the myocardium. Typically, the APD is longer in the septum, as this region depolarises first and depolarises last, and shorter in the free wall [Murgatroyd and Krahn, 2002]. Sub-endocardial APD is also generally longer than sub-epicardial APD.

In a clinical routine, cardiac electrophysiology is assessed non-invasively with electrocardiograms (ECG). The ECG is obtained by placing on the torso of the patient two or more skin electrodes that measure the tiny electrical signal produced by the heart. The ECG signal of healthy subjects is characterised by well-defined patterns that correspond to precise cardiac events (Figure 2.4, right panel). The ECG signal starts with the *P wave*, which is the manifestation of the initial atrial depolarisation triggered by the sinoatrial nodes. The pause at the atrioventricular nodes is reflected by a short return to the rest state of the ECG signal, the *PR interval*. When the electrical wave spreads over the ventricles, it generates the *QRS complex* of the ECG. Finally, the *T wave* corresponds to the repolarisation of the myocardium. Contrary to the QRS complex, the shape and length of the T wave depend on the cardiac rhythm. The higher the heart rate, the shorter the T wave. Analysing the shape, trigger times and lengths of the ECG waves provides crucial insights into the cardiac electrophysiology of a patient. ECG is often sufficient to detect bundle branch blocks or severe arrhythmia. Yet, slight variations in cardiac conductivity cannot be detected by the ECG. Their quantitative evaluation requires invasive endocardial mappings obtained through catheter.

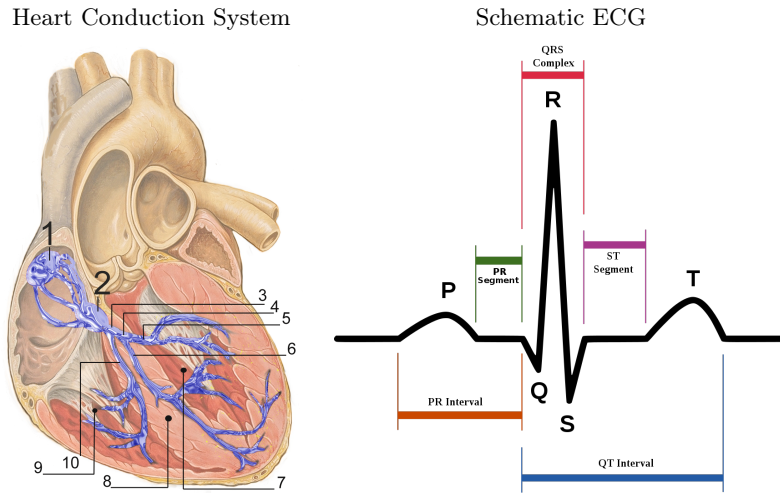


Figure 2.4: *Left panel:* Heart conduction system. the electrical activity of the heart is triggered by the sinuatrial nodes (1) and then the atrioventricular nodes (2). It is transported by the left bundle (5) and the right bundle (10) branches and finally transmitted to the myocardium (8) through the Purkinje fibres (9). *Right panel:* Simplified ECG with the main electrical waves (*Images from Wikipedia*).

Because the electrical wave controls the cardiac contraction, its healthiness is crucial for an efficient cardiac function. The vast majority of sudden deaths are due to local alterations of the cardiac electrophysiology. The APD can be locally altered by lesions, after a myocardial infarct for instance or due to chronic diseases or surgery sequels. Such adverse events can indeed damage cardiac cells, reducing their electrical conductivity, thus hampering the cardiac synchrony. Regions of non-excitable cells can for instance yield wave re-entry, the main source of cardiac fibrillation. Modelling cardiac electrophysiology is therefore a very active research field to predict complex and life-threatening cardiac abnormalities.




Figure 2.5: Ion exchanges at the surface of the cell membrane that generate the action potential (Figure from [Marbán, 2002]).




Figure 2.6: Drawing of a myocyte and its constituent ((©Penn State Altoona).

2.1.3.2 Cardiac Biomechanics

A cardiac cell is made up of several *myofibrils*, which are themselves a succession of *sarcomeres*, the contractile unit of the cardiac muscle (Figure 2.6). A sarcomere is delimited by two Z-discs, between which interleaved proteins of constant length, termed *filaments*, are stacked: *actin* and *myosin* (Figure 2.7). Actin is a thin filament that is partially connected to the Z-discs. Myosin is a thick filament, indirectly connected to the Z-discs by a very thin protein, *titin*. Myosin is characterised by myosin-binding proteins, the *myosin heads*, that are tightly connected to the myosin filament and that can also bind to the actin. These bindings are the engines of the sarcomere contraction.

When the calcium ions enter into the sarcomere, the myosin heads are “enabled” and connect to the actin filaments (Figure 2.8). They progressively slide towards the Z-discs, outwards, which results in the shortening of the sarcomere. The space between the myosin filaments and the Z-discs, called I band, shortens and the H zone, the space between neighbouring actin filaments disappears. When the calcium ions are released back from the cell, the myosin-bindings are disabled: The myosin heads unbind from the actin and the sarcomere comes back to its rest state.

2.1.4 The Cardiac Cycle

The cardiac cycle consists of a succession of four different phases as reported in Figure 2.9. When the ventricles start contracting, the atrioventricular valves close (first heart noise) and the ventricular pressures quickly rise. It is the *isovolumetric contraction* phase. As soon as the ventricular pressures become higher than the arterial pressures (80 *mmHg* for the aorta and 15 *mmHg* for the pulmonary artery in normal hearts), the arterial valves open and the blood is ejected. This is the *ejection* phase. During this stage, the ventricular pressures continue rising as the ventricles continue contracting and then decrease as the myocardium starts relaxing. When the ventricular pressures fall below the arterial pressures, the arterial valves close (second heart noise) and the isovolumetric *relaxation* phase starts. The pressures continue decreasing until they become lower than the atrial pressures (12 *mmHg* for




Figure 2.7: Sarcomere constitution and sliding model of cell contraction (*Images from Wikipedia*).




Figure 2.8: Cardiac contraction model. **A-** The sarcomere is activated, the myosin heads bind to the actin filament. **B-** The myosin heads fold and **C-** the myosin slides along the actin: the sarcomere contracts. **D-** When the cell is deactivated, the myosin heads unbind and the sarcomere returns to its rest position (*Images from Wikipedia*).

the left atrium and 5 mmHg for the right atrium). At this point, the atrioventricular valves open and blood flows into the ventricles. This is the *filling* phase. During the initial stages of the filling, the blood passively flows into the ventricles because of the pressure gradient between the cavities. As the sinoatrial nodes are stimulated, which starts a new cardiac cycle, the atria contract and further pump the blood into the ventricles, with a small increase of the atrial pressures: this is the active filling.

The contraction phases (contraction and ejection) constitute the cardiac *systole*.




Figure 2.9: Cycle of the left ventricle. The right ventricle goes through an identical cycle but with lower pressures (*Images from Wikipedia*).

Similarly, the relaxation phases (relaxation and filling) constitute the cardiac *diastole*. The end-systole time frame corresponds to the time when the heart is fully contracted whereas at the end-diastole time, the heart is maximally filled.

The arterial pressures also vary during the cardiac cycle. During ejection, the arterial pressures quickly increase, followed by a progressive decrease. This phenomenon, known as *Windkessel effect* [Frank, 1899] is due to the compliance of the arteries. Intuitively, the arterial walls dilate during ejection to accommodate the cardiac pulse (Figure 2.10). The artery stores the excess of blood that cannot be injected into the circulatory system, which has a finite capacity. As a result, the arterial pressures increase as the arteries dilate. When the ejection strength decreases, the excess of blood that has been stored by the arteries is progressively released to the circulatory system. In other words, the arteries transform the discontinuous blood flow pumped by the heart into the steady flow of the circulatory system.

2.2 Imaging the Heart

Assessing the cardiac function of patients requires visualising the myocardium to detect functional abnormalities and lesions. Several imaging modalities are available in clinics, each one with its advantages and limitations. This section provides an overview of the most common imaging modalities. It is not an exhaustive review but rather a brief presentation of the images that are used in the remainder of




Figure 2.10: The Windkessel effect. During systole, the artery dilates to accommodate the cardiac pulse. When the valves close, the blood is released progressively in the circulatory system (©Urban and Fischer 2003, Roche Lexikon Medizin, 5. Aufl.)

this manuscript. The reader is referred to [Liang and Lauterbur, 1999; Chambers, 2001; Shah et al., 2005; Teske et al., 2007; Topol and Califf, 2007] or other related references for further details.

Three main cardiac imaging techniques are used for diagnosis in cardiology: echocardiography, cardiac magnetic resonance imaging (CMR) and cardiac computed tomography (CT). Nowadays, all of them provide 3D images of the beating heart although the clinical evaluation of cardiac anatomy and function is still based on 2D views wisely selected by the radiologist. Four standard 2D views have been defined (Figure 2.11) according to their orientation with respect to the *long axis* of the left ventricle (the axis that passes through the apex and the mitral valve). The short axis views are orthogonal to the long-axis (Figure 2.11, bottom row). They display the two ventricles and are used to assess the radial deformations (heart thickening) and the circumferential displacements (heart twisting). The four-chamber view is used when the four cavities must be assessed. The two chamber view enables to assess the left atrium and ventricle only, while the long-axis view is used to analyse the aortic valves. For regional assessment, the left myocardium is divided into 3 planes: the *base* close to the valve, the *mid-plane* and the *apical plane*, which are subdivided into 16 regions as defined by the American Heart Association, with the apical cap being a standalone region (the 17th region) [Cerqueira et al., 2002] (Figure 2.11).

2.2.1 Echocardiography

2.2.1.1 Overview

Echocardiography is the most common imaging device in clinics thanks to its ubiquity, ease of use and low cost. Echocardiography is based on ultrasound waves. Like sonar in submarines, a probe sends multidirectional ultrasound waves through the body. At tissue interfaces, the waves are reflected, generating echoes that are captured by the probe and used to reconstruct the 2D images. Echocardiography thus provides images of tissue interfaces.




Figure 2.11: Main 2D views of the heart with AHA subdivision of the left ventricle (Image from [Lang et al., 2005], AHA subdivision from [Cerqueira et al., 2002])

In a clinical routine, echocardiography is used to evaluate the cardiac anatomy and function of patients [Topol and Califf, 2007] (Figure 2.12). The large majority of machines render 2D slices of the beating heart along the direction defined by the cardiologist. Usually, the ultrasound probe is placed on the thorax of the patient but in some cases the cardiologist may introduce the probe in patient oesophagus to visualise myocardium regions that may be hidden by the ribs. In the last decade, 3D echocardiography machines have been introduced into the clinical environment to provide a full view of the heart and to reduce rater variability. The current 3D techniques consist in stitching 2D slices together to reconstruct 3D volumes. These techniques may require several heart beats to reconstruct one 3D volume, resulting in low frame-rates. This limitation is being relieved by new ultrasound probes that achieve real-time full-volume 3D echocardiography imaging and capture the 3D heart in one heart-beat, paving the way to consistent and comprehensive assessment of cardiac anatomy and function [Zamorano, 2009].

Echocardiography also enables to evaluate the 1D cardiac motion and blood flows through Doppler imaging. Spatial velocity along the ultrasound probe is estimated using the Doppler effect of travelling waves. 1D blood flows can be estimated, enabling the quantification of valve regurgitations or septal defects (Figure 2.12). Similarly, unidirectional tissue velocities is used to estimate the strain rate (the variation of velocities between two spatial points) and, after time integration, the strain along the ultrasound beam (Tissue Doppler Imaging, TDI [Teske et al., 2007]). By wisely positioning the ultrasound probe the clinician can have a 1D estimate of the longitudinal or the radial myocardium strain.

Recently, 2D strain imaging has been developed to estimate the 2D strain of

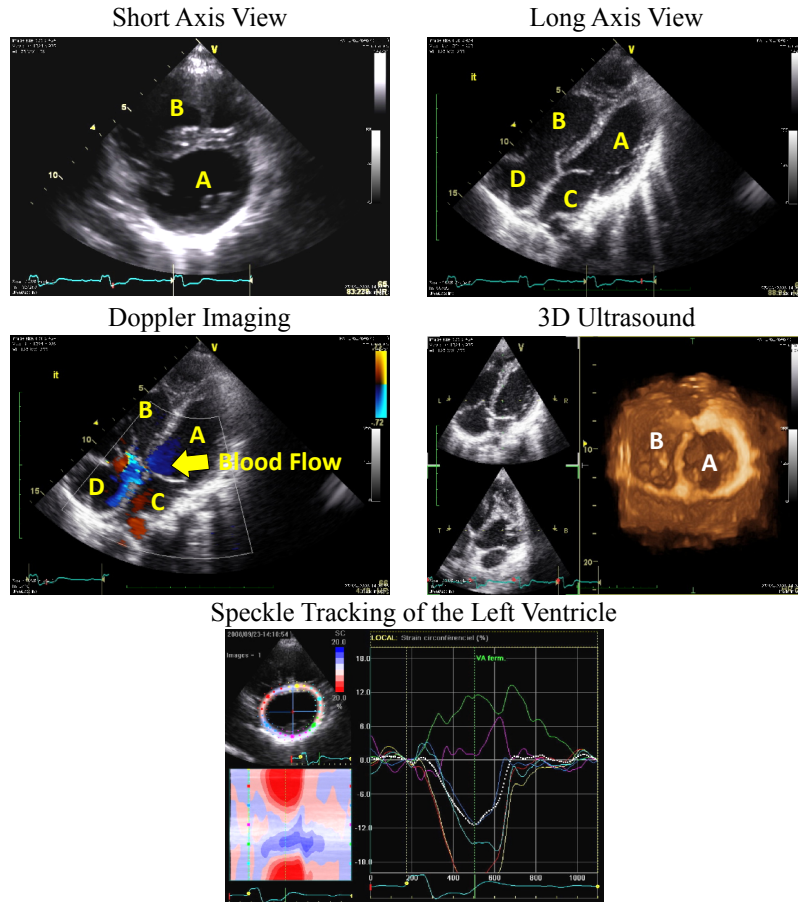


Figure 2.12: Echocardiography of a patient with repaired tetralogy of Fallot. **A**- Left ventricle. **B**- Right ventricle. **C**- Left atrium. **D**- Right atrium (*Images from APHP, Necker-Enfants Malades, Paris*).

the myocardium. The underlying principle is to track the ultrasound noise, called speckle, using an optical flow algorithm. Indeed, ultrasound speckle is consistent over time as it is generated by the diffraction of the ultrasound waves by the tissues. It hence depends on the tissue properties. The displacements recovered by the speckle tracking are then used to estimate the 2D myocardium strain [Ledesma-Carbayo et al., 2002]. This technique has received a great interest from the cardiologists since it is more comprehensive and accurate than TDI [Teske et al., 2007]. However, it still suffers from large inter rater variability [Chung et al., 2008]. The recent real-time full-volume 3D machines extend this technique to 3D strain measurements, aiming at alleviating the rater variability issue and providing 3D assessment of the heart [Zamorano, 2009].

2.2.1.2 Advantages and Current Limitations

The main advantages of echocardiography are:

- + Non-invasive
- + Low cost
- + Ubiquity
- + Good spatial resolution
- + High frame rate
- + Enable to quantify cardiac deformations and blood flows

and its current limitations:

- Low tissue contrast
- High noise and image artefacts
- Blind regions, masked by highly reflective tissues like bones
- Large high rater variability [Chung et al., 2008]

The advantages of echocardiography make this imaging modality the first choice for quick assessment of the cardiac function. However, because of the inter-rater variability and of the poor image quality, more advanced imaging techniques are required for more accurate evaluations.

2.2.2 Cardiac Magnetic Resonance Images

2.2.2.1 Overview

Cardiac Magnetic Resonance Imaging (MRI) is gaining a leading position in cardiology, in particular for diagnosis and evaluation of congenital heart defects where MRI is considered as the gold standard [Helbing and De Roos, 2000; Krishnamurthy, 2008]. In simple terms, MR scanners use strong magnetic fields to image biological tissues. Within the MR scanner, a strong and constant magnetic field of 1T to 3T (and more for research), according to the machine, aligns the spin of the hydrogen atoms. An additional magnetic field, called pulse, is overlaid to re-orient the aligned spins. When the pulse is disabled, the spins come back to their initial orientation along the reference magnetic field in an amount of time that depends on the biological tissue. MRI thus consists in measuring this time. The spatial sampling is obtained by applying gradients to the magnetic field, meaning that the field varies over space. We refer the reader to [Liang and Lauterbur, 1999] for further details.

Tissue response to the magnetic fields depends on the length, amplitude and frequency of the pulse fields. In other words, different pulse sequences yield different tissue contrast. The radiologist can thus choose which tissue to visualise or to remove from the images. For instance, fat can be explicitly removed from the images by using fat-removal sequences. For cardiology, special acquisition techniques have been developed to accommodate the cardiac motion. First, the acquisition is gated, i.e. synchronised, to the ECG to acquire images at a specific instant of the cardiac cycle. This enables to acquire a 3D image of the heart over several heart beats. Two gating strategies are commonly used: *prospective gating* and *retrospective gating*. The prospective strategy consists in acquiring the images at a very specific instant of the cardiac cycle defined by the ECG. Conversely, the retrospective strategy consists in first acquiring all the data, independently of the instant of the cardiac cycle.

The data is subsequently reorganised according to the recorded ECG. Another difficulty of cardiac MRI is the respiratory motion, which can seriously affect the image quality by blurring the images. Breath-hold strategies have been designed to minimise motion artefacts. The subject holds its breath for a short time, during which the radiologist acquires a small portion of the image. To acquire a full view of the heart, several breath holds are required, which may introduce slice misalignments when the subject holds its breath at different position (see Section 3.2.2.1). Respiratory navigators can also be used to synchronise the acquisition with the respiratory motion.

In cardiology, the state-of-the-art pulse sequence is the Steady State Free Precession (SSFP), which is an improved T1 sequence especially tailored to the cardiac tissues [Liang and Lauterbur, 1999; Shah et al., 2005] (Figure 2.13). When SSFP sequence is used alone, one gets the anatomical cardiac cine MRI (cMRI) widely used in clinical settings to visualise the cardiac anatomy and the apparent motion. Isometric 3D images of the heart at end-diastole are acquired by gating the acquisition at the end-diastole time instant (usually identified as the R wave of the ECG). Dynamic sequences of the beating heart are also available by stacking short-axis or long-axis 2D+t slices together. The underlying principle is to acquire at each breath hold images of the cardiac motion observed in a specific 2D plane. The 2D+t planes are then stacked to reconstruct the 3D+t volume (Figure 2.14). Such an approach provides good in plane spatial resolution (in the order of magnitude of the millimetre) and temporal frame rates (about 40 frames or more with the most recent machines). However, the slice thickness is usually large (about 1 cm) and the 3D volume may suffer from slice misalignments due to inconsistencies in the breath hold.

The versatility of MRI has allowed the development of pulse sequences for imaging the cardiac function and tissue integrity. While standard SSFP MRI only exhibit the cardiac anatomy and the apparent motion of the myocardium, tagged MRI aims at evaluating the complete cardiac deformation. First introduced by [Zerhouni et al., 1988], tagged MRI consists in marking the tissues, in particular the myocardium, with a grid by using strong magnetic fields (Figure 2.13). SSFP pulse sequences are then employed to image the marked beating heart over time. During post-processing, the tag grids are tracked and the estimated displacements are used to compute the cardiac deformations [McVeigh, 1998; Moore et al., 2000]. Nowadays, tagged MRI is considered the gold standard for assessing the circumferential and the longitudinal cardiac motion, which is hardly visible in anatomical cine MRI. More sophisticated sequences have been developed in the last decade to measure the cardiac deformation, for instance displacement encoding (DENSE encoding) [Aletras et al., 1999; Kim et al., 2004; Feng et al., 2009], velocity encoding for blood flows and myocardium motion [Van Dijk, 1984; Bogren and Buonocore, 1999] and finally direct strain encoding [Osman et al., 2001; Youssef et al., 2008].

Tissue integrity is evaluated on MRI by imaging how contrast products like gadolinium are fixed by the biological tissues. This approach is commonly used in patients with heart failures to identify and quantify myocardium scars that may

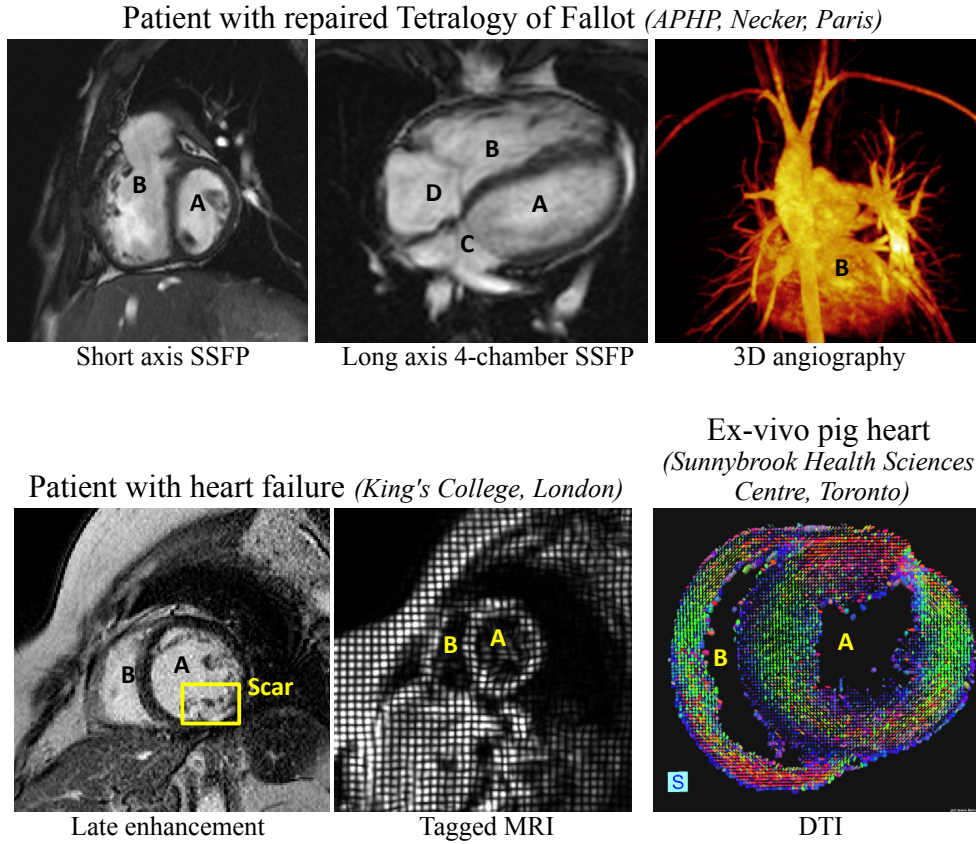


Figure 2.13: Example of MR images. **A**- Left ventricle, **B**- Right ventricle, **C**- Left atrium, **D**- Right atrium. Angiography colours are artificial to highlight the vessels. DTI colours encode fibre directions.

hamper the cardiac function. A similar technique is employed to image the coronary arteries and the great arteries. The contrast product makes the blood very bright in the images, facilitating the visualisation of small vessels. This modality is referred to as MR angiogram (MRA).

A last MR modality that can have a great impact in cardiology is diffusion tensor imaging (DTI). DTI measures the diffusivity of proton particles in anisotropic media. Its most common application is brain imaging, where it is used to quantify the direction of the axon bundles in the white matter [Le Bihan et al., 1986; Basser et al., 1994]. The diffusion signal is measured along several direction, called gradients (at least 6 directions are needed). Then, a diffusion tensor is estimated, which encodes the three main directions of diffusion in the 3D space. Tracking algorithms can finally be applied to track the fibres through space. In cardiology, numerous DTI studies have been carried out to image the 3D orientation of myocytes in *ex-vivo* hearts [Geerts et al., 2002; Helm et al., 2005a], yielding spatially consistent 3D characterisations of fibre orientations in contrast to dissection studies that do not preserve the cardiac anatomy. However, *in-vivo* applications are still challenging due

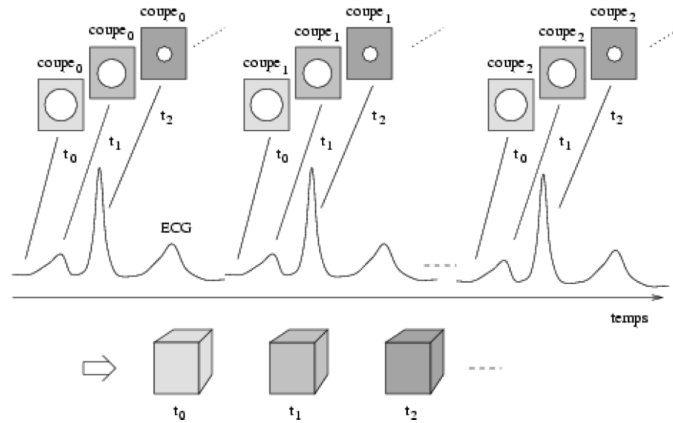


Figure 2.14: ECG gated MRI acquisition of 3D+t images (Image from [Sermesant, 2003])

to the cardiac motion that fades the diffusion signal. Intense research is striving to design efficient sequences for imaging *in-vivo* fibre orientation in patients. Some of them correct the diffusion signal using myocardium 3D strain measured with velocity encoding [Tseng et al., 2000; Chen et al., 2005]. Nonetheless, the acquisition is time consuming and requires long breath holds, which make this modality challenging to apply in clinical routine [Wu et al., 2009]. To tackle this limitation, [Toussaint et al., 2010] propose to acquire few 2D slices at wise cardiac positions and orientations, and then to extrapolate the diffusion signal elsewhere using a computational model of fibre orientation.

2.2.2.2 Advantages and Current Limitations

The main advantages of cardiac MRI are:

- + Non-invasive
- + Good spatial resolution according to the MR modalities
- + High frame rate
- + High image contrast
- + Little rater variability
- + Enable to quantify cardiac deformations and blood flows
- + High versatility, the radiologist can fine tune the sequences to get the optimal contrast

and its current limitations:

- High cost
- Not comfortable for the patient
- Metal sensitive
- Prone to slice misalignments and image phantoms

Despite its elevated commercial and operational cost, MRI is becoming the gold standard in cardiology, in particular in paediatrics. It significantly reduces rater variability and its high contrast makes the assessment of the cardiac anatomy and function more accurate than with echocardiography. However, MRI is not always available. Furthermore, patients may feel uncomfortable inside the machine because of claustrophobia or just cannot have MRI scanner because of metal implants. In some cases, complete anaesthesia with controlled respiratory stops are required, in particular in very young patients.

2.2.3 Cardiac Computed Tomography

2.2.3.1 Overview

The underlying idea of computed tomography (CT) scan is to acquire X-ray images along several directions and to recombine them on computers to reconstruct the 3D image [Topol and Califf, 2007; Schoenhagen et al., 2005]. As for CMR, the acquisition is ECG gated to accommodate the cardiac motion. With the advance of multi-slice CT scans (up to 256 slices can be acquired simultaneously in recent machines), the heart can be imaged during the entire cardiac cycle, the spatial and temporal resolution depending on the number of slices the machine can acquire simultaneously.

CT is used to visualise the cardiac anatomy and its apparent motion (Figure 2.15). CT images benefit from higher spatial resolution than MR and the 3D images are usually isotropic. Like some MR sequences, CT images are quantitative. The grey level values are proportional to the mass density. Furthermore, CT is not affected by metallic objects: it can be used after the implant of pacemakers or any other device. However, the patient is irradiated by a large amount of X-rays to account for the cardiac motion, which can be dangerous for the patient as X-rays are carcinogenic. That is why CT imaging is avoided in paediatrics as children are more sensitive than adults to X-rays. Current research is striving to reduce the amount of irradiation to acceptable levels while keeping good image quality.




Figure 2.15: Different views of a 3D CT image of the heart. **A**- Left ventricle, **B**- Right ventricle, **C**- Left atrium, **D**- Right atrium, **E**- Pulmonary artery, **F**- Aorta.

2.2.3.2 Advantages and Limitations

The main advantages of cardiac CT are:

- + Excellent isotropic spatial resolution
- + Good image contrast
- + No rater variability

and its current limitations:

- Invasive, the patient is irradiated by hazardous x-ray beams
- Image cardiac anatomy only

During this thesis, only cMRI of the beating heart were available for processing as we dealt with children and teenagers. Echocardiography was also performed on these patients but these images could not be processed for technical reason. Because of that, the developments described later on in this manuscript focus on cMRI although they are general enough to be applied on other modalities.

2.3 Tetralogy of Fallot: Physiopathology and Clinical Challenges

Tetralogy of Fallot (ToF) is a severe congenital heart defect that affects mainly the right ventricle (RV) of the heart. It is the most common defect among the cyanotic congenital heart diseases (better known as “blue-baby” syndrom) and it accounts for 5 to 6% of all congenital heart diseases, with an incidence of 1 out of 2500 births [Hoffman and Kaplan, 2002]. ToF was described in 1672 by Niels Stensen, in 1773 by Edward Sandifort but the name of the pathology was given by Etienne-Louis Arthur Fallot in 1888. It is characterised by four major defects (Figure 2.16):

Pulmonary Stenosis The right ventricle outflow tract (RVOT) and the pulmonary artery are narrower than in normal subjects due to overgrown cardiac tissues proximal to the pulmonary valves. Because of the stenosis, the right ventricle pressure is higher than in normal to accommodate the abnormal resistance towards the lungs.

Ventricular Septal Defect (VSD) A hole in the inter-ventricular septum creates a blood shunt between the two ventricles. The VSD is generally in the upper part of the septum, just below the arteries. It can be relatively large.

Overriding Aorta The root of the aorta and the aorta valves are placed directly on top of the VSD. The aorta is bi-ventricular, meaning that its root is connected to the left ventricle and to the right ventricle.

Right Ventricular Hypertrophy. The right ventricle is much thicker than in normal subjects, mainly due to the pulmonary stenosis.




Figure 2.16: Structure and function of the heart in tetralogy of Fallot subjects. Due to the defects, oxygen-depleted blood comes back to the systemic system. (Images by Mariana Ruiz Villareal, Public Domain)

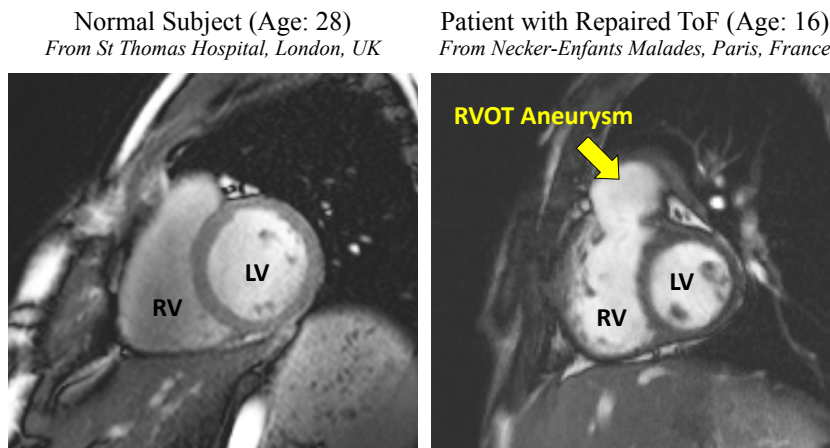


Figure 2.17: *Left panel:* MRI of a normal heart. *Right panel:* MRI of a patient with repaired tetralogy of Fallot. Observe the dilated right ventricle (RV) compared to the left ventricle (LV) in ToF in contrast with the normal heart. The patient also has an RVOT aneurysm due to the corrective patch.

As a consequence of these defects, oxygen-depleted blood returns to the systemic system without being re-oxygenated into the lungs. The “blue” aspect of ToF babies is the physiological symptom of this poor oxygenation, which, in the most severe cases, can lead to death. ToF infants thus require paliative surgical repair in their early months. Performed the first time by C. Walton Lillehei at the University of Minnesota in 1954, ToF complete repair consists in closing the VSD with a Gore-Tex patch or homograft and in relieving the pulmonary stenosis by resecting the

stenosis to reduce the pressure overload. The RVOT or the pulmonary artery may be enlarged by placing a ventricular patch, a trans-annular patch or a homograft to further reduce the stenosis.

Nowadays, ToF repair is relatively standard, with very low perioperative mortality ($< 5\%$) [Murphy et al., 1993]. Nonetheless, pulmonary valves may be damaged by the surgery which, together with the enlargement of the RVOT, may cause chronic pulmonary regurgitations. During child growth, the regurgitations are usually asymptomatic. However, they yield severe long-term complications such as extreme RV dilation, RVOT aneurysm (Figure 2.17, right panel), arrhythmias and RV function impairment with consequences on the LV pump efficiency due to an abnormal motion of the inter-ventricular septum [Oosterhof et al., 2006]. In addition, residual pulmonary stenosis may remain, further impairing the RV function. If nothing is done to control these complications, the risks of life-threatening events increases and life expectancy is reduced [Nollert et al., 1997].

The clinical challenges raised by ToF are therefore shifted from the initial repair of the congenital malformation to the long-term management of these patients. Thanks to the low mortality and morbidity of the initial repair, the population with repaired ToF is steadily increasing, requiring more clinical efforts to regularly follow them and detect when their cardiac function becomes too impaired. To decrease the risks of sudden death, pulmonary valve replacement is considered to be an effective treatment. It has been shown that the RV can recover an almost normal shape and function after valve replacement [Therrien et al., 2005; Geva, 2006]. Because most of the repaired ToF patients are asymptomatic, cardiologists tend to wait as long as possible to avoid multiple re-interventions due to the limited lifespan of the implanted valves. However, if the valves are implanted too late, the RV does not recover a normal shape and function anymore: it is definitively damaged [Therrien et al., 2000; Oosterhof, 2006]. **There is therefore a need of quantitative and reproducible metrics to comprehensively evaluate the cardiac function in these patients.** Unfortunately, defining the best timing for re-intervention is controversial. The most used criterion is based on the RV end-diastole volume (EDV) but it is not clear whether this indicator represents all the risks of sudden death and impaired function [Gatzoulis et al., 1995; Geva, 2006]. **There is thus a need of methods to model the cardiac shape and function in the future to predict possible cardiac failures that would justify immediate valve replacement.**

In addition to the timing for re-intervention, there is controversy on which PVR strategy to apply for a given patient. Two PVR techniques are becoming prevalent (Figure 2.18). On the one hand, the cardiologist replaces the pulmonary valves through surgery or, more recently, using catheters (the procedure is then called percutaneous PVR (PPVR) [Khambadkone et al., 2005]). Note that PPVR is possible only if the RVOT diameter is lower than 22 mm due to the limited diameter of the artificial valves [Schievano et al., 2007]. With this strategy, the RV progressively remodels on its own to recover a normal cardiac function. On the other hand, a recent surgical technique consists in reducing the RVOT, replacing the pulmonary valves and remodelling the RV directly [del Nido, 2006]. The surgeon not only replaces

the valves but also intentionally resects the regions of the RV myocardium that are impaired by fibrosis or scars, to reduce RV volume and improve its function directly. However, postoperative effects of these techniques upon the RV function are difficult to predict as they mostly depend on the patient physiopathology. Choosing the appropriate therapy for a given patient remains a clinical challenge. **There is hence a need for personalised therapy planning.**




Figure 2.18: Pulmonary valve replacement strategies. *Left panel:* Pulmonary valves are replaced with RVOT reduction if too dilated. RV scars and fibrosis far from the valve insertion may be left. *Right panel:* Valve replacement and direct RV volume reduction. The RV is resected till the apex to remove any scar tissue (*Image modified from [del Nido, 2006]*).

Quantifying and predicting the RV remodelling in repaired ToF patients is therefore crucial for the clinical management and therapy planning of these patients. How much the myocardium of a patient is impaired? Will his right ventricle deteriorate in the near future? Is now the best timing for pulmonary valve replacement? If yes, which valve replacement strategy should one apply for a specific patient? What will be the effects of the valve implant on his cardiac function? Nowadays, no clinical criteria provide clear-cut answers to these questions. The cardiologist relies on thumb rules and its personal experience. The problem is all the more challenged by the extreme variability of ToF physiopathology: one patient may have a unique pathology course. Finally, because patients are children, invasive assessment of the cardiac function are not recommended (no CT images, no endocardial electrophysiological maps, no pressure catheter, etc.). This constraint further challenges the diagnosis. **The works done in this thesis are motivated by these clinical difficulties. The idea is to study how computational models can provide the cardiologists with hints on the cardiac function in repaired ToF despite the sparse clinical data, how it will evolve in the future and how**

it will be affected by valve replacement. In the following, tetralogy of Fallot (ToF) always refer to the condition *after* the initial repair.

Prerequisite: Semi-Automatic Delineation of the Heart

Contents

3.1 Motivation	33
3.2 A Pipeline for the Dynamic Segmentation of the Myocardium	35
3.2.1 Overview of the Segmentation Pipeline	35
3.2.2 Image Preprocessing	36
3.2.3 Interactive 3D Segmentation of Myocardium Boundaries using Interpolating Implicit Surfaces	40
3.2.4 Smooth Propagation of the Segmentation over the Cardiac Se- quence	45
3.2.5 Reconstruction of the Beating Myocardium	50
3.3 Results on Tetralogy of Fallot	51
3.3.1 Segmentation of the Myocardium on cMRI of Patients with Repaired Tetralogy of Fallot	51
3.3.2 Quantification of the Cardiac Function from Segmentation	54
3.3.3 Examples on Other Pathologies and Imaging Modalities	54
3.4 Discussion	55

3.1 Motivation

The methods presented later in this manuscript rely on segmentations of the patient myocardium. The algorithm for estimating the cardiac function in medical images requires a mask of the bi-ventricular myocardium (Chapter 4, page 61 and Chapter 5, page 89). The statistical analyses of the right ventricle shape are performed directly on the 3D surfaces of the right ventricle endocardia (Chapter 6, page 107 and Chapter 7, page 133). The electromechanical models of the heart are computed on 3D volume meshes that represent the myocardium of the patient and calibrated using the variation of blood pool volumes computed from the segmentations, among other data (Chapter 8, page 155 and Chapter 9, page 193). Segmenting the heart is therefore crucial, although not the main topic of the thesis. In some studies, no fully-automatic and off-the-shelf methods were at our disposal.

The task of cardiac segmentation was all the more challenged when dealing with repaired tetralogy of Fallot patients. For these patients only clinical short-axis cMRI with standard image quality and large slice thickness were available. Additionally, the cardiac anatomy of these patients was extremely variable, with various degrees of right ventricle dilation and, in some cases, with a large aneurysm of the right ventricle outflow tract. These features make the automated segmentation of the heart in this population very challenging.

We thus developed an integrated segmentation pipeline that combines specific image processing algorithms to cope with the variability in imaging modalities and cardiac morphology encountered during this thesis. Nevertheless, each element of the pipeline was independent from the others such that the whole framework could be adapted to other modalities, pathologies or advances in image processing. The aim was not to develop a “better than state-of-the-art algorithm” but rather a multi-purpose tool that can be used to start the forthcoming analyses and be adapted to other applications.

State of the Art

Delineating the myocardium in cardiac images is a fundamental step of the clinical evaluation of the cardiac anatomy and function. However, the manual delineation of the myocardium is a time-consuming and tedious task, sensitive to rater expertise. In the recent years, a large amount of studies proposed solutions for the automatic segmentation of the heart in medical images. If few studies relied exclusively on image information [Sonka and Fitzpatrick, 2000; Jolly, 2006; Lempitsky et al., 2009; Cousty et al., 2010], the large majority of methods use higher-level models to make the segmentation robust to the often poor quality of medical images [Frangi et al., 2001]. A standard technique consists in using deformable models, as first proposed by [McInerney and Terzopoulos, 1995] for cardiac segmentation, to cope with the ill-defined anatomical boundaries. The recent progresses are leading to push-button full-heart segmentation algorithms, like those proposed by [Zheng et al., 2008; Ecabert et al., 2008].

While some methods drive the deformable model using prior knowledge about cardiac biomechanics [Sermesant et al., 2003; Bistoquet et al., 2007; Zhu et al., 2010], the heart is usually segmented on medical images by deforming atlases or using shape models like active shape models (ASM) [Cootes et al., 1995], like in [van Assen et al., 2006], or active appearance models (AAM) [Cootes et al., 1998], like in [Mitchell et al., 2002; Andreopoulos and Tsotsos, 2008; Zhang et al., 2010b]. In these approaches, the deformable model is guided towards the cardiac boundaries under a shape constraint learned from a training set of data. Advanced boundary detectors can then be implemented to improve the training stage [Peters et al., 2010].

Model-based methods present some limitations that make their use challenging in tetralogy of Fallot patients. All reported studies showed good segmentation accuracy. However, they are all designed for a specific imaging modality (echocardiography [Noble and Boukerroui, 2006], MRI [Andreopoulos and Tsotsos, 2008] or

CT [Ecabert et al., 2008; Zheng et al., 2008]) and often tested on normal subjects. Because they heavily rely on prior knowledge, they cannot be applied directly to other imaging modalities or pathologies without relearning the underlying model. In [Zhang et al., 2010b], the authors developed an AAM model to segment the bi-ventricular myocardium in repaired ToF patients. Yet, that model cannot be applied directly on healthy subjects or adults with heart failure.

Another limitation is the sensitivity of model-based segmentation algorithms to initialisation and to local minima. To cope with the initialisation problem, some authors proposed to register an atlas of the heart to the image to process [Lorenzo-Valdes et al., 2004], to use generalised Hough transform to localise the heart in the images [Pednekar et al., 2006; Ecabert et al., 2008] or to employ machine learning object detection techniques to automatically detect reliable fiducials [Zheng et al., 2008; Lu et al., 2009; Zheng et al., 2009]. The risk of finding local minima is reduced by integrating additional constraints like temporal consistency [Montagnat and Delingette, 2005; Lynch et al., 2008] or coupled endocardium/epicardium segmentation [Kaus et al., 2004; Zhu et al., 2010]. Multi-resolution schemes have also been proposed [Montagnat and Delingette, 1998]. In [Ecabert et al., 2008] for instance, the authors successively detect the heart, rigidly align an atlas to it, which is deformed using piecewise affine transformation and finally refined using a shape-constrained parametric deformable model. Fast and efficient machine learning algorithms based on probabilistic boosting trees [Tu, 2005] have also been proposed to automatically find the optimal segmentation, yielding efficient segmentations. In [Zheng et al., 2008] for instance, the heart is first detected using marginal space learning. A 4-chamber model of the heart is then automatically positioned and refined to match the observed myocardium boundaries.

Nevertheless, the vast majority of the solutions relies on assumptions that may be violated by the extreme anatomy encountered in ToF patients. For this reason, we developed an interactive approach that can be applied to any situation, although more user inputs are required to get accurate delineations.

3.2 A Pipeline for the Dynamic Segmentation of the Myocardium

3.2.1 Overview of the Segmentation Pipeline

The segmentation of the bi-ventricular myocardium in dynamic cMRI is achieved in four main steps (Figure 3.1). First, the raw images are prepared for processing (Section 3.2.2). This step corrects the image artefacts and the large slice thickness that characterise cMRI. Next, the user manually delineates the boundaries of the left ventricle (LV), of the right ventricle (RV) and of the epicardium on the first frame of the preprocessed cardiac sequence (Section 3.2.3). This is achieved with an interactive 3D segmentation tool based on implicit variational surfaces. The manual delineations are then automatically propagated over the cardiac sequence by using diffeomorphic demons [Vercauteren et al., 2009], an efficient non-linear image

registration algorithm (see Chapter 4, page 61 for details), and simplex-based deformable model [Delingette, 1999] (Section 3.2.4). Finally, the dynamic myocardium is automatically reconstructed from the propagated delineations (Section 3.2.5).

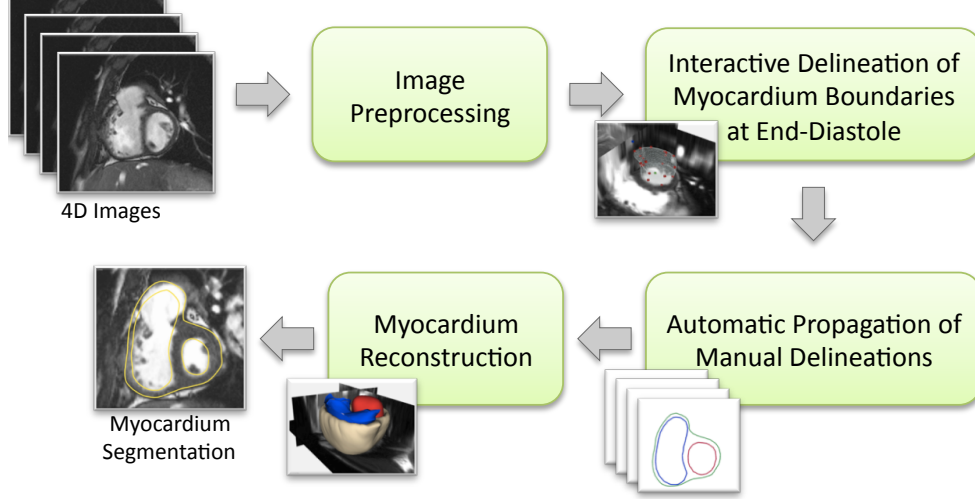


Figure 3.1: Main steps of the segmentation pipeline.

Notations In the following, we denote by I_k , $k = 0 \dots N - 1$ the N frames of the cardiac sequence. We assume that I_0 is the reference end-diastole time frame. M_{l_k} is the binary mask of the structure l (LV, RV, Epicardium or Myocardium) at the frame k . \mathcal{T}_{l_k} is the 3D triangulation of the structure l at the frame k . \mathbf{x} is the spatial coordinate vector $\mathbf{x} = (x, y, z)^T$.

3.2.2 Image Preprocessing

The purpose of this step is to prepare the images for the segmentation of the beating myocardium and further processing like the tracking of the myocardium. The idea is to correct the most common artefacts of standard cMRI: slice misalignment, large slice thickness and low contrast.

3.2.2.1 Semi-Automatic Correction of Slice Misalignment

The first important image artefact we need to deal with is slice misalignment due to inconsistent breath-holds (Figure 3.2, left panel). Automatically correcting slice misalignment is an active research topic in both image acquisition and image processing. The problem is ill-posed as one needs to replace the misaligned slices at their correct position in the 3D space, using the 3D information estimated from the other slices (which can also be misaligned). A common approach consists in using 2D-to-3D image registration algorithms [Chandler et al., 2008]. Others propose to use several views simultaneously (short axis, long axis, 4-chamber, etc.) to guide

the slice correction [Elen et al., 2010]. However, such images are not always available and the patients may have moved during the different acquisitions. During this thesis we had to deal with these artefacts, although their automatic correction was not the main topic of research. We thus simplified the problem and used existing tools.

To correct slice misalignment we assumed that a misaligned slice is only translated or rotated with respect to the image volume: the displacement is only 2D. Correcting that slice thus consists in estimating the 2D transformation that replaces it back into the volume. This is achieved by registering the misaligned slice with the previous slice in the 3D image using a 2D-to-2D image registration algorithm. Of course, assuming 2D displacements does not always hold as the misalignment is most probably three-dimensional. The automatic correction may thus be suboptimal. In case of failure, the correction can be manually adjusted using Isis¹, an interactive volume reconstruction software developed at Asclepios. It has to be noted however that during this thesis, no manual refinement has been needed.

Because several slices may be misaligned in the volume, the 2D rigid transformations are estimated recursively, from bottom to top. Let's denote n the number of slices of the 2D stack, S_0 the first slice, S_{n-1} the last slice and S_a the first misaligned slice. We first estimate the 2D transformation that aligns S_a to S_{a-1} , which we denote R_a . To this end, we use the block matching rigid registration algorithm proposed by [Ourselin et al., 2001]. This algorithm relies on a dense displacement field that minimises block matching differences between the two images to estimate the parameters of the global rigid transformation. This approach is well-suited for our purposes as it takes into account local displacements everywhere in the images. Information provided by other structures like the right ventricle, the liver or the ribs automatically helps the algorithm to preserve the longitudinal curvature of the heart. To increase the convergence speed and the robustness of the algorithm, a multi-resolution scheme with 5 levels is used. Once estimated, the transformation R_a is applied to the misaligned slice and we continue on the following misaligned slices to correct the entire volume (Algorithm 1). We finally apply the transformations R_a 's to the respective slices of the other time frames to correct the whole cardiac sequence consistently. Figure 3.2, right panel, shows a result of automated slice correction.

3.2.2.2 Isotropic Resampling of Images

Large slice thickness may reduce the accuracy of the interactive delineation and of the myocardium tracking. To avoid these limitations, we resample all the images of the corrected cardiac sequence to have isotropic voxel size. To this end, we take the smallest spacing and then resample the images along the other dimensions using trilinear interpolation. Let $d_x \times d_y \times d_z$ and $v_x \times v_y \times v_z$ be the dimensions and voxel spacings of the original image. We assume without loss of generality that the smallest spacing is $s = v_x$. Since the real dimensions in *mm* of the image must not

¹<http://www-sop.inria.fr/asclepios/software.php>

Algorithm 1 Automatic Correction of 2D Slice Misalignment**Require:** Cardiac sequence $I_{f=0\dots N-1}$ (N number of time frames)**Require:** List of misaligned slices in increasing order A

- 1: $\{Stage\ 1\}$
- 2: Split I_0 into n 2D slices $S_{i=1\dots n}$
- 3: **for all** $a \in A$ **do**
- 4: Estimate the 2D rigid transformation R_a that registers S_a to S_{a-1} using block-matching rigid registration algorithm [Ourselin et al., 2001].
- 5: Update the misaligned slice $S_a \leftarrow R_a \circ S_a$
- 6: $\{Stage\ 2\}$
- 7: **for all** $f \in 1 \dots N - 1$ **do**
- 8: Split I_f into n 2D slices $S_{i=1\dots n}$
- 9: **for all** $a \in A$ **do**
- 10: Apply R_a to S_a , $S_a \leftarrow R_a \circ S_a$
- 11: Rebuild the volume I_f from the n corrected slices $S_{i=1\dots n}$
- 12: **return** Corrected volumes I_f




Figure 3.2: Automatic correction of slice misalignment using slice-by-slice block-matching 2D rigid registration. One can see how the second slice has been corrected (yellow arrow).

change after resampling, the dimensions of the resampled image write:

$$d'_x = d_x \quad d'_y = \lceil d_y v_y / s \rceil \quad d'_z = \lceil d_z v_z / s \rceil \quad (3.1)$$

where $\lceil \cdot \rceil$ denotes the ceiling function. We then scan the voxels of the new isotropic image and compute their grey level intensity by using trilinear interpolation. An example of isotropic resampling is reported in Figure 3.3.

3.2.2.3 Histogram-Based Improvement of Image Contrast

The last step consists in improving the contrast of the images and in normalising the grey level intensities with respect to a fixed scale. Each image of the cardiac sequence is processed independently in order to have constant image intensities. The contrast is automatically enhanced by clamping the tails of the grey level histogram. Let l_k and L_k be the grey level intensities of the 1^{st} and 99^{th} quantiles of the histogram of the image I_k . Voxels whose intensity is lower than l_k or higher than

L_k are modified such that their intensity becomes l_k or L_k respectively. In other words, we automatically window the grey level intensities by saturating the very dark and very light voxels. Afterwards, the grey levels are normalised by scaling them between a fixed range, $[0 - 255]$ in our experiments, according to the scaling equation: $I_k(\mathbf{x}) = 255(I_k(\mathbf{x}) - l_k)/(L_k - l_k)$. Figure 3.4 illustrates the effects of this processing.




Figure 3.3: Image resampling to get isometric voxel spacing

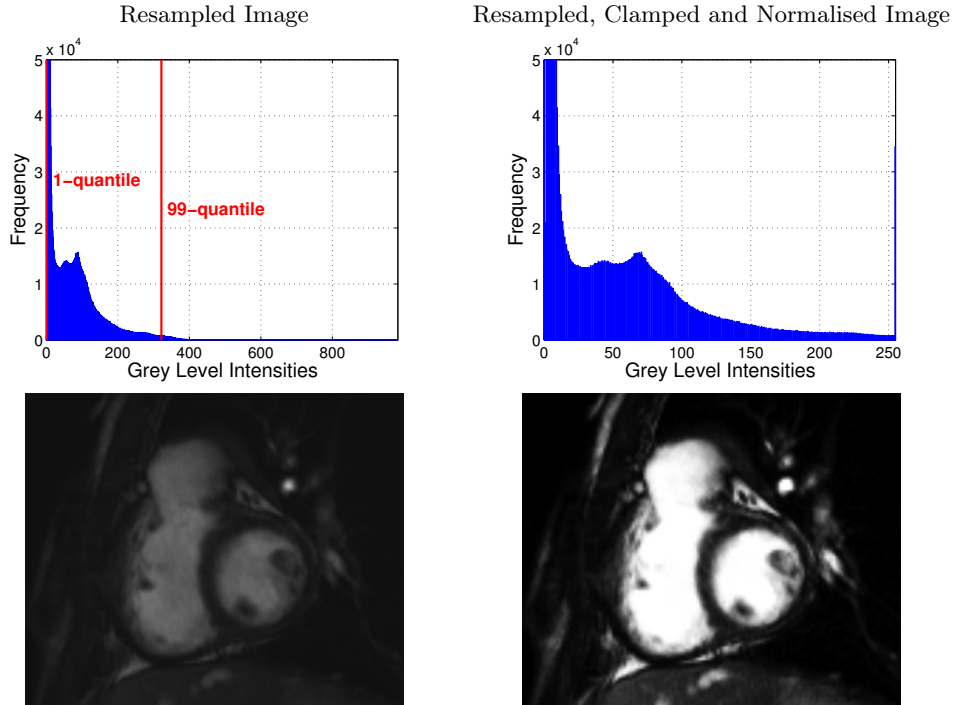


Figure 3.4: Image clamping and normalisation to get a consistent range of grey level intensities over the cardiac cycle. In these images, the contrast has been set to the minimum and maximum grey levels to highlight the differences in intensity dynamics.

3.2.3 Interactive 3D Segmentation of Myocardium Boundaries using Interpolating Implicit Surfaces

The next step of the segmentation pipeline is the delineation of the boundaries of the heart. For that matter, we developed an interactive 3D surface modeller to delineate the left endocardium (LV), the right endocardium (RV) and the bi-ventricular epicardium on the first frame of the cardiac sequence. The idea is to sculpt 3D surfaces by placing control points inside, on and outside the boundaries to delineate (Figure 3.5). Our approach relies on interpolating implicit surfaces, also known as variational implicit surfaces [Turk and O’Brien, 2002]. The control points are strong constraints of the implicit function. In the current implementation, no image information is used during the delineation: the tool relies entirely on the user interactions.

This approach has been preferred to the standard slice-by-slice painting as it naturally handles the 3D nature of the heart. The user works directly in the 3D space, without switching between the different 2D views. The segmentation is spatially consistent and easier to achieve. In addition, our approach automatically provides smooth meshes of the delineated boundaries and an accurate mask of the inner volume of the surface. In the following we detail the underlying principles.

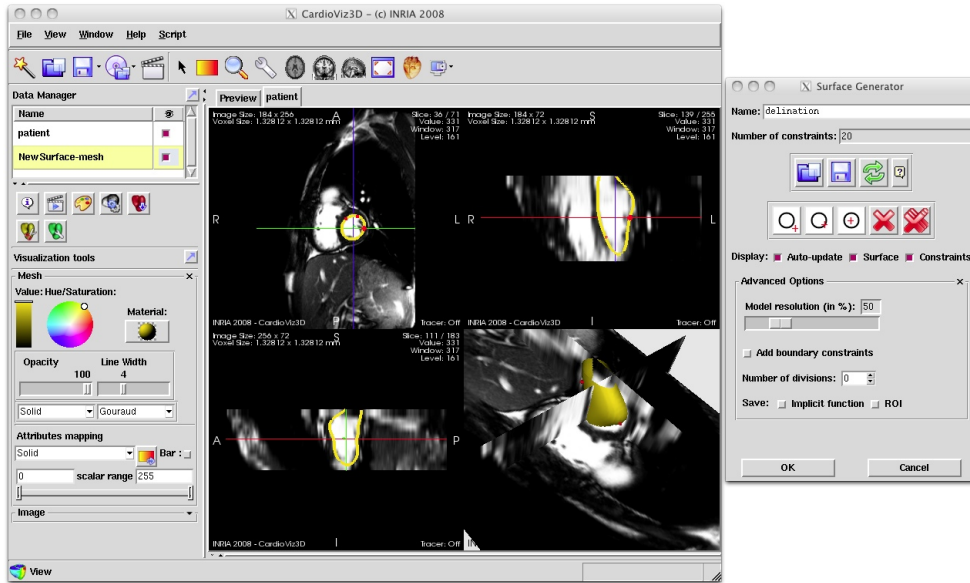


Figure 3.5: Interface of the interactive 3D segmentation tool.

3.2.3.1 Interpolating Implicit Surfaces

An implicit surface \mathcal{S} is defined as the 0-level set of a scalar-valued function $f(\mathbf{x})$ of the image space $\Omega \in \mathbb{R}^3$ [Bloomenthal and Bajaj, 1997]. For instance, a circle of radius one is the 0-level set of the implicit function $f(\mathbf{x}) = 1 - \|\mathbf{x}\|$ (Figure 3.6). Proposed by [Turk and O’Brien, 2002], interpolating implicit surfaces interpolate a

set of control points while minimising the curvature of the surface. They are smooth by construction and can represent 3D objects of any topology. Let \mathbf{c}_i , $i = 1 \dots m$, be the spatial positions of m control points, or constraints, defined by the user and $h_i \in \mathbb{R}$ their scalar value. The interpolating surface defined by the \mathbf{c}_i 's is the 0-level set of the implicit function f that optimises the constrained problem:

$$\begin{cases} f = \operatorname{argmin}_f \left(\int_{\Omega} \frac{\partial^2 f}{\partial x^2}(\mathbf{x}) + \frac{\partial^2 f}{\partial x \partial y}(\mathbf{x}) + \frac{\partial^2 f}{\partial y^2}(\mathbf{x}) d\mathbf{x} \right) \\ h_i = f(\mathbf{c}_i), \quad \forall i = 1 \dots m \end{cases} \quad (3.2)$$




Figure 3.6: A circle of radius 1 is the 0-level set of the implicit function $f(\mathbf{x}) = 1 - \|\mathbf{x}\|$

Equation 3.2 is known as the scattered data interpolation problem [Wendland, 2005]. A large literature in the computer graphics community is devoted to the efficient resolution of this equation (see [Turk and O'Brien, 2002; Wendland, 2005] and references therein). An elegant solution consists in using functions that are linear combinations of radial basis functions (RBF) centred on the control points [Schaback, 1995; Turk and O'Brien, 1999a]. RBF are radially symmetric functions about a single point, called centre. The most common RBF for surface reconstruction are the functions $r(\mathbf{x}) = \|\mathbf{x}\|^2 \log(\|\mathbf{x}\|)$ in 2D and $r(\mathbf{x}) = \|\mathbf{x}\|^3$ in 3D [Schaback, 1995; Wendland, 2005]. More advanced RBF have been proposed to reconstruct surfaces from very dense scattered points [Morse et al., 2001], 4D interpolation [Turk and O'Brien, 1999b] or even anisotropic distributions [Dinh et al., 2001]. But for our application, where few control points are used (< 200), the standard 3D RBF $r(\mathbf{x}) = \|\mathbf{x}\|^3$ is largely sufficient [Schaback, 1995]. With this framework, the interpolating implicit function is of the form:

$$f(\mathbf{x}) = \sum_{j=1}^m \lambda_j r(\mathbf{x} - \mathbf{c}_j) + P(\mathbf{x}) \quad (3.3)$$

In this equation, the λ_j 's are weight factors and the polynomial $P(\mathbf{x}) = p_0 + p_1x + p_2y + p_3z$ accounts for the linear and constant portions of f . A minimiser of Equation 3.2 that writes as in Equation 3.3 is calculated by finding the weights λ_j and the coefficients p_i such that f verifies all the constraints [Turk and O'Brien, 2002;

Wendland, 2005]:

$$h_i = f(\mathbf{c}_i) = \sum_{j=1}^m \lambda_j r(\mathbf{c}_i - \mathbf{c}_j) + P(\mathbf{c}_i), \forall i$$

with the conditional constraints:

$$\begin{aligned} \sum_{j=1}^m \lambda_j &= 0, & \sum_{j=1}^m \lambda_j c_{jx} &= 0, \\ \sum_{j=1}^m \lambda_j c_{jy} &= 0, & \sum_{j=1}^m \lambda_j c_{jz} &= 0 \end{aligned}$$

We have $m + 4$ linear equations and $m + 4$ variables (the λ_j 's and the p_i 's), which are gathered into the linear system:

$$\begin{bmatrix} r_{11} & r_{12} & \dots & r_{1m} & 1 & c_{1x} & c_{1y} & c_{1z} \\ r_{21} & r_{22} & \dots & r_{2m} & 1 & c_{2x} & c_{2y} & c_{2z} \\ \vdots & \vdots & & \vdots & \vdots & \vdots & \vdots & \vdots \\ r_{m1} & r_{m2} & \dots & r_{mm} & 1 & c_{mx} & c_{my} & c_{mz} \\ 1 & 1 & \dots & 1 & 0 & 0 & 0 & 0 \\ c_{1x} & c_{2x} & \dots & c_{mx} & 0 & 0 & 0 & 0 \\ c_{1y} & c_{2y} & \dots & c_{my} & 0 & 0 & 0 & 0 \\ c_{1z} & c_{2z} & \dots & c_{mz} & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \\ \vdots \\ \lambda_m \\ p_0 \\ p_1 \\ p_2 \\ p_3 \end{bmatrix} = \begin{bmatrix} h_1 \\ h_2 \\ \vdots \\ h_m \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad (3.4)$$

In this equation, r_{ij} is the value at \mathbf{c}_i of the RBF centred in \mathbf{c}_j , $r_{ij} = r(\mathbf{c}_i - \mathbf{c}_j)$.

When all the control points are distinct, the submatrix of the r_{ij} 's is definite positive and the Equation 3.4 has one unique solution. If few constraints are used, as in our case, direct solvers based on LU decomposition for instance can be used to compute the weights of the interpolating implicit function $f(\mathbf{x})$ that satisfies all the constraints. Then, the interpolating implicit function f is evaluated at every point of a sampling grid Γ of the image space Ω . Finally, the marching cubes algorithm [Lorensen and Cline, 1987] is applied to compute the 3D triangulated mesh \mathcal{T} of the 0-level set of f . The interpolation algorithm is summarised in Algorithm 2.

Algorithm 2 Interpolating Implicit Surfaces

Require: Spatial position of m control points \mathbf{c}_i and their value h_i , $i = 1 \dots m$

- 1: Compute the implicit function parameters λ_i , p_0 , p_1 , p_2 and p_3 (Equation 3.4)
 - 2: Evaluate the implicit function f on the sampling grid Γ
 - 3: Compute the interpolating triangulated mesh \mathcal{T} of the 0-level set of f (marching cubes algorithm [Lorensen and Cline, 1987]).
-

From a computational point of view, the complexity of the interpolation algorithm only depends on the number of constraints m and on the resolution of the sampling grid on which the implicit function f is evaluated. The computational complexity of the LU solver increases in $O(m^3)$ as m increases. In parallel, the

computation time related to the evaluation of the implicit function f at every point of the sampling grid Γ increases linearly in m . The resolution of the sampling grid is also critical as the implicit function must be evaluated on the entire grid to extract the implicit surface. Yet, these limitations are not observed in practice since the number of control points that are placed by the user is low ($m < 200$) in contrast with the typical data sets processed by the computer graphics community ($m > 10000$) [Turk and O'brien, 2002; Morse et al., 2001].

3.2.3.2 Interactive 3D Segmentation

The interpolating implicit surfaces are the heart of the interactive surface modeller. The tool allows placing three types of points in the 3D image space (Figure 3.7):

- on points:** They identify the voxels through which the interpolating surface must pass. Their value is 0.
- inside points:** They identify the inner region of the surface (the blood pools for instance). Their value is fixed to $+1$.
- outside points:** They identify the outer region of the surface (the background). Their value is fixed to -1 .




Figure 3.7: The implicit surface is interactively sculpted by placing control points inside, on and outside the boundary to delineate. Pairs of on/outside points can be used to control the surface normal.

The interpolating surface is defined as the 0-level set of the implicit function f computed according to the Algorithm 2. Each time the user adds, deletes or moves a control points, the interpolation algorithm is called and the 0-level set triangulation \mathcal{T} is updated in real-time. The mesh related to the 0-level set is immediately visible in the 3D space and in all the 2D views of the image. To increase the accuracy of the segmentation and to control the speed of the algorithm, the sampling grid Γ is isotropic and embeds the image space Ω . The resolution of Γ is computed from the smallest voxel spacing of the image. It can be made coarser, to speed up the algorithm when images are large but at the price of a lower accuracy, and conversely. This strategy enables the user to accurately segment images with large slice thicknesses.

A useful side-product of the method is its ability to provide automatically an accurate mask of the segmented surface. By construction, the interpolating implicit function f is positive inside the surface and negative outside. A binary mask of the inner region is therefore obtained by a simple thresholding (Figure 3.8). This feature is particularly useful for propagating the segmentation to the other frames of the cardiac sequence, as we shall see in Section 3.2.4.

At the end of that stage, one has the 3D triangulations of the LV, RV and epicardium at the end-diastole time frame (\mathcal{T}_{LV_0} , \mathcal{T}_{RV_0} and \mathcal{T}_{Epi_0}) and their related binary masks (M_{LV_0} , M_{RV_0} and M_{Epi_0}) (Figure 3.9).




Figure 3.8: *Left panel:* Example of an implicit function. The interpolating implicit surface is defined by the 0-level set. *Right panel:* The mask of the interpolating implicit surface is obtained by thresholding the implicit function.

Guidelines for Fast and Intuitive Surface Modelling

1. Starts by placing an *inside point* within the structure to segment
2. Sculpt the 3D surface by positioning *on points*
3. To control the curvature of the surface, prefer using *inside* and *outside points*
4. To control the surface normal, place pairs of *on point* and *outside point* very close to each other. The segment defined by the two points define the surface normal (Figure 3.7)
5. During the sculpting, reduce the resolution of the sampling grid Γ if necessary to maximise the reactivity of the tool.
6. At the end of the process, set the resolution of the sampling grid Γ according to the required resolution of the final segmentation.

These guidelines are consistent with the strategies listed in [Turk and O'brien, 2002]. However, in [Turk and O'brien, 2002], the different types of constraints are used separately. In our application, the user can use all of them simultaneously to precisely delineate the cardiac boundaries as illustrated in Figure 3.10².

²A movie illustrating the interactive segmentation tool in action is available at: <http://www-sop.inria.fr/asclepios/software/CardioViz3D/howto/>

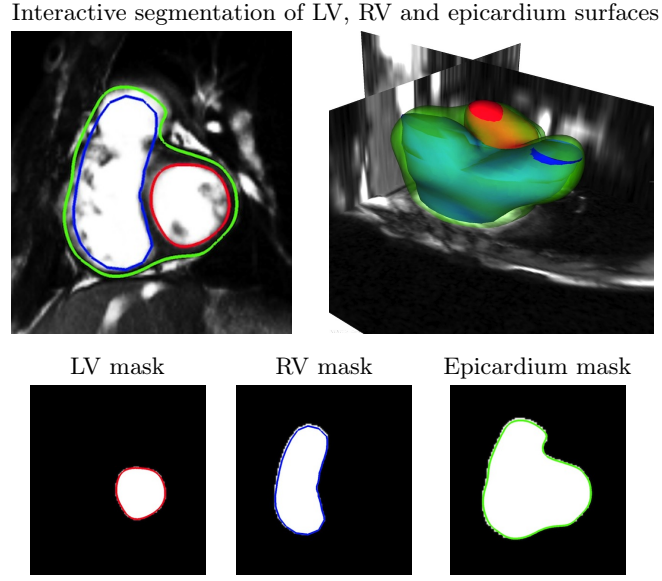


Figure 3.9: Interactive delineation of left endocardium (*in red*), right endocardium (*in blue*) and epicardium (*in green*). The surfaces are sculpted in 3D using the interactive 3D segmentation tool. The binary masks are automatically computed by thresholding the interpolating implicit function.

3.2.4 Smooth Propagation of the Segmentation over the Cardiac Sequence

Once the 3D surfaces of the LV, RV and epicardium are defined, we automatically propagate them over the cardiac cycle (Figure 3.1). A standard strategy consists in estimating the cardiac motion from the images by using non-linear image registration. The segmentations are then warped with the estimated deformations and propagated over the entire sequence [Perperidis et al., 2005; Bistoquet et al., 2007]. However, the quality of the dynamic segmentation highly depends on the quality of the motion estimation. Noisy deformations yield irregular segmentations, with holes and bumps. To cope with this issue, we developed an *ad-hoc* two-step propagation approach:

1. The heart is tracked using the diffeomorphic demons, a non-linear registration algorithm [Vercauteren et al., 2009] (Chapter 4). Demons yields dense deformation fields that map the first frame of the sequence to the following frames. The LV, RV and epicardium masks are then warped using the estimated deformations.
2. The binary masks are iteratively meshed using simplex deformable models [Delingette, 1999] to ensure smooth and temporally consistent segmentations. As a side product, point-to-point correspondence over the temporal sequence is obtained.




Figure 3.10: Interactive 3D segmentation of the left ventricle endocardium. From top to bottom, the user adds control points inside (*green points*), on (*red points*) and outside (*blue points*) the left ventricle endocardium to interactively sculpt the 3D mesh (19 points used in this example).

3.2.4.1 Step 1: Cardiac Motion Estimation

The transformation ϕ_k that match the initial frame I_0 to the frame I_k is estimated with the diffeomorphic demons [Vercauteren et al., 2009]. We refer the reader to Chapter 4, which provides a detailed description of that algorithm along with some contributions for cardiac motion estimation. The transformations ϕ_k are recursively estimated to increase registration accuracy and to minimise error propagation. Assume that the transformation ϕ_k that registers I_0 to I_k is known. We first compute the frame-by-frame transformation $\Phi_{k \rightarrow k+1}$ that matches I_k to I_{k+1} . We then estimate the transformation ϕ_{k+1} by initialising the demons algorithm with the composed transformation $\phi_k \circ \Phi_{k \rightarrow k+1}$. This strategy is more stable than frame-by-frame propagation and more robust to large deformations than directly registering I_0 to I_k . Furthermore, it can be partially parallelised by computing all the frame-by-frame registrations on a cluster of computers. The overall computational time is therefore not much longer than direct tracking. In the following, the ϕ_k 's are known.

3.2.4.2 Step 2: Segmentation Propagation using Simplex Deformable Models

To propagate the manual delineations of the myocardium over the cardiac sequence, one could directly deform the 3D surfaces with the previously computed transformations. Each vertex of the mesh would be transported according to the displacement estimated at its spatial position. However, this approach is highly sensitive to the quality of the transformation and often yields irregular and bumpy meshes. We tackle this limitation by propagating instead the binary masks of the surfaces M_{LV_k} , M_{RV_k} and M_{Epi_k} (Figure 3.15 and 3.16). The dynamic 3D surfaces are then computed by iteratively segmenting the masks with 2D simplex deformable models, which regularise the boundaries and ensure temporal consistency.

About Deformable Models

Algorithms based on deformable models are efficient techniques to segment objects in images. Appeared in [Kass et al., 1988], they are curves or surfaces that evolve within the image according to forces that ensure their smoothness and drive them towards the boundaries of the object to delineate. Two main types of deformable models can be distinguished: parametric deformable models [Terzopoulos, 1984; Kass et al., 1988] and geometric deformable models [Caselles et al., 1997; Malladi et al., 1995]. We refer the reader to [Montagnat et al., 2001] for a detailed review. Parametric deformable models, or snakes, rely on a parametric representation of the surface to deform it. The convergence is fast, which makes the method suitable for real-time applications. However, the driving forces must derive from potential energies. Furthermore, the model parameterisation can degenerate (point accumulation for instance), thus requiring regular remeshing steps. Geometric deformable models rely on the curve evolution theory and on the level set method [Sethian et al., 1999]. In this approach, the contours are defined as the 0-level set of

an implicit function as with the interpolating implicit surfaces. They can be driven by forces that do not necessarily derive from a potential energy and, as they are parameter-free, can handle topology changes easily. Furthermore, the regularisation is based on curvature flows, which yields smoother results than parametric models.

First introduced by [Delingette, 1999], simplex deformable models have the advantages of both worlds. A simplex model is a parametric model represented by the dual of triangulation. To each face of a triangulation corresponds a vertex of a simplex mesh. Hence, every vertex of a closed simplex mesh has exactly 3 neighbours (Figure 3.11). Similar to snakes, simplex models are driven by an *external force*, which moves the vertices along the surface normal, proportionally to the distance to the closest point of interest [Kimia et al., 1992]. Similar to the geometric models, simplex models are regularised using *internal force* based on surface curvature, which ensures mesh smoothness. Overall, simplex deformable models present four key advantages over the snakes:

1. the external force does not necessarily derive from potential energies
2. the surface curvature is simple to compute, resulting in efficient internal forces based on curvature that do not necessarily derive from a potential energy,
3. the distribution of the vertices can be controlled during the surface evolution for increased accuracy. No remeshing is needed.
4. the regularisation preserves the shape of the mesh, there is no “shrinking-effect”

The reader is referred to [Delingette, 1999] for further details.




Figure 3.11: Duality between triangulation (plain lines, circles) and 2D simplex meshes (dashed lines, squares). To each face of the triangulation corresponds a vertex of the simplex mesh, which has exactly 3 neighbours. (Image Courtesy of [Delingette, 1999])

A Simplex Deformable Model for Mask Regularisation

In our pipeline, we segmented the propagated masks using a standard simplex model. The internal force minimises the surface curvature but preserves its global area [Delingette, 1999]. We also employ an isotropic face constraint (all the faces of the simple mesh should have the same area) to get as regular meshes as possible. The external force is based on image isocontours (Figure 3.12). Assume that the background intensity is 0 and the foreground grey level is 255. The external force at a given vertex P of the simplex mesh is oriented along the simplex normal and proportional to the distance between P and the closest point M on the 128 isocontour.




Figure 3.12: Each vertex P of the simplex deformable model is driven towards the boundaries of the object to segment by an external force \mathbf{f}_e perpendicular to the simplex model and proportional to the distance PM .

Propagation of Segmentation over the Cardiac Sequence

Let M_{LV_k} , M_{RV_k} and M_{Epi_k} be the masks propagated over the cardiac sequence using the cardiac deformations estimated during the first stage. We first convert the manual end-diastole triangulations \mathcal{T}_{LV_0} , \mathcal{T}_{RV_0} and \mathcal{T}_{Epi_0} into simplex meshes Σ_{LV_0} , Σ_{RV_0} and Σ_{Epi_0} . Let Σ_0 be one of these simplex meshes without loss of generality. The segmentation of the mask M_{k+1} is initialised using the simplex model Σ_k that results from the segmentation of the mask M_k . To preserve point correspondence, the parameters of the simplex evolution are set such that the simplex model moves slowly and smoothly towards the object boundaries (large internal force and low external force). The rigidity of the model is strong to regularise holes and bumps.

The first three columns of Figure 3.15 and Figure 3.16 illustrate the dynamic segmentation of the left and right endocardia and the epicardium in a patient with repaired ToF. From these images one can see how the simplex deformable model regularises the propagated masks. The surfaces are smoother than the propagated masks, which are particularly noisy at end-systole when the papillary muscles visually “merge” with the myocardium, thus misleading the registration algorithm.

3.2.5 Reconstruction of the Beating Myocardium

We finally reconstruct the dynamic myocardium from the smooth surfaces. To that end, we rasterise the surfaces into binary masks whose voxel values are either zero or one. Myocardium reconstruction is then easily performed by applying the arithmetic formula on the voxels:

$$M_{myo_k}(\mathbf{x}) = M_{Epi_k}(\mathbf{x}) - [M_{LV_k}(\mathbf{x}) + M_{RV_k}(\mathbf{x})] \quad (3.5)$$

Intuitively, we remove from the epicardium mask the voxels that are laying within the endocardia.

However, applying directly this formula may yield holes in the myocardium if the endocardial surfaces or the epicardium have not been propagated properly. To tackle this issue, we reconstruct the myocardium such that its thickness is at least 3mm, the thickness of a normal right ventricle wall in children. During our experiments, we observed that the propagation of the endocardial surfaces is much more reliable than the propagation of the epicardium because of the high contrast between the cardiac muscle and the blood. We thus keep the endocardia and modify the epicardium such that the thickness of the reconstructed myocardium is at least 3mm. An Euclidian distance map is computed from the boundaries of the RV mask M_{RV_k} by using Chamfer algorithm [Borgefors, 1996]. Let d_{RV_k} be this map. Inside the mask, $d_{RV_k} = 0$. Outside the mask, $d_{RV_k}(\mathbf{x}) > 0$. A voxel belongs to the myocardium if it belongs to the epicardium mask or if its distance to the boundaries of the RV mask is lower than 3mm. The reconstruction formula thus writes:

$$M_{Myo_k}(\mathbf{x}) = [M_{Epi_k}(\mathbf{x}) \vee (d_{RV_k} \leq 3)] - [M_{LV_k}(\mathbf{x}) + M_{RV_k}(\mathbf{x})] \quad (3.6)$$

Figure 3.13 illustrates that arithmetic operation on the images.

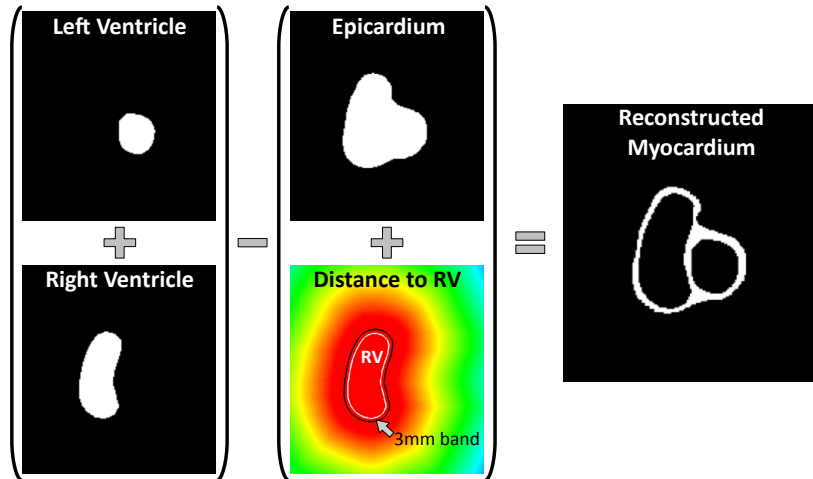


Figure 3.13: The myocardium is reconstructed from the LV, RV and epicardium masks. 3-mm myocardium thickness is ensured through Chamfer distance maps from the RV mask.

Myocardium reconstruction is performed for all the frames k of the cardiac sequence. The user can also use a valve plane to cut the bi-ventricular myocardium at the base for simulation purposes for instance. This can be achieved using arithmetic on images as previously: $M_{Myo_k}(\mathbf{x}) = [M_{Epi_k}(\mathbf{x}) \mid (d_{RV_k} \leq 3)] - [M_{LV_k}(\mathbf{x}) + M_{RV_k}(\mathbf{x})] * (M_{valve})$, where M_{valve} is a binary image whose voxel grey level is zero above the valve plane and one below. Afterwards, the triangulated mesh \mathcal{T}_{Myo} related to the myocardium is extracted from the first frame using meshing algorithms such as marching cubes or CGAL³. Finally, the mesh is propagated over the cardiac sequence using the previous simplex deformable model to ensure mesh consistency and point correspondence.

3.3 Results on Tetralogy of Fallot

3.3.1 Segmentation of the Myocardium on cMRI of Patients with Repaired Tetralogy of Fallot

Figures 3.14, 3.15 and 3.16 illustrate the dynamic segmentation of the bi-ventricular myocardium of a patient with repaired tetralogy of Fallot. One can identify from the 3D representation of patient's heart the aneurysm of the right ventricle outflow tract (RVOT).

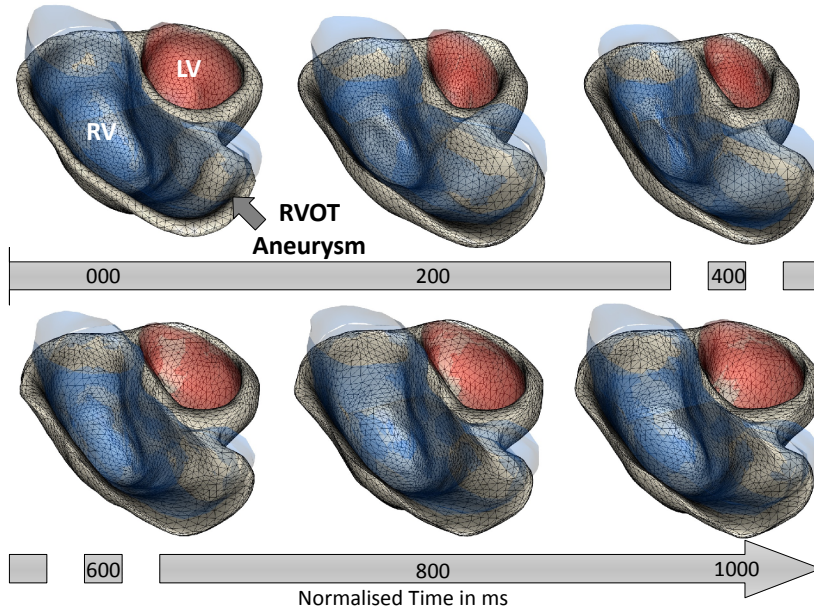


Figure 3.14: Example of dynamic 3D segmentation of the myocardium (*brown*), the left ventricle (*LV, in red*) and the right ventricle (*RV, in blue*). Observe the aneurysm at the right ventricle outflow tract (RVOT) often visible in patients with repaired tetralogy of Fallot.

³<http://www.cgal.org>

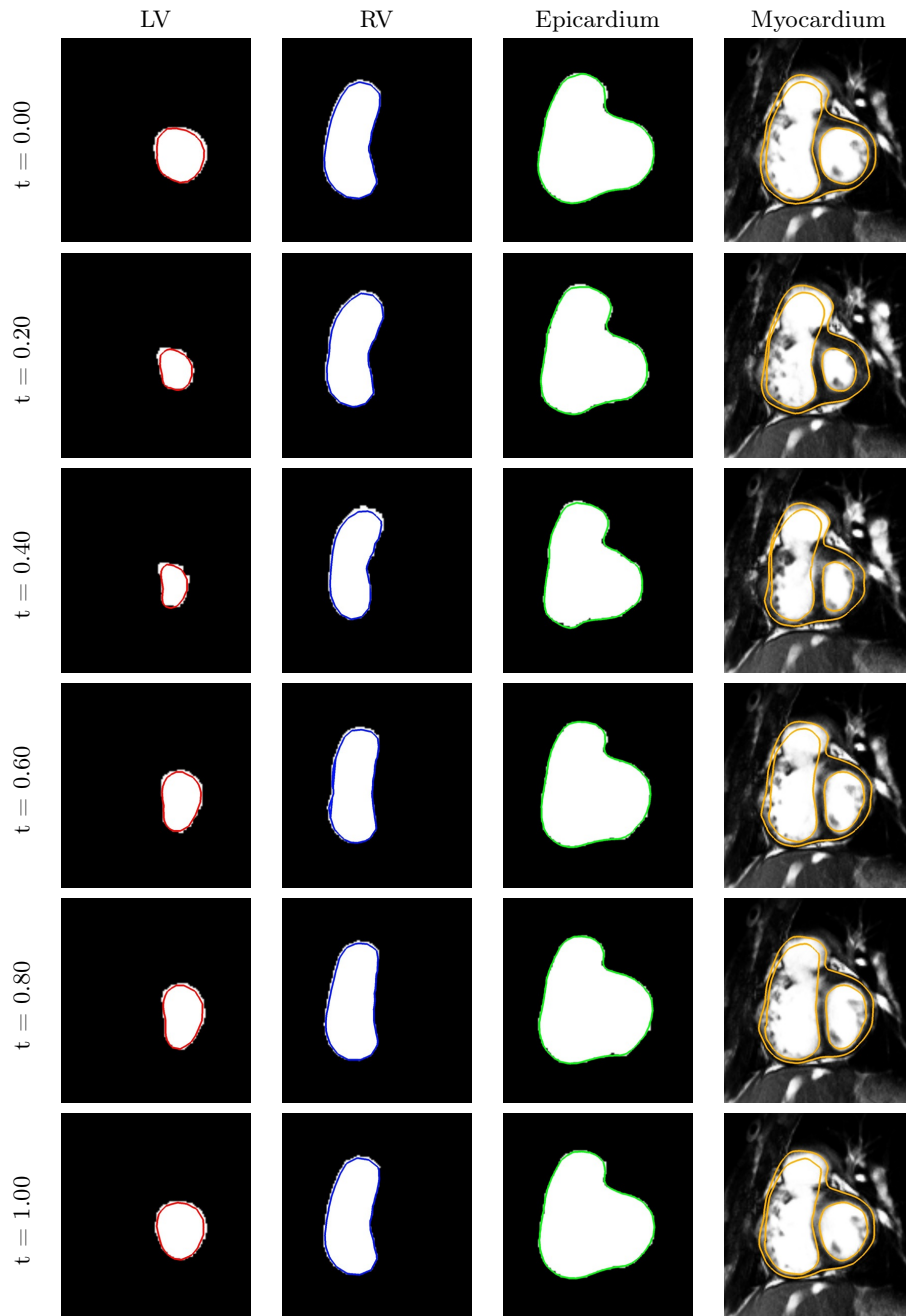


Figure 3.15: Short axis view of the dynamic myocardium segmentation in a patient with repaired tetralogy of Fallot. The time is normalised to 1 s. The first free columns show the propagated masks and the final simplex deformable models that regularise them. The last column shows the reconstructed myocardium segmentation. The segmentation is visually satisfying, in particular when considering the challenges raised by cMRI and ToF patients.




Figure 3.16: Long axis view of the dynamic myocardium segmentation in a patient with repaired tetralogy of Fallot. The time is normalised to 1 s. In this view the impacts of the simplex regularisation are more evident. In particular, at end-systole ($t = 0.20$ s and $t = 0.40$ s) the simplex deformable model smooths the bumps and fills the holes of the propagated mask. In that patient the myocardium was cut at the base using a user-defined valve plane (*see text for details*)

3.3.2 Quantification of the Cardiac Function from Segmentation

The dynamic segmentation of the myocardium enables to compute key clinical parameters about the cardiac function. For instance, one can easily compute the blood pool volumes and how they vary over time (Figure 3.17, left panel). From the volumes at end-diastole (ED) and end-systole (ES), one computes the ejection fraction (EF) of each ventricle, a key clinical parameter that quantifies the cardiac pump efficiency. EF is calculated as:

$$EF = (Volume(ED) - Volume(ES)) / Volume(ED) * 100 \quad (3.7)$$

In addition to volume quantification, the dynamic segmentation of the myocardium enables to assess the radial displacements of the muscle over time. Such a measurement is useful to detect cardiac asynchrony and abnormal radial motion patterns. In virtue to the point correspondences provided by the smooth iterative simplex-based propagation, we can estimate the displacement of any vertex of the mesh with respect to the end-diastolic position of reference. The displacements are then projected along the surface normal at the vertex position to get the radial displacement. Figure 3.17, right panel, illustrates the radial displacements of the heart in our patient. The dyskinetic motion of the RVOT aneurysm is immediately identified. At end-systole, when the heart is contracted (inward motion highlighted in red), the outflow tract goes outwards (in blue).

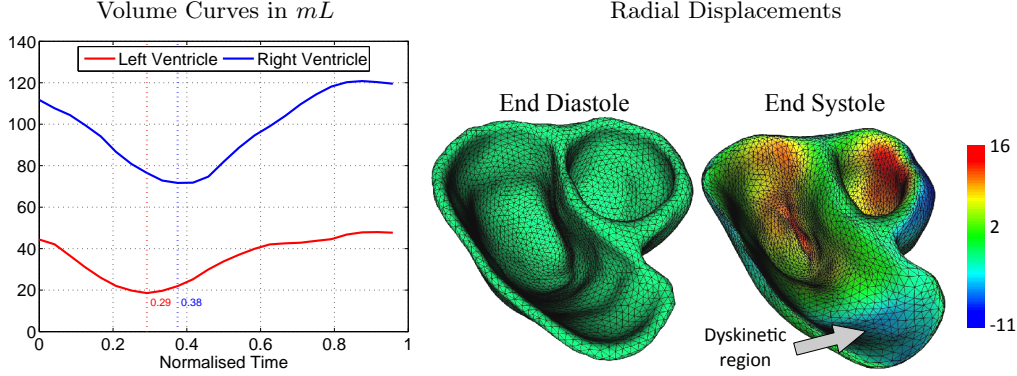


Figure 3.17: *Left panel*: Volume curves computed from the dynamic segmentation. From these curves we quantify a normal LV ejection fraction (61%) but a below-normal RV value (41%). We also observe a left/right dyssynchrony (vertical bars, about 100 ms shift between LV and RV end-systole peaks). *Right panel*: Radial displacements. One can see the dyskinetic motion of the right ventricle outflow tract, which dilates (in blue) when the heart contracts (in red).

3.3.3 Examples on Other Pathologies and Imaging Modalities

This segmentation pipeline has been successfully applied to other pathologies and imaging modality. In particular, it has been applied on CT images of patients with

heart failure [Peyrat et al., 2009] (Figure 3.18, A), on MR images of patients with heart failure [Billet et al., 2009; Sermesant et al., 2009] (Figure 3.18, B), on healthy subjects [Billet et al., 2008; Mansi et al., 2009d] (Figure 3.18, C) and even on animals [Chabiniok et al., 2009] (Figure 3.18, D).

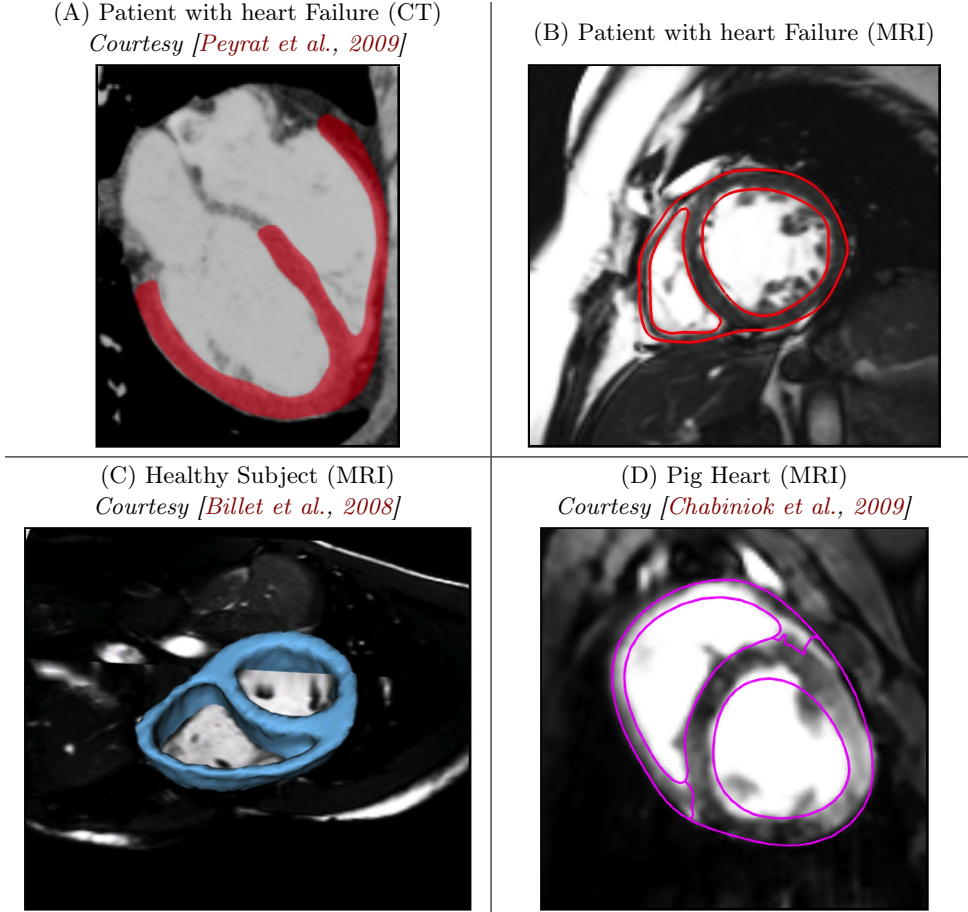


Figure 3.18: Other applications of the proposed segmentation pipeline

3.4 Discussion

Most of the contributions of this thesis rely on 3D surfaces of the patients' cardiac anatomy. Developing a segmentation pipeline was therefore needed although not the main topic of research. We decided to develop a modular framework that combines specific algorithms that already exist. The choice of every component was guided by the difficulties of the problem. In particular, we wanted a tool that could cope with any imaging modality, but most especially cMRI as these images were the only available during this thesis. We wanted to quickly segment the myocardium of any subject, but with particular focus on tetralogy of Fallot patients, whose right ventricle anatomy challenged most of the state-of-the-art algorithms. Because no automatic tools that were available could cope with such a variability, we implemented

an interactive 3D surface modeller and made it available to the community as part of the free software CardioViz3D⁴ [Toussaint et al., 2008]. The delineated surfaces were then propagated automatically using non-linear image registration and simplex deformable models to ensure spatial and temporal consistency. Qualitatively, the results were satisfying and adapted to the subsequent analyses. Unfortunately, the lack of ground truth prevented quantitative evaluations, in particular regarding the automatic propagation. Finally, the modularity of the framework enabled us to integrate results from other segmentation algorithms when available.

The interactive surface modeller is probably the most important contribution of this section. We use it for the interactive segmentation of anatomical structures, as an interactive sketcher [Mao et al., 2009]. This contrasts with previous approaches that use variational implicit surfaces to reconstruct surfaces from sparse anatomical data [Yoo et al., 2001]. As it has been shown in Section 3.3.3, the surface modeller has been employed with success by other users [Billet et al., 2008; Toussaint et al., 2008; Sermesant et al., 2009; Billet et al., 2009; Peyrat et al., 2009] and [Chabiniok et al., 2009]. Working directly in the 3D space, it yields segmentations that are smooth and spatially consistent. The left ventricle is particularly easy to segment thanks to its ellipsoidal shape that suits perfectly the implicit surface framework.

The modeller has room for improvements. At the present time, it relies exclusively on user inputs, no image information is used. An interesting development thus consists in adding image information to interactively guide the sculpted surface to the closest boundaries. An idea would be to transform, at the end of the modelling process, the interpolating implicit surface into a simplex deformable model. The model would then be driven towards image boundaries but constrained to pass through the *on* control points placed by the user [Delingette et al., 2001]. Another limitation of the tool raises from the use of 3D thin-plate RBF, which are easy to implement and fast to compute but have global support. Because they have non-null values everywhere in space, each constraint has an influence everywhere in the image. This can yield undesired behaviours during the surface modelling. The shape of the surface can globally change when adding a single point. One can alleviate this limitation by using more complex RBF with compact support, as proposed by [Schaback, 1995; Morse et al., 2001]. The influence of each constraint would become local, the user thus controls the local shape of the surface perfectly. Additionally, despite the increased complexity of the RBF, the computational efficiency would be improved as the evaluation of the implicit function at a given voxel of the sampling grid would require evaluating only the few neighbouring RBF.

Other interesting developments consist in improving the propagation of the manual delineation over the cardiac sequence. In our pipeline, no prior knowledge is used about the cardiac motion. The masks are directly propagated, without constraints. A first idea consists in using the incompressible demons algorithm we present in the Chapter 4 to improve the accuracy of the recovered deformations. Alternatively, reconstructing the myocardium from the propagated surfaces could also be done under the incompressibility constraint, like in [Zhu et al., 2010], which could yield

⁴<http://www-sop.inria.fr/asclepios/software/CardioViz3D/>

more accurate results than our *ad-hoc* minimal-thickness rule. The tracking algorithm can also be improved by enforcing temporal constraints as in [Declerck et al., 1997; Perperidis et al., 2004; Delhay et al., 2007; Yang et al., 2008b; Ionasec et al., 2009, 2010]. The simplex deformable model could also be constrained temporally [Montagnat and Delingette, 2005]. Nevertheless, the more complex the algorithm is, the more priors are integrated and the less general the algorithm becomes.

Part II

DIAGNOSIS: QUANTIFICATION OF CARDIAC DEFORMATION ON ANATOMICAL IMAGES

Efficient Elastic Incompressible Image Registration

Contents

4.1 Motivation	62
4.2 Background: Log-Domain Diffeomorphic Demons	65
4.3 From Gaussian Regularisation to Elastic-Like Regularisation	67
4.3.1 Insights into LogDemons Gaussian Regularisation	67
4.3.2 Elastic-Like LogDemons	68
4.4 Incompressible LogDemons	70
4.5 Experiments and Results	73
4.5.1 Incompressibility of Deformations Parameterised by Divergence-Free Velocities	73
4.5.2 Global Incompressibility Recovery Using Elastic Regularisation	74
4.5.3 Local Incompressibility Recovery Using Volume-Preserving Constraint	77
4.6 Discussion	80
4.7 Appendices	83

In Chapter 3 we proposed a pipeline for delineating the myocardium from medical images. The segmentation gives us hints on the heart shape and apparent motion. Yet, the heart is a moving organ and quantifying its intrinsic three-dimensional deformation is crucial for a comprehensive evaluation of its condition [Moore et al., 2000]. Imaging modalities have been developed to measure the cardiac deformation but when they are not available, little can be provided to the clinicians to assist their diagnosis. Unfortunately, this situation is common in paediatrics.

This chapter and Chapter 5 tackle the objective of estimating the cardiac deformation from dynamic anatomical images like standard cMRI. To cope with the scarce textures and the low image quality, we constrain a non-linear image registration algorithm based on demons approach to provide elastic and incompressible deformations, two major features of the myocardium [Glass et al., 1991]. This leads us to closely investigate demons regularisation. In the following sections, we present the proposed methodology. How it applies to cardiac deformation recovery is detailed in Chapter 5.

4.1 Motivation

Tissue tracking in sequences of medical images is an important task in many clinical applications, either for disease diagnosis or therapy guidance. However there is no easy way to achieve it, even with user input. A standard approach is to use non-linear image registration to estimate dense spatial transformations between images. For instance, cardiac motion is commonly estimated by non-linearly co-registering the frames of a cardiac sequence, yielding a dense displacement field that quantifies myocardium motion [Rueckert et al., 2006b; Bistoquet et al., 2008].

In practice, non-linear image registration is performed by minimising a dissimilarity criterion between the images to register, up to a regularisation term that models prior knowledge about the spatial transformations. For clinical applications, it is often required to estimate smooth one-to-one mappings for consistent analyses. In addition, the transformations must be compatible with the properties of the tissue to track, such as elasticity and incompressibility. This is all the more important if the estimated deformations are used to analyse anatomical changes between different time points [Ashburner et al., 1998]. Nevertheless, adding these constraints to image registration algorithms often yields additional computational complexity.

State of the Art

Diffeomorphic Non-Linear Image Registration

With the recent advances in computational anatomy, mathematical frameworks based on diffeomorphic deformations have been developed to estimate one-to-one differentiable mappings between two images. Recent algorithms parameterise the transformations by velocity fields according to the Lagrange transport equation [Arnold, 1989]. When the velocity fields vary over time, large diffeomorphic deformations can be estimated using the Large Deformation Diffeomorphic Metric Mappings (LDDMM) [Miller et al., 2002b; Beg et al., 2005]. However, complex partial differential equations must be integrated over time, resulting in a large computational cost. To tackle this limitation, recent works relied on *stationary* velocity fields, which are efficiently integrated through exponential maps [Arsigny et al., 2006a; Bossa et al., 2007; Hernandez et al., 2009].

Among the methods based on stationary velocities, *logDemons* [Vercauteren et al., 2008] is an efficient non-linear registration algorithm based on the demons alternate optimisation [Thirion, 1998]. Two images are registered by alternating an optimisation step, which updates the stationary velocity field in a voxel-wise manner, and a Gaussian smoothing step, which models a diffusion motion. LogDemons algorithm is appealing as it ensures diffeomorphic mappings, it enables one to work on velocities and transformations simultaneously, and its complexity is linear in the number of voxels. However, the diffusion prior may not be appropriate for tracking biological tissues as it has no physical meaning. Mathematical justification of the demons optimisation step has been provided [Cachier et al., 2003; Vercauteren et al., 2009] but the theoretical foundations of the Gaussian regularisation still has to be

consolidated [Pennec et al., 1999; Modersitzki, 2004; Cahill et al., 2009].

Elastic Non-Linear Image Registration

First introduced by [Broit, 1981], elastic registration algorithms consist in regularising the transformations by using the linear elasticity equation, also known as the Lamé equation. Nonetheless, these regularisers are suitable for small displacements only and can yield discontinuities in the derivatives of the regularised deformations, and thus in the strains [Modersitzki, 2004]. Techniques based on smooth elastic-body splines have been developed [Rueckert et al., 1999; Sorzano et al., 2005]. Yet, they are computationally demanding and diffeomorphic mappings are ensured through *ad-hoc* penalisation of the registration energy [Rueckert et al., 2006a; DeCraene et al., 2009], which makes the computation of the inverse transformation critical. [Cachier and Ayache, 2004] demonstrated that the linear Lamé equation is actually a specific first order isotropic differential quadratic form (IDQF) of the transformation. High order IDQF can be designed, resulting in elastic-like regularisation of any order of smoothness. [Cachier and Ayache, 2004] also conjectured a separable elastic-like vector filter that behaves like an IDQF of infinite order. This filter has been used in demons algorithms in place of the Gaussian kernel [Cachier and Ayache, 2004; Mansi et al., 2009d], resulting in efficient and stable elastic-like registration. Yet, its link with IDQF energies is still not clear.

Incompressible Non-Linear Image Registration

Incompressible deformations cannot be recovered with elastic regularisers alone, explicit constraints must be added. A transformation is locally volume-preserving if its Jacobian determinant equals one. This constraint is non-linear and requires *ad-hoc* numerical schemes that are computationally demanding. Jacobian determinant can be directly constrained [Haber and Modersitzki, 2004] or introduced as a penalty term in the registration energy [Tanner et al., 2002; Rohlfing et al., 2003]. In [Bischoff et al., 2008], the authors proposed to use the linear approximation of that constraint, i.e. the divergence of the displacements is null. However, volume drifts appear when deformations become large, which the authors controlled by penalising the energy functional.

When estimating fluid motion, incompressibility is satisfied if the velocity of the fluid is divergence-free. Building up on this observation, countless optical flow techniques [Horn and Schunck, 1981] based on the continuity equation and the divergence-free constraint have been developed to estimate incompressible fluid motion from 3D images (see [Heitz et al., 2009] and reference therein). [Song and Leahy, 1991] and [Gorce et al., 1997] applied this approach to estimate 3D cardiac velocity from 4D CT images. [Cuzol et al., 2007] combined the optical flow algorithm with the Helmholtz decomposition to estimate 2D fluid motion parameterised by divergence-free and curl-free parameter maps. [Saddi et al., 2007] constrained a fluid registration algorithm to be incompressible by projecting the update velocities onto the space of divergence-free vector field using Helmholtz decomposition. All

these techniques showed satisfying results and demonstrated that incompressibility constraints can improve the estimation of incompressible fluid motion. However, the fluid model may not be suitable for tracking elastic biological tissues: the incompressibility condition on the transformation is usually preferred for biological applications.

Interestingly, one can demonstrate that diffeomorphic transformations parameterised by divergence-free velocity fields through the transport equation are incompressible [Evans, 1998]. [Hinkle et al., 2009] for instance used this property to guide image reconstruction of incompressible organ with LDDMM and divergence-free time varying velocity fields.

Model-Based Non-Linear Image Registration

For tracking biological tissues, some authors proposed to guide non-linear registration algorithms with biomechanical models. These approaches have been successfully applied in cardiac motion estimation [Papademetris et al., 2000; Sinusas et al., 2001; Veress et al., 2005; Phatak et al., 2009; Sundar et al., 2009a] and brain shift estimation [Ferrant et al., 2001; Clatz et al., 2005]. However, the underlying models often rely on physical parameters that are difficult to determine for a given patient. Besides, such models may not apply anymore in pathological cases. Finally, they all require meshing the space domain in order to solve complex partial differential equations. For all these reasons we prefer here a purely image-driven algorithm.

This chapter describes a consistent and efficient framework for elastic incompressible non-linear registration based on logDemons algorithm. The registration is entirely performed in the log-domain. The constraint is strongly ensured and is applied directly in the demons minimisation space. After a brief introduction to the log-Demons algorithm (Section 4.2), we propose a mathematical justification of demons Gaussian regularisation that enables to adapt the algorithm to other transformation models (Section 4.3). We then replace that regulariser with multi-order IDQF whose minimiser is exactly computed with [Cachier and Ayache, 2004] elastic-like vector filter. Finally, strong incompressibility is ensured by constraining the stationary velocity fields that parameterise the transformations to be divergence-free (Section 4.4). Our method, hereafter termed *iLogDemons*, is mathematically consistent, all its elements and parameters are controlled. The proposed demons framework results in the following advantages with respect to previous techniques: i) the elastic regulariser and the incompressibility constraint are linear, yielding low computational overhead, ii) they are rigorously integrated in the demons energy functional, yielding closed form minimisers that can be easily enabled/disabled by the user and applied to subdomains of the images, and iii) incompressibility constraint is strongly enforced: no volume drifts appear. Section 4.5 reports results on synthetic datasets with known ground truth.

4.2 Background: Log-Domain Diffeomorphic Demons

Proposed by [Vercauteren et al., 2008], log-domain diffeomorphic demons algorithm, hereafter termed *logDemons*, is an efficient non-linear registration algorithm based on demons approach [Thirion, 1998]. Given a reference image R and a template image T , logDemons estimates a dense spatial transformation ϕ that best aligns T to R . This is achieved by alternating an *optimisation step*, which updates the transformation in a voxel-wise manner, and a *regularisation step*, which traditionally consists in a Gaussian smoothing. In [Cachier et al., 2003], the authors justified the demons algorithm by the alternate minimisation of the energy functional:

$$\mathcal{E}(\phi, \phi_c) = \frac{1}{\sigma_i^2} \|R - T \circ \phi_c\|_{L_2}^2 + \frac{1}{\sigma_x^2} \text{dist}(\phi_c, \phi) + \frac{1}{\sigma_d^2} \|\nabla \phi\|_{L_2}^2 \quad (4.1)$$

In this equation, ϕ is the dense spatial transformation to estimate and ϕ_c is an intermediate transformation, called *correspondences*, that matches the two images without considering the regularity of the transformation. The first term of Equation 4.1 is the similarity criterion or data term. It measures how R and $T \circ \phi_c$ are similar. σ_i^2 relates to the noise in the images. The last term of Equation 4.1 is the regulariser whose strength is controlled by σ_d^2 . It ensures the spatial smoothness of the transformation ϕ , here by penalising large gradients, and models prior knowledge about the transformation to recover. The second term of Equation 4.1 couples the correspondences ϕ_c with the smooth transformation ϕ . It unifies in a common mathematical framework the optimisation step, which amounts to minimising $\mathcal{E}(\phi, \phi_c)$ with respect to ϕ_c , and the regularisation step, which consists in minimising $\mathcal{E}(\phi, \phi_c)$ with respect to ϕ .

In logDemons, the registration energy $\mathcal{E}(\phi, \phi_c)$ is adapted to estimate spatial transformations that are parameterised by *stationary* velocity fields \mathbf{v} through the exponential map $\phi = \exp(\mathbf{v})$. Such transformations belong to the subspace of diffeomorphisms \mathbb{G} generated by the one-parameter subgroups of diffeomorphisms. The tangent space of velocities V is the *log-domain*. As \mathbf{v} is stationary, the exponential map $\exp(\mathbf{v})$ is efficiently computed using a “scaling-and-squaring” algorithm [Argigny et al., 2006a] (see appendices). Alternatively, ϕ can be defined as the solution of the Lagrange transport equation: $\partial \phi(\mathbf{x}, t) / \partial t = \mathbf{v}(\phi(\mathbf{x}, t))$, $\phi(\mathbf{x}, 0) = \mathbf{x}$. With $\phi = \exp(\mathbf{v})$ and $\phi_c = \exp(\mathbf{v}_c)$, logDemons energy writes in the log-domain:

$$\begin{aligned} \mathcal{E}(\mathbf{v}, \mathbf{v}_c) = & \frac{1}{\sigma_i^2} \|R - T \circ \exp(\mathbf{v}_c)\|_{L_2}^2 + \\ & \frac{1}{\sigma_x^2} \|\log(\exp(-\mathbf{v}) \circ \exp(\mathbf{v}_c))\|_{L_2}^2 + \frac{1}{\sigma_d^2} \|\nabla \mathbf{v}\|_{L_2}^2 \end{aligned} \quad (4.2)$$

During the optimisation step, we minimise $\mathcal{E}(\mathbf{v}, \mathbf{v}_c)$ with respect to \mathbf{v}_c . The correspondence field ϕ_c is modelled with the diffeomorphic update rule $\phi_c = \exp(\mathbf{v}_c) = \phi \circ \exp(\delta \mathbf{v})$, where ϕ is the current estimate of the transformation to recover and $\delta \mathbf{v}$ is an unknown small update velocity field, the so-called demons force. Gauss-Newton

algorithm and an efficient second-order minimisation (ESM) scheme yield:

$$\delta \mathbf{v}(\mathbf{x}) = -\frac{R(\mathbf{x}) - T \circ \phi(\mathbf{x})}{\|J(\mathbf{x})\|^2 + \sigma_i^2/\sigma_x^2} J(\mathbf{x}) \quad (4.3)$$

where $J(\mathbf{x})$ is the symmetric gradient $J(\mathbf{x}) = (\nabla R(\mathbf{x}) + \nabla(T \circ \phi)(\mathbf{x}))/2$. In virtue of the diffeomorphic update rule $\phi_c = \phi \circ \exp(\delta \mathbf{v}) = \exp(\mathbf{v}) \circ \exp(\delta \mathbf{v})$, we can apply Baker-Campbell-Hausdorff (BCH) formula to estimate the correspondence velocity \mathbf{v}_c without computing the logarithm of the updated correspondences ϕ_c . Indeed, BCH formula gives an approximation of the velocity field $\mathbf{v}_c = Z(\mathbf{v}, \delta \mathbf{v})$ such that $\exp(\mathbf{v}_c) = \exp(\mathbf{v}) \circ \exp(\delta \mathbf{v})$. As shown in [Vercauteren et al., 2008], the first order approximation is sufficient for image registration purposes. We thus have:

$$\begin{aligned} \mathbf{v}_c &= Z(\mathbf{v}, \delta \mathbf{v}) \\ &= \mathbf{v} + \delta \mathbf{v} + 1/2[\mathbf{v}, \delta \mathbf{v}] + 1/12[\mathbf{v}, [\mathbf{v}, \delta \mathbf{v}]] + O(\|\delta \mathbf{v}\|^2) \end{aligned} \quad (4.4)$$

In the previous equation, the Lie bracket $[\cdot, \cdot]$ is defined by $[\mathbf{v}, \delta \mathbf{v}] = (\nabla \mathbf{v})\delta \mathbf{v} - (\nabla \delta \mathbf{v})\mathbf{v}$. Although it is not clear whether theoretically the space \mathbb{G} is a BCH-Lie group [Glockner, 2006], BCH composition of diffeomorphisms of \mathbb{G} has experimentally shown promising results in terms of image registration and statistics on diffeomorphisms [Bossa et al., 2007; Vercauteren et al., 2008].

Once \mathbf{v}_c is calculated, Equation 4.2 is minimised with respect to \mathbf{v} . This step is performed by smoothing \mathbf{v}_c with a Gaussian kernel G_σ . Next section investigates how this smoothing relates to $\mathcal{E}(\mathbf{v}, \mathbf{v}_c)$. The main steps of the logDemons algorithm are reported in the pseudo-code (Algorithm 3).

Algorithm 3 LogDemons Registration

Require: Stationary velocity field \mathbf{v}^0 . {Usually $\mathbf{v}^0 = \mathbf{0}$, i.e. $\phi^0 = \text{Id}$ }.

- 1: **loop** {over n until convergence}
 - 2: Compute the update velocity: $\delta \mathbf{v}^n$ (Equation 4.3).
 - 3: Fluid-like registration: $\delta \mathbf{v}^n \leftarrow G_{\sigma_f} \star \delta \mathbf{v}^n$, G_{σ_f} is a Gaussian kernel.
 - 4: Update the velocity field: $\mathbf{v}^n \leftarrow Z(\mathbf{v}^{n-1}, \delta \mathbf{v}^n)$ (Equation 4.4).
 - 5: Diffusion-like registration: $\mathbf{v}^n \leftarrow G_\sigma \star \mathbf{v}^n$, G_σ is a Gaussian kernel.
 - 6: Update the warped image $T \circ \phi^n = T \circ \exp(\mathbf{v}^n)$.
 - 7: **return** \mathbf{v} , $\phi = \exp(\mathbf{v})$ and $\phi^{-1} = \exp(-\mathbf{v})$.
-

About LogDemons Parameters LogDemons is controlled by four parameters: the image noise σ_i^2 , the uncertainty on the correspondences σ_x^2 and the regularisation strengths σ_f^2 and σ^2 . The noise in the images is estimated at every voxel by $\sigma_i^2(\mathbf{x}) = |R(\mathbf{x}) - T \circ \phi(\mathbf{x})|^2$ [Cachier et al., 1999; Vercauteren et al., 2009]. As in Thirion demons, such an estimator normalises $\delta \mathbf{v}$ to prevent too large updates that would hamper the stability of the algorithm. In particular, building on [Cachier et al., 1999], Vercauteren et al. demonstrated that the maximum amplitude of the update velocity $\delta \mathbf{v}$ is upper bounded by $\sigma_x/2$ [Vercauteren et al., 2009]. More global

noise estimators fail to limit the update velocities, which can become large and ultimately yield non diffeomorphic transformations. This behaviour was confirmed experimentally using Gaussian based noise estimated. σ_f^2 controls the strength of the fluid-like regularisation. In practice, $\sigma_f^2 = 0.5$ is recommended. Finally, σ^2 controls the strength of the regularisation, as we shall discuss in the next section. It has to be stressed that here, the three parameters σ_x^2 , σ_f^2 and σ^2 are explicitly decoupled as in [Cachier et al., 2003; Vercauteren et al., 2008] in contrast to other formulations of demons where σ_x^2 is implicit.

4.3 From Demons Gaussian Regularisation to Elastic-Like Regularisation

4.3.1 Insights into LogDemons Gaussian Regularisation

A consistent mathematical formulation of logDemons regularisation is required to adapt the algorithm to other transformation models. In scale-space theory, one can demonstrate that the Gaussian smoothing is the solution to the Tikhonov estimation problem with equal weighting of the spatial derivatives in the Taylor series sense [Nielsen et al., 1994]. We thus replace the first-order logDemons regulariser $\|\nabla \mathbf{v}\|^2 / \sigma_d^2$ by the infinite sum Tikhonov regulariser defined by:

$$\mathcal{R}_{dif}(\mathbf{v}) = \int_{\Omega} \sum_{k=1}^{+\infty} \frac{\partial_{i_1 \dots i_k} \mathbf{v}_{i_{k+1}} \partial_{i_1 \dots i_k} \mathbf{v}_{i_{k+1}}}{\sigma_x^2 \sigma_d^{2k} k!} \quad (4.5)$$

In this equation, Ω is the image domain and $\partial_{i_k \dots i_l}$ denotes the composition of spatial derivatives $\partial_{i_k} \dots \partial_{i_l}$. A simplified Einstein notation convention is used: Indices that are repeated twice in a product are summed all over their range (e.g. $\mathbf{v}_i \mathbf{v}_i = \mathbf{v}_1^2 + \mathbf{v}_2^2 + \mathbf{v}_3^2$ if $\mathbf{v} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$). \mathcal{R}_{dif} has been divided by σ_x^2 to simplify the integration of the regulariser with demons coupling term. The regularisation weight σ_d^2 is function of the derivative orders to preserve the shape of the impulse response related to the regulariser [Nielsen et al., 1994]. Minimising the registration energy Equation 4.2, modified using Equation 4.5, with respect to \mathbf{v} amounts to minimising the regularisation energy:

$$\mathcal{E}_{reg}(\mathbf{v}) = \frac{1}{\sigma_x^2} \|\log(\exp(-\mathbf{v}) \circ \exp(\mathbf{v}_c))\|_{L_2}^2 + \int_{\Omega} \sum_{k=1}^{+\infty} \frac{\partial_{i_1 \dots i_k} \mathbf{v}_{i_{k+1}} \partial_{i_1 \dots i_k} \mathbf{v}_{i_{k+1}}}{\sigma_x^2 \sigma_d^{2k} k!} \quad (4.6)$$

By linearising the first term using the zeroth order approximation of BCH formula (Equation 4.4), $\log(\exp(-\mathbf{v}) \circ \exp(\mathbf{v}_c)) \approx \mathbf{v}_c - \mathbf{v}$, Equation 4.6 becomes a Tikhonov estimation problem:

$$\mathcal{E}_{reg}(\mathbf{v}) = \frac{1}{\sigma_x^2} \|\mathbf{v}_c - \mathbf{v}\|_{L_2}^2 + \int_{\Omega} \sum_{k=1}^{+\infty} \frac{\partial_{i_1 \dots i_k} \mathbf{v}_{i_{k+1}} \partial_{i_1 \dots i_k} \mathbf{v}_{i_{k+1}}}{\sigma_x^2 \sigma_d^{2k} k!} \quad (4.7)$$

It follows from the optimal condition $\partial_{\mathbf{v}} \mathcal{E}_{reg}(\mathbf{v}) = 0$:

$$\mathbf{v} + \sum_{k=1}^{+\infty} \frac{(-1)^k \Delta^k \mathbf{v}}{\sigma_d^{2k} k!} = \mathbf{v}_c$$

where Δ is the Laplacian operator. We solve this equation in the Fourier domain. Let $\hat{\mathbf{v}}(\mathbf{w}) = \mathfrak{F}(\mathbf{v}(\mathbf{x}))$ be the Fourier transform of the velocity field $\mathbf{v}(\mathbf{x})$, \mathbf{w} is the frequency variable. We have:

$$\sum_{k=0}^{\infty} \frac{(\mathbf{w}^T \mathbf{w})^k}{\sigma_d^{2k} k!} \hat{\mathbf{v}}(\mathbf{w}) = \exp\left(\frac{\mathbf{w}^T \mathbf{w}}{\sigma_d^2}\right) \hat{\mathbf{v}}(\mathbf{w}) = \hat{\mathbf{v}}_c(\mathbf{w}) \quad (4.8)$$

where the exponential appears from its Taylor series. Let $\sigma^2 = 2/\sigma_d^2$. The velocity \mathbf{v} is obtained by smoothing \mathbf{v}_c with the Gaussian kernel

$$G_{\sigma} = 1/\sqrt{2\pi\sigma^2}^d \exp(-\mathbf{x}^T \mathbf{x}/(2\sigma^2))$$

d being the image dimension. We retrieve the demons regularisation. We verify that the higher σ_d^2 , the lower σ^2 and the less regularised the transformation: The width σ^2 of the Gaussian kernel corresponds to the strength of the regularisation σ_d^2 . It has to be noted that the derivation of Equation 4.7 with respect to \mathbf{v} implies that all the spatial derivatives of \mathbf{v} vanish at the boundaries $\partial\Omega$ of Ω . Gaussian smoothing must be performed accordingly by extending the image periodically for instance or by ensuring that the moving structure stays far from the image boundaries.

Scale-space theory provides another interesting insight into the algorithm. As shown in [Nielsen et al., 1994], the Gaussian smoothing, which is now the exact minimiser of the regularisation energy $\mathcal{E}_{reg}(\mathbf{v})$, is also the solution of the heat equation. This further justifies why logDemons, and demons algorithms in general, are considered as diffusion registration methods [Pennec et al., 1999; Modersitzki, 2004; Cahill et al., 2009].

4.3.2 Elastic-Like LogDemons

The previous developments enable us to directly integrate an elastic regularisation in the logDemons framework. To preserve demons computational efficiency, we propose a regulariser based on multi-order isotropic differential quadratic forms (IDQF) whose minimiser is exactly computed using the separable elastic-like kernel filter proposed by [Cachier and Ayache, 2004]. With the simplified Einstein convention, the k^{th} -order IDQF of a vector field \mathbf{v} is defined by:

$$Q_{el}^k(\mathbf{v}) = \alpha_k \partial_{i_1 \dots i_k} \mathbf{v}_{i_{k+1}} \partial_{i_1 \dots i_k} \mathbf{v}_{i_{k+1}} + \beta_k \partial_{i_1 \dots i_k} \mathbf{v}_{i_{k+1}} \partial_{i_{k+1} i_2 \dots i_k} \mathbf{v}_{i_1} \quad (4.9)$$

α_k and β_k are scalar coefficients of \mathbb{R} , $\alpha_k \geq 0$ and $\beta_k \geq -\alpha_k$ to ensure the positiveness of Q_{el}^k . With $\alpha_k = 1/k!$ we recognise in the first term of Q_{el}^k the k^{th} term of the previous Tikhonov regulariser (Equation 4.5). Elasticity is thus obtained through the second term of Q_{el}^k . We define the elastic regularisation as:

$$\mathcal{E}_{reg}(\mathbf{v}) = \frac{1}{\sigma_x^2} \|\mathbf{v}_c - \mathbf{v}\|_{L_2}^2 + \int_{\Omega} \sum_{k=1}^{+\infty} \frac{Q_{el}^k(\mathbf{v})}{\sigma_d^{2k} \sigma_x^2} \quad (4.10)$$

From the functional derivatives:

$$\begin{aligned}\partial_{\mathbf{v}}(\partial_{i_1 \dots i_k} \mathbf{v}_{i_{k+1}} \partial_{i_1 \dots i_k} \mathbf{v}_{i_{k+1}}) &= (-1)^k \Delta^k \mathbf{v} \\ \partial_{\mathbf{v}}(\partial_{i_1 \dots i_k} \mathbf{v}_{i_{k+1}} \partial_{i_{k+1} i_2 \dots i_k} \mathbf{v}_{i_1}) &= (-1)^k \Delta^{k-1} \nabla \nabla^T \mathbf{v}\end{aligned}$$

It follows the optimal condition:

$$\mathbf{v} + \sum_{k=1}^{\infty} \frac{(-1)^k}{\sigma_t^{2k}} \left[\alpha_k \Delta^k \mathbf{v} + \beta_k \Delta^{k-1} \nabla \nabla^T \mathbf{v} \right] = \mathbf{v}_c \quad (4.11)$$

which is solved in the Fourier domain. Note that when $k = 1$, the regulariser becomes the first-order Lamé elastic equation. According to the identities

$$\begin{aligned}\mathfrak{F}(\Delta^k \mathbf{v}(\mathbf{x})) &= (-1)^k (\mathbf{w}^T \mathbf{w})^k \hat{\mathbf{v}}(\mathbf{w}) \\ \mathfrak{F}(\Delta^{k-1} \nabla \nabla^T \mathbf{v}(\mathbf{x})) &= (-1)^k (\mathbf{w}^T \mathbf{w})^{k-1} \mathbf{w} \mathbf{w}^T \hat{\mathbf{v}}(\mathbf{w})\end{aligned}$$

Equation 4.11 is transformed as:

$$\left[\underbrace{\left(1 + \sum_{k=1}^{\infty} \frac{\alpha_k (\mathbf{w}^T \mathbf{w})^k}{\sigma_d^{2k}} \right)}_A \text{Id} + \underbrace{\left(\sum_{k=1}^{\infty} \frac{\beta_k (\mathbf{w}^T \mathbf{w})^{k-1}}{\sigma_d^{2k}} \right)}_B \mathbf{w} \mathbf{w}^T \right] \hat{\mathbf{v}}(\mathbf{w}) = \hat{\mathbf{v}}_c(\mathbf{w})$$

Since A and B are scalars, we can apply Sherman-Morrison inversion formula, which yields the closed form solution:

$$\hat{\mathbf{v}}(\mathbf{w}) = \underbrace{\left[\frac{1}{A} \text{Id} - \frac{1}{A} \left(\frac{B}{A + B \mathbf{w}^T \mathbf{w}} \right) \mathbf{w} \mathbf{w}^T \right]}_M \hat{\mathbf{v}}_c(\mathbf{w})$$

The optimal velocity field \mathbf{v} is therefore obtained by filtering the correspondence velocity \mathbf{v}_c with the filter M in the Fourier domain. Computational efficiency is greatly improved by choosing the coefficients α_k and β_k such that M is separable. If $\alpha_k = 1/k!$, A is the Gaussian kernel found in the previous section. One can demonstrate that if β_k is defined by (see appendices, page 84):

$$\begin{cases} \beta_0 = 0, \\ \beta_k = \sum_{i=1}^k \gamma^i \sigma_d^{2i} / (k-i)!, \quad \forall k \geq 1 \end{cases}$$

and $\gamma \in \mathbb{R}$, then the second term of the filter M is proportional to the Hessian of the Gaussian kernel $\exp(\mathbf{w}^T \mathbf{w} / \sigma_d^2)$. With $\sigma^2 = 2/\sigma_d^2$ and $\gamma = \sigma^2 \kappa / (\kappa + 1)$ we retrieve the elastic-like separable vector filter proposed by [Cachier and Ayache, 2004]:

$$\mathbf{v} = \left(G_{\sigma} \text{Id} + \frac{\sigma^2 \kappa}{1 + \kappa} \mathcal{H} G_{\sigma} \right) \star \mathbf{v}_c = G_{\sigma, \kappa} \star \mathbf{v}_c \quad (4.12)$$

$\mathcal{H}G_\sigma(\mathbf{x})$ is the Hessian of the Gaussian kernel G_σ and $G_{\sigma,\kappa}$ is the elastic-like vector filter. When $\kappa = 0$, $G_{\sigma,\kappa=0}$ is the Gaussian filter and the elastic regularisation energy (Equation 4.10) is exactly the diffusion energy (Equation 4.7). It is therefore straightforward to switch between regulariser.

As for diffusion regularisation, σ^2 controls the strength of the regularisation. The *elastic* parameter κ behaves like the Poisson ratio ν of the theory of elasticity by controlling the cross-effects of the smoothing between the vector components. In [Cachier and Ayache, 2004], the authors showed that the higher κ , the more incompressible the deformation. This property still holds here even though $G_{\sigma,\kappa}$ acts on velocities and not on deformations. Indeed, the stationary velocities \mathbf{v} are parameters of the deformations, their norm is directly related to the length of the deformations ($\int_0^1 \|\mathbf{v}(t)\|_V^2 dt^{1/2} = \|\mathbf{v}(t=0)\|_V^2$, V is the space of velocities). Smoothing $\mathbf{v}(t=0)$ thus amounts to smoothing ϕ . This is very different from fluid registration which regularises the infinitesimal increments i.e., the instant velocities $\mathbf{v}(t)$. However, elastic-like regularisation may not be sufficient to recover locally incompressible deformations. κ only controls incompressibility at a global scale as it is a global parameter. Furthermore, perfect incompressibility would be reached only when $\kappa \rightarrow \infty$. Hard constraints must be used when strong incompressibility is required.

4.4 Incompressible LogDemons

A transformation ϕ is locally incompressible if its Jacobian determinant $|\nabla\phi|$ equals one. This non-linear constraint however is computationally demanding. For diffeomorphic transformations one can show that the condition on fluid motion holds. Integrating divergence-free velocities over time yields incompressible deformations (see appendices, page 84). Making logDemons incompressible is thus achieved by constraining the velocity field \mathbf{v} to be divergence-free. This only alters the regularisation step as the optimisation stage minimises $\mathcal{E}(\mathbf{v}, \mathbf{v}_c)$ with respect to the correspondence velocity \mathbf{v}_c , which is not constrained.

Helmholtz decomposition states that any velocity \mathbf{v} that vanishes at infinity can be uniquely decomposed into the sum of a divergence-free field and a curl-free field. Using variational calculus and Lagrangian multipliers, Simard and Mailloux demonstrated that the Helmholtz decomposition projects \mathbf{v} onto the space of divergence-free vector field in the L_2 -norm sense [Simard and Mailloux, 1988]. We employ here a similar technique to integrate the divergence free constraint in the registration energy.

We want to minimise Equation 4.10 under the divergence-free constraint $\nabla \cdot \mathbf{v} = 0$. Let p be the Lagrangian multiplier associated to the constraint. p is a scalar field of the Sobolev space $H^1(\Omega) = \{p \in L^2(\Omega) \mid \forall i \partial p / \partial x_i \in L^2(\Omega)\}$ that vanishes at infinity ($p \in H_0^1$). Intuitively, p is the pressure field that generates the sinks and the sources of the velocity field. Minimisers of the energy Equation 4.10 (or Equation 4.7

if $\kappa = 0$) under the divergence-free constraint are optima of the Lagrange function:

$$\mathcal{P}_{reg}(\mathbf{v}, p) = \frac{1}{\sigma_x^2} \|\mathbf{v}_c - \mathbf{v}\|_{L_2}^2 + \int_{\Omega} \sum_{k=1}^{+\infty} \frac{Q_{el}^k}{\sigma_x^2 \sigma_d^{2k}} - \frac{2}{\sigma_x^2} \int_{\Omega} p \nabla \cdot \mathbf{v} \quad (4.13)$$

The Gâteaux derivatives of the previous equation yields two optimal conditions:

$$\nabla \cdot \mathbf{v} = 0 \quad (4.14)$$

$$\mathbf{v} + \sum_{k=1}^{\infty} \frac{(-1)^k}{\sigma_d^{2k}} (\alpha_k \Delta^k \mathbf{v} + \beta_k \Delta^{k-1} \nabla \nabla^T \mathbf{v}) = \mathbf{v}_c - \nabla p \quad (4.15)$$

with $p = 0$ at the boundaries $\partial\Omega$ of the image domain. The optimal velocity field \mathbf{v} is therefore computed by smoothing the right hand side of Equation 4.15, $\mathbf{g} = \mathbf{v}_c - \nabla p$, with the kernel $G_{\sigma, \kappa}$. To compute \mathbf{g} , we take the divergence of Equation 4.15. Knowing that under the optimal condition we $\nabla \cdot \mathbf{v} = 0$, we get the Poisson equation under 0-Dirichlet boundary conditions:

$$\Delta p = \nabla \cdot \mathbf{v}_c \in H^{-1}(\Omega) \quad (4.16)$$

where $H^{-1}(\Omega)$ is the dual of $H_0^1(\Omega)$ ¹. p can thus be computed independently of \mathbf{v} by solving Equation 4.16. This is exactly the Helmholtz decomposition of \mathbf{v}_c . $\mathbf{g} = \mathbf{v}_c - \nabla p$ is the L_2 projection of \mathbf{v}_c to the space of divergence-free vector fields, as $\nabla \cdot \mathbf{g} = \nabla \cdot \mathbf{v}_c - \Delta p = 0$. ∇p is the orthogonal curl-free component. Ensuring divergence-free velocity fields thus consists in projecting the correspondence velocity onto the space of divergence-free vector fields and smoothing the result.

With this approach, the incompressibility constraint can be applied within a sub-domain $\Gamma \subset \Omega$ only by defining $p \in H_0^1(\Gamma)$, $p = 0$ on Ω/Γ . This may be useful for tracking incompressible tissues localised in space, like the cardiac muscle. However, particular care must be taken at the domain boundaries $\partial\Gamma$. Although Gaussian smoothing theoretically preserves vector field divergence, in practice unconstrained velocities close to $\partial\Gamma$ may leak inside the incompressible domain due to the Gaussian convolution, ultimately resulting in volume drifts. Yet, Gaussian filter and vector derivatives commute for well-designed filters such as Deriche recursive filters [Deriche, 1993]. We therefore replace the theoretical “project-and-smooth” strategy by a “smooth-and-project” approach that preserves the divergence close to $\partial\Gamma$. To further limit numerical instabilities, a smooth domain transition is implemented in a narrow band around Γ by diffusing the pressure field p using heat-transfer equation [Evans, 1998]. The main steps of the proposed algorithm, henceforth termed *iLogDemons*, are summarised in (Algorithm 4). Note that the incompressibility constraint can be easily disabled by skipping the 6th and 7th steps of the algorithm.

¹In the continuous domain, one would need p to be twice differentiable in order to compute the Laplacian. H_0^1 is sufficient in practice as Equation 4.16 is solved on a discrete grid, the Laplacian being computed on discrete points which can be seen as delta Dirac. The Sobolev space of distributions $H^{-1}(\Omega)$ is thus enough to calculate Δp .

Algorithm 4 iLogDemos: Incompressible Elastic LogDemos Registration

Require: Stationary velocity field \mathbf{v}^0 . {Usually $\mathbf{v}^0 = \mathbf{0}$ i.e. $\phi^0 = \text{Id}$ }.

- 1: **loop** {over n until convergence}
 - 2: Compute the update velocity: $\delta\mathbf{v}^n$ (Equation 4.3).
 - 3: Fluid-like regularisation: $\delta\mathbf{v}^n \leftarrow G_{\sigma_f} \star \delta\mathbf{v}^n$, G_{σ_f} is a Gaussian kernel.
 - 4: Update the correspondence velocity: $\mathbf{v}^n \leftarrow Z(\mathbf{v}^{n-1}, \delta\mathbf{v}^n)$ (Equation 4.4).
 - 5: Elastic-like regularisation: $\mathbf{v}^n \leftarrow G_{\sigma, \kappa} \star \mathbf{v}^n$ (Equation 4.12).
 - 6: Solve: $\Delta p = \nabla \cdot \mathbf{v}^n$ with 0-Dirichlet boundary conditions (Equation 4.19).
 - 7: Project the velocity field: $\mathbf{v}^n \leftarrow \mathbf{v}^n - \nabla p$.
 - 8: Update the warped image $T \circ \phi^n = T \circ \exp(\mathbf{v}^n)$.
 - 9: **return** \mathbf{v} , $\phi = \exp(\mathbf{v})$ and $\phi^{-1} = \exp(-\mathbf{v})$.
-

About Divergence-Free Update Velocity One could also constrain the correspondence field ϕ_c to be incompressible in order to find the optimal image matching that satisfies the constraint (see appendices, page 86). When the transformation ϕ is incompressible, it follows from the diffeomorphic update rule $\phi_c \leftarrow \phi \circ \exp(\delta\mathbf{v})$ that ϕ_c is incompressible if $\delta\mathbf{v}$ is divergence-free. Yet, from a theoretical perspective, adding such a constraint to the iLogDemos would have little effect on the result. In theory, constraining \mathbf{v} to be divergence-free amounts to projecting \mathbf{v}_c to the space of divergence-free vector fields. ϕ_c is therefore incompressible. Since the composition of two continuous incompressible fields is incompressible, $\exp(\delta\mathbf{v})$ is also incompressible and $\delta\mathbf{v}$ is divergence-free. When the zeroth order BCH approximation is used to compute \mathbf{v}_c , the linearity of the projector yields the same conclusion. The two approaches are hence equivalent. Yet, small differences may arise in practice due to the numerical approximations (scaling-and-squaring integration, numerical accuracy of the composition, etc.). We will experimentally evaluate when it is necessary to use this additional constraint.

Numerical Implementation The algorithm has been implemented using ITK and the open source implementation of the log-domain demons [Dru and Vercauteren, 2009]. The Poisson Equation 4.16 is discretised on the image grid using finite difference schemes [Simard and Mailloux, 1988] as the incompressible domain Γ may be of irregular shape. Its resolution is performed using PETSc library (see appendices, page 87).

Algorithm Complexity Despite the additional constraints, the complexity of the algorithm remains reasonable. Demons update velocity is computed at each voxel. The elastic-like filter is computed using Gaussian convolutions, no significant overhead is added to the original Gaussian filtering. The complexity of the divergence-free projector directly depends on the number of voxels of the incompressible domain Γ . It requires computing i) the divergence of the velocity field, ii) the gradient of the pressure field p , and iii) solving a linear system with $n \times n$ elements, where n is the number of voxels of the incompressible domain. The di-

vergence and gradient operators are linear in the number of voxels. The Poisson Equation 4.16 is solved at each iteration using iterative solvers like GMRES [Saad, 2003]. To speed up the process, we build the linear matrix and precondition it only once, at the very beginning of the algorithm, as the matrix is constant. However the system resolution can be time consuming if the domain is large. Should the incompressibility be ensured over the entire image domain Ω , more efficient Fourier techniques can be preferred [Hinkle et al., 2009].

4.5 Experiments and Results

Three experiments were performed on synthetically generated datasets to evaluate how much iLogDemons improves the recovery of incompressible deformations with respect to the original logDemons.

1. As a preliminary experiment, we verified that deformations parameterised by divergence-free stationary velocity fields are actually incompressible.
2. Then, we tested the ability of the elastic regularisation alone to estimate random incompressible deformations.
3. Finally, we evaluated the incompressibility constraint on large analytic incompressible transformations.

In the following, *logDemons* refers to the unconstrained logDemons algorithm, either with diffusion or elastic-like regularisation (Algorithm 3). *iLogDemons* refers to the proposed incompressible logDemons algorithm, where the velocities \mathbf{v} are constrained to be divergence-free (Algorithm 4). We also evaluated the fully constrained iLogDemons, where both the update velocities $\delta\mathbf{v}$ and the velocity \mathbf{v} are divergence-free. This algorithm is called *i²LogDemons*.

4.5.1 Incompressibility of Deformations Parameterised by Divergence-Free Velocities

We first verified that “scaling-and-squaring” divergence-free velocity fields (Algorithm 5, page 83) yields incompressible deformations. One hundred 3D velocity fields were generated on a $1mm$ -isotropic grid. Each voxel was assigned a random velocity according to a Gaussian distribution with high standard deviation (SD) ($SD = 5000 mm/s$) to get large deformations (large displacements and large strains). The resulting fields were smoothed with a Gaussian kernel ($SD = 3 mm$) and normalised to get a L_2 -norm of $2 mm$. On the one hand, we integrated these velocities to get deformations of reference (mean \pm SD of deformation amplitude: $1.91 \pm 0.99 mm$). On the other hand, we projected the velocities to the space of divergence-free vector fields by using Helmholtz decomposition and integrated the resulting vector fields, which theoretically yields incompressible deformations (average amplitude: $1.65 \pm 0.79 mm$). This experiment showed that despite the numerical approximations, the deformations parameterised by divergence-free velocity fields

are nearly incompressible. Their Jacobian determinant was close to the incompressibility condition ($|\nabla\phi| = 0.99 \pm 0.03$) in contrast to the Jacobian determinant of the reference fields ($|\nabla\phi| = 1 \pm 0.53$) (Figure 4.1).




Figure 4.1: Scaling-and-squaring integration of stationary velocity fields. a) 2D slice of a random 3D velocity field \mathbf{v} . b) Jacobian determinant of the deformation parameterised by \mathbf{v} c) Jacobian determinant of the deformation parameterised by the projection of \mathbf{v} onto the space of divergence-free velocities. One can clearly see that integrating divergence-free velocities yields almost incompressible deformations.

4.5.2 Global Incompressibility Recovery Using Elastic Regularisation

Elastic regulariser theoretically provides more incompressible deformation fields, controlled by the global parameter κ . Here, we experimentally test how much this feature alone (without incompressibility constraint) can help in recovering random volume-preserving deformation fields.

4.5.2.1 Illustration on Translated Cubes

We first tested the elastic-like regularisation on a toy example to have an intuition of the results. Two translated black-and-white small cubes were co-registered using diffusion ($\sigma^2 = 1, \kappa = 0$) and elastic-like ($\sigma^2 = 1, \kappa = 0.5$) regularisation. As illustrated in Figure 4.2, the elastic-like regularisation yielded deformations globally more incompressible as it distributed the smoothing across deformation components. The compressions around the cube were decreased, propagated over the image domain.

4.5.2.2 Quantitative Evaluation on Random Incompressible Deformations

The impact of the elastic-like regularisation on the estimation of incompressible deformations was quantified on synthetic data sets generated as follows. A 3D isotropic Steady-State Free Precession (SSFP) MR image of the heart ($53 \times 60 \times 60$

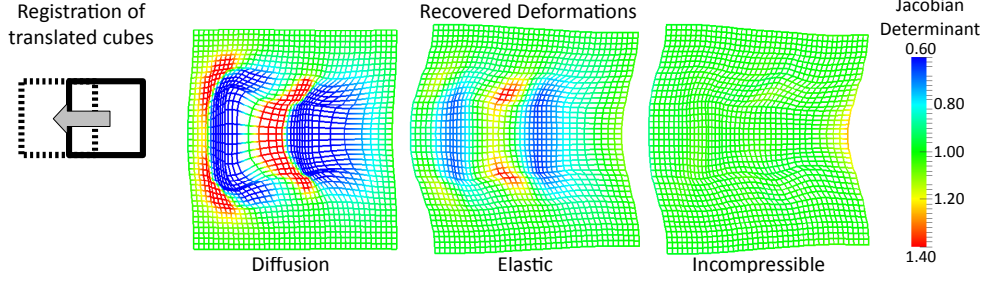


Figure 4.2: Registration of two translated cubes using diffusion logDemons, elastic-like logDemons and iLogDemons. At similar grey level RMSE, elastic-like regularisation yielded a stiffer deformations but only iLogDemons provided an incompressible deformation close to the true translation.

slices, 1 mm^3 voxel spacing), henceforth called test image, was warped by 50 random incompressible deformation fields. Warped and test images were then altered with slight Gaussian noise ($\text{SD}=3$, range of grey level intensities: $[0, 198]$) (Figure 4.3A-B).

The random incompressible deformation fields were generated by integrating divergence-free velocity fields. Each voxel was assigned a random velocity according to a Gaussian distribution with high standard deviation to get large displacements and large strains ($\text{SD}=5000 \text{ mm/s}$). The resulting velocity field was smoothed with a Gaussian kernel ($\text{SD}=3 \text{ mm}$) and its L_2 -norm was normalised to 2 mm . We then made it divergence-free using Helmholtz decomposition and integrated the result with the “scaling-and-squaring” algorithm to get the final incompressible deformation (Figure 4.3C). The average L_2 -norm of the deformation fields was $1.83 \pm 0.77 \text{ mm}$ (mean \pm standard deviation SD) and, as expected, their Jacobian determinant was 0.99 ± 0.03 , close to the incompressibility condition (Figure 4.3D).

We registered the 50 warped images to the test image with and without elastic regularisation. The following registration parameters were used, $\sigma_f^2 = 1$, $\sigma^2 = 1$ and $\sigma_x = 1$. No multi-resolution scheme was used as we aimed at comparing two methods rather than pure performance. The number of demons iterations was fixed to 50. Several elastic parameter values were tested $\kappa = \{0, 0.1, 0.5, 1, 2, 10, 100\}$. Registration accuracy was measured using the distance to the true deformation field (DTF) and the relative mean squared error of image intensities (RMSE) defined by:

$$\begin{aligned} \text{DTF}(\phi, \phi_{ref}) &= \|\phi - \phi_{ref}\|_{L_2} \\ \text{RMSE}(T, R \circ \phi) &= \|T - R \circ \phi\|^2 / \|T - R\|^2 \end{aligned}$$

For both indices, the lower the value, the better. Variations in registration performances were quantified using the coefficients of variation $\nu = \text{sd}/\text{mean}$ of RMSE, DTF and Jacobian determinant. Low ν values mean little impact of the elastic regularisation on a particular metric.

The results showed that deformation field recovery and image matching accuracy did not change significantly by increasing κ ($\nu_{\text{DTF}} \approx 1.6\%$, $\nu_{\text{RMSE}} \approx 8.5\%$,

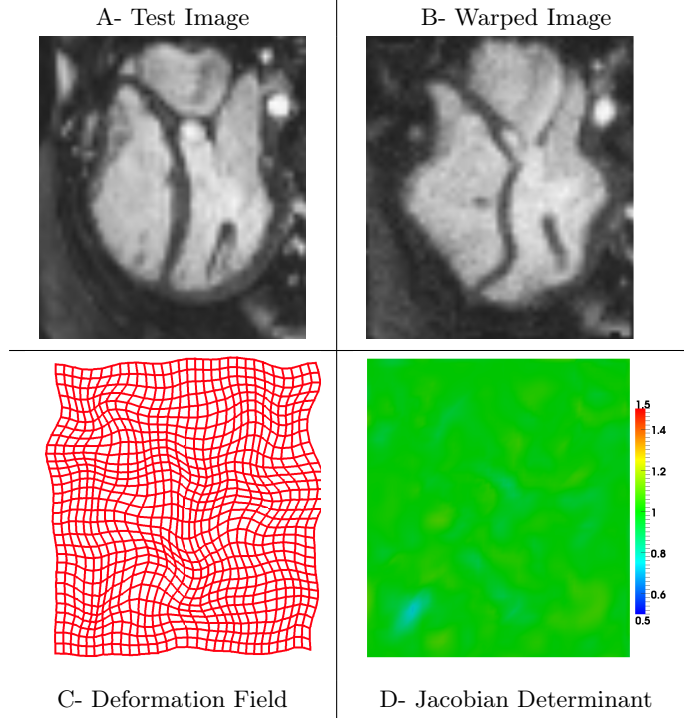


Figure 4.3: Synthetic 3D image warped with a random incompressible deformation field (here represented by a warped grid).




Figure 4.4: Effect of elastic-like regularisation on registration performances. These curves show that elastic regularisation, controlled by the parameter κ , does not affect registration accuracy (low variation of distance to true field, *left panel*) while it significantly decreases Jacobian determinant standard deviation (*right panel*). The deformation is globally more incompressible.

Figure 4.4). However, increasing κ largely reduced the standard deviation of the Jacobian determinant ($\nu_{\text{std}(\text{Jac.})} \approx 36\%$) while its mean was close to one ($\nu_{\text{mean}(\text{Jac.})} \approx 0.18\%$). The elastic regularisation thus improved global incompressibility of the deformation. However, it did not change the local accuracy of the registration. Strong

constraints are needed to better recover locally incompressible deformations.

4.5.3 Local Incompressibility Recovery Using Volume-Preserving Constraint

We now evaluate how much the incompressibility constraint, without elastic regularisation, can recover locally but strong volume-preserving deformation fields.

Illustration on Translated Cubes

As for the elastic regularisation, we first tested the incompressibility constraint on the small cubes to get an intuition of the results. Elastic regularisation was disabled ($\sigma^2 = 1$, $\kappa = 0$) and the incompressibility constraint was turned on. As one can see in Figure 4.2, the incompressibility constraint also prevented the compressions around the cubes. Qualitatively, the true global translation was better recovered.

Quantitative Evaluation on Analytic Incompressible Whirls

We quantified the previous qualitative observation on synthetic data generated by warping the test image with analytic whirl transformations (Figure 4.5). We decided not to use the previous synthetic dataset to avoid any bias as the deformations were generated using divergence free velocities, like the proposed constraint. Furthermore, analytic whirls enable to work on much larger but still volume-preserving deformations.

Eight volume preserving whirl transformations were created as in [Saddi et al., 2007]. The voxel O at the centre of the image domain was the centre of the whirl. All the voxels P that were outside the sphere of radius R and centred in O did not move. The voxels P inside the sphere were rotated with respect to O with an angle $\alpha(P) = \alpha_0(1 - \text{dist}(P, O)/R)^2$. The strength of the deformation was controlled by the whirl angle α_0 , spanning from 10° to 80° . Within the whirl domain, the L_2 -norm varied from 0.52 mm to 4.78 mm but the Jacobian determinant remained close to one (worst value: $|\phi_{\alpha_0=80^\circ}| = 1 \pm 0.04$).

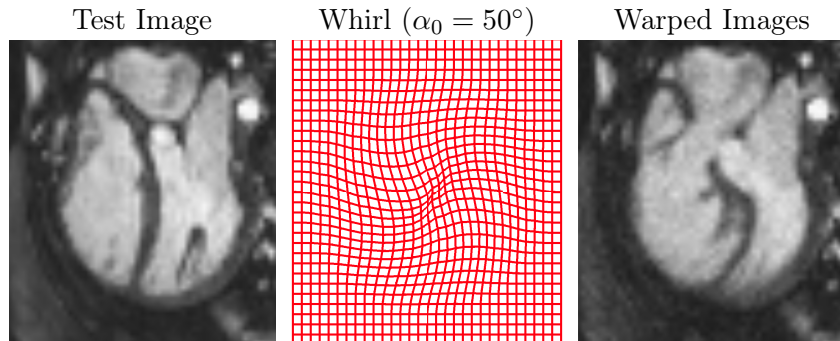


Figure 4.5: Synthetic 3D image warped with an analytic whirl transformation (here represented by a warped grid).

The 8 images T warped using the whirl transformations ϕ_{α_0} were registered to the test image R using LogDemons and iLogDemons ($\sigma_x = 1$, $\sigma^2 = 1$, $\sigma_f^2 = 1$, $\kappa = 0$, number of iterations fixed to 150 to ensure convergence at any whirl angle). RMSE, Jacobian determinant and DTF are reported in Figure 4.6. As expected, the deformation fields estimated with iLogDemons were almost incompressible. Jacobian determinants were always equal to 1 ± 0.02 independently of the strength of the whirl to recover. Image matching accuracy was not affected by the incompressibility constraint, showing only 0.6% decrease. The higher RMSE at small whirl angles is due to the relative nature of that metric. In those cases, the images are already fairly close to each other. Slight image matching errors yield larger RMSE. Most importantly, iLogDemons significantly improved the accuracy of the recovered deformation fields. Means and standard deviations of DTF were systematically lower (average improvements of 29% and 36% respectively). The larger the deformation, the more significant the improvement while RMSE stayed comparable. This experiment demonstrated the importance of the transformation model. As illustrated in Figure 4.7, regions with homogeneous grey levels provided few information to accurately estimate the whirl. With the iLogDemons, the incompressibility constraint helped the algorithm by ensuring that the estimated deformation is of the same type as the true field. This feature is particularly interesting for clinical applications, where deformations must be reliably estimated from ill-textured images.

We also investigated whether the performances were improved using i²LogDemons, which enforces the update velocity field to be divergence-free. Results, reported in Figure 4.6, were in agreement with the theoretical considerations (Section 4.4). Relevant differences only appeared at large deformations ($\alpha_0 \geq 70^\circ$). Finally, it should be noted that all these observations continue to hold with $\sigma_x = 2, 4$ and on random incompressible fields (experiments not reported here), supporting robustness to parameters.

Robustness of iLogDemons with Respect to the Slice Thickness

With a view to the estimation of myocardium strain from clinical cardiac images with poor through-plane resolution, we tested the robustness of the iLogDemons with respect to image slice thickness. To this end, reference and warped images of the whirl data sets were degraded by artificially increasing the slice thickness along the z-axis (Figure 4.8). Every N consecutive slices were grouped together and averaged to simulate partial volume effect. In-plane resolution was preserved. The resulting images were resampled to get $1mm$ -isotropic voxels. Four datasets were generated with slice thicknesses spanning from $1mm$ to $10mm$. The registration parameters were $\sigma_f^2 = 1$, $\sigma^2 = 1$ and $\sigma_x = 1$. Results are reported in Figure 4.9.

Not surprisingly, increasing the slice thickness decreased the overall registration accuracy as less image information was available: RMSE and DTF steadily increased. LogDemons yielded better image matching (lower RMSE) but the recovered deformation fields were less accurate than those estimated using iLogDemons

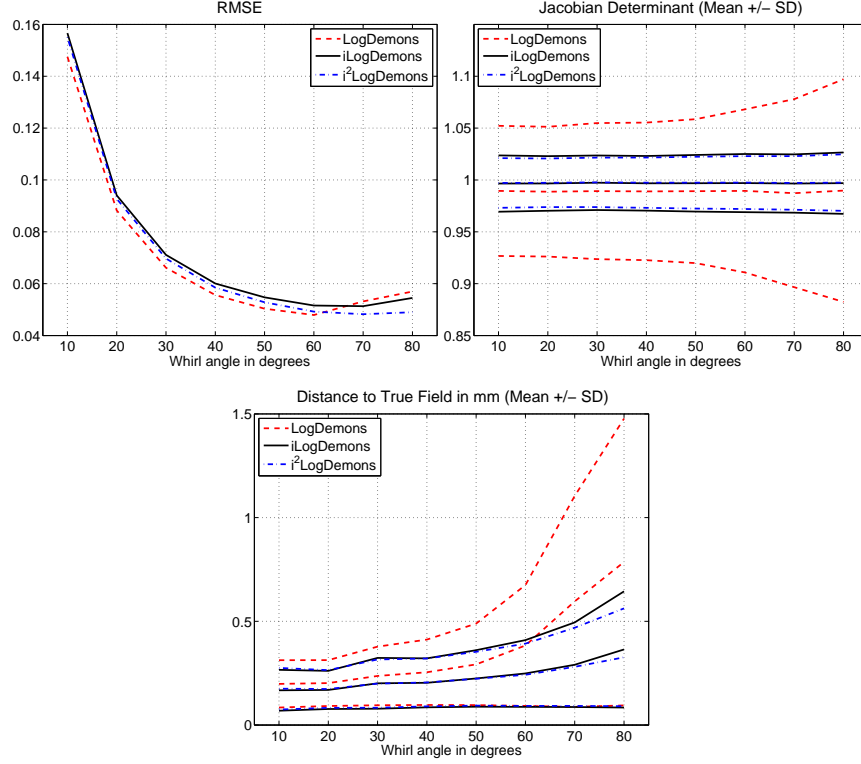


Figure 4.6: Results of the registration of whirl datasets. With similar image matching performances (similar RMSE of image intensities, *top-left panel*), iLogDemos provided more incompressible deformations (Jacobian determinant closer to one, *top-right panel*) and outperformed LogDemos in terms of deformation field accuracy (lower DTF, *bottom panel*). Constraining the update velocity to be divergence-free (i^2 LogDemos) did not improve the results significantly.

(higher DTF). Incompressibility constraints helped the algorithm to recover the incompressible whirl. This experiment further confirmed that i^2 Logdemons did not improve registration accuracy with respect to iLogDemos, even with large slice thickness. Figure 4.10 illustrates these findings on a particular case. Far from image gradients, the deformation estimated by logDemos were weak and erroneous (arrow A). Even worse, near strong image gradients, the deformation field can be wrong (arrow B). Thanks to the incompressibility constraint, iLogDemos alleviated these pitfalls and recovered a plausible through-plane motion (arrow C). These results motivate the use of iLogDemos to estimate the motion of incompressible organs in medical images.

To conclude, the experiments on synthetic data showed that i) elastic regularisation does provide globally more incompressible deformation fields but does not significantly improve the local accuracy of the registration; ii) the proposed strong incompressibility constraint provides almost incompressible deformations and do improve the recovery of volume preserving deformations. Both contributions are




Figure 4.7: Streamlines of true and estimated whirl deformations (whirl angle $\alpha_0 = 60^\circ$). Colours encode deformation amplitude in mm. iLogDemons better estimated the whirl transformation in regions with poor texture (*yellow arrow*), providing more accurate motion and deformation amplitude.

therefore complementary and can be used jointly to obtain smooth, globally and locally incompressible deformation fields.

4.6 Discussion

This chapter presented a method for elastic incompressible diffeomorphic registration based on the logDemons algorithm proposed by [Vercauteren et al., 2008]. We first established that logDemons Gaussian regularisation minimises an infinite order Tikhonov regulariser. Our formulation constitutes a well-posed formulation of demons algorithm with controlled parameters. An important theoretical condition on the coupling term $\|\log(\phi^{-1} \circ \phi_c)\|^2$ appeared. One must be able to linearise this term such that the regularisation energy is written as a least-square problem to justify the Gaussian regularisation.

Equipped with a closed form expression of demons regularisation, we adapted it to elastic registration by replacing the Tikhonov regulariser with the infinite sum of isotropic differential quadratic forms whose minimiser is exactly computed through convolution with the separable elastic-like vector filter proposed by [Cachier and Ayache, 2004].

We then enforced the algorithm to provide incompressible deformations by constraining the search space of stationary velocities to the space of divergence-free




Figure 4.8: Synthetic images with user-defined slice thickness

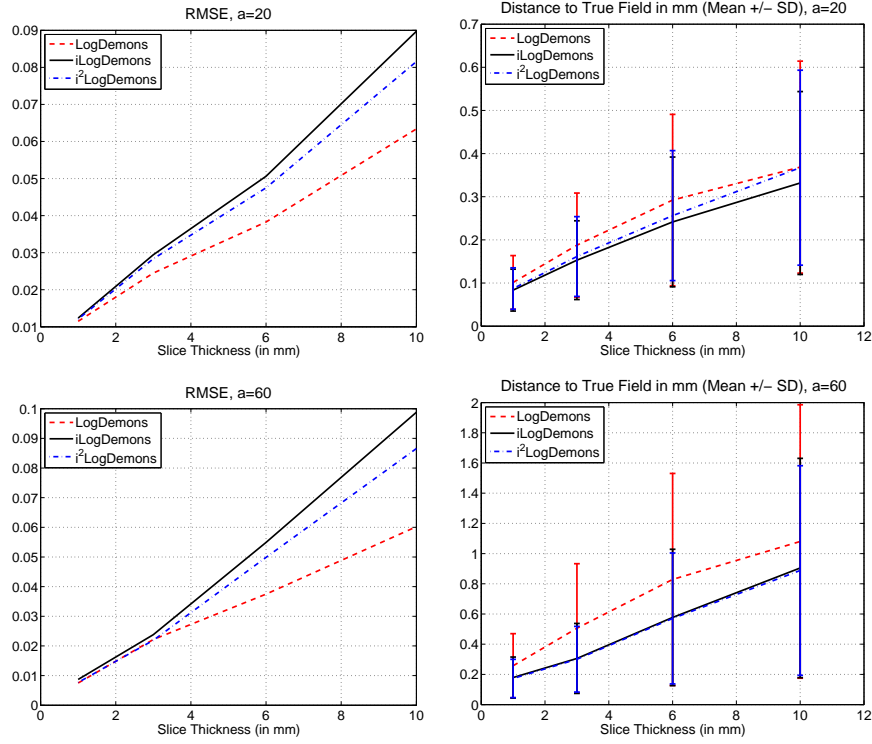


Figure 4.9: RMSE and DTF with respect to the slice thickness for two different whirl angles. Despite higher RMSE, estimated deformation fields were systematically more accurate using incompressible constraints (iLogDemons) than using log-Demons. One can also see that constraining the update velocity to be divergence-free (i^2 LogDemons) did not improve the results in terms of DTF.

vector fields. In practice, this is achieved by adding a new term to the deformation field estimated by the logDemons. The constraint can therefore be enabled/disabled by the user, no *ad-hoc* minimisation schemes being required. Compared with traditional methods, our approach is well posed, provides diffeomorphic transformations, introduces only one extra parameter, the elastic parameter κ , and can be applied




Figure 4.10: Streamlines of 3D whirl transformations ($\alpha_0 = 60^\circ$) recovered from an image with 6mm-thick slices. Contrary to logDemons, which failed to recover the whirl transformation in homogeneous regions (A) and was misguided by strong gradients (B), iLogDemons improved the recovery of the through-plane whirl deformations (C) despite the lack of image information. Colours encode the amplitude of the through-plane z-direction in mm.

within a localised area of the image only.

The synthetic experiments demonstrated that the proposed elastic-like regulariser provides elastic deformations with infinite order of smoothness. Contrary to more traditional approaches based on the linear Lamé equation, our method relies on separable vector filters that can be implemented using efficient Gaussian filters [Deriche, 1993]. As it relies on a kernel, our technique may recall spline-based elastic registration algorithms [Sorzano et al., 2005]. However, the transformation provided by the iLogDemons are diffeomorphic and estimated everywhere in the image.

The synthetic experiments also supported the proposed incompressibility constraint. We showed that deformations parameterised by stationary divergence-free velocities are nearly incompressible despite the approximations and the numerical accuracy. The linearity of the divergence allows efficient implementation of the constraint, in contrast to previous approaches based on the non-linear determinant constraint which was ensured either by *ad-hoc* numerical schemes [Haber and Modersitzki, 2004] or energy penalisation [Tanner et al., 2002; Rohlfing et al., 2003; Bistoguet et al., 2008]. In all our experiments, the Jacobian determinant of the estimated deformations remained close to one independently of the strength of the deformations. We also showed i²LogDemons, which also constrains the update velocities to be divergence-free, does not significantly improve the recovery of incompressible deformations.

Some theoretical aspects still need to be consolidated. When the incompressibility constraint is applied within a limited subdomain of the image, smooth transitions are ensured by artificially diffusing the Lagrangian pressure field p . Intuitively, this technique would be equivalent to using a mask of the incompressible domain with smooth transitions. Experiments supported this approach as no numerical instabilities appeared. However, a rigorous formulation that explicitly integrates the mask into the registration energy functional would enable more efficient numerical

schemes.

A second theoretical challenge to investigate is the fluid-like Gaussian regularisation, or better said viscous regularisation, of the update velocities $\delta \mathbf{v}$ (third step of Algorithm 3, page 66 and Algorithm 4, page 72). Intuitively, this intermediate smoothing controls the regularity of the log-domain, which must be smooth enough to ensure that the integrated transformations are diffeomorphisms [DoCarmo, 1992]. Although this step is optional, its use greatly contributes to the stability of the algorithm: A slight smoothing of the update velocities is recommended. It would be interesting however to develop a theoretical proof of this intuition to implement more sophisticated regularisation schemes. A possible direction would be to endow the space of velocity fields with a kernel norm, like in [Hernandez et al., 2009].

Numerically, our approach is relatively simple to implement as it is based on Gaussian filters and it requires solving a linear system with constant stiffness matrix. In this work the Poisson equation was solved in the space domain to be able to constrain incompressibility in regions of arbitrary shapes, which would have been difficult to achieve with Fourier techniques. Nevertheless, elastic-like divergence-free filters implemented in the Fourier domain are more efficient for whole-domain incompressibility constraint, as in [Hinkle et al., 2009]. As additional future directions, it would be interesting to integrate our approach into registration methods based on time-varying velocity fields [Beg et al., 2005; Hinkle et al., 2009]. Also, more advanced regularisation schemes could be integrated like the locally adaptive regularisers proposed by [Stefanescu et al., 2004; Cahill et al., 2009].

4.7 Appendices

Time-Integration of Stationary Velocity Fields

[Arsigny et al., 2006a] devised an efficient way to compute the exponential map of a stationary velocity field $\phi = \exp(\mathbf{v})$ by observing that, in virtue of the properties of one-parameter subgroups ($t \mapsto \exp(t\mathbf{v})$), we have $\forall n \in \mathbb{N}, \exp(\mathbf{v}) = \exp(\mathbf{v}/n)^n$. If n is large enough, $\exp(\mathbf{v}/n)$ can be approximated by \mathbf{v}/n (*scaling*), which is then composed $\log_2(n)$ times to get the exponential (*squaring*). The pseudo-code of the algorithm is:

Algorithm 5 Scaling-and-Squaring Algorithm

Require: Velocity field \mathbf{v}

- 1: Choose n such that $\|\mathbf{v}/2^n\| \leq 0.5$
 - 2: Explicit first-order integration: $\mathbf{u} \leftarrow \mathbf{v}/2^n$
 - 3: **loop** $\{n \text{ times}\}$
 - 4: $\mathbf{u} \leftarrow \mathbf{u} \circ \mathbf{u}$
 - 5: **return** Displacement field \mathbf{u}
-

On the Equivalence Between Divergence-Free Velocities and Incompressible Diffeomorphic Transformations

In fluid dynamics, a fluid is said incompressible if the divergence of its velocity \mathbf{v} is null, $\nabla \cdot \mathbf{v} = 0$ [Evans, 1998]. Interestingly, deformations of the one-parameter subgroup \mathbb{G} are incompressible if the velocity field that parameterise them is divergence-free. As transformations $\phi \in \mathbb{G}$ are solutions to the Lagrangian ODE

$$\begin{cases} \frac{\partial \phi(\mathbf{x}, t)}{\partial t} = \mathbf{v}(\phi(\mathbf{x}, t)), \\ \phi(\mathbf{x}, 0) = \mathbf{x} \end{cases} \quad (4.17)$$

they are incompressible if and only if at each time instant t , $\det(\nabla \phi(\mathbf{x}, t)) = 1$. This relation holds for $t = 0$ as $\det(\nabla \phi(\mathbf{x}, 0)) = \det(\nabla \mathbf{x}) = 1$. For $t > 0$, the variation of volume due to ϕ is given by the time-derivative of the determinant:

$$\frac{\partial \det(\nabla \phi(\mathbf{x}, t))}{\partial t} = \det(\nabla \phi(\mathbf{x}, t)) \operatorname{tr} \left(\nabla \phi(\mathbf{x}, t)^{-1} \frac{\partial(\nabla \phi(\mathbf{x}, t))}{\partial t} \right)$$

Swapping time and spatial derivatives yields:

$$\frac{\partial \det(\nabla \phi(\mathbf{x}, t))}{\partial t} = \det(\nabla \phi(\mathbf{x}, t)) \operatorname{tr}(\nabla \phi(\mathbf{x}, t)^{-1} \nabla \mathbf{v}(\phi(t, \mathbf{x})))$$

It follows from the chain rule on $\nabla \mathbf{v}(\phi(t, \mathbf{x}))$:

$$\frac{\partial \det(\nabla \phi(\mathbf{x}, t))}{\partial t} = \det(\nabla \phi(\mathbf{x}, t)) \nabla \cdot \mathbf{v}(\mathbf{x})$$

When the velocity \mathbf{v} is divergence-free, the previous equation is homogeneous and its solution is constant. Since $\det(\nabla \phi(\mathbf{x}, 0)) = 1$, we have $\det(\nabla \phi(\mathbf{x}, t)) = 1$ for all $t > 0$. The two constraints are equivalent. In simple words, $\nabla \cdot \mathbf{v}(x) = 0$ is the Eulerian condition of incompressibility, which focuses on a specific location \mathbf{x} of space. Conversely, $\det(\nabla \phi(\mathbf{x}, t)) = 1$ is the Lagrangian condition of the constraint, which tracks the material point as it moves within the domain. The Lagrangian frame and the Eulerian frame are bridged by the transport Equation 4.17.

Conditions on IDQF Parameters to Get Separable Vector Filters

As described in Section 4.3.2, logDemons elastic-like regularisation is obtained by solving in the Fourier domain the optimal condition:

$$\left[\underbrace{\left(1 + \sum_{k=1}^{\infty} \frac{\alpha_k (\mathbf{w}^T \mathbf{w})^k}{\sigma_d^{2k}} \right)}_A \operatorname{Id} + \underbrace{\left(\sum_{k=1}^{\infty} \frac{\beta_k (\mathbf{w}^T \mathbf{w})^{k-1}}{\sigma_d^{2k}} \right)}_B \mathbf{w} \mathbf{w}^T \right] \hat{\mathbf{v}}(\mathbf{w}) = \hat{\mathbf{v}}_c(\mathbf{w})$$

Sherman-Morrison inversion Lemma gives:

$$\hat{\mathbf{v}}(\mathbf{w}) = \underbrace{\left[\frac{1}{A} \operatorname{Id} - \frac{1}{A} \left(\frac{B}{A + B \mathbf{w}^T \mathbf{w}} \right) \mathbf{w} \mathbf{w}^T \right]}_M \hat{\mathbf{v}}_c(\mathbf{w})$$

We seek α_k and β_k such that the filter M is separable to preserve demons computational efficiency. $\alpha_k = 1/k!$ yields $A = \exp(\mathbf{w}^T \mathbf{w} / \sigma_t^2)$. As a result, if $B/(A + B\mathbf{w}^T \mathbf{w})$ is a scalar $\gamma \in \mathbb{R}$, the inverse Fourier transform of the second term of M is the Hessian of $\mathfrak{F}^{-1}(\gamma \exp(-\mathbf{w}^T \mathbf{w} / \sigma_d^2))$. The idea thus consists in finding the coefficients β_k such that $B/(A + B\mathbf{w}^T \mathbf{w}) = \gamma$. With $\beta_0 = 0$, this writes:

$$\sum_{k=0}^{\infty} \frac{\beta_{k+1}}{\sigma_d^{2(k+1)}} (\mathbf{w}^T \mathbf{w})^k = \sum_{k=0}^{\infty} \gamma \left(\frac{1}{\sigma_d^{2k} k!} + \frac{\beta_k}{\sigma_d^{2k}} \right) (\mathbf{w}^T \mathbf{w})^k$$

This equation defines a recursive relationship between the β_k 's:

$$\begin{cases} \beta_0 = 0 \\ \beta_{k+1} = \gamma \left(\frac{\sigma_d^2}{k!} + \beta_k \right) \end{cases} \quad \forall k \geq 1$$

from which we deduce the closed form:

$$\beta_k = \sum_{i=1}^k \frac{\gamma^i \sigma_d^{2i}}{(k-i)!}$$

The proof of this relationship is achieved by recurrence. The previous formula is verified for $k=1$. We assume it to be true for k and we verify it still holds for $k+1$. To this end, we replace β_k by the conjectured formula in the recursive expression of β_{k+1} :

$$\begin{aligned} \beta_{k+1} &= \sigma_d^2 \gamma \left(\frac{1}{k!} + \sum_{i=1}^k \frac{\gamma^i \sigma_d^{2i}}{(k-i)!} \right) = \frac{\gamma \sigma_d^2}{k!} + \sum_{i=1}^k \frac{\gamma^{i+1} \sigma_d^{2(i+1)}}{(k-i)!} \\ &= \frac{\gamma \sigma_d^2}{k!} + \sum_{i=2}^{k+1} \frac{\gamma^i \sigma_d^{2i}}{(k-(i-1))!} = \sum_{i=1}^{k+1} \frac{\gamma^i \sigma_d^{2i}}{(k+1-i)!} \end{aligned}$$

which proves the result. With these coefficients, the filter M writes:

$$\hat{\mathbf{v}}(\mathbf{w}) = (\exp(-\mathbf{w}^T \mathbf{w} / \sigma_d^2) \text{Id} - \gamma \mathbf{w} \mathbf{w}^T \exp(-\mathbf{w}^T \mathbf{w} / \sigma_d^2)) \star \hat{\mathbf{v}}_c(\mathbf{w})$$

which becomes in the space domain (Δ is the Hessian operator and d is the dimension of the image domain Ω)

$$\begin{aligned} \mathbf{v}(\mathbf{x}) &= \sqrt{\sigma_d^2 / (4\pi)}^d \exp(-4\sigma_d^2 \mathbf{x}^T \mathbf{x}) \star \mathbf{v}_c(\mathbf{x}) \\ &\quad + \gamma \Delta \left(\sqrt{\sigma_d^2 / (4\pi)}^d \exp(-4\sigma_d^2 \mathbf{x}^T \mathbf{x}) \right) \star \mathbf{v}_c(\mathbf{x}) \end{aligned}$$

Defining $\sigma^2 = 2/\sigma_d^2$ and $\gamma = \sigma^2 \kappa / (1 + \kappa)$ yields Equation 4.12.

Demons Optimisation under the Divergence-Free Constraint

Divergence-free update velocities are computed by minimising the constrained optimisation energy:

$$\begin{cases} \mathcal{E}_{cor}(\delta \mathbf{v}) = \frac{1}{\sigma_i^2} \|R - T \circ \phi \circ \exp(\delta \mathbf{v})\|_{L_2}^2 + \frac{1}{\sigma_x^2} \|\log(\phi^{-1} \circ \phi \circ \exp(\delta \mathbf{v}))\|_{L_2}^2 \\ \nabla \cdot \delta \mathbf{v} = 0 \end{cases}$$

Let $p(\mathbf{x})$ be a scalar field belonging to the Sobolev space $H_0^1(\Omega)$ that vanishes at infinity. The Lagrangian function $\mathcal{P}_{corr}(\delta \mathbf{v}, p)$ related to the constrained correspondence energy $\mathcal{E}_{corr}(\delta \mathbf{v})$ writes (the Lagrangian multiplier $p(\mathbf{x})$ has been multiplied by 2 to simplify calculations):

$$\begin{aligned} \mathcal{P}_{corr}(\delta \mathbf{v}, p) &= \frac{1}{\sigma_i^2} \int_{\Omega} \|R(\mathbf{x}) - T \circ \phi \circ \exp \delta \mathbf{v}(\mathbf{x})\|^2 d\mathbf{x} \\ &+ \frac{1}{\sigma_x^2} \int_{\Omega} \|\delta \mathbf{v}(\mathbf{x})\|^2 d\mathbf{x} - 2 \int_{\Omega} p(\mathbf{x}) \nabla \cdot \delta \mathbf{v}(\mathbf{x}) d\mathbf{x} \end{aligned}$$

Differentiating $\mathcal{P}_{corr}(\delta \mathbf{v}, p)$ with respect to p yields the constraint $\nabla \cdot \delta \mathbf{v} = 0$. To differentiate $\mathcal{P}_{corr}(\delta \mathbf{v}, p)$ with respect to $\delta \mathbf{v}$, we linearise the similarity criterion as in [Vercauteren et al., 2009]:

$$\|R(\mathbf{x}) - T \circ \phi \circ \exp \delta \mathbf{v}(\mathbf{x})\|^2 \approx \|R(\mathbf{x}) - T \circ \phi(\mathbf{x}) + J(\mathbf{x})^T \delta \mathbf{v}(\mathbf{x})\|^2$$

where $J(\mathbf{x})$ is the symmetric gradient $J(\mathbf{x}) = (\nabla R(\mathbf{x}) + \nabla(T \circ \phi)(\mathbf{x}))/2$. It follows the linear least square problem:

$$\begin{aligned} \mathcal{P}_{corr}(\delta \mathbf{v}, p) &= \frac{1}{\sigma_i^2} \int_{\Omega} \|R(\mathbf{x}) - T \circ \phi(\mathbf{x}) + J(\mathbf{x})^T \delta \mathbf{v}(\mathbf{x})\|^2 d\mathbf{x} \\ &+ \frac{1}{\sigma_x^2} \int_{\Omega} \|\delta \mathbf{v}(\mathbf{x})\|^2 d\mathbf{x} - 2 \int_{\Omega} p(\mathbf{x}) \nabla \cdot \delta \mathbf{v}(\mathbf{x}) d\mathbf{x} \end{aligned}$$

The optimal condition $\partial_{\delta \mathbf{v}} \mathcal{P}_{corr}(\delta \mathbf{v}, p) = 0$ thus writes:

$$\underbrace{\left(J(\mathbf{x}) J(\mathbf{x})^T + \frac{\sigma_i^2}{\sigma_x^2} \text{Id} \right)}_{D(\mathbf{x})} \delta \mathbf{v}(\mathbf{x}) = -J(\mathbf{x}) \left(R(\mathbf{x}) - T \circ \phi(\mathbf{x}) \right) - \sigma_i^2 \nabla p(\mathbf{x})$$

As the tensor $D(\mathbf{x})$ is always invertible, we can calculate the optimal divergence-free update velocity field:

$$\delta \mathbf{v}^*(\mathbf{x}) = - \underbrace{\left(R(\mathbf{x}) - T \circ \phi(\mathbf{x}) \right) D(\mathbf{x})^{-1} J(\mathbf{x})}_{\delta \mathbf{v}(\mathbf{x})} - \sigma_i^2 D(\mathbf{x})^{-1} \nabla p(\mathbf{x}) \quad (4.18)$$

The first term of the previous equation is exactly the logDemons update velocity field $\delta \mathbf{v}$ (Equation 4.3). The scalar field p is calculated by solving the Poisson

equation under 0-Dirichlet boundary conditions that results from the divergence of the previous equation:

$$\nabla \cdot (\sigma_i^2 D(\mathbf{x})^{-1} \nabla p(\mathbf{x})) = \nabla \cdot \delta \mathbf{v}(\mathbf{x}) \quad (4.19)$$

Because the tensor $D(\mathbf{x})$ is updated at each iteration, the operator $\nabla \cdot (\sigma_i^2 D^{-1} \nabla)$ is not constant. The matrix of the related linear system must therefore be built and pre-conditioned at each time step, which can be computationally demanding if the domain Ω is large. Furthermore, $D(\mathbf{x})$ is computed at every voxel of the image domain Ω independently. The resulting tensor field can therefore be noisy, likely yielding numerical instabilities. To alleviate this limitation, D is smoothed using Log-Euclidian techniques [Arsigny et al., 2006b]: each component of $\log(D)$ is smoothed with a Gaussian kernel G_{σ_D} and the result is exponentiated to get a smooth tensor field. We fixed σ_D equal to the strength σ_f of the fluid regularisation.

Numerical Implementation of Poisson Equations

Let \mathbf{f} be a vector field defined over a rectangular image domain Ω in \mathbb{R}^d . The dimension of Ω along the i^{th} axis is s_i and the voxel spacing v_i . We want to solve the Poisson equation with 0-Dirichlet boundary conditions:

$$\nabla \cdot (M \nabla p) = \nabla \cdot \mathbf{f}$$

This equation is solved by discretising the spatial derivatives, which results in the matrix form

$$Kp = R \quad (4.20)$$

The matrix K , so-called stiffness matrix of the linear system, is a $n \times n$ matrix where n is the total number of voxels of Ω . R is the $n \times 1$ right-hand side vector. The stiffness matrix K is sparse symmetric positive if M is symmetric positive and if divergence and gradient operators are conjugate. The latter property is satisfied when the gradient and divergence operators are discretised using backward and forward first-order finite difference schemes respectively.

For convenience we define the mapping $\mu : \mathbb{N}^d \rightarrow \mathbb{N}$ that associates a unique index ι to each voxel of the image. ι^{th} row of the right-hand side vector R is therefore the divergence of \mathbf{f} at the ι^{th} voxel of the image. Let $(\mathbf{e}_i)_{i=1\dots d}$ be the canonical coordinate system of the image domain Ω and $M(\mathbf{x}) = (m_{i,j}(\mathbf{x}))$. With the forward and backward first order finite difference schemes for the divergence and gradient operators respectively we have:

$$\begin{aligned} \nabla p(\mathbf{x}) &= \sum_{i=1}^d \frac{p(\mathbf{x}) - p(\mathbf{x} - \mathbf{e}_i)}{v_i} \mathbf{e}_i \\ M(\mathbf{x}) \nabla p(\mathbf{x}) &= \sum_{j=1}^d \left(\sum_{i=1}^d m_{j,i}(\mathbf{x}) \frac{p(\mathbf{x}) - p(\mathbf{x} - \mathbf{e}_i)}{v_i} \right) \mathbf{e}_j \end{aligned}$$

and finally:

$$\begin{aligned} \nabla \cdot M(\mathbf{x}) \nabla p(\mathbf{x}) = \sum_{j=1}^d \sum_{i=1}^d \left(m_{j,i}(\mathbf{x} + \mathbf{e}_j) \frac{p(\mathbf{x} + \mathbf{e}_j) - p(\mathbf{x} + \mathbf{e}_j - \mathbf{e}_i)}{v_j v_i} \right. \\ \left. - m_{j,i}(\mathbf{x}) \frac{p(\mathbf{x}) - p(\mathbf{x} - \mathbf{e}_i)}{v_j v_i} \right) \end{aligned} \quad (4.21)$$

When $M = \text{Id}_d$ the Laplacian operator writes:

$$\Delta p(\mathbf{x}) = \sum_{i=1}^d \left(\frac{p(\mathbf{x} - \mathbf{e}_i) + p(\mathbf{x} + \mathbf{e}_i)}{v_i^2} - \frac{2p(\mathbf{x})}{v_i^2} \right)$$

We use these expressions to populate the stiffness matrix K . Assume that the ι^{th} voxel of the image is at the physical point \mathbf{x} ($\mu(\mathbf{x}) = \iota$). The ι^{th} row of K is populated according to Equation 4.21 as follows:

$$\begin{aligned} K_{\iota,\iota} &= - \sum_{j=1}^d \sum_{i=1}^d m_{j,i}(\mathbf{x}) / v_j v_i \\ \forall j, \quad K_{\iota,\mu(\mathbf{x}+\mathbf{e}_j)} &= \sum_{i=1}^d m_{j,i}(\mathbf{x} + \mathbf{e}_j) / v_j v_i \\ \forall i, \quad K_{\iota,\mu(\mathbf{x}-\mathbf{e}_i)} &= \sum_{j=1}^d m_{j,i}(\mathbf{x} - \mathbf{e}_i) / v_j v_i \\ \forall (j, i), \quad K_{\iota,\mu(\mathbf{x}+\mathbf{e}_j-\mathbf{e}_i)} &= -m_{j,i}(\mathbf{x} + \mathbf{e}_j - \mathbf{e}_i) / v_j v_i \end{aligned}$$

If $M = \text{Id}_d$, these relations become:

$$\begin{aligned} K_{\iota,\iota} &= - \sum_{i=1}^d 2/v_i^2 \\ K_{\iota,\mu(\mathbf{x}-\mathbf{e}_i)} &= K_{\iota,\mu(\mathbf{x}+\mathbf{e}_i)} = 1/v_i^2 \quad \forall i = 1 \dots d \end{aligned}$$

When a voxel is outside or on the boundary of the domain Ω , the related factor is not reported in the stiffness matrix K as p equals 0 at that voxel (0-Dirichlet boundary conditions).

Estimation of 3D Myocardium Strain from Cine MRI

Contents

5.1	Motivation	89
5.2	Cardiac Tracking Algorithm and Strain Calculation	91
5.3	Experiments and Results	93
5.3.1	LV Myocardium Strain in Patients with Heart Failure	93
5.3.2	LV Myocardium Strain in a ToF Patient	98
5.3.3	Computation Time	102
5.4	Discussion	102

In the following we investigate the clinical question that motivated the development of the iLogDemons: Can we estimate the myocardium deformation from standard anatomical cMRI? As the heart is an elastic near-incompressible tissue [Glass et al., 1991], it is reasonable to use these priors to enhance the accuracy of the estimated displacements. In this chapter we test that idea on real cases.

5.1 Motivation

Analysing cardiac motion provides crucial insights into the condition of the cardiac function. Myocardium strain can be quantitatively assessed with tagged MRI (tMRI) or ultrasound modalities like 2D-strain based on speckle tracking (recently 3D-strain is becoming available). On the one hand, tMRI is nowadays considered the gold standard for myocardium motion assessment [Zerhouni et al., 1988; McVeigh, 1998]. However, that MR protocol is still experimental and is not readily available in paediatrics. On the other hand, 2D speckle tracking is available in clinical environment. Nonetheless, it only provides a partial, two-dimensional evaluation of the myocardium motion [Teske et al., 2007]. Estimating the myocardium strain from clinical gated cMRI constitutes an attractive alternative. cMRI is widely available and yields detailed 3D anatomical images of the beating heart with constant image quality over time. Yet, accurately estimating the 3D cardiac motion from these images is challenging. Due to the lack of consistent texture within the myocardium, only the apparent radial motion of the heart can be assessed. Furthermore, due to acquisition limitations, slice thickness of cMRI is usually very large, of the order of the centimetre. Prior knowledge about cardiac dynamics is thus required.

Regardless of the difficulty, various methods have been proposed to estimate the deformation of the myocardium from cMRI. Some approaches drive an image registration algorithm with biomechanical finite element models that simulate the passive properties of the myocardium [Papademetris et al., 2000; Veress et al., 2005; Schaerer et al., 2010]. These methods successfully estimate the left ventricle myocardium displacement. Nonetheless, they rely on specific biomechanical parameters that are difficult to determine for a given patient. Besides, such models may not apply anymore in pathological cases. Finally, they require meshing the space domain to solve complex partial differential equations, which may be difficult to achieve. Purely image-driven methods have thus been proposed. In [Delhay et al., 2007], the authors estimate the cardiac motion using a spatio-temporal model based on Free-Form Deformation (FFD) [Rueckert et al., 1999] and Kalman filters, under the assumption of periodic motion. Yet, the proposed approach may not provide deformations consistent with the biomechanical properties of the myocardium. In [Feng et al., 2008], the authors use a B-spline FFD algorithm regularised with myocardium contours that are manually delineated on all the time frames. The results are consistent with tagged MRI but the method requires manual contours for all the time frames, which may hinder its application in clinical settings.

During the cardiac cycle, it has been reported that the volume of the heart muscle does not vary significantly (about 5% of volume variation [Glass et al., 1991]). It is therefore reasonable to use incompressibility constraints to estimate the cardiac motion, as in [Bistoquet et al., 2008]. In that work, the authors present an FFD-like algorithm that ensures myocardium near-incompressibility. To that end, the dense deformation field is interpolated using divergence-free radial basis functions. However, divergence-free displacements are incompressible only when the deformations are small, which may not be the case at the end-systole when the heart is fully contracted. An *ad-hoc* penalisation is used to control volume drifts. Nonetheless, constraining the displacements to be divergence-free may not reflect the true motion as incompressible deformations are not necessarily divergence-free.

In the following we apply iLogDemons to estimate the 3D strain of the left-ventricular (LV) myocardium from standard anatomical cMRI of the heart. Contrary to [Bistoquet et al., 2008], here we employ a strong incompressibility constraint, completed with an elastic-like regularisation to estimate physiologically plausible deformations. The algorithm is tested on two adults with heart failure and one teenager with repaired ToF. The results are compared with those obtained using the original logDemons, to quantify the added value of our method, and the ground truth provided by tagged MRI (in the adult patients) and echocardiography 2D-strain (in the young patient).

5.2 Cardiac Tracking Algorithm and Strain Calculation

Tracking Protocol

The displacement of the heart muscle is estimated by tracking the organ over the cardiac cycle using iLogDemos. We recall that iLogDemos estimates a deformation $\phi = \text{Id} + \mathbf{u}$ that maps two images, \mathbf{u} is the displacement field. ϕ is parameterised by a stationary velocity field \mathbf{v} through the exponential map $\phi = \exp(\mathbf{v})$. In iLogDemos, \mathbf{v} is divergence-free. An important observation is that demons algorithm, like other common registration algorithms, provides resampling fields. In other words, the 3D displacement \mathbf{u} of structures that move between the images I_0 and I_n is calculated by registering I_n to I_0 as illustrated in Figure 5.1.




Figure 5.1: Registering the moving image to the fixed image yields a displacement field $\mathbf{u}(\mathbf{x})$ that maps each voxel \mathbf{x} of the fixed image to a point in the moving image. The field $\mathbf{u}(\mathbf{x})$ is called the resampling field.

In that way, estimating the deformation of the heart amounts to “stabilising” the cardiac sequence: All the frames are registered back to a frame of reference, usually the end-diastole time frame identified by the R-wave peak. At that instant of the cardiac cycle, the heart is nearly at rest position, i.e. almost relaxed. From a computational point of view, it is also easier to estimate compressions than expansions due to the fixed-size sampling grid. To minimise error propagation and ensure temporal consistency, the stabilisation is performed recursively as depicted in Figure 5.2. Without loss of generality, let I_0 be the reference frame and I_n the subsequent frames. Assume now that the deformation $\phi_{I_{n-1} \rightarrow I_0}$ and thus the velocity field $\mathbf{v}_{I_{n-1} \rightarrow I_0}$, are known. We first estimate the frame-by-frame transformation parameterised by the velocity $\mathbf{v}_{I_n \rightarrow I_{n-1}}$ to take advantage of the small inter-frame cardiac deformation. We then concatenate the resulting transformation to $\phi_{I_{n-1} \rightarrow I_0}$, which amounts to composing the velocity fields using BCH formula, $\mathbf{v}_{concat} = Z(\mathbf{v}_{I_{n-1} \rightarrow I_0}, \mathbf{v}_{I_n \rightarrow I_{n-1}})$. The deformation parameterised by \mathbf{v}_{concat} would normally map the current frame I_n to the reference frame I_0 but to further reduce

registration errors, we refine the transformation by registering I_n to I_0 with initial velocity \mathbf{v}_{concat} .




Figure 5.2: Recursive tracking algorithm. Knowing the velocity $\mathbf{v}_{I_{n-1} \rightarrow I_0}$ (green): i) Estimate $\mathbf{v}_{I_n \rightarrow I_{n-1}}$ (blue). ii) Concatenate $\mathbf{v}_{I_n \rightarrow I_{n-1}}$ and $\mathbf{v}_{I_{n-1} \rightarrow I_0}$ (grey) using BCH formula. iii) Estimate $\mathbf{v}_{I_n \rightarrow I_0}$ using ii) as initialisation (red).

In cardiac sequences, other structures surrounding the heart may have their volume vary. Ventricular blood pools in particular are not incompressible, the isovolumetric phase being usually too short to be visible in standard clinical dynamic images. Considering them as incompressible may yield “locking”, i.e. the myocardium remains frozen as the entire volume of the heart cannot change. To cope with that, we apply the incompressibility constraint only within the myocardium (Figure 5.3). Besides, this greatly decreases the computation time as the Poisson equations related to the constraint are solved on a subspace of the image only. It has to be noted that owing to the backward stabilisation strategy, the myocardial region must be defined only on the reference frame in contrast to previous approaches [Papademetris et al., 2000; Feng et al., 2008]. This can be achieved using the interactive surface modeller presented in Chapter 3 for instance.




Figure 5.3: Short-axis cMRI of patient 1. Incompressibility is ensured only within the myocardium (outlined in yellow). Note the lack of consistent texture within the myocardium even in the in-plane image (*left panel*) and the coarse through plane resolution (*right panel*).

Strain Calculation

The 3D myocardium strain at the spatial position \mathbf{x} is computed from the estimated displacement $\mathbf{u}(\mathbf{x})$ using the 3D Lagrangian finite strain tensor:

$$\mathbf{E}(\mathbf{x}) = \frac{1}{2} [\nabla \mathbf{u}(\mathbf{x}) + \nabla \mathbf{u}^T(\mathbf{x}) + \nabla \mathbf{u}^T(\mathbf{x}) \nabla \mathbf{u}(\mathbf{x})]$$

The strain tensor is projected onto a local prolate coordinate system to get the radial, circumferential and longitudinal indices commonly used in clinics [Moore et al., 2000] (Figure 5.4, left panel). The local prolate coordinate system is defined as follows. At each material point of the myocardium, the radial axis \mathbf{e}_r is the outward normal of the epicardial surface. The circumferential axis is the cross product of the left ventricle (LV) long axis \mathbf{e}_{LA} with \mathbf{e}_r : $\mathbf{e}_c = \mathbf{e}_{LA} \wedge \mathbf{e}_r$. \mathbf{e}_c lies in the transverse plane of the heart. Finally, the longitudinal axis is the cross product of the radial axis with the circumferential axis: $\mathbf{e}_l = \mathbf{e}_c \wedge \mathbf{e}_r$. \mathbf{e}_l lies in the transverse plane and is tangential to the epicardial surface. The 3D strain $\mathbf{E}(\mathbf{x})$ in the direction $\mathbf{v}(\mathbf{x})$ is given by $E_{\mathbf{v}}(\mathbf{x}) = \mathbf{v}(\mathbf{x})^T \mathbf{E}(\mathbf{x}) \mathbf{v}(\mathbf{x})$. Finally, the myocardium is automatically subdivided into 17 AHA zones for regional assessment [Cerqueira et al., 2002] (Figure 5.4, right panel and Section 2.2, page 19).




Figure 5.4: *Left panel*: Prolate coordinate system. The radial direction \mathbf{e}_r is given by the outward normal of the epicardium surface. The circumferential direction \mathbf{e}_c is computed at each point of the myocardium by $\mathbf{e}_c = \mathbf{e}_{LA} \wedge \mathbf{e}_r$, where \mathbf{e}_{LA} is the long-axis direction. The longitudinal direction \mathbf{e}_l is finally obtained by $\mathbf{e}_l = \mathbf{e}_c \wedge \mathbf{e}_r$. *Right panel*: AHA subdivision of the myocardium (see Section 2.2, page 19).

5.3 Experiments and Results

5.3.1 Left Ventricle Myocardium Strain in Patients with Heart Failure

5.3.1.1 Patient Selection and Image Preparation

We first tested our algorithm on two patients with heart failure recruited at Saint Thomas Hospital, King's College London, U.K. Anatomical cMRI were acquired in the short axis view with multiple breath-holds (Achieva, Philips Medical System,

30 time frames, 1.4 mm^2 isotropic in-plane resolution, 10 mm slice thickness). For the first patient, 3D tagged MR images (tMRI) were acquired during the same exam (CSPAMM encoding, 23 time frames, 1.0 mm^3 isotropic resolution, tag size $\approx 7\text{ mm}$, Figure 5.5). No manual tracking of the tag grids was available since this task is extremely difficult due to the 3D nature of the motion. For the second patient, 2D tMRI were acquired in the short axis view (23 time frames, 1.1 mm^2 in-plane resolution, 18 mm slice thickness, tag size $\approx 6\text{ mm}$). For this case, manual tracking of the tag grids was performed by an expert in the short axis view. All images fully covered both ventricles. The two cMRI and the 3D tMRI presented no slice misalignments. The tMRI were spatially and temporally aligned to the cMRI using DICOM information. For the incompressibility constraint, we manually delineated the myocardium on the end-diastole time frame using the interactive surface modeller (Chapter 3). Finally, the dense transformations were defined on an isotropic sampling grid adjusted to the image dimensions to cope with the large slice thickness that can introduce high frequencies in the transformations, resulting in numerical instabilities and lower registration accuracy. In practice, this amounts to linearly resampling the cMRI to get isotropic voxels.




Figure 5.5: 3D tMRI of patient 1.

The heart was tracked using both logDemons and iLogDemons for comparison. Registration parameters were fixed: $\sigma_x = 1$, $\sigma^2 = 2$ and $\sigma_f^2 = 0.5$. For the iLogDemons only, elastic regularisation was applied everywhere in the image with $\kappa = 1$. A 2-level multi-resolution scheme was used and the registration was automatically stopped as soon as RMSE stopped decreasing.

5.3.1.2 Comparison with 3D tMRI

In a first stage, we estimated the cardiac motion of patient 1 by tracking the heart in the 3D tMRI using both iLogDemons and logDemons. We verified that iLogDemons preserved the volume of the myocardium below the values reported in the literature during the entire cardiac cycle (average volume variation: 2%, maximum volume variation: 6%) contrary to logDemons (average volume variation: 26%, maximum volume variation: 32%) (Figure 5.6, top-left panel). We also observed that the

incompressibility constraint reduced the deviations of the estimated displacements throughout the cardiac cycle despite the simple tracking procedure.




Figure 5.6: Jacobian determinant of myocardium deformation. Curves represent mean (plain lines) \pm standard deviation (dashed lines). Incompressibility constraint significantly decreased myocardium volume variations during the cardiac cycle and controlled the volume deviation at the end of the cycle.




Figure 5.7: Virtual grids warped using the displacements estimated from tMRI using logDemons (*in red*) and iLogDemons (in green). Visually, no significant differences were noticed within the myocardium, which was confirmed by the low L_2 -distance between the fields ($1.7 \pm 0.71mm$)

For visual assessment, we applied the estimated deformations on virtual planes manually positioned at end-diastole (Figure 5.8). Realistic deformations consistent with the tag grids were obtained with both algorithms (Figure 5.7). The similar performances between logDemons and iLogDemons, quantified by the low L_2 -distance between recovered deformation fields (Table 5.1), is justified by the fact that the tag grids provided enough texture information within the myocardium to guide the registration. As no ground truth was available, we considered the displacements estimated on the 3D tMRI using iLogDemons as reference.

We then estimated the 3D motion of the heart from the cMRI and compared the results with the reference tMRI deformation. Visual assessment of the warped virtual planes showed that the incompressibility constraint did help in recovering longitudinal and circumferential displacements despite the large slice thickness and the lack of texture features within the myocardium (Figure 5.8, blue and red curves). The average and standard deviation of the L_2 -distance between the displacements estimated from cMRI and the reference tMRI displacements improved by 12% and 22% respectively (Table 5.1). The longitudinal and circumferential strains computed over the entire myocardium confirmed this finding (Figure 5.9). Values obtained using iLogDemons were much closer to the reference (86% of improvement for radial strain, 89% for circumferential strain and 65% for longitudinal strain). Similar degrees of improvement were obtained with respect to the motion estimated from tMRI using logDemons. The amplitude of the radial strain was more plausible and the temporal variations of the circumferential and longitudinal strains were consistent with the clinical literature [Moore et al., 2000]. Note that logDemons exhibited a wrong lengthening in both longitudinal and circumferential directions at the beginning of the cardiac contraction. Furthermore, the variability in strain measurements is significantly reduced using iLogDemons, which suggests that the estimated motion is globally more consistent.

Table 5.1: L_2 -distances averaged over the cardiac cycle between estimated displacements. Values to be compared with the tag size: $6mm$. Tracking cardiac motion on tMRI with logDemons and iLogDemons yielded globally little differences. When tracking the heart on cMRI, iLogDemons improved the results thanks to the incompressibility constraint and the elastic-like regularisation, which cope with the lack of myocardial texture and the large slice thickness.

Method	L_2 -distance (mean \pm sd, max)
iLogDemons (<i>tMRI</i>)	reference
logDemons (<i>tMRI</i>)	$1.7 \pm 0.71mm, 3.2mm$
logDemons (<i>cMRI</i>)	$3.2 \pm 0.92mm, 4.6mm$
iLogDemons (<i>cMRI</i>)	$2.8 \pm 0.72mm, 4.0mm$

5.3.1.3 Comparison with Manual Tracking

We then estimated the cardiac motion of the second patient using logDemons and iLogDemons on cMRI and compared the results with manual tracking of 2D tag grids. As with the previous patient, iLogDemons controlled the variations of myocardium volume over the cardiac cycle compared to logDemons (average volume variation: 7% and maximum volume variation: 10% for iLogDemons; average volume variation: 42% and maximum volume variation: 54% for logDemons, see Figure 5.10).

Point-wise motion comparison between cMRI and manual tracking was not possible due to a non-perfect tMRI-to-cMRI alignment because of tMRI slice misalign-

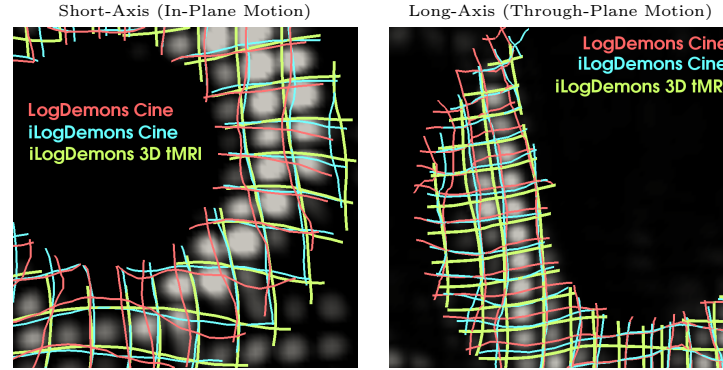


Figure 5.8: Close-up of the tMRI of patient 1 at end-systole. The virtual tag planes were warped with the deformation estimated on the tMRI using iLogDemons (*green*) and with those estimated on the cMRI using logDemons (*red*) and iLogDemons (*blue*). From the blue and red planes one can see that iLogDemons better estimated myocardium motion even in cMRI.

ment and patient motion between scans. We thus compared the regional displacements averaged over the 12 basal and mid-ventricular AHA zones (Figure 5.4, right panel). For fair comparison, we transformed the 3D displacements estimated from the cMRI to apparent 2D displacements by warping the short-axis displacements (XY-plane) along the through plane motion (Z-direction).

The results showed that iLogDemons, in this patient, improved the accuracy of the recovered motion. The amplitude of the radial displacements was closer to the ground truth (global error with respect to tag displacements from $2.4 \pm 2.4mm$ to $1.2 \pm 1.6mm$, about the voxel size, Figure 5.11). One can observe that logDemons already estimated realistic radial motion patterns, yet over-estimated, as the cardiac motion visible in cMRI is mainly radial to the LV boundaries (we mainly see the thickening of the heart). Both LogDemons and iLogDemons recovered negative radial displacements, i.e. towards the centre of the LV blood pool (Figure 5.4), for most of the LV regions. The two algorithms also detected the abnormal dyskinetic motion of the 3rd and 9th zones of the inter-ventricular septum (Figure 5.4). Their displacement was positive, i.e. towards the right ventricle. This abnormal pattern may be due to an electrical asynchrony or to a lesion.

The circumferential displacements provided by iLogDemons were also more realistic than those estimated using logDemons, yet still under-estimated (Figure 5.12). The sign of the estimated circumferential displacements were more consistent with the ground truth estimated from tMRIs. For instance, contrary to logDemons, iLogDemons correctly estimated the counter-clockwise motion of the 3rd and 9th zones (positive values, arrow A in Figure 5.12) and the clockwise motion of the 1st and 7th zones (negative values, arrow B in Figure 5.12). Quantitatively, the overall error decreased by 34%, from $3.5 \pm 2.0mm$ to $2.3 \pm 1.7mm$. In conclusion, the incompressibility constraint assisted the registration algorithm by redistributing the apparent radial displacement across the other directions to preserve myocardium volume.

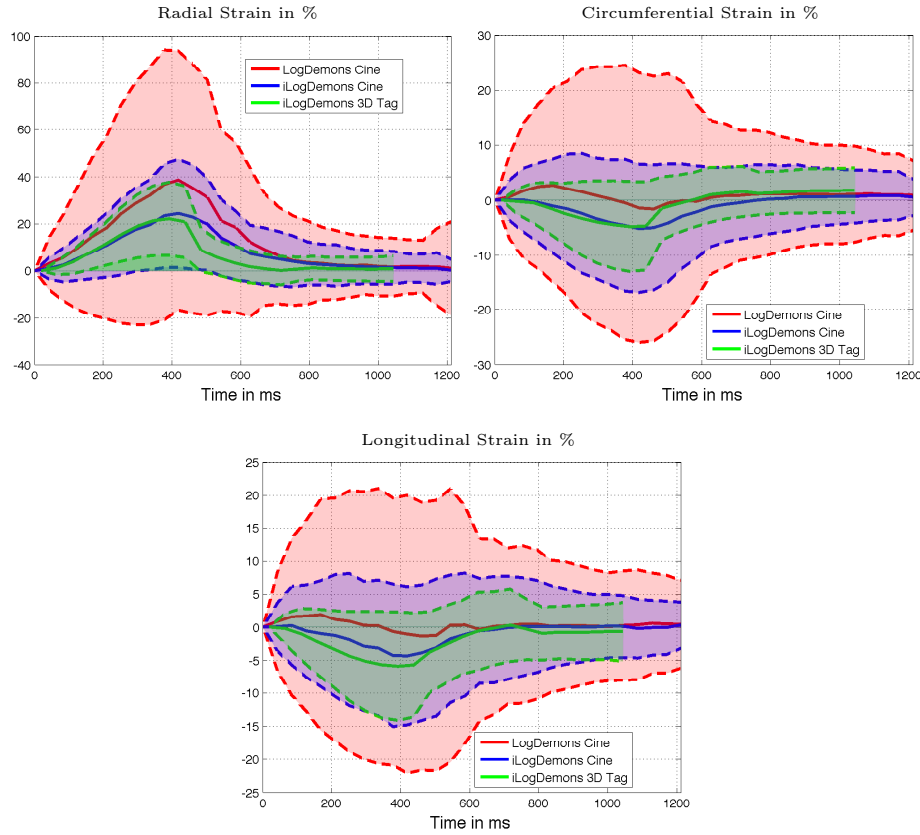


Figure 5.9: Myocardium strains computed from cMRI and 3D tMRI of patient 1. Mean and standard deviation computed over the entire left ventricle. iLogDemons better estimated circumferential and longitudinal strains despite the lack of image information and the large slice thickness.

5.3.2 Left Ventricle Myocardium Strain in a ToF Patient

5.3.2.1 Patient Selection and Image Preparation

In a second experiment we evaluated the iLogDemons on a patient (age=10) with repaired ToF. The estimated myocardium strains were compared with echocardiography 2D-strain measurements.

Anatomical SSFP cMRI of the heart were acquired in the short-axis view covering the entirety of both ventricles (10 9.6 mm-thick slices; $1.02 \times 1.02 \text{ mm}^2$ in-plane resolution; 25 frames) using a 1.5T MR scanner (Avanto, Siemens Medical Systems, Erlangen, Germany). No long-axis cMRI were available. Visual inspection of the images revealed no slice misalignment. As in the previous experiments, the images were made isotropic to get well-defined dense deformation fields. Circumferential 2D-strain measurements were performed in the short-axis view (frame-rate: 80 fps) using Automatic Functional Imaging (Vivid7, General Electrics, Vingmed Ultrasound) as described in [Teske et al., 2007].




Figure 5.10: Jacobian determinant of myocardium deformation. Curves represent mean (plain lines) \pm standard deviation (dashed lines). As for the first patient, the incompressibility constraint significantly decreased myocardium volume variations during the cardiac cycle.




Figure 5.11: Basal and mid-ventricular radial displacements of patient 2 (in mm, averaged per zone). Compared with the radial displacements measured by an expert on tMRI (*left panels*), the displacements estimated on cMRI using iLogDemons (*mid panel*) had more realistic amplitudes compared with those estimated with logDemons (*right panels*). Note that both algorithms recovered realistic radial motion patterns over the cardiac cycle as the image gradients of the cMRI are sufficient to estimate the thickening of the heart.

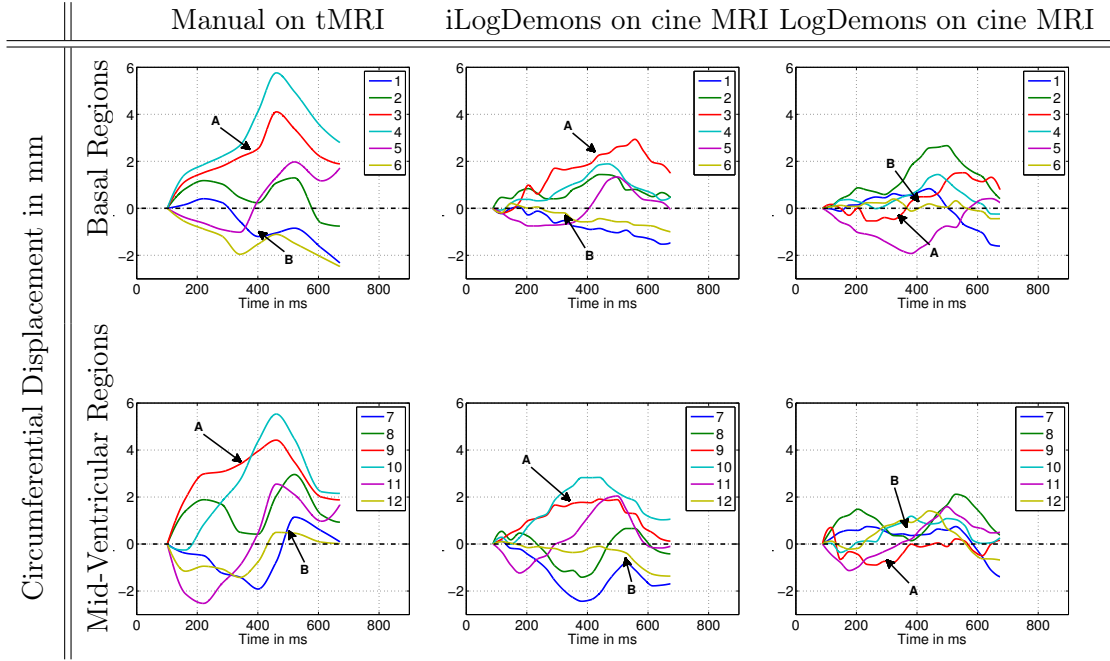


Figure 5.12: Basal and mid-ventricular circumferential displacements of patient 2 (in mm, averaged per zone). Limited by the lack of consistent texture information, the circumferential displacements of the heart estimated on cMRI using both logDemons (*right panel*) and iLogDemons (*mid panel*) were globally under-estimated when compared to manual tracking on tMRI (*left panel*). Yet, iLogDemons motions presented more realistic patterns, as highlighted by the directions of the displacements of zones 3-9 (A) and zones 1-7 (B) for instance (positive displacement: counter-clockwise; negative displacement: clockwise).

5.3.2.2 Comparison with Echocardiography 2D-Strain

The heart was tracked on the cMRI using iLogDemons. Registration parameters were kept identical to the previous experiments: $\sigma_x = 1$, $\sigma^2 = 2$, $\sigma_f^2 = 0.5$, and $\kappa = 1$. A 2-level multi-resolution scheme was also used and the registration was automatically stopped as soon as RMSE stopped decreasing. For fair comparisons, we transformed the estimated 3D displacements into 2D apparent displacements by warping the short axis displacements (XY-plane) along the through-plane motion (Z-axis).

Despite of the different nature of the measurements, results provided by the iLogDemons were fairly similar to the measured 2D-strains. As one can see in Figure 5.13, estimated regional circumferential strains exhibited variation patterns similar to the 2D strain measurements. Nonetheless, a more thorough validation would be necessary to quantify this result. Furthermore, care should be taken when comparing with 2D-strain as this measurement is only partial (only the 2D deformations are measured) and prone to rater variability [Chung et al., 2008].





Figure 5.13: Circumferential strain throughout the cardiac cycle. *Top panel*: 2D-strain measurement. *Bottom panel*: cMRI estimation with iLogDemons. Colours between the two images correspond to the same myocardium regions. iLogDemons recovered plausible strain patterns from cMRI. In particular, it correctly identified the abnormal motion of the anterior zone (*in cyan*).

5.3.3 Computation Time

For the first adult patient, the frame-by-frame registration took on average 130 *s* with logDemons (4.8 *s* per iterations), 300 *s* with iLogDemons (10.9 *s* per iterations) and 1310 *s* with i²LogDemons (48.6 *s* per iterations) on a MacPro 2 × 3.2GHz Quad-Core Intel Xeon, 16GB of RAM but mono-core execution. The incompressibility was ensured on about 97000 voxels ($\approx 4\%$ of the image size $171 \times 61 \times 83$). iLogDemons was about 2 times slower than logDemons but with still reasonable computational time although the algorithm was not at all optimised. i²LogDemons was too computationally intensive while yielding no significant improvements due to the repeated building, preconditioning and resolution of the linear system (the results are not reported here for the sake of clarity). In iLogDemons, the linear system was built and preconditioned only once, at the beginning of the registration. Similar computation times were obtained for the other two patients. In the current implementation, iLogDemons already appears to be faster than other methods [Tanner et al., 2002; Haber and Modersitzki, 2004; Saddi et al., 2007], although direct comparison is not evident due to the different experimental set-up. Fourier-based methods [Hinkle et al., 2009] or optimised multi-grid system resolution [Saddi et al., 2007] will most probably further improve our method.

5.4 Discussion

This chapter demonstrated the effectiveness of iLogDemons algorithm to recover cardiac motion from short-axis gated cMRI. Of course, we do not claim that tracking the heart on standard clinical images using iLogDemons should replace cardiac motion estimation from tMRI or more recent imaging modalities, which, must be preferred whenever they are available. We only see iLogDemons as a complementary tool that could estimate the myocardium strain when only anatomical images are at the clinician disposal, for retrospective studies or in paediatrics for instance.

During our experiments, we observed that enforcing incompressible elasticity significantly enhanced the realism and the accuracy of the estimated displacements and strains in two adults with heart failure and one young patient with repaired ToF. Recovered deformation fields were closer to those computed automatically or manually from tMRI and 2D-strain. The physiological constraints reoriented the displacement vectors within the myocardium, which enabled the recovery of longitudinal and circumferential motions despite the poor myocardium texture information and the large slice thickness. We also observed that iLogDemons slightly improved the estimation of the deformation from tMRI, making the tracking more robust to tag fading.

The possibility of applying the incompressibility constraint in a limited domain of the image has been crucial for this application, as only the myocardium is incompressible. Blood pool volume must vary to ensure correct registration, otherwise locking may happen. Moreover, as the myocardial region is relatively small, little computational overhead is added to the logDemons algorithm.

Despite the promising results, some limitations have still to be tackled for that particular application. In this thesis, we used isotropic elastic regularisation. Because the cardiac muscle is anisotropic, that transformation model may not be ideal. Yet, designing efficient anisotropic smoothing is not straightforward and, in that case, it is even more challenged by the spatial variation of the cardiac anisotropy. We thus decided to use isotropic filters for the sake of efficiency but thanks to the proposed regularisation framework, more advanced anisotropic regularisation schemes could be investigated in the future.

Another limitation is the proposed tracking strategy, which is rather simple. The focus of our work was on the improvement of the frame-by-frame registration algorithm. However, the tracking can also be enhanced to take into account the specificities of the cardiac motion. We could for instance integrate temporal priors like in [Ledesma-Carbayo et al., 2002; Perperidis et al., 2004; Delhay et al., 2007; DeCraene et al., 2009] or use more sophisticated 4D registration techniques [Sundar et al., 2009b]. These improvements will most probably increase the accuracy of the estimated deformations.

From a clinical point of view, an effort of validation must be conducted on both tMRI and cMRI. In that study we relied on manual tracking of the tMRI but this task is tedious, prone to rater variability or even impossible for 3D tMRI. Comparison with other well-established tools would help in evaluating our method [Young et al., 1995; Chandrashekara et al., 2004; Rueckert et al., 2006b; Arts et al., 2010]. We are currently acquiring and working on larger datasets of patients with heart failure. Finally, iLogDemons can be applied on other cardiac imaging modalities, in particular 3D ultrasound and CT images.

Part III

PROGNOSIS: IDENTIFICATION AND MODELLING OF PATHOLOGICAL HEART GROWTH

Statistical Modelling of the Heart Shape and its Growth

Contents

6.1 Motivation	108
6.2 Unbiased Template of Shapes Using Currents	109
6.2.1 Background and State of the Art	109
6.2.2 Data Preparation	112
6.2.3 Unbiased and Non-Parametric Statistical Models on Surfaces	112
6.2.4 Results on Pulmonary Arteries	116
6.2.5 Towards Statistical Analyses of Shapes	117
6.3 Identification of Pathological Heart Shape Patterns	118
6.3.1 Background and State of the Art	118
6.3.2 Model Reduction on Principal Components	119
6.3.3 Results on Pulmonary Arteries	121
6.3.4 Exploring Factors between Shape and Clinical Features	122
6.4 Estimation of a Generative Model of Heart Growth	125
6.4.1 Background and State of the Art	125
6.4.2 Model Reduction on Partial Least Squares Components	127
6.4.3 A Generative Model of Heart Growth	129

Clinical prognosis of patients with chronic disease is often required to plan therapies. Understanding how a pathology will alter the heart is thus of primary importance. For instance, predicting the heart shape in repaired tetralogy of Fallot (ToF) is crucial for planning valve replacement.

However, predicting the remodelling of the heart is tremendously challenging, even in healthy subjects. In that context, we tackled that question from a statistical point of view. Inspired from strategies seen in neuroscience and functional brain imaging, the idea is to detect in a population of patients the patterns of the heart shape that are relevant to the disease under study. Cross-sectional analyses are then applied to model the observed heart remodelling. This chapter describes the methodology of the approach, based on group-wise statistical analyses. Results specific to ToF patients are detailed in Chapter 7.

6.1 Motivation

In recent years, increasing evidence has been reported on the tight relationship between heart shape and function [Narula et al., 2008]. The heart naturally remodels itself during growth: It becomes larger and thicker so as to adapt its pump efficacy to the growing body. Similarly, pathologies can seriously affect the heart. Patients with heart failures usually have dilated, thin and akinetic myocardium [Remme and Swedberg, 2001]. Congenital heart diseases and corrective surgery can have tremendous long-term sequels on the heart [Oosterhof et al., 2006]. Analysing how the heart shape is altered by pathologies and how it evolves over time can provide precious insights into pathological mechanisms and their time course, eventually resulting in quantitative metrics that can be used for disease diagnostic and therapy planning. However, the extreme variability in heart shape and the 3D nature of the problem make this task dauntingly complex. So far, only global indicators, like volume, cardiac output, etc. have been used to quantify pathological remodelling [Therrien et al., 2005; Oosterhof, 2006]. Furthermore, the large variety of biological phenomena involved in cardiac remodelling makes prohibitively difficult the construction of universal direct models of myocardium growth. To date, proposed models focus on specific remodelling phenomena only, like cardiac fibre realignment or myocardial thickening for instance [Rodriguez et al., 1994; Geerts-Ossevoort, 2002; Kroon et al., 2007, 2009].

Recent advances in image processing now enable computational analyses of shape in large populations of patients. Tools have been developed to identify abnormal patterns in patients compared to controls [Ashburner et al., 1998] and to study organ remodelling over time [Thompson et al., 2000], with particular emphasis on brain structures. These studies are now attainable thanks to consistent mathematical frameworks and complex registration algorithms. It is now possible to attack the problem of 3D heart remodelling from a statistical point of view. This strategy is commonly known as computational anatomy [Grenander and Miller, 1998; Pennec, 2008].

Two categories of approaches for shape analysis can be distinguished: *parametric* methods and *non-parametric* methods. Although both categories share the same philosophy, i.e. analysing the principal components of shapes descriptors; they are fundamentally different in the way shapes are computationally represented.

Parametric methods consist in studying the variability of anatomical landmark positions among a population [Bookstein, 1986; Kendall, 1989]. If consistent landmarks can be identified in a population of shapes, the variation of their spatial position indicates the variation of the shapes. Performing principal component analysis (PCA) on the landmark positions provides the observed shape variability. This model is often called point distribution model (PDM). [Cootes et al., 1995] for instance relied on PDM to guide a deformable model for image segmentation using shape priors. Although very powerful, PDM suffer from a major limitation: they require accurate and consistent landmarks. Solutions to that issue have been proposed, like the Minimum Description Length strategy [Davies et al., 2001] or

Spherical Harmonic parameterisation [Styner et al., 2004]. However, the uncertainty on the landmark positioning may include bias in the subsequent statistical analyses.

Non-parametric approaches study how a representative template, or atlas, of the population deforms within this population. The idea is not to look at the shapes *per se* but at the transformations that deform the template to the observations [Thompson, 1917; Grenander and Miller, 1998; Rueckert et al., 2001]. The main advantage of these methods is that no mesh correspondences are needed, thanks to well-posed correspondence-free registration algorithms. However, these approaches rely on an ideal template that best represents the population of study, which can be difficult to estimate [Guimond et al., 2000; Joshi et al., 2004; Allasonniere et al., 2007; Durrleman et al., 2008].

As a first step towards the personalised characterisation of 3D pathological patterns of the heart shape and their evolution in patients, we propose to carry out group-wise analyses relating shape to clinical features. The overall aim is two-fold. First, we exhibit shape patterns relevant to the severity of a pathology by relating the shape descriptors to clinical indices. Second, we propose to build a generative atlas of heart remodelling through a cross-sectional regression analysis of heart shapes.

Figure 6.1 outlines the main steps of the proposed approach. We first estimate an unbiased template of the observed heart shapes and the related deformations (Section 6.2). We then perform two independent analyses. On the one hand, we relate the observed shapes to clinical indices that quantify the severity of a pathology (Section 6.3). To that end, we compute shape variability using principal component analysis (PCA) and project the observations onto the PCA subspace to reduce model dimensionality and increase statistical significance. The resulting low-dimensional shape representation is related to the clinical indices using standard univariate and multivariate designs. On the other hand, we estimate a generative model of heart growth from the observations (Section 6.4). To that end, we project the observations onto a subspace estimated by partial least squares (PLS) between deformations and a clinical feature that represents patient growth, like age or body surface area. PLS is more suitable for that purpose than PCA as it automatically extracts modes that are simultaneously relevant to shape and the clinical feature. The generative model of growth is then obtained using canonical correlation analysis (CCA). The following sections describe the methods of the analyses, illustrated on a small dataset of pulmonary arteries. Chapter 7 reports the results specific to tetralogy of Fallot.

6.2 Unbiased Template of Shapes Using Currents

6.2.1 Background and State of the Art

The first step of the analysis consists in estimating an ideal template of the heart shape from a cohort of subjects. Let a population of N patients whose hearts have been segmented from medical images (using the pipeline described in Chapter 3 for

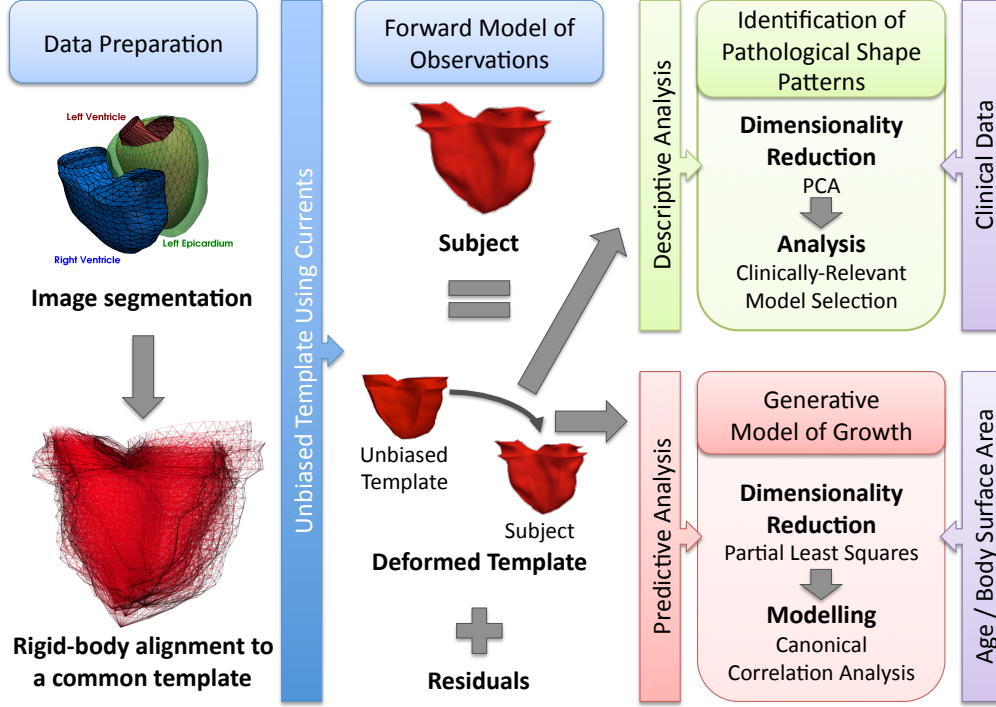


Figure 6.1: Framework for statistical analyses and modelling of the heart shape based on currents, exploratory statistics and regression models.

instance). The cardiac anatomy of a patient is represented by a surface S^i , the superscript i stands for the patient index (Figure 6.2). We assume that the surfaces have been rigidly aligned to a common space. In the following, we call *shapes*, or *observations*, the computational representations of the surfaces S^i of the patients and denote them by \mathcal{T}^i . They belong to the *space of observations*. Note that the \mathcal{T}^i 's are usually different from the theoretical surfaces S^i . In the large majority of applications, the \mathcal{T}^i 's are triangulated surface meshes that approximate the S^i 's. We call *template*, or *atlas*, the reference shape from which the deformations are computed. That template must be as unbiased as possible from the population used to estimate it. The atlas belongs to the ideal *space of shapes*.

In non-parametric shape representations, the observations can be modelled in two different ways, resulting in two different constructions of the ideal template. The *backward* approach models the template as the average of the deformed observations plus some residuals that account for variabilities in the space of shapes [Guimond et al., 2000; Joshi et al., 2004]. Let \mathcal{T}^i be the observed heart shape of the patient i , $\bar{\mathcal{T}}$ the ideal template, ε^i the residuals, ϕ^i the transformation that maps \mathcal{T}^i to $\bar{\mathcal{T}}$ and $*$ the action of the transformation on the shape. The backward model writes:

$$\bar{\mathcal{T}} = \phi^i * \mathcal{T}^i + \varepsilon^i$$




Figure 6.2: Triangulated surfaces of the cardiac anatomy of a patient with repaired tetralogy of Fallot.

Such a template is efficiently computed from the observations using iterative registration/averaging techniques [Guimond et al., 2000]. However, the parameters of the observations cannot be easily estimated from the data. In particular, how the residuals are mapped back to the space of observations is not evident. Moreover, separating shape variations and noise due to data artefacts is not obvious with this framework. These uncertainties make this a challenging approach for the statistical analyses of shapes.

The *forward* approach reverses the point of view by representing the observations as deformations of the template plus some residuals [Allasonniere et al., 2007; Durrleman et al., 2008], which writes

$$\mathcal{T}^i = \phi_*^i \bar{\mathcal{T}} + \varepsilon^i$$

In this framework, the parameters are easier to identify. The transformation ϕ^i maps the ideal template to the observation, it encodes most of the shape information that is represented by the template. The residuals are defined in the space of the observations: they correspond to the shape features that are not captured by the template and the deformations, such as topology changes, shape outliers due to image artefacts, etc. Model parameters can therefore be faithfully estimated from images and clinical data. Nevertheless, this framework requires more complex mathematical tools to represent the shapes, deformations and residuals in a consistent way.

In our analyses, we rely on the forward approach as we aim to statistically relate shapes to clinical features. More precisely we use the method proposed by [Durrleman et al., 2009a], which is particularly suitable for our purposes as i) it relies on currents, an elegant non-parametric representation of shapes that defines vector spaces; ii) the template and the deformations are computed simultaneously and consistently; and iii) the parameters are well-defined, the amount of shape information to analyse is fully controlled.

6.2.2 Data Preparation

The very first step of the analysis consists in extracting 3D surfaces of the heart from medical images. As the method is non-parametric, no point correspondences are necessary, any image segmentation algorithm can be used to segment the boundaries of the heart. Once the surfaces have been delineated, we align them in a common space to remove positioning effects in the subsequent statistical analyses. If point correspondences are available, standard least square rigid registration can be employed [Mosier, 1939; Arun et al., 1987]. Otherwise, non-parametric rigid registration can be used [Jian and Vemuri, 2005; Hufnagel et al., 2007]. In the following, we assume the meshes are already aligned.

6.2.3 Unbiased and Non-Parametric Statistical Models on Surfaces

This section briefly outlines the forward model of shapes for the sake of completeness. The reader is referred to [Durrleman et al., 2009a; Durrleman, 2010] for further details.

6.2.3.1 Non-Parametric Representation of Shapes with Currents

Currents are used to represent the surfaces \mathcal{S}^i and the residuals ε^i . Before going into the mathematical details, let's consider a simple analogy illustrated in Figure 6.3. In the recent years, 3D scanners have been developed to digitalise material objects. These machines acquire the geometry of an object by probing its surface with laser beams. The diffraction of the beams on the surface is captured by cameras and the resulting signal is used to reconstruct the object on the computer. Likewise, currents characterise shapes by probing them using varying vector fields $\omega \in W$ (the scanner laser beams). The shape is characterised by how it integrates these vector fields (the scanner camera).




Figure 6.3: 3D scanners digitalise objects by probing their shape using varying laser beams (*left panel*). Currents work in the same way. They characterise a shape by probing it using varying vector fields (*right panel*)

Mathematically, a current is a continuous linear mapping $\mathcal{L}_W(\omega)$ from a vector space W to \mathbb{R} , i.e. it is an application that integrates vector fields. The current of a surface \mathcal{S} is the flux of a test vector field $\omega \in W$ across that surface. The

shape \mathcal{T} of the surface S is uniquely characterised by the variations of the flux as the test vector field varies. The core element that makes this framework possible is to choose as vector space of the test vector fields W a vector space generated by a kernel, typically the Gaussian kernel $K_W(\mathbf{x}, \mathbf{y}) = \exp(-\|\mathbf{x} - \mathbf{y}\|^2 / \lambda_W^2)$ (W is a reproducible kernel Hilbert space (r.k.h.s.)). W is the dense span of basis vector fields of the form $\omega(\mathbf{x}) = K_W(\mathbf{x}_a, \mathbf{x}) \mathbf{a}$, where the vectors \mathbf{a} are given and fixed at the spatial positions \mathbf{x}_a . The kernel K_W defines an inner product in W that is easily computed by $\langle \omega(\cdot), \nu(\cdot) \rangle_W = \langle K_W(\mathbf{x}_a, \cdot) \mathbf{a}, K_W(\mathbf{x}_b, \cdot) \mathbf{b} \rangle_W = \mathbf{a}^T K_W(\mathbf{x}_a, \mathbf{x}_b) \mathbf{b}$, where $\omega(\cdot) = K_W(\mathbf{x}_a, \cdot) \mathbf{a}$ and $\nu(\cdot) = K_W(\mathbf{x}_b, \cdot) \mathbf{b}$ are two vector fields of W . A consequence of these properties is that the space of currents W^* , which is the dual of W , is the dense span of the dual representations of the basis vectors $\omega(\cdot)$, called Dirac delta currents $\delta_{\mathbf{x}_a}^{\mathbf{a}}(\omega)$ and defined by:

$$\delta_{\mathbf{x}_a}^{\mathbf{a}}(\omega) = \langle K_W(\mathbf{x}_a, \cdot) \mathbf{a}, \omega(\cdot) \rangle_W = \mathbf{a}^T \omega(\mathbf{x}_a) \quad (6.1)$$

Intuitively, a Dirac delta current is an infinitesimal vector \mathbf{a} that is concentrated at the spatial position \mathbf{x}_a . In that way, the current \mathcal{T} of a surface S can be decomposed into an infinite set of Dirac delta currents defined at each point of the surface and oriented along the surface normal. Computationally, the surfaces are represented by discrete triangulated meshes: Their current representation is the finite sum

$$\mathcal{T}(\omega) = \sum_k \delta_{\mathbf{x}_k}^{\mathbf{a}_k}(\omega) \quad (6.2)$$

where \mathbf{x}_k are the barycentres of the mesh faces and \mathbf{a}_k their normal (Figure 6.4, left panel). The vector field ω dual of the current $\mathcal{T}(\omega)$ is the spatial convolution of every normal vector \mathbf{a}_k with the kernel K_W , $\omega(\mathbf{x}) = \sum_k K_W(\mathbf{x}, \mathbf{x}_k) \mathbf{a}_k$. In [Durrleman et al., 2009a], the authors propose an efficient greedy algorithm to approximate the current representation of a surface by a minimal yet optimal set of Dirac delta currents (Figure 6.4, right panel). Computational complexity is further decreased by using FFT-based kernel convolution techniques.

As currents are linear applications, they define a vector space on shapes denoted W^* . The sum of two currents is the union of their Dirac delta currents, i.e. the union of the two surfaces. Likewise, scaling a current amounts to scaling the amplitude of the Dirac delta currents. In addition to that, the space of currents W^* is equipped by construction with the inner-product $\langle \delta_{\mathbf{x}_a}^{\mathbf{a}}(\omega), \delta_{\mathbf{x}_b}^{\mathbf{b}}(\omega) \rangle_{W^*} = \langle K_W(\mathbf{x}_a, \cdot) \mathbf{a}, K_W(\mathbf{x}_b, \cdot) \mathbf{b} \rangle_W = \mathbf{a}^T K_W(\mathbf{x}_a, \mathbf{x}_b) \mathbf{b}$. The distance between two shapes can therefore be computed as the W -norm of the difference of their currents. The space of currents thus enables one to compute mean, standard deviations and other descriptive statistics on shapes. Finally, discrete Gaussian variables in the space of currents can be defined to model the residuals ε^i . In practice, we place a random vector \mathbf{a}_p that follows a normal distribution at each point \mathbf{x}_p of a 3D grid that embeds the shape \mathcal{T} .

We conclude this section with an important remark about the width λ_W of the kernel K_W that controls the resolution of the current representation. Basically, the larger λ_W , the coarser the resolution and the less “accurate” the representation. This

Original Shape (1476 delta currents) Compressed Shape (281 delta currents)




Figure 6.4: The Dirac delta currents of a triangulated mesh are the normal vectors of every face, centred at the face barycentres. A greedy algorithm reduces the amount of delta currents needed to represent the shape while preserving the accuracy of the representation.




Figure 6.5: The size of the kernel K_W enables one to choose the level of details of the current representation. Large λ_W discards subtle shape features (A) that would be recovered using smaller λ_W (B) (*Images courtesy of S. Durrleman*)

parameter thus sets the level of shape details we want to study. As it is illustrated in Figure 6.5; small λ_W values enable one to capture little differences between surfaces, whereas large λ_W discard them.

6.2.3.2 Surface Registration Using LDDMM on Currents

The smooth diffeomorphic transformation ϕ^i that registers the template \bar{T} to the observations \mathcal{T}^i is estimated using the Large Deformation Diffeomorphic Metric Mappings (LDDMM) [Dupuis et al., 1998; Miller et al., 2002a; Beg et al., 2005] on currents [Vaillant and Glaunes, 2005; Glaunes, 2005]. Diffeomorphic transformations are crucial to ensure one-to-one mapping between the two shapes. No material loss nor topology changes are allowed, thus guaranteeing the consistency of the analysis.

Similar to the iLogDemons (Chapter 4), ϕ^i is calculated by integrating the Lagrangian transport equation $\partial\phi^i(\mathbf{x}, t)/\partial t = \mathbf{v}^i(\phi^i(\mathbf{x}, t), t)$, $\phi^i(\mathbf{x}, t = 0) = \mathbf{x}$. However, the velocity field \mathbf{v}^i are now time-dependent, yet completely determined by the initial velocity field $\mathbf{v}^i(\mathbf{x}, t = 0)$, denoted $\mathbf{v}_0^i(\mathbf{x})$ to simplify the notations. The initial velocity $\mathbf{v}_0^i(\mathbf{x})$ belongs to a r.k.h.s. V generated by the Gaussian kernel $K_V(\mathbf{x}, \mathbf{y}) = \exp(-\|\mathbf{x} - \mathbf{y}\|^2/\lambda_V^2)$, which defines an inner product of V . The \mathbf{v}_0^i 's are therefore defined as the convolution of sets of *moment* vectors β^i by the kernel K_V . Intuitively, the moment vectors contain the initial kinetic energy that is necessary

to cover the geodesic path ϕ^i . For discrete surfaces, the moments are defined at the locations \mathbf{x}_k of the Dirac delta currents $\delta_{\mathbf{x}_k}^{\mathbf{a}_k}$ and the initial velocity field satisfies $\mathbf{v}_0^i(\mathbf{x}) = \sum_k K_V(\mathbf{x}_k, \mathbf{x}) \beta_k^i$. The inner product between two velocities \mathbf{v}_0^i and \mathbf{v}_0^j thus writes:

$$\langle \mathbf{v}_0^i, \mathbf{v}_0^j \rangle_V = \sum_{k,l} \beta_k^{iT} K_V(\mathbf{x}_k, \mathbf{x}_l) \beta_l^j$$

The transformation ϕ^i is estimated by minimising the registration energy defined on the current:

$$\mathcal{E}(\mathbf{v}) = \|\phi^i(t=1)_* \bar{\mathcal{T}} - \mathcal{T}^i\|_{W^*}^2 + \nu \int_0^1 \|\mathbf{v}(t)\|_V dt \quad (6.3)$$

In this equation, ν is a weight parameter that controls the strength of the regularisation and $\phi^i(t=1)_* \bar{\mathcal{T}}$ is the action of the diffeomorphism at time $t=1$ on the template $\bar{\mathcal{T}}$. How this energy is minimised is detailed in [Glaunes, 2005; Vialant and Glaunes, 2005]. What is important to consider for our application is the effect of the kernel K_V on the estimated deformations. The matching criterion is regularised by minimising the length along the geodesic diffeomorphism, which is computed by integrating the V -norm of the velocity field $\mathbf{v}(t)$ over time. Hence, the kernel K_V controls the smoothness of the velocity \mathbf{v}_0 , and indirectly of the transformation ϕ . Intuitively, λ_V^2 controls the size of the spatial region that is deformed consistently. When λ_V^2 is large, wide spatial regions are deformed in a coherent way, and conversely. One can thus control the scale of shape deformation to analyse. For studying global shape differences, large λ_V^2 are suggested, and conversely.

6.2.3.3 Unbiased Template of Shapes


The template $\bar{\mathcal{T}}$ and the deformations ϕ^i are estimated jointly and consistently by means of an alternate two-step strategy. The idea is to consider the forward model as a Bayesian problem where the shapes \mathcal{T}^i are the observations, the template $\bar{\mathcal{T}}$ is the unknown, the transformations ϕ^i that register $\bar{\mathcal{T}}$ to the observations are hidden variables and the residuals ε^i are random Gaussian variables represented by currents. To compute $\bar{\mathcal{T}}$ and the ϕ^i 's simultaneously, we solve the approximated minimisation problem:

$$(\bar{\mathcal{T}}, \phi^i) = \operatorname{argmin}_{\bar{\mathcal{T}}, \mathbf{v}_0^i} \left\{ \|\phi^i(t=1)_* \bar{\mathcal{T}} - \mathcal{T}^i\|_{W^*}^2 + \tau \sum_{i=1}^N \|\mathbf{v}_0^i\|_V^2 \right\} \quad (6.4)$$

where τ weights the geodesic regularisation and \mathbf{v}_0^i is the initial velocity field that parameterises the transformation ϕ^i . This equation is similar to the registration energy (6.3), except that here the template $\bar{\mathcal{T}}$ is also unknown. The first term is the distance between two currents, the transformed template $\phi_*^i \bar{\mathcal{T}}$ and the target subject \mathcal{T}^i . The second term is the geodesic regularisation, applied to all the transformations to estimate.

[Durrleman et al., 2009a] proposed an alternate minimisation of that energy (Figure 6.6). We first assume the template $\bar{\mathcal{T}}$ is known. Minimising Equation 6.4

with respect to the ϕ^i 's amount to registering the present estimate of the template \bar{T} to each patient independently. Once the ϕ^i 's are calculated, Equation 6.4 is minimised with respect to the template \bar{T} . This is achieved using a gradient descent method, see [Durrleman et al., 2009a] for details. Intuitively, this second stage updates the template from the transformations ϕ^i to reduce the overall registration error. The new template minimises the registration errors for all the patient at the same time, i.e. it is more centred with respect to the population. The algorithm is initialised with the mean current of all the observations and it is iterated until convergence. In all our experiments, we fixed $\tau = 10^{-3}$. Only the more intuitive kernel sizes λ_V^2 and λ_W^2 were adjusted.

A- Estimate the Transformations ϕ^i 's

B- Centre the Template




Figure 6.6: Joint estimation of template and transformations. Given a template, the transformations that map it to the shapes are first estimated (A). The template is then centred to minimise the overall registration error (B).

A parallel implementation of the algorithm has been developed to process large amounts of subjects on clusters of computers (Algorithm 6). A scheduler script controls the execution of the algorithm. It dispatches the template-to-patient registrations to all the available computers as these steps are independent from each other (Figure 6.6, Step A). It then waits for the computers to perform the registrations, after which it centres the template (Figure 6.6, Step B) and loop until convergence. As a result, the computation time required to estimate the template minimally depends on the number of subjects.

6.2.4 Results on Pulmonary Arteries

As an illustration, we computed the average template of the pulmonary artery in five patients with repaired tetralogy of Fallot. In these patients, the pulmonary artery is severely altered due to chronic pulmonary regurgitations subsequent to the initial cardiac repair (see Section 2.3, page 28). Of course, no clinical conclusions can be drawn from this experiment due to the small size of the population. We just aim to illustrate the methodology.

Algorithm 6 Parallel Template and Deformations Estimation**Require:** N linearly registered currents \mathcal{T}^i .

- 1: Compute the initial template $\bar{\mathcal{T}}^0$, the average of the N currents.
- 2: **loop** {over n until convergence, $\bar{\mathcal{T}}^n = \bar{\mathcal{T}}^{n-1}$ }
- 3: **Dispatch:** estimate the transformation ϕ^{i^n} that registers $\bar{\mathcal{T}}^{n-1}$ to $\mathcal{T}^i \forall i$.
- 4: **Wait:** all the registrations are done
- 5: **Update:** compute the new template $\bar{\mathcal{T}}^n$ from the deformations ϕ^{i^n} and the previous template $\bar{\mathcal{T}}^{n-1}$
- 6: **return** Final template $\bar{\mathcal{T}}^n$ and deformations ϕ^{i^n} .

For the five selected patients, we segmented the pulmonary artery from MR angiography through image thresholding and manual corrections. The artery was segmented from the pulmonary valve annulus, which connects the artery to the right ventricle, to about 1 – 2 cm after the pulmonary branches, which go towards the lungs. Figure 6.7 illustrates the segmentations. From the segmented meshes, we computed the unbiased template, illustrated in Figure 6.8




Figure 6.7: Pulmonary arteries of five patients suffering from repaired tetralogy of Fallot. One can see the different patterns of dilation, with different degrees of bulging of the pulmonary trunk.

6.2.5 Towards Statistical Analyses of Shapes

Equipped with a consistent framework for representing the shapes of a population of patients, we can now relate the observed shape features to pathology indices. As we have seen in the previous sections, the shape information is captured by either the deformations ϕ^i 's or the residuals ε^i 's. The amount of information present in each component is controlled by the user through the parameters λ_V and λ_W . The ϕ_i 's encode the features that are represented by the template $\bar{\mathcal{T}}$ whereas the residuals ε^i capture all the remaining characteristics. The question of which information to




Figure 6.8: Template of the pulmonary artery estimated from a population of five patients with repaired ToF. Observations are represented in wireframe.

use is therefore crucial. In our application, we analyse the deformations ϕ^i only as we mainly focus on the regional changes of the heart anatomy due to the pathology. Besides, heart segmentation is prone to local inconsistencies due to image artefacts, which may challenge the analysis of the residuals.

The major difficulty of statistical analyses of shapes is the very large dimensionality of the data to analyse. The deformations ϕ^i 's are studied by considering the initial velocities \mathbf{v}_0^i 's, which can be parameterised by thousands of moments β^i (one at each point of the current representation). A large number of patients would thus be necessary to get statistical significance. In addition to that, neighbouring moments are highly correlated by construction, hence resulting in collinear shape descriptors that may violate assumptions of the most usual statistical tests. To tackle these issues, we propose to project the initial velocities \mathbf{v}_0^i on optimal subspaces. The idea is to reduce the number of variables that are necessary to describe an observation. Two subspace techniques are investigated. To explore shape patterns relevant to pathologies, we rely on the traditional Principal Component Analysis (PCA), which extracts components that best explain the shape variance of the population. To predict the shape from an external variable, we instead use partial least squares (PLS), which estimates the optimal subspace that best explains both the observed shape variance and its covariance with the external variable.

6.3 Identification of Pathological Heart Shape Patterns

6.3.1 Background and State of the Art

The first aim of the analysis is to identify shape patterns that relates to pathologies in a population of patients. Numerous studies investigated the shape variability observed in populations. The large majority of them compare groups of patients with controls to identify differences due to the pathology (see [Cates et al., 2008])

and references therein for instance).

In this thesis we explore an alternative idea, which is to correlate the heart shape variability with clinical parameters that quantify the severity of the pathology. The hypothesis is that such correlations may reveal quantitative 3D patterns of the pathological morphology. Few related strategies can be found in the literature. In [Querol et al., 2006], the authors use PCA to relate the main modes of variation of bone shapes with biomechanical properties simulated using finite element methods. This approach enabled them to identify variations of bone stress due to changes in bone anatomy. In [Boisvert et al., 2008], the authors performed PCA on shape descriptors of scoliotic spines to classify patients with different pathological spine morphology. Resulting shape modes were consistent with the clinical classification of reference. More recently, Blanc *et al.* proposed to mix shape descriptors and explicit anatomical constraints representative of bones morphology to get anatomically consistent shape modes [Blanc et al., 2009]. The obtained statistical model was able to generate realistic bone shapes by discarding variations due to noise for instance.

In this work we apply standard statistical design to select the modes of variation that are relevant to the studied pathology. Figure 6.9 presents the different steps of the method. We first compute deformation modes using PCA. The modes represent the shape variations observed in our population. We then project each patient to a subspace of modes to reduce the dimensionality of the problem. Each patient is represented by a shape vector constituted of few elements. We finally apply standard statistical analyses between shape vectors and clinical indices to explore shape patterns that may relate to the pathology.

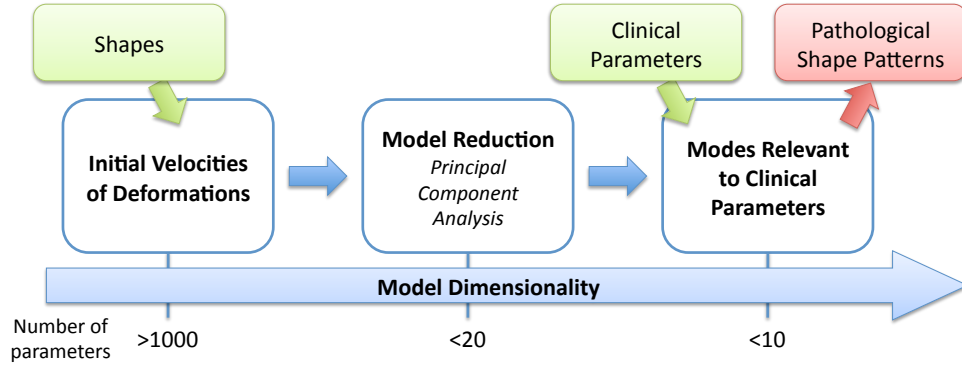


Figure 6.9: Shape analysis pipeline. Pathological shape patterns are identified by relating PCA deformation modes to clinical parameters. Model dimensionality is reduced using PCA to increase statistical significance.

6.3.2 Model Reduction on Principal Components

Principal Component Analysis (PCA) is applied on the initial velocity fields \mathbf{v}_0^i to extract the main deformation modes observed in the population. PCA finds basis

vectors, the *modes*, of the space of variables (here the initial velocities) that best explain their variance. The modes \mathbf{p} are calculated by solving the eigenvalue problem $\Sigma \mathbf{p} = \mu \mathbf{p}$, where the elements $\sigma_{i,j}$ of the covariance matrix Σ are calculated in the kernel space V by:

$$\sigma_{ij} = \langle \mathbf{v}_0^i - \bar{\mathbf{v}}_0, \mathbf{v}_0^j - \bar{\mathbf{v}}_0 \rangle_V = \sum_{\mathbf{x}_k, \mathbf{x}_l} (\beta_{\mathbf{x}_k}^i - \bar{\beta}_{\mathbf{x}_k}) K_S(\mathbf{x}_k, \mathbf{x}_l) (\beta_{\mathbf{x}_l}^j - \bar{\beta}_{\mathbf{x}_l}) \quad (6.5)$$

$\bar{\mathbf{v}}_0$ is the mean initial velocity field computed over the population and the superscripts i and j denote two different subjects. The principal components are obtained by computing the spectral decomposition $\Sigma = P M P^T$. M is the diagonal matrix of the eigenvalues σ^m , or variances, sorted in decreasing order and P is the orthonormal matrix (in the L_2 -norm sense) of the eigenvectors \mathbf{p}^m . The m^{th} loading \mathbf{l}^m of the PCA decomposition is given by the formula:

$$\mathbf{l}^m = \sum_i \mathbf{p}^m[i] (\beta^i - \bar{\beta}) \quad (6.6)$$

In this equation, $\mathbf{p}^m[i]$ is the i^{th} element of the m^{th} eigenvector of Σ , $\beta^i = (\beta_{\mathbf{x}_1}^i \dots \beta_{\mathbf{x}_r}^i)^T$ is the $r \times 3$ matrix that gathers the r moment vectors of patient i , and $\bar{\beta}$ is the average moment. By construction, the loadings \mathbf{l}^m belong to the space of moments. We observe that the V -norm of the initial velocity field parameterised by \mathbf{l}^m , denoted $\mathbf{v}_{0|m}$, equals the m^{th} variance σ^m :

$$\begin{aligned} \|\mathbf{v}_{0|m}\|_V^2 &= \mathbf{l}^{mT} K_S \mathbf{l}^m = \left(\sum_i \mathbf{p}^m[i] (\beta^i - \bar{\beta}) \right) K_S \left(\sum_i \mathbf{p}^m[i] (\beta^i - \bar{\beta}) \right) \\ &= \sum_{i,j} \mathbf{p}^m[i] \mathbf{p}^m[j] \underbrace{(\beta^i - \bar{\beta}) K_S (\beta^j - \bar{\beta})}_{\sigma_{i,j}} \\ &= \mathbf{p}^{mT} \Sigma \mathbf{p}^m = \sigma^m \end{aligned}$$

As a result, the variability captured by the m^{th} deformation mode between $[-k\sigma^m; +k\sigma^m]$ is visualised by deforming the template \bar{T} with the deformations ϕ^{-m} and ϕ^{+m} parameterised by the moments $\bar{\beta} - k\mathbf{l}^m$ and $\bar{\beta} + k\mathbf{l}^m$ respectively.

By construction there is at most $N - 1$ PCA modes, where N is the number of observations. Using all of them fully describe the observations but usually only a subspace is considered by choosing $p < N - 1$ components. A standard approach is to choose the p first modes that explain at least $\pi\%$ of the total variance: $\sum_{m=1}^p \sigma^m$. Each subject is then projected onto the selected PCA subspace and is represented by a unique shape vector $\mathbf{s}^i = \{s^{i,m}\}_{m=1..p}$ whose elements are the scalar product between the initial velocity \mathbf{v}_0^i of the patient and the velocity of the m^{th} mode:

$$\begin{aligned} s^{i,m} &= \langle \mathbf{v}_0^i - \bar{\mathbf{v}}_0, \mathbf{v}_{0|m} \rangle_V \\ &= \sum_{\mathbf{x}_k, \mathbf{x}_l} [\beta_{\mathbf{x}_k}^i - \bar{\beta}] K_S(\mathbf{x}_k, \mathbf{x}_l) \mathbf{l}_{\mathbf{x}_l}^m \end{aligned}$$

The $s^{i,m}$'s quantify the amount of variability along the m^{th} mode present in patient deformation.

In conclusion, we have reduced the amount of data needed to represent the shape of a patient by two or more orders of magnitude. The precision of the representation is controlled by the number of components of the PCA subspace. However, it has to be observed that modes with low variances may still be relevant to external clinical parameters. For instance, a mode that captures a local bulging is probably more related to the pathology than a global scaling of the shape although this bulging is not very visible in the population. Consequently, we will be fairly conservative in the selection of the PCA subspace and select the modes based on their relationship with the clinical parameters of interest and not their variance as described in the following sections.

6.3.3 Results on Pulmonary Arteries

As an illustration, we consider the previous set of five pulmonary arteries of tetralogy of Fallot patients. From the estimated template and deformations, we computed the first three PCA deformation modes and applied them to the template to visualise the shape variations. As one can see from Figure 6.10, the deformation modes captured different degrees of elongation of the pulmonary branches and bulging or bending of the pulmonary trunk. The idea in the following sections is to correlate these modes with clinical parameters in order to identify features relevant to the pathology.




Figure 6.10: Template and first three deformation modes of pulmonary artery computed from five ToF patients. The deformation modes captured different degrees of elongation of the pulmonary branches and bulging or bending of the trunk.

6.3.4 Exploring Factors between Shape and Clinical Features

Deformation modes relevant to the pathology are explored by statistically relating the shape vectors \mathbf{s}^i to clinical parameters. Two categories of parameters can be distinguished. On the one hand, *ordinal* parameters are qualitative or quantitative indices that classify patients by ordered levels. Pulmonary regurgitation for instance can be qualitatively classified as none, trace, mild or severe. Patients can be quantitatively classified by their age. On the other hand, *continuous* parameters are obtained by direct measurements of physical entities like flows, volumes or surfaces. Because of the different nature of the parameters, two approaches are needed.

In the following sections, a statistical nomenclature is adopted. The PCA shape vectors \mathbf{s}^i are gathered into an $N \times p$ matrix of variables, $\mathbf{X} = (\mathbf{s}^1 \dots \mathbf{s}^N)^T$ (Figure 6.11). N is the number of patients in the population and p the number of elements of the shape vectors, i.e. the number of components of the PCA subspace. The columns of \mathbf{X} , which correspond to a specific mode, are called *predictors* and are seen as different variables of a single observation. Similarly, the q clinical parameters under study are gathered into an $N \times q$ matrix \mathbf{Y} , called *response* variables. In this work, only one clinical parameter is studied at a time ($q = 1$) to have full control on the statistical analyses and to maximise significance.




Figure 6.11: Matrices of the clinical variable \mathbf{Y} and the shape descriptors \mathbf{X} .

6.3.4.1 Relating Shape to Ordinal Parameters

Ordinal parameters classify patients in different groups. Deformation modes are selected based on their ability to separate these groups. When the number of observations is not sufficient to convert ordinal parameters into continuous numbers, non-parametric statistical tools based on ranks must be employed. These approaches are robust to non-normal errors and resistant to outliers. The analysis is performed in two steps. First, we test if groups of patients can be distinguished. In the affirmative, post-hoc pair-wise tests are applied to detect which groups actually differ. In our studies we prefer univariate analyses to multivariate designs to maximise statistical power and control the statistical tests, although inter-component effects

may exist.

Rank-based Kruskal-Wallis one-way analysis of variance is used to test if a mode can distinguish groups of patients [Saporta, 2006]. This analysis is the non-parametric variant of the one-way analysis of variance (ANOVA) in what it detects if two groups of patients have different median ranks. Let \mathbf{x}^m be the m^{th} predictor (m^{th} column of \mathbf{X}) and Y be the ordinal clinical parameter. Kruskal-Wallis analysis consists in ordering the elements of \mathbf{x}^m in increasing order, replacing their value by their rank and grouping the observations according to the clinical levels defined in Y . The analysis then tests if median ranks can be distinguished although it does not say which groups differ.

A post-hoc two-way rank-based Wilconxon test is required to detect which groups actually differ, as ANOVA requires post-hoc two-way Student's t-test. Let A and B be two groups of patients identified by Y , and $s_A^{i,m}$ and $s_B^{j,m}$ the shape descriptors along the m^{th} mode of the i^{th} patient of group A and the j^{th} patient of group B . This test compares the groups A and B by analysing the distribution of the sorted pairs $(s_A^{i,m}, s_B^{j,m})$. If the two groups come from a same distribution, the number of pairs that verify $s_A^{i,m} > s_B^{j,m}$ should be pretty close to the number of pairs with $s_A^{i,m} < s_B^{j,m}$. This test is rejected if this equality is not verified. As multiple comparisons are performed (levels are compared pair-wise), the probability of rejecting the hypothesis just by chance increases significantly. Bonferroni correction for multiple tests is thus used, which consists in dividing the chosen level of significance (the maximum p -value) by the total number of tests.

These two steps are complementary since finding an effect between a mode and the investigated clinical parameter does not necessarily mean that this mode is able to differentiate levels. In particular, when the variance of each group is sufficiently high to mask the distance between their mean, no level can be differentiated. It has to be noted that finding no effects does not mean that there is no effect between a mode and the clinical parameter. With more patients for instance, new findings may be obtained. At the end of the process, we come up with a subset of modes that can distinguish different levels of the clinical parameters under study. Visualising these modes may reveal anatomical patterns related to the pathology.

6.3.4.2 Relating Shape to Continuous Parameters

A similar approach can be applied to relate the deformation modes to continuous clinical parameters by replacing the Kruskal-Wallis analysis by ANOVA and the post-hoc Wilconxon test by Student's t-test. However, a regression approach is preferred in order to account for cross-component effects and to detect the direction of correlations. The idea is to perform a multiple linear regression between the shape descriptors \mathbf{s}^i and the clinical parameter Y , which is usually known as Principal Component Regression [Massy, 1965; Rosipal and Krämer, 2006]. Mathematically, the multiple linear regression writes:

$$Y = a_0 + \sum_{l=1}^p a_l \mathbf{x}^m = a_0 + \mathbf{a}^T \mathbf{X} \quad (6.7)$$

In the previous equation, a_0 is the intercept and \mathbf{a} is the column vector of the regression coefficients a_l . The a_l 's reflect the conditional correlations between the predictors and the response. Their sign shows the direction of the correlation, information that we cannot get with univariate analyses. Their value quantifies the amount of information that a single predictor adds to the model when all the remainder predictors are already included. This information is difficult to interpret as it depends on the other modes. Yet, one can assess whether a predictor is significant to the overall linear model by testing the null hypothesis H_0 that $a_l = 0$ using Student's t -test on the ratio between a_l and its standard error [Saporta, 2006]. This information is provided by any statistical software. We can therefore discard the modes that are not significant to the linear model, and thus not relevant to the response variable, through step-wise manual approaches or exhaustive criterion-based automatic techniques.

Step-wise mode selection consists in iteratively removing the less significant predictors from the linear model until model overall significance stops increasing (i.e. the p -value related to the null hypothesis that the estimated linear model is a constant model, stops decreasing). The remaining modes are those that are relevant to the response variable Y . Although this approach showed promising results in previous studies [Mansi et al., 2009c], the selected modes may be unstable with respect to the observations. Different modes may be selected if new patients are added into the database. We shall prefer more robust and automatic modes selection based on criterion evaluation.

Two major criteria are available in the literature: the Bayesian Information Criterion (BIC), also known as Schwartz criterion [Schwartz, 1978] and the Akaike Information Criterion (AIC) [Akaike, 1974]. The idea is to maximise the log-likelihood of the model according to its predictors (seen as parameters of a probabilistic model) while penalising the amount of parameters of the model. Combination of predictors are exhaustively tested and the model which minimises the criterion is chosen.

Let $L(\Theta)$ be the likelihood of the linear model whose parameters are Θ , namely the predictors X . BIC and AIC write:

$$BIC = -2 \ln L(\Theta) + \ln(N)p \quad (6.8)$$

$$AIC = -2 \ln L(\Theta) + 2p \quad (6.9)$$

Apparently very similar, these two criteria often yield different models. When the number of observations is high ($N > 8$), BIC penalisation on the number of parameters is higher than AIC penalisation. This criterion thus favours models with as less parameters as possible. Coming from Bayes theory, BIC tries to find the most probable model given the observations. It can be demonstrated that when the number of observations N is very large, the probability to find the “true” underlying model tends to one. BIC should therefore be used to design models that best explain the data, although the definition of a “true” model is far from being clear. AIC penalisation on the number of parameters is less severe. AIC has been developed from the Kullback-Leibler divergence, which measures how a given distribution differs from a reference distribution. In simple words, AIC estimates how much information is lost

by using the reduced model instead of the theoretical distribution. It thus favours models whose distributions parameterised by the selected predictors are as close as possible to the unknown theoretical distribution. Models with highest likelihood, and therefore with higher predictive power, will be preferred. Nonetheless, AIC may never find the “true” underlying model even for large N . In our experiments, reported in Chapter 7, the modes that are selected by the three strategies (step-wise backward selection, AIC and BIC) are compared.

6.4 Estimation of a Generative Model of Heart Growth

6.4.1 Background and State of the Art

In the previous sections we described tools to identify shape characteristics that are relevant to a pathology. Can we go beyond shape description and predict a heart shape from clinical data? Can we apply statistical methods to derive an atlas of the cardiac remodelling observed in a population?

Ideally, one would use longitudinal data to study organ remodelling, as in [Thompson et al., 2000; Beg et al., 2004; Qiu et al., 2008; Aljabar et al., 2008; Durrleman et al., 2009b]. However, such datasets are challenging to acquire, entailing long-term projects with rigorous protocols. In this thesis, too few longitudinal datasets were available to perform longitudinal analyses. We therefore relied on cross-sectional regression to estimate an average generative model of the heart growth. The underlying idea is to consider each subject as a time-point realisation of a hidden event and to model how these realisations “evolve” over time through regression analyses (Figure 6.12).




Figure 6.12: Cross-sectional regression of shapes. Each subject is associated to a point in time (his age for instance). A regression model is derived from the temporal data.

Cardiac remodelling over time is investigated by regressing the initial velocities \mathbf{v}_0^i that parameterise the deformations ϕ^i with an index that represents patient growth. Patient age is a typical choice but in paediatrics it may not reflect growth precisely, as teenagers can enter into puberty at different ages. For this population,

clinicians prefer using body surface area (BSA), a quantitative and continuous index of body physiology, calculated from subject height and weight, which correlates with age. Among the several formula available in the literature, Dubois formula [DuBois and DuBois, 1915] is a standard choice:

$$BSA(m^2) = 0.007184 \times weight(kg)^{0.425} \times height(cm)^{0.725} \quad (6.10)$$

As in the previous section, the dimensionality of the problem must be reduced to get statistically significant regressions. Furthermore, we need to remove any collinearity of the predictors to get reliable regression models. A standard strategy consists in regressing on principal components (Section 6.3.4). Although this approach yielded promising results on 18 tetralogy of Fallot patients [Mansi et al., 2009c], it may not be optimal as PCA decomposition does not consider the effect of body surface area. In the following, we propose instead to project the observations onto a subspace estimated with Partial Least Squares (PLS) [Wold, 1966; Rosipal and Krämer, 2006]. Given two sets of variables X and Y , PLS computes the components that span the spaces of variables which best explain the variance of X , the variance of Y and the covariance between X and Y . Furthermore, the PLS components are such that the regression $Y = f(X)$ is optimal. PLS is therefore suited for exhibiting deformation modes that best relate to an external parameter. Furthermore, PLS is much more efficient than standard regression when the number of predictors is much higher than the number of observations. Thanks to these advantages, PLS has been largely applied in computational chemistry [Wold et al., 2001] and more recently in functional brain imaging [Worsley, 1997; McIntosh and Lobaugh, 2004; Ji et al., 2009], where it is used to predict activation patterns related to cognitive activity. In medical imaging, [Ablitt et al., 2004] relied on PLS to predict cardiac motion from respiratory signals to improve cardiac image acquisition. More recently, [Rao et al., 2008] used PLS and canonical correlation analysis to predict the shape of a brain structure from neighbouring anatomies in view of automatic brain segmentation. In both approaches PLS showed good prediction power.

Contrary to PCA, PLS is not symmetric as it extracts modes that optimise the regression $Y = f(X)$. In some cases one would prefer canonical correlation analysis (CCA), which is symmetric. However, contrary to CCA, PLS can be applied to datasets that have much less observations than variables, like in our case, as it does not require inverting large matrices. Ideally, we would like to use PLS to have a model that predicts the heart shape from BSA, $shape = f(BSA)$. Yet, performing this regression directly is not possible due to the large number of variables we would need to predict (the deformation parameters) compared to the predictors (the BSA). We thus revert the point of view and apply PLS with the predictive model $BSA = f(shape)$ to select the optimal subspace relevant to both shape and BSA, as illustrated in Figure 6.13. The observations are then projected onto the PLS subspace, resulting in a compact shape representation of the observations. Finally, CCA is applied to estimate the generative model of heart growth.

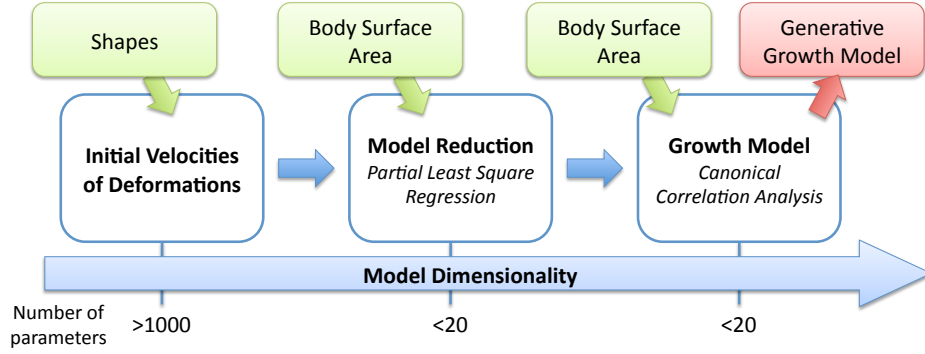


Figure 6.13: Generative model of the heart growth. Partial least squares is used to exhibit deformation modes that maximise shape variance and shape covariance with BSA. Canonical correlation analysis is used to build the generative model on the selected PLS modes.

6.4.2 Model Reduction on Partial Least Squares Components

PLS combines PCA and linear regression to predict response variables Y from a set of predictor variables X . Contrary to principal components regression, PLS automatically finds an orthonormal basis of the space of predictors X that is also relevant to Y . The basis vectors, called *scores*, *components* or *modes*, maximise the variances of both sets of variables and their covariance. In simple words, PLS does PCA, model selection and linear regression at the same time to provide the optimal result [Höskuldsson, 1988]. Mathematically, PLS consists in estimating the normalised weight vectors \mathbf{r} and \mathbf{s} that verify:

$$\max_{|\mathbf{r}|=|\mathbf{s}|=1} \text{cov}(\mathbf{X}\mathbf{r}, \mathbf{Y}\mathbf{s}) = \max_{|\mathbf{r}|=|\mathbf{s}|=1} \text{var}(\mathbf{X}\mathbf{r}) \text{corr}(\mathbf{X}\mathbf{r}, \mathbf{Y}\mathbf{s})^2 \text{var}(\mathbf{Y}\mathbf{s}) \quad (6.11)$$

under the constraint of optimal regression between the predictors X and the response Y . For comparison, PCA solves the optimisation problem $\max_{|\mathbf{r}|=1} \text{var}(\mathbf{X}\mathbf{r})$.

Body Surface Area
(Response)

$$Y = \begin{bmatrix} y^1 \\ \vdots \\ y^i \\ \vdots \\ y^N \end{bmatrix} \begin{matrix} 1 \\ \vdots \\ i \text{ } i^{\text{th}} \text{ patient} \\ \vdots \\ N \end{matrix}$$

Deformation Moments (Predictors)

$$X = \begin{bmatrix} 1 & \dots & \beta_{jx}^1 & \beta_{jy}^1 & \beta_{jz}^1 & \dots & 3r \\ \vdots & & \vdots & \vdots & \vdots & & \vdots \\ \beta_{1x}^i & \dots & \beta_{jx}^i & \beta_{jy}^i & \beta_{jz}^i & \dots & \beta_{px}^i \\ \vdots & & \vdots & \vdots & \vdots & & \vdots \\ \dots & & \beta_{jx}^N & \beta_{jy}^N & \beta_{jz}^N & \dots & \end{bmatrix} \begin{matrix} 1 \\ \vdots \\ i \text{ } i^{\text{th}} \text{ patient} \\ \vdots \\ N \end{matrix}$$

*xyz component of
jth moment*

Figure 6.14: Matrices of BSA, Y , and centred moment vectors that parameterise the deformations, X .

PLS deformation modes that are relevant to BSA are computed as follows. The predictor variables X are the r moments β^i that parameterise the initial velocity fields \mathbf{v}_0^i , gathered into a $N \times 3r$ matrix. The response variable Y is the $N \times 1$ matrix of the BSA values. N is the number of observations (patients) (Figure 6.14). We define X_c and Y_c the centred variables $X_c = X - \bar{X}$ and $Y_c = Y - \bar{Y}$ respectively. They are modelled by the linear relationships:

$$X_c = TP^T + E \quad (6.12)$$

$$Y_c = UQ^T + F \quad (6.13)$$

These equations are similar to the PCA decomposition: T and U are the matrices of PLS modes (equivalent to the PCA modes); P and Q are the matrices of loadings; and E and F are residuals matrices that are null if the ranks of T and U equal those of X_c and Y_c respectively. T and U are orthonormal, i.e. $T^T T = U^T U = \text{Id}$, but contrary to PCA, the loadings are not necessarily orthogonal. The PLS modes also verify the linear regression equation:

$$U = TD + G \quad (6.14)$$

where D is a diagonal weight matrix and G is the matrix of residuals. Note that PLS asymetry is due to the previous equation. By integrating (6.14) into (6.13), we get $Y_c = TC^T + F^*$. We can then express the PLS modes T in function of X_c , $T = (X_c - E)P^{T+}$, where P^{T+} is the Moore-Penrose pseudo inverse of P^T . The PLS regression model thus write:

$$Y_c^* = X_c B \quad (6.15)$$

where Y_c^* are the predicted response variables Y_c and B is the matrix of regression coefficients given by $B = P^{T+} D Q^T$.

The optimal PLS modes and loadings are computed iteratively, as illustrated in Algorithm 7. The idea is to find, for a given set of variables X_c and Y_c , the unitary weight vectors \mathbf{r} and \mathbf{s} that maximise the PLS criterion (6.11). It can be demonstrated that \mathbf{r} and \mathbf{s} are the first eigenvector of the squared covariance matrices $X_c^T Y_c Y_c^T X_c$ and $Y_c^T X_c X_c^T Y_c$ respectively [Höskuldsson, 1988]. The PLS modes \mathbf{t} and \mathbf{u} are then linear combination of the variables X_c and Y_c weighted by \mathbf{r} and \mathbf{s} respectively. To find the following modes, we remove from the data the variance explained by \mathbf{r} and \mathbf{s} . For the variables X_c , this step, called *deflation*, consists in subtracting from X_c the matrix $\mathbf{t} \mathbf{p}^T$ (cf. Equation 6.12). The same approach is applied to the variables Y_c although here we have to take into account the regression. According to Equation 6.13 and Equation 6.14, Y_c is deflated by subtracting $\mathbf{t}[\mathbf{t}^T Y] / (\mathbf{t}^T \mathbf{t})$. The algorithm loops until all the PLS modes are extracted.

One can demonstrate that the PLS modes are automatically ordered by decreasing variance, like PCA, and covariance between the variables [Höskuldsson, 1988]. PLS modes are therefore also ordered by decreasing correlation, as we shall see in the results on tetralogy of Fallot reported in Chapter 7. We can thus select a subset of PLS modes by taking the first p PLS modes that simultaneously explain $\pi_X\%$ of variance of X_c and $\pi_Y\%$ of variance of Y_c .

Algorithm 7 Partial Least Squares Space Decomposition**Require:** X, Y, p number of PLS modes to extract, $p \leq N - 1$.

- 1: $X_c^0 \leftarrow X - \bar{X}$
- 2: $Y_c^0 \leftarrow Y - \bar{Y}$
- 3: **for** $n = 1$ to p **do**
- 4: $\mathbf{r}^n \leftarrow$ first eigenvector of $X_c^{nT} Y_c^n Y_c^{nT} X_c^n$ *{weight vector of X_c^n }*
- 5: $\mathbf{r}^n \leftarrow \mathbf{r}^n / \|\mathbf{r}^n\|$ *{normalisation}*
- 6: $\mathbf{t}^n \leftarrow X_c^n \mathbf{r}^n$ *{ n^{th} PLS mode of X }*
- 7: $\mathbf{s}^n \leftarrow Y_c^{nT} \mathbf{t}^n / (\mathbf{t}^{nT} \mathbf{t}^n)$ *{weight vector of Y_c^n }*
- 8: $\mathbf{s}^n \leftarrow \mathbf{s}^n / \|\mathbf{s}^n\|$ *{normalisation}*
- 9: $\mathbf{u}^n \leftarrow Y_c^n \mathbf{s}^n$ *{ n^{th} PLS mode of Y }*
- 10: $\mathbf{p}^n \leftarrow X_c^{nT} \mathbf{t}^n / (\mathbf{t}^{nT} \mathbf{t}^n)$ *{ n^{th} loading of X }*
- 11: $\mathbf{q}^n \leftarrow Y_c^{nT} \mathbf{u}^n / (\mathbf{u}^{nT} \mathbf{u}^n)$ *{ n^{th} loading of Y }*
- 12: $X_c^{n+1} \leftarrow X_c^n - \mathbf{t}^n \mathbf{p}^{nT}$ *{deflation}*
- 13: $Y_c^{n+1} \leftarrow Y_c^n - \mathbf{t}^n [\mathbf{t}^{nT} Y_c^n / (\mathbf{t}^{nT} \mathbf{t}^n)]$ *{deflation with regression}*
- 14: **return** $T = (\mathbf{t}^n)_{n=1 \dots p}$, $P = (\mathbf{p}^n)_{n=1 \dots p}$, $U = (\mathbf{u}^n)_{n=1 \dots p}$, $Q = (\mathbf{q}^n)_{n=1 \dots p}$

As in the previous section, the patients are represented by PLS shape vectors $\mathbf{t}^i = \{t^{m,i}\}_{m=1 \dots p}$, whose elements $t^{m,i}$ are the scalar products $\langle \mathbf{b}^i - \bar{\mathbf{b}}, \mathbf{p}^m \rangle$, \mathbf{b}^i is the column-wise deformation moments of the patient i , $\mathbf{b}^i = (\beta_{1x}^i, \beta_{1y}^i, \beta_{1z}^i, \dots, \beta_{px}^i, \beta_{py}^i, \beta_{pz}^i)^T$, $\bar{\mathbf{b}}$ is the average moment vector computed over the population and \mathbf{p}^m is the m^{th} PLS loading.

6.4.3 A Generative Model of Heart Growth

The PLS modes constitute an optimal subspace for relating shape to BSA. In our framework, they optimise the linear regression that predicts BSA from the deformation modes, $BSA = f(shape)$. Yet, our aim is the converse, to predict the shape given a BSA. We thus need to know how much each PLS mode varies when BSA varies. To that end, we employ canonical correlation analysis (CCA) on the PLS shape vectors \mathbf{t}^i 's. CCA generalises the notion of scalar correlation coefficients to sets of variables [Hotelling, 1936]. Basically, CCA finds bases of the space of variables of maximum correlation. The basis vectors quantify how much a variable of one set correlates with a variable of the second set. Mathematically, CCA calculates the components \mathbf{r} and \mathbf{s} that maximise the correlation between the two sets of variables:

$$\max_{|\mathbf{r}|=|\mathbf{s}|=1} \text{corr}(\mathbf{X}\mathbf{r}, \mathbf{Y}\mathbf{s})^2 \quad (6.16)$$

where \mathbf{Y} is the one-column matrix of population BSA's and \mathbf{X} is the matrix that gathers the p PLS shape descriptors \mathbf{t}^i (Figure 6.15). Variance and covariance matrices are defined by:

$$V_{XX} = \frac{1}{N-1} \mathbf{X}^T \mathbf{X} \quad V_{YY} = \frac{1}{N-1} \mathbf{Y}^T \mathbf{Y} \quad V_{XY} = \frac{1}{N-1} \mathbf{X}^T \mathbf{Y}$$




Figure 6.15: Matrices of BSA's Y and PLS shape vectors X .

The matrix $\Gamma = V_{XX}^{-1/2} V_{XY} V_{YY}^{-1/2}$ can be seen as a multi-variate generalisation of the uni-dimensional correlation coefficient $r = \sigma_{xy} / \sqrt{\sigma_{xx} \sigma_{yy}}$. The sought correlations are obtained by SVD decomposition of Γ :

$$\Gamma = ASB^T$$

S is the diagonal matrix of the correlation coefficients between correlation vectors and A and B are rotation matrices of correlation vectors, i.e. $A^T A = B^T B = \text{Id}$. In our application, Y is a one-column matrix. Hence, S has only one non-null coefficient R , which is the overall correlation between the PLS shape vectors X and BSA. B is a scalar equal to ± 1 that determines the direction of BSA correlation. The elements of the first correlation vector of A , denoted by ρ , relate to the amplitude and direction of correlations of each predictor, namely each PLS mode, when Y varies along the direction defined by the sign of B . In other words, when BSA varies by 1, the k^{th} predictor varies by $BR\rho[k]$. We can therefore compute a generative average model of heart growth by artificially increasing BSA and deforming the template \bar{T} with the growth deformation Φ parameterised by the moments $\mu = B \cdot R \cdot \sum_k \rho[k] \mathbf{p}^k$, where \mathbf{p}^k is the k^{th} PLS loading.

Test of Significance Bartlett-Lawley test can be used to test if the correlations R and ρ are statistically different from 0 [Fujikoshi and Veitch, 1979]. The underlying principle consists in iteratively testing the rank of the matrix Γ . In the general case, let γ_k be the eigenvalues of Γ sorted in decreasing order. If $\text{rank}(\Gamma) = 0$, the predictors X and the response Y are not correlated. We assume that Γ is of rank $k - 1$. We can then test the hypothesis H_0 that the k^{th} eigenvalue is null, i.e. Γ is of rank $k - 1$, by calculating the statistic:

$$L(\Gamma) = - \left(N - 1 - k - 0.5(p + q + 1) + \sum_{i=1}^k \gamma_i^{-2} \right) \ln \left(\prod_{k+1}^{\min(q,p)} (1 - \gamma_i^2) \right)$$

In this equation, p and q are the number of predictors and responses respectively. Asymptotically, $L(\Gamma)$ follows a χ^2 distribution with $(p - k + 1)(q - k + 1)$ degrees

of freedom, yielding a p -value that reflects the probability to reject H_0 . In our case, $q = 1$ (only one clinical parameter, the BSA). The rank of Γ can hence be only 0 or 1. $L(\Gamma)$ writes:

$$L(\Gamma) = - (N - 2 - 0.5(p + 2) + \gamma_1^{-2})$$

— \diamond —

In the next chapter, this framework is applied on the right ventricle of patients with repaired ToF to identify pathological patterns of the RV shape and to estimate a generative model of heart remodelling.

Identification and Modelling of Right Ventricular Growth in Tetralogy of Fallot

Contents

7.1 Clinical Motivation	134
7.2 Patient Selection and Data Preparation	134
7.2.1 Patient Selection	134
7.2.2 Data Preparation	135
7.2.3 Statistical Atlas of the Right Ventricle in ToF	136
7.3 Exploring RV Shape Patterns Related to Regurgitations . .	137
7.4 Generative Statistical Model of the RV Growth in ToF . .	141
7.4.1 Partial Least Square Regression of Shapes	141
7.4.2 Generative Model of RV Growth	142
7.4.3 Comparison with PCA Regression	144
7.4.4 Validation of the Model Generalisation	148
7.5 Discussion	150

In Chapter 6 we proposed a framework to analyse and model the heart shape in a population of patients and to relate it to specific clinical features. In this chapter we apply it to the specific case of repaired tetralogy of Fallot (ToF). Two key clinical questions are investigated. On the one hand, it is crucial to understand the morphological changes of the right ventricle (RV) shape due to the disease. Identifying and quantifying such abnormal patterns could constitute new shape-based features for clinical decision support. On the other hand, predicting how the RV will evolve in the future is one of the key question cardiologists need to answer to decide the optimal time for pulmonary valve replacement. As a first step towards these aims, we apply the group-wise methods previously described on tetralogy of Fallot data. Exhibited pathological shape patterns and cardiac remodelling were consistent with observations reported in the clinical literature. The analysis also demonstrated that partial least squares provides more realistic and predictive models than standard PCA approaches.

7.1 Motivation: What are the Anatomical Evidences of Tetralogy of Fallot?

Contrary to the left ventricle, whose shape under pathological conditions is well documented (see [Zhong et al., 2010] and references therein), RV anatomy is complex and varies tremendously among ToF patients. Several studies investigated possible correlations between clinical parameters in ToF [Therrien et al., 2005; Geva, 2006; Samyn et al., 2007; Frigiola et al., 2008; Bodhey et al., 2008]. However, few works have quantified the anatomical alterations of the RV due to the disease [Geva, 2006; Sheehan et al., 2007; Zhang et al., 2010b]. In [Sheehan et al., 2007], the authors measured the most striking differences in RV shape compared to normals, which enabled to identify abnormal RV remodelling in ToF. Yet, only one-dimensional indices were considered in that study despite the availability of 3D segmentations. Complex 3D deformations may thus have been overlooked. In [Zhang et al., 2010b], the authors presented a 4D active appearance model of the beating heart based on point distribution models to segment the RV in MRI. New shape-based indices were proposed to classify patients from controls, achieving very good classification rates. Nevertheless, the authors did not correlate their model with clinical features of ToF.

The clinical challenges raised by ToF encourage us to apply image-based shape analysis techniques to model the anatomical alterations of the RV due to pathological factors. In Chapter 6, we presented a framework based on a forward model of shapes that is suitable for statistical analyses. That framework enables to correlate 3D shapes with clinical parameters of interest and to build, from a population of patients, a statistical model of anatomical growth. Starting from the assumption that pulmonary regurgitations is an important factor of RV dysfunction [Frigiola et al., 2004; Geva, 2006], we first investigate their effect on the RV anatomy (Section 7.3). To that end, we relate the 3D RV anatomy to pulmonary and tricuspid regurgitations as quantified from echocardiography and MRI evaluations. The idea is to identify quantitative 3D shape features of RV integrity relevant to the regurgitation severity. In a second stage, we build a statistical model of RV remodelling by using the partial least squares (PLS) - canonical correlation analysis (CCA) method (see Section 6.4, page 125). The idea is to investigate the average remodelling observed in a population. Such a model could help in understanding the progression of the disease (Section 7.4). Before detailing these analyses, we present in the next section the investigated population and the construction of the ideal template.

7.2 Patient Selection and Data Preparation

7.2.1 Patient Selection

32 young ToF patients (19 males and 13 females) were selected according to the following criteria:

1. Patient body surface area, ultrasound evaluation of pulmonary and tricuspid regurgitations, cine MR images (cMRI) and pulmonary regurgitation fractions

measured from flow MRI were available,

2. Patient age was from 10 to 30 years,
3. Pulmonary regurgitation fraction was higher than 10%,
4. Patients have undergone no valve replacement

Eight out of 32 patients had an aneurysm of the right ventricle outflow tract (RVOT). Body-surface area (BSA) was computed for each patient using Dubois formula [DuBois and DuBois, 1915]. Table 7.1 reports the mean and standard deviations (SD) of these parameters.

7.2.2 Data Preparation

Echocardiography

Colour Doppler ultrasound (sweep speeds: 50-100 mm/s) was used to quantify tricuspid (TriReg) and pulmonary (TPVReg, TPV for trans-pulmonary valve) regurgitations according to the guidelines reported in [Pongiglione and Trocchio, 2006]. From the regurgitant jet flows identified by colour Doppler imaging, five levels of TriReg and TPVReg were identified: none, trace, mild, moderate and severe regurgitations. Because of the small number of patients, levels were grouped as follows to enhance the statistical power of the analysis.

- None, Trace and Mild TPVReg were grouped together as Mild level; Moderate and Severe were grouped as Severe level, thus resulting in two different groups: Mild and Severe
- Mild, Moderate and Severe TriReg were grouped together as Moderate level, thus resulting in three different TriReg groups: None, Trace and Mild.

MR Image Preparation

Steady-State Free Precession (SSFP) cMRI of the heart were acquired with 1.5T MR scanners (Avanto, Siemens at GOSH and Necker; Achieva Philips at OPBG). Images were acquired in the short-axis view covering entirely both ventricles (10-15 slices; isotropic in-plane resolution: $1.1 \times 1.1 \text{ mm}$ to $1.7 \times 1.7 \text{ mm}$; slice thickness: 5 mm to 10 mm ; 25 to 40 phases). Pulmonary regurgitation fractions (PRF), defined as the percentage of backward blood flow (blood that comes back to the RV) over the outward flow (blood that goes to the lungs), were estimated using 2D+t flow MRI acquired at the proximal pulmonary artery section. In 25 patients, both end-diastole volume (EDV) and end-systole volume (ESV) were calculated from manual segmentations of the RV endocardium at end-diastole and end-systole. Stroke volumes ($SV = EDV - ESV$) and pulmonary regurgitant volumes ($PRV = (EDV - ESV) \times PRF$) were then derived (Table 7.1).

Table 7.1: Mean and standard deviation (SD) of clinical parameters in 32 ToF patients. Values with * are computed on a subset of 25 patients.

Parameters	Mean \pm SD
Age	16.1 ± 4.1 year
Body Surface Area (BSA)	1.53 ± 0.35 m^2
End-Diastole Volume (EDV)	211 ± 90 mL^*
End-Systole Volume (ESV)	116 ± 52 mL^*
Pulmonary Regurgitation Fraction (PRF)	$39 \pm 11\%$
Pulmonary Regurgitation Volume (PRV)	38 ± 24 mL^*
RVOT Aneurysm	8/32 patients

Surface Mesh Preparation

In this study we focused on the RV shape at end-diastole, when the anatomical features of the pathology are the most evident [Sheehan et al., 2007]. The RV endocardium at end-diastole was segmented on the MRI cardiac sequence using the methods proposed by [Zheng et al., 2008; Yang et al., 2008b]. In simple words, the RV endocardium was delineated by fitting an anatomically accurate geometrical model. Its position, orientation and scale in the end-diastole image was determined automatically using marginal space learning. Then, local boundaries were estimated by training a probabilistic boosting tree classifier with steerable features. The mesh was finally resampled in local anatomical coordinates to guarantee vertex correspondence. We could hence use a standard least-square method to rigidly align the patients to a common space to reduce positioning effects in the shape analyses (Figure 7.1, left panel) (the method first appeared in [Mosier, 1939] but is better known in computer vision from [Arun et al., 1987]). In non-reported experiments, we observed that small perturbations in the rigid-body alignment does not affect the statistical findings significantly, thanks to the unbiased estimation of the template.

7.2.3 Statistical Atlas of the Right Ventricle in ToF

As it has been described in Chapter 6, estimating the unbiased RV template requires setting two parameters: λ_V , which defines the “stiffness” of the non-linear diffeomorphic deformations (the higher λ_V , the more rigid the transformations); and λ_W , which characterises the resolution of the currents representation (low λ_W values enable to analyse subtle shape features). As we were mainly interested in the regional alterations of the RV due to ToF (dilation, valve enlargement, regional bulging), the “rigidity” parameter of the transformation, λ_V , was set to 30 mm , about the diameter of the pulmonary annulus. The currents resolution was set to $\lambda_W = 10\text{ mm}$ to have good mesh matching while discarding features due to image artefacts (small bumps, cuts, etc.). Besides, lower λ_W values would have been inappropriate as the image slice thickness was $\approx 10\text{ mm}$.

Five iterations of the alternate minimisation were needed to reach convergence. The resulting template was well centred (standardised mean $\nu = \text{mean}/\text{sd} = 0.2$).

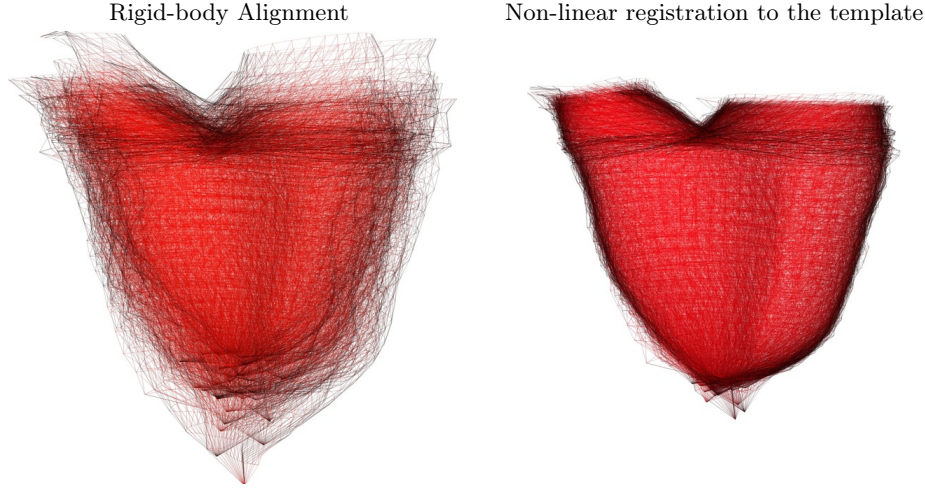


Figure 7.1: 3D RV meshes of 32 young ToF patients. *Left panel:* The meshes were rigidly registered to a representative patient of the dataset. Observe the large variability in shape. *Right panel:* The meshes are registered back to the unbiased template using the non-linear deformations estimated during the template creation.

Figure 7.1, right panel, shows the meshes of all the patients matched back onto the unbiased template. As one can see, most of the shape variability visible in the linear registration (Figure 7.1, left panel) was captured by the diffeomorphic template-to-subject transformations. Remaining local differences were mostly due to segmentation artefacts, thus not relevant for our analyses.

Interestingly, the age of the closest patient to the template, in terms of the norm in the space of currents W , was 16 and his BSA was $1.64 m^2$. These indices were fairly close to the observed averages (Table 7.1), which suggested that in our population, the mean shape was consistent with the mean age and BSA.

7.3 Exploring Right Ventricle Shape Patterns Related to Regurgitations in Tetralogy of Fallot

Equipped with the template and the deformations that map it to the patients, we now investigate the relationships between shapes and clinical indices that quantify regurgitation severity. Shape modes relevant to these indices may identify pathological morphological patterns. To that end, we first performed principal component analysis (PCA) on the deformations as described in Section 6.3.2, page 119. The first 14 deformation modes were selected, representing more than 90% of the spectral energy (Figure 7.2). Then, the patient shape vectors \mathbf{s}^i were computed by projecting the initial velocities \mathbf{v}_0^i onto the PCA subspace (the \mathbf{s}^i 's are 14-element vectors). For the following exploratory analyses, the statistical tests were performed on the PCA shape vectors \mathbf{s}^i as described in Section 6.3. The level of significance was set at $p < 0.1$ and multiple comparisons were corrected using Bonferroni adjustment.

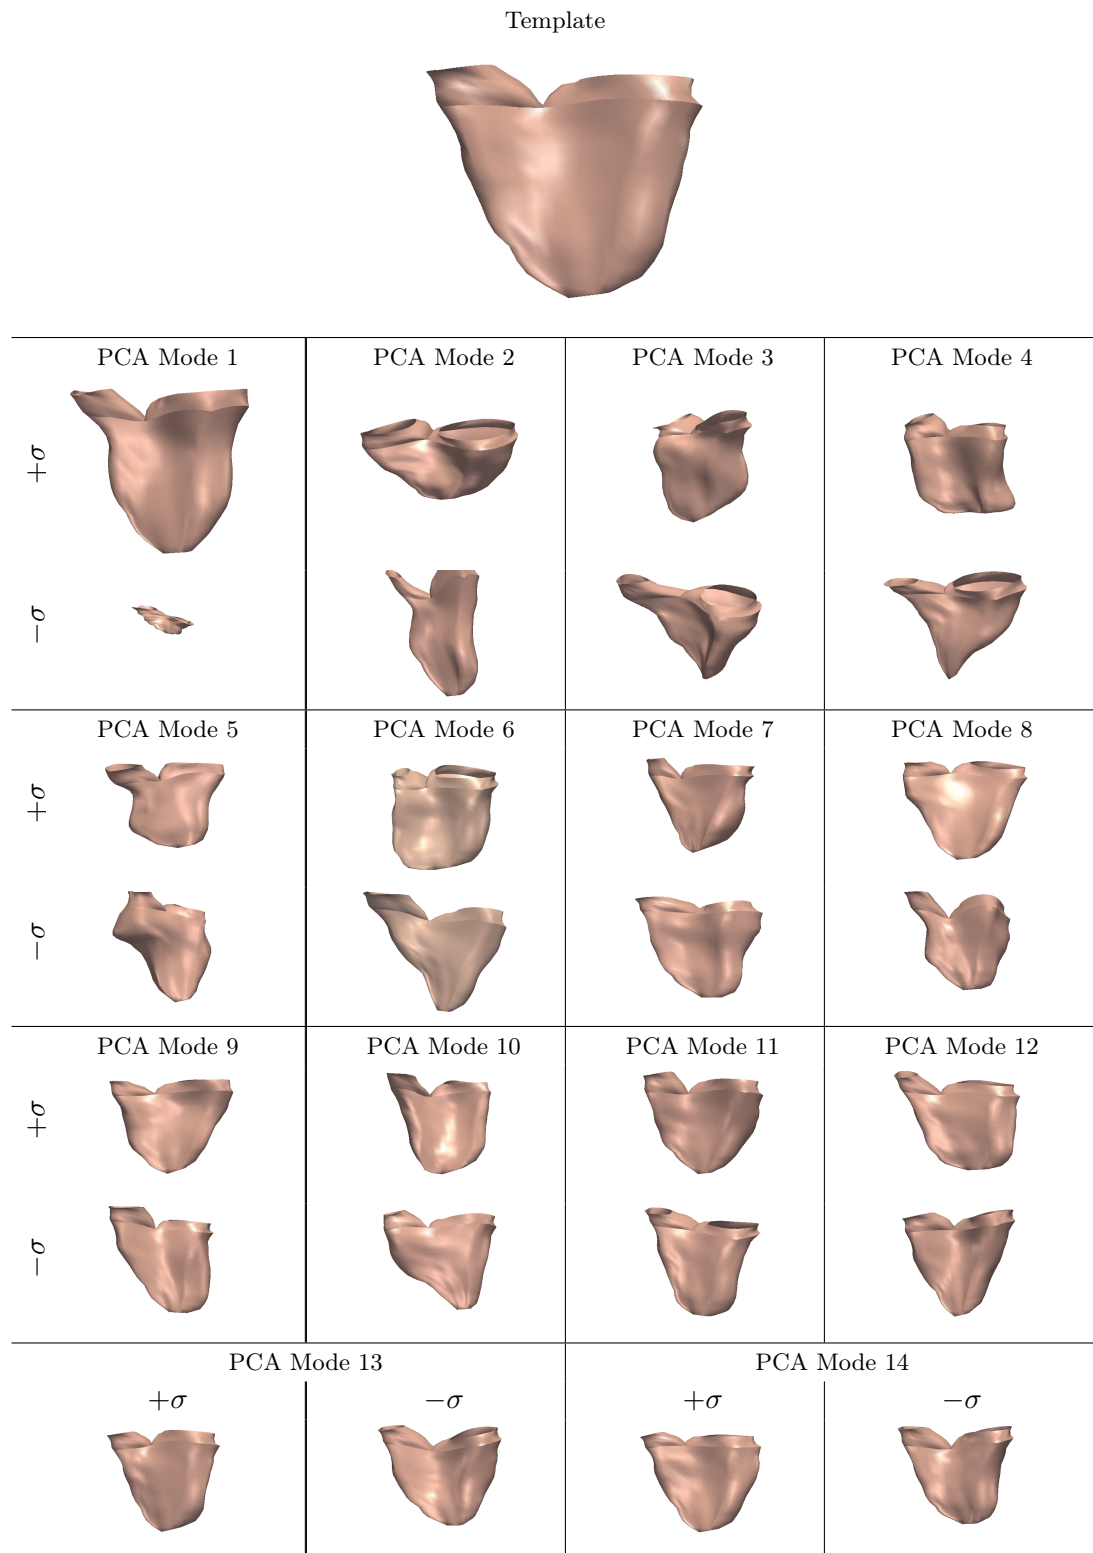


Figure 7.2: 14 first PCA deformation modes from a population of 32 patients suffering from repaired tetralogy of Fallot (90% of variance).

Correlating Shape with Colour Doppler Evaluation of Regurgitations

Kruskal-Wallis analysis between shape vectors \mathbf{s}^i and transpulmonary regurgitation (TPVReg) revealed a significant effect of pulmonary regurgitations on two deformation modes, 5 and 7 ($p < 0.05$). Post-hoc pair-wise Wilconxon test confirmed the finding, both modes separated the mild and severe levels ($p < 0.05$). Visually, mode 5 captured a bulging of the RV basis. The bulging was associated with strong deformations of the valves and the apparition of an aneurysm on the outflow tract. Mode 7 separated regional bulging of the free wall: basal bulging on the one hand and apical bulging on the other hand (Figure 7.2).

Kruskal-Wallis analysis showed a significant effect of tricuspid regurgitations (TriReg) on deformation mode 13 ($p < 0.1$). According to pair-wise Wilconxon test, this mode separated two TriReg levels: none versus mild ($p < 0.1$). As one can see from Figure 7.3, it captured a deformation of the tricuspid annulus, from circular to ovoid, and of the RV inflow tract (base and apex).

It has to be note however that these analyses did not provide the direction of the correlations. We did not know the evolution of the mode when the severity of the regurgitation increased.




Figure 7.3: Deformation mode related to tricuspid regurgitations (TriReg). The mode captured a deformation of the tricuspid valve and of the RV inflow tract.

Correlating Shape with MRI Pulmonary Regurgitation Volume

We then tested the shape vectors \mathbf{s}^i with respect to pulmonary regurgitation volume (PRV) as measured from MRI. Because of missing data, that analysis was performed on 25 patients only. As PRV is a quantitative and continue parameter, we analysed the shape vectors by linear regression ($\text{PRV} = a_0 + \sum_{l=1}^{14} a_l \mathbf{s}[l]$), which yielded good and significant correlation ($R^2 = 0.77$, $p = 0.08$). BIC and backward model reduction selected the same deformation modes (1, 4, 5, 7, 8 and 12), with still good model fit ($R^2 = 0.66$, $p < 0.005$), whereas AIC criterion kept two more deformation modes (6 and 11, $R^2 = 0.71$, $p < 0.005$). Table 7.2 reports the regression coefficients a_l . Interestingly, among the selected modes, we retrieved the deformation modes that were found relevant to TPVReg (mode 5, $p < 0.05$, mode 7, $p < 0.01$). This result suggested a good stability of our analysis. The sign of the regression coefficients indicated the direction of correlation (Figure 7.4). Mode 5 was negatively correlated, suggesting that when PRV increases, the 5th deformation modes goes towards $-\sigma$, i.e. the RV base bulges and an RVOT aneurysm appears.

The RV basal bulging was already reported by [Sheehan et al., 2007]. Similarly, the positive coefficient associated to mode 7 indicated that the RV apex dilates when PRV increases ($+\sigma$). This is consistent with the findings reported by [Bodhey et al., 2008]. We finally noticed that the first mode, which mainly captured an overall RV dilation, was positively correlated with PRV (the RV dilates as PRV increases), suggesting cross-effects between PRV and RV growth as we shall see in the next section.




Figure 7.4: Deformations related to increasing pulmonary regurgitation volumes (PRV). Mode 5 suggested basal bulging and apparition of RVOT aneurysm as PRV increases. Mode 7 captured localised dilations of the RV apex. These two modes were also relevant to TPVReg, confirming the consistency of the analysis.

Table 7.2: Regression coefficients a_l between PCA shape vectors \mathbf{s} and pulmonary regurgitation volume (PRV). Significant coefficients are highlighted in bold ($p < 0.1$). Coefficients were fairly constant after model reduction, confirming the stability of the analysis. The sign indicates the direction of the modes for increasing PRV.

	PCA-All		PCA-BIC		PCA-AIC	
	$R^2 = 0.78, p = 0.076$		$R^2 = 0.66, p = 0.001$		$R^2 = 0.72, p = 0.003$	
	Coef.	t -values	Coef.	t -values	Coef.	t -values
a_0	43.5	10.3	41.94	12.48	43.13	12.91
a_1	$5.98 \cdot 10^{-3}$	2.38	$7.57 \cdot 10^{-3}$	4.11	$7.15 \cdot 10^{-3}$	3.95
a_2	$6.05 \cdot 10^{-3}$	0.84				
a_3	$6.84 \cdot 10^{-3}$	0.54				
a_4	0.017	0.95	0.020	2.09	0.022	2.35
a_5	-0.021	-1.97	-0.017	-2.22	-0.017	-2.33
a_6	-0.017	-1.27			-0.015	-1.34
a_7	0.034	1.55	0.040	3.07	0.042	3.2
a_8	0.025	1.55	0.023	1.64	0.022	1.64
a_9	-0.014	-0.74				
a_{10}	0.023	1.15				
a_{11}	0.050	1.59			0.03	1.24
a_{12}	-0.045	-1.75	-0.046	-2.09	-0.048	-2.25
a_{13}	0.019	0.51				
a_{14}	-0.035	-0.94				

Correlating Shape with MRI Pulmonary Regurgitation Fraction

We finally tested the shape vectors \mathbf{s}^i against the pulmonary regurgitation fractions (PRF) measured from MRI, although this index is not very representative of the RV preload as it is relative to the output blood flow [Grosse-Wortmann and Redington, 2009]. Linear regression between shape vectors and PRF ($\text{PRF} = a_0 + \sum_{l=1}^{14} a_l \mathbf{s}[l]$) had a low and not significant fit ($R^2 = 0.4$, $p = 0.81$). Only mode 5 was found significant to the model ($p = 0.06$) with negative coefficient, as with PRV, which further supported the stability of the analysis. However that mode was not kept by AIC, BIC and backward model reduction. The poor correlation between shape and PRF may be related to the low relevance of the PRF index for assessing the severity of RV impairment.

7.4 Generative Statistical Model of the RV Growth in ToF

We then investigated the heart growth observed in our population using the PLS method presented in Section 6.4, page 125. The results were compared to PCA regression to quantify the added value of PLS.

7.4.1 Partial Least Square Regression of Shapes

Patient growth was quantified by body surface area (BSA) index. In our dataset, BSA correlates with age ($R^2 > 0.5$, $p < 0.001$). We first reduced the dimensionality of the problem by projecting each individual onto the subspace of partial least squares (PLS) modes (Algorithm 7, Chapter 6). We selected the 7 first modes, which represented 98% of the observed BSA variability and 66% of the observed shape variability in the population (Figure 7.5). It has to be noted that in that analysis, we were mainly interested in the shape information that is relevant to BSA and not to the total shape variability. This contrasts with PCA, which explains only the observed shape variability independently of external variables. From the selected 7 modes, we computed the PLS shape vectors \mathbf{t}^i of all the patients by projecting their initial velocity \mathbf{v}_0^i onto the PLS subspace.

By construction, PLS already provided a regression between deformation moments and BSA. However, here PLS was used for space decomposition and model reduction. Similar to PCA space decomposition, we estimated a linear regression model between PLS shape vectors and BSA, $\text{BSA} = a_0 + \sum_{l=1}^7 a_l \mathbf{t}[l]$, to get the directions of the correlations between the PLS modes and BSA. Table 7.3 reports the regression coefficients a_l , the related t -values and the overall model significance. The fit was very strong ($R^2 = 0.85$, $p < 10^{-5}$), with low residual standard error ($\sigma_{\varepsilon_{PLS}} = 0.16 \text{ m}^2$). We used that regression to remove the uncertainty on the direction of the PLS mode. We chose to orient them such that the correlations were positive (PLS mode towards $+\sigma$ when BSA increases). Linear model reduction (AIC, BIC and backward strategy) kept all the 7 components: PLS automatically




Figure 7.5: Cumulative percentage of variance explained by PLS modes. 7 components simultaneously explained 98% of BSA variance and 66% of shape variance.

Table 7.3: Coefficients a_i of the linear regression between PLS shape vectors and BSA ($R^2 = 0.85$, $p < 10^{-5}$). All the coefficients were found very significant ($p < 0.005$). PLS automatically extracted deformation modes relevant to BSA. Coefficient sign was used to specify PLS mode direction: positive when BSA increases.

	a_0	a_1	a_2	a_3	a_4	a_5	a_6	a_7
Coef. $\times 10^{-4}$	$1.52 \cdot 10^4$	9.2	10.0	19.2	13.4	12.4	11.6	3.5
t -values	50.7	10.6	9.16	6.87	6.56	5.52	4.65	3.16

found the modes that were pertinent to BSA. Furthermore, the t -values of the regression coefficients steadily decreased, which confirmed that PLS extracts modes with decreasing covariance with BSA.

Figure 7.6 illustrates the exhibited PLS modes. Visually, PLS mode 1 captured an overall RV dilation. The second mode was very similar to the 5th PCA mode that was found relevant to pulmonary regurgitations (Figures 7.2 and 7.4): the RV base bulges and an aneurysm appears in the RVOT. This observation further supported the hypothesised relationship between the growing heart in ToF and the severity of pulmonary regurgitations, which impact the RVOT (bulging and apparition of aneurysm). Mode 3 exhibited a significant bulging of the RV apex, while mode 4 showed a deformation of the apical shape towards the pulmonary artery. Mode 5 and 7 captured a clear elongation of the RVOT which, along with mode 2, may capture the RVOT aneurysm. Finally, mode 6 exhibited a strong enlargement of the tricuspid valve and pulmonary annulus.

7.4.2 Generative Model of RV Growth

We then computed the generative model of RV growth using Canonical Correlation Analysis (CCA) on the PLS shape vectors as described in Section 6.4, page 125. Overall correlation between PLS shape vectors \mathbf{t}^i and BSA was $R = 0.92$, confirming the strong correlation between the PLS modes and BSA. Bartlett-Lawley test of

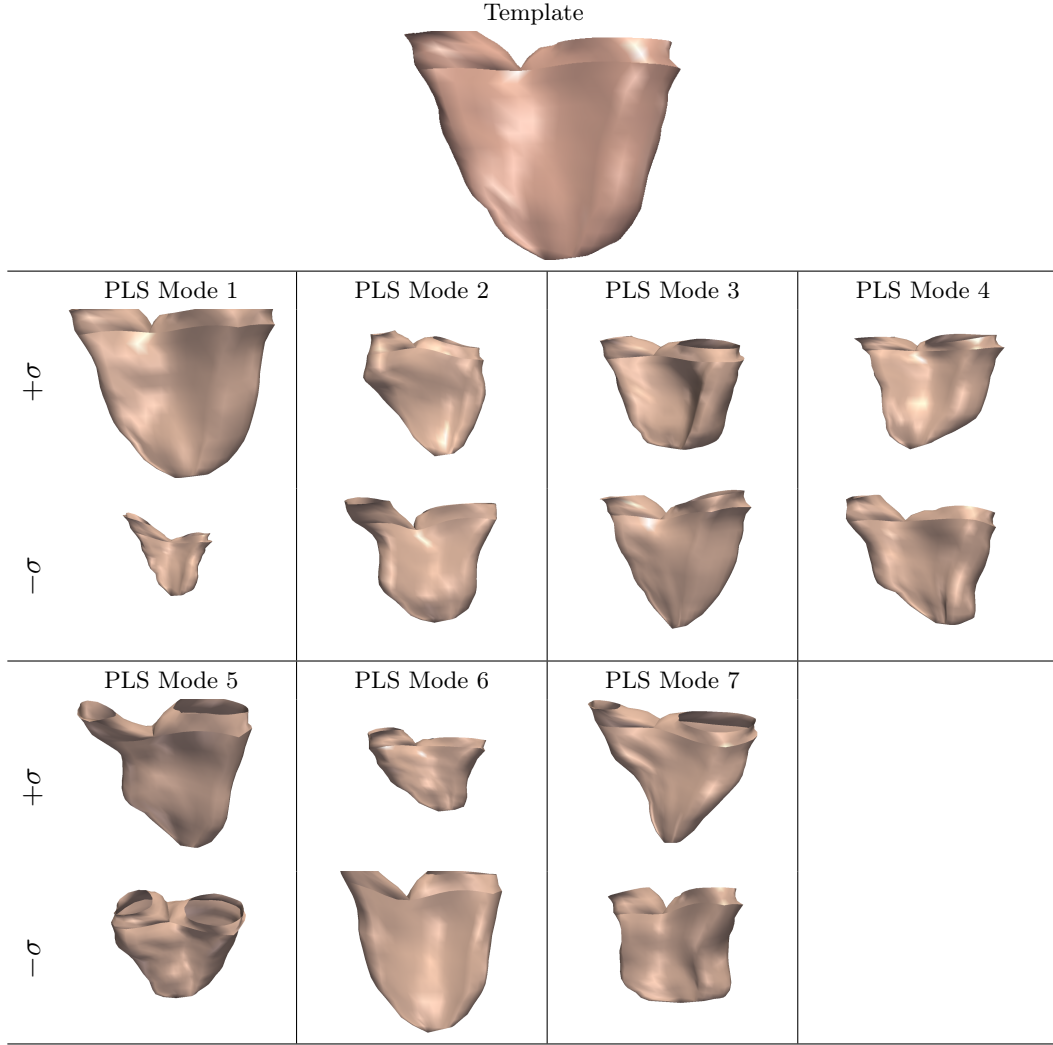


Figure 7.6: Seven first PLS loadings computed with respect to body surface area from a population of 32 patients with repaired tetralogy of Fallot. The 7 loadings explained 98% of the covariance between shape and BSA and 66% of the shape variability observed in the population. Shape evolves towards $+\sigma$ when BSA increases.

correlation significance confirmed these results with a returned probability value lower than $p < 10^{-6}$. The canonical correlation coefficients were

$$\rho_{PLS} = \{0.74, 0.45, 0.34, 0.34, 0.13, 0.07, 0.02\},$$

as illustrated in Figure 7.7, left panel. All the correlation coefficients were positive since the direction of the PLS modes was explicitly chosen such that the regression coefficients were positive. Two important observations can be done from the correlation coefficients ρ_{PLS} (Figure 7.7, left panel). First, the $\rho_{PLS}[m]$'s steadily decreased as a consequence of the PLS algorithm: the PLS modes were automatically ordered by decreasing correlation. Second, the decay of the correlation coefficients




Figure 7.7: CCA correlation coefficients for the three RV growth model: PLS, PCA-BIC, PCA-AIC. *Black line:* fitted exponential. PLS model provided better correlation with body surface area, required less modes and their individual correlation with BSA decreased much faster than for PCA models.

was exponential. In particular, the “time” constant τ of the fitted exponential was equal to 1.75, which means that 1.75 modes were necessary to divide the correlation coefficients by $e \approx 2.72$. After 7 modes, the correlation was already of 0.02.

We finally computed the generative model of RV growth from the ρ_{PLS} , as detailed in Section 6.4.3, page 129. Visually, the obtained RV growth model was found realistic (Figure 7.8). As BSA increased, RV volume increased while the RV free-wall, the tricuspid valve and the RVOT dilated. In proportion, the RV apex dilated much more than the RV inlet and outlet, which was consistent with the observations reported by [Bodhey et al., 2008]. We also observed a progressive deformation of the RVOT: an aneurysm appeared at the late stages of the growth model. Our model thus suggested that the aneurysm appears later during patient life, as a consequence of the cross-effects between growth and regurgitations identified by the second PLS mode (Figure 7.6), which was very similar to the 5th PCA mode (Figure 7.2) that was found relevant to pulmonary regurgitations (Section 7.3). However, the septum of the statistical model of growth became more concave as BSA increased. More precisely, the dilation of the RV inlet folded the septum close to the tricuspid valve. This feature was rather counter intuitive with respect to what is usually observed in ToF patients. Normally, the RV in these patients is rounder than in normals due to increased RV pressure subsequent to myocardium impairment [Geva, 2006; Sheehan et al., 2007]. By analysing the patients of the study, we observed that they did have relatively concave septum, in particular in young adults. This pattern may be particular to our population. Yet, the septum may still be less concave than in normal. Comparison with controls would help in quantifying this feature.

7.4.3 Comparison with PCA Regression

To quantify the added value of the PLS method in our application, we compared the RV growth model obtained on the PLS subspace with models computed on PCA subspaces like in [Mansi et al., 2009c] (Section 6.3.4).

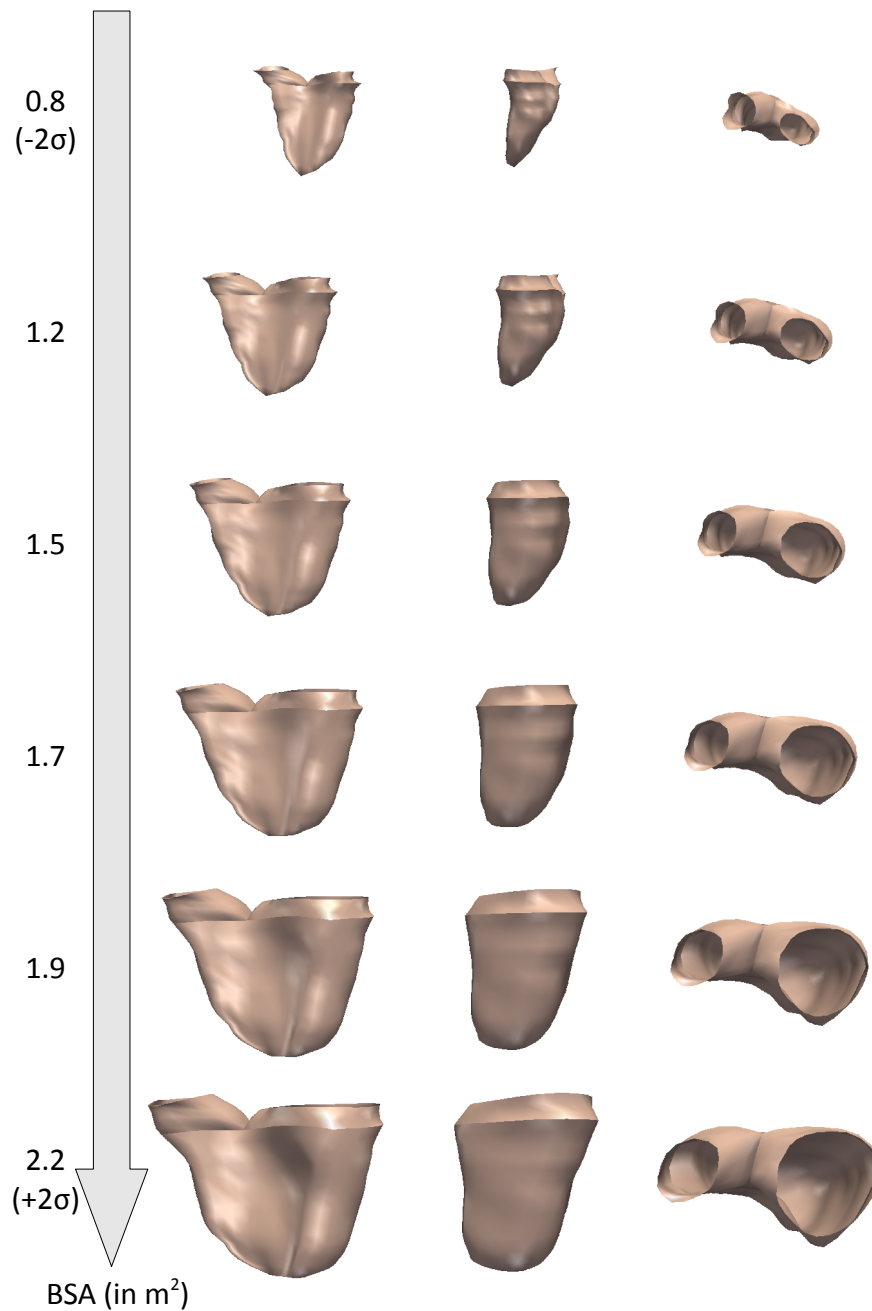


Figure 7.8: Statistical atlas of RV remodelling computed from a population of 32 ToF patients using partial least squares (PLS) subspace and canonical correlation analysis (CCA). While BSA increases, RV globally enlarge. The apex and valves dilate. Simultaneously, the RV free wall becomes rounder and the septum more concave due to the dilation of the RV inlet. At high BSA, an aneurysm appears in the RVOT.

Table 7.4: Linear regression coefficients a_l between PCA shape vectors and BSA. Coefficient sign of the selected modes was used to specify mode direction: positive when BSA increases. Significant coefficients are highlighted in bold ($p < 0.1$). Coefficients were fairly constant after model reduction, confirming the stability of the statistical tests. BIC and backward model reduction yielded the same reduced models.

	PCA-All		PCA-BIC		PCA-AIC	
	$R^2 = 0.71, p = 0.07$		$R^2 = 0.60, p = 0.002$		$R^2 = 0.69, p = 0.004$	
	Coef.	t -values	Coef.	t -values	Coef.	t -values
a_0	$1.52 \cdot 10^0$	28.6	$1.53 \cdot 10^0$	31.7	$1.52 \cdot 10^0$	32.8
a_1	$7.4 \cdot 10^{-5}$	2.4	$7.9 \cdot 10^{-5}$	2.9	$7.3 \cdot 10^{-5}$	-2.7
a_2	$2.7 \cdot 10^{-4}$	2.8	$2.3 \cdot 10^{-4}$	2.9	$2.9 \cdot 10^{-4}$	-3.5
a_3	$8.2 \cdot 10^{-6}$	0.1				
a_4	$6.5 \cdot 10^{-5}$	0.6				
a_5	$-5.8 \cdot 10^{-5}$	-0.5				
a_6	$4.1 \cdot 10^{-4}$	2.1	$3.8 \cdot 10^{-4}$	2.2	$4.1 \cdot 10^{-4}$	2.5
a_7	$-5.7 \cdot 10^{-5}$	0.3				
a_8	$2.5 \cdot 10^{-4}$	1.3			$2.5 \cdot 10^{-4}$	1.4
a_9	$5.8 \cdot 10^{-4}$	2.2	$5.0 \cdot 10^{-4}$	2.1	$6.0 \cdot 10^{-4}$	2.6
a_{10}	$-1.4 \cdot 10^{-4}$	-0.6				
a_{11}	$4.2 \cdot 10^{-4}$	1.3			$4.4 \cdot 10^{-4}$	1.5
a_{12}	$3.7 \cdot 10^{-4}$	1.0			$4.0 \cdot 10^{-4}$	1.2
a_{13}	$1.4 \cdot 10^{-3}$	2.5	$1.2 \cdot 10^{-3}$	2.6	$1.4 \cdot 10^{-3}$	3.0
a_{14}	$7.3 \cdot 10^{-4}$	1.5	$7.6 \cdot 10^{-4}$	1.7	$7.4 \cdot 10^{-4}$	1.7

We determined the linear regression between the PCA shape vectors \mathbf{s}^i and BSA. The modes relevant to the linear model were selected using backward model reduction, AIC and BIC criteria for comparison. Six deformation modes were selected with the backward model reduction strategy and the BIC criterion (1, 2, 6, 9, 13 and 14), henceforth termed *PCA-BIC model*. These six modes represented about 51% of the observed shape variability. Three more modes (8, 11 and 12) were added according to the AIC criterion. That model, called *PCA-AIC* from now on, represented about 59% of the shape variability. Table 7.4 reports the overall significance of each model, the regression coefficients and their t -values.

Thanks to the three additional modes, the PCA-AIC model had a better prediction power as highlighted by the higher model fit ($R^2 = 0.69$) and lower residual errors ($\sigma_{\varepsilon_{PCA-AIC}} = 0.24 m^2$) than the PCA-BIC model ($R^2 = 0.60$, $\sigma_{\varepsilon_{PCA-BIC}} = 0.25 m^2$). This slight increase in performance, which was theoretically expected, was achieved at the price of more variables. Nonetheless, both PCA models were less fitted to BSA and less accurate than the PLS model ($R^2 = 0.85$, $\sigma_{\varepsilon_{PLS}} = 0.16 m^2$). The added value of PLS was even more evident on the 95% confidence intervals of the linear models, which quantify the confidence on the fitted values according to the confidence on the regression coefficients a_l . The confidence interval of the PLS model

was $0.11 m^2$, from 2.5 to 3 times narrower than PCA-BIC ($0.26 m^2$) and PCA-AIC ($0.29 m^2$) confidence intervals. Finally, although only seven PLS modes were used, they explained more shape variability (66%) than the two PCA models (51% and 59% for PCA-BIC and PCA-AIC respectively), with greater BSA prediction power. We can therefore conclude that the PLS model was more accurate, more precise and required less components than PCA-based models while still representing more shape and BSA variability.

We also computed the RV growth model using CCA on the PCA-AIC and PCA-BIC modes and visually compared them to the PLS RV growth model. For the PCA-BIC model, overall correlation coefficient with BSA was $R = 0.77$. The correlation vector of the PCA-BIC deformation modes was

$$\rho_{PCA-BIC} = \{0.55, 0.47, 0.35, 0.29, 0.39, 0.34\},$$

(Bartlett-Lawley correlation significance test $p < 0.005$). For the PCA-AIC model, overall correlation with BSA was slightly improved ($R = 0.83$) thanks to the additional variables, with the correlation vector

$$\rho_{PCA-AIC} = \{0.50, 0.48, 0.33, 0.20, 0.29, 0.16, 0.07, 0.39, 0.31\},$$

(Bartlett-Lawley correlation significance test $p < 0.005$). Yet, both PCA models were less correlated with BSA than the PLS model ($R = 0.92$), suggesting a poorer representation of the visible anatomical changes that occur during growth.

We also observed that the individual correlations of the PCA modes with BSA did not depend on their variance. Modes with low variance, like mode 13 for instance, had higher correlation with BSA than first modes, like mode 6 (Figure 7.7, mid and right panel). This observation confirmed that PCA modes, contrary to PLS, can be relevant to external clinical parameters even if they explain little shape variability. It is therefore crucial to keep as much PCA modes as possible in the analysis. Artificial cut-off can discard shape information related to the clinical parameter of interest. For this reason we also analysed more PCA modes until 98% of the spectral energy was explained and before model over-fitting. No other PCA deformation modes were found relevant to BSA.

We also investigated the correlation decay by reordering the PCA modes according to their correlation to BSA. As illustrated in Figure 7.7, the correlation coefficients did not decrease as fast as in the PLS coefficients (time constants of fitted exponentials: $\tau_{PCA-BIC} = 8.33$, $\tau_{PCA-AIC} = 4.76$ and $\tau_{PLS} = 1.75$). More modes were required to explain less correlation than in PLS model. These experiments confirmed another advantage of PLS over PCA decomposition: PLS automatically determines the minimum yet optimal number of components that simultaneously explain shape variability, BSA variability and their covariance, ordered by decreasing order of covariance and correlation.

Finally, we visually compared the PLS and PCA growth models. PCA-AIC and PCA-BIC models were very similar, with a maximum point-to-point distance of $1.6 mm$ (about the in-plane resolution of the original images) over the entire

growth model. The additional PCA modes kept by the AIC criterion did not added significant information on the remodelling of the RV shape. Compared to the PLS model however the PCA growth models significantly differed in two different regions (Figure 7.9). First, the PCA models showed a more elongated RVOT at low BSA, elongation that progressively disappeared at larger BSA. PLS model exhibited the exact contrary: the RVOT was first normal and then deformed to be more elongated, with a clear enlargement of the pulmonary annulus and with the apparition of an aneurysm at the late stages. Second, the PCA models captured a dilated RV apex at low BSA, which disappeared, relatively to the basal bulging, as BSA increased. Again, PLS model exhibited the inverse phenomenon: the apex was relatively normal at low BSA, without bulging of the apical free-wall. Then, it progressively dilated along with the basal bulging to yield a “rounder” free wall. The differences in these two regions was quantified by large point-to-point distances between the two models ($\approx 9\text{ mm}$ at the pulmonary valve annulus and $\approx 7\text{ mm}$ at the apex). The anatomical changes captured by the PLS growth model were more consistent with observations reported in the literature than the PCA models [Sheehan et al., 2007; Bodhey et al., 2008].

7.4.4 Validation of the Model Generalisation

Generalising the statistical model of RV remodelling is crucial for patient management and therapy planning. We thus tested the robustness of our model on seven new patients who matched the selection criteria (mean age \pm SD = 20 ± 5). Ideally one would like to predict the RV shape of a patient for a given BSA. However, this challenging task is still subject to intense research as it involves acquiring large databases of longitudinal data and developing tools based on complex mathematical theories, like parallel transport on Riemannian manifolds [DoCarmo, 1992]. In this work, we rather tested our model by predicting the BSA of the seven test patients from their RV shape. Although this application has little clinical relevance, it enables to evaluate the ability of the model to represent new patients.

The unbiased template estimated in Section 7.2.3 was registered to each test patient. The resulting deformations ϕ^i were then projected onto the PLS subspace to get the shape vectors \mathbf{t}^i of the new patients. For comparison, we also computed the PCA shape vectors $\mathbf{s}_{PCA-BIC}^i$ and $\mathbf{s}_{PCA-AIC}^i$ by projecting the deformations ϕ^i onto the PCA-BIC and PCA-AIC subspaces. We finally estimated the BSA from the shape vectors using the previously estimated PLS, PCA-BIC and PCA-AIC linear models. 95% prediction intervals were also computed for each model to quantify the intervals into which the predicted values have 95% of chance to fall. The narrower the interval, the more precise the model. Results are reported in Table 7.5.

For all the models (PLS, PCA-AIC and PCA-BIC), the predicted BSA values compared successfully with the measured values, the average error was below the BSA standard deviation of the training population (Table 7.1). The PLS model provided better predictions in 4 patients out of 7 (1,3,4 and 7), as showed in Figure 7.10, left panel. On average about 33% and 25% of improvements were obtained

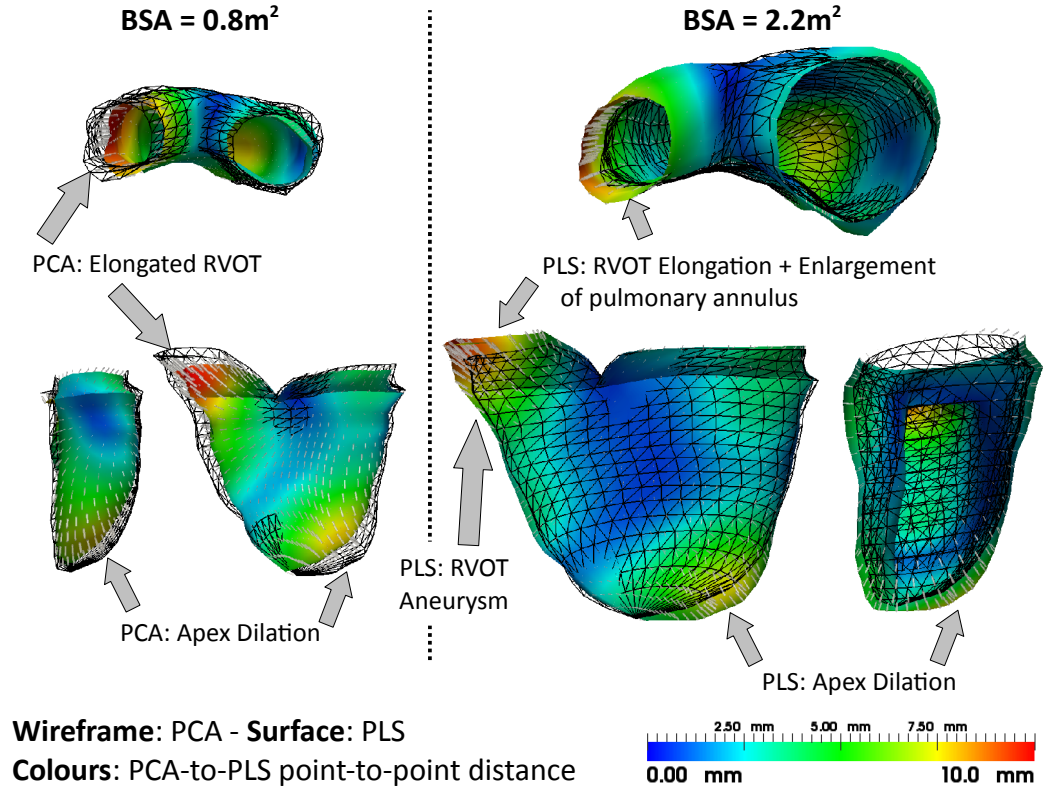


Figure 7.9: Point-to-point distance between PLS RV growth model and PCA-AIC model. The main differences were at the outflow tract and apex. The PLS model showed a more realistic growth pattern than PCA models according to observations reported in the literature (*See text for details*).

Table 7.5: Average prediction errors and 95% prediction intervals for the PLS, PCA-AIC and PCA-BIC models. PLS provided the best predictions both in terms of accuracy and precision.

Model	Mean Error	95% Prediction Interval
PLS	$0.18 m^2$	$0.35 m^2$
PCA-AIC	$0.27 m^2$	$0.59 m^2$
PCA-BIC	$0.24 m^2$	$0.60 m^2$

with respect to PCA-AIC and PCA-BIC respectively. Moreover, the PLS model was much more precise than the PCA models. PLS 95% prediction interval was about 40% narrower than for PCA models, about 1SD of the measured BSA in the training population. These results confirmed that the PLS modes better captured the anatomical changes relevant to patient growth under tetralogy of Fallot, with encouraging generalisation.

It should be noted however that the BSA of patients 5 and 7 could not be

predicted by any model, and in particular by the PLS model. A possible reason for this is the high BSA of these patients. They are thus at the tail of the population distribution (Figure 7.10, right panel), the model is less precise for these subjects. Adding more patients with similar BSA would probably improve the results.

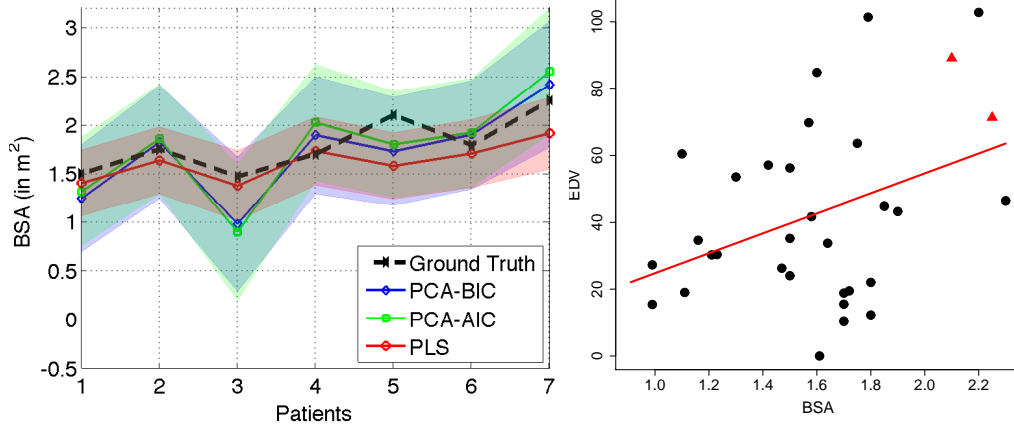


Figure 7.10: *Left panel*: Measured BSA (*in black*) and estimated values using PLS, PCA-AIC and PCA-BIC models. 95% prediction interval are represented in filled patches. One can see from the prediction intervals that the PLS model yielded more precise predictions (bellow population BSA standard deviation). PLS fit was better for 4 patients (1, 3, 4 and 7) out of 7. *Right panel*: Population distribution with respect to BSA and RV end diastolic volume (EDV). Patients 5 and 7 are in red. They are clearly at the limit of the distribution, which may explain why the model is less representative for these patients.

7.5 Discussion

In this chapter, we studied the correlations between the shape of the right ventricle at end-diastole and clinical parameters to identify pathological shape patterns and quantify right ventricle remodelling. The shape at end-diastole was considered as it is the time when the effects of the pathology on the anatomy are the most evident [Sheehan et al., 2007]. An unbiased template was estimated from a population of 32 ToF patients using currents representation and LDDMM registration algorithm. Multivariate statistical analyses on the deformations highlighted shape patterns related to the severity of the regurgitations and provided a generative model of the observed RV remodelling. The findings were found clinically pertinent as they exhibited realistic changes in RV anatomy previously reported in the literature [Sheehan et al., 2007; Bodhey et al., 2008]. To the best of our knowledge, that study constitutes a first attempt at correlating the 3D shape of the right ventricle with clinical measurements in ToF. These analyses may yield quantitative image-based indices about RV anatomy and remodelling in ToF.

The effects of regurgitation severity measured from colour Doppler ultrasound were analysed on a component-by-component basis to preserve the statistical power of the tests due to the ordinal nature of the ultrasound data. The groups were not sufficiently populated to apply more comprehensive statistics. However, despite this limitation, selected PCA deformation modes were consistent with those found by analysing the shape with pulmonary regurgitation fraction and regurgitation volumes measured from MRI. These latter analyses also provided the direction of correlation, further specifying the shape evolution as pulmonary regurgitation increases. In particular, we identified that the base bulges, the apex dilates and an aneurysm appears in the RVOT at the late stage of growth. These findings were consistent with reported observations in the literature. The selected modes can therefore constitute quantitative features of regurgitation severity.

The effects of growth on the RV shape was modelled using partial least squares space decomposition. We compared the results with the more classical PCA decomposition and showed that the PLS model was visually more consistent with reported observations and quantitatively more predictive and precise than PCA models. It also required less components to explain more shape and BSA variability than PCA approaches. PLS thus constitutes an ideal tool for this kind of analyses.

Several technical and clinical questions still remain open. From a technical point of view, it would be interesting to apply the statistical model of RV remodelling to predict the RV shape of a given patient, going from group analysis to individual analysis. That work would indeed have a tremendous clinical impact. To that end, parallel transport algorithms would be necessary to “transport” the deformation associated to the growth model to the patient RV anatomy. However, this task is not trivial as it requires algorithms tailored to the Riemannian manifold of diffeomorphisms [DoCarmo, 1992; Qiu et al., 2008]. As a first step, one could use a simplified parallel transport, as the method proposed by [Rao et al., 2002]. Similar techniques could be used to investigate the 4D cardiac motion in order to correlate shape and function with pathology and growth. Parallel transport would then be used to transport the patient cardiac motion to the template space [Rao et al., 2002; Beg et al., 2004]. An alternative for the 4D analysis would be to use the spatio-temporal model developed by [Durrleman et al., 2009b] although more complex models that take into account the cardiac biomechanics would be preferable in order to constrain the statistical analyses. A possible direction would be to use manifold learning for instance (see [Hamm et al., 2010] and references therein for further details).

Another possible improvement is to better integrate the PLS algorithm into the currents and LDDMM frameworks. In this work, we applied the standard PLS decomposition algorithm on the deformation moments. However, we do know that the moments are parameters of velocity fields that belong to a Gaussian kernel space. Kernel PLS [Rosipal and Trejo, 2002] would therefore be a natural way to handle these data, like we did for the PCA decomposition. Such an approach is expected to provide better predictions and precision. Similarly, the linear models used in this study may be too restrictive. It is indeed acknowledged that teenager growth follows a rather “sigmoidal” trend. Based on this observation, one could fit

polynomial or sigmoidal models to the data. A more elegant way to cope with this would be to perform kernel regression and kernel canonical correlation analysis by taking advantage of the kernel space of the initial velocities.

From a statistical point of view, the main limitation of our study is the absence of stability analysis of the findings with respect to the training population. Although obtained results were qualitatively consistent with those published in [Mansi et al., 2009c] and those obtained in non-reported experiments on different datasets, more quantitative evaluations would be necessary to confirm the presented results. A first task towards this aim is to gather much more patients in order to get enough statistical power and, above all, to be able to perform cross-validation. Three-fold leave-one out analyses could then be carried out as follows. The patient set would be divided into three groups of equal size. The first group of patients would be used to estimate the unbiased template of RV shape. We could then compare the resulting templates and quantify their variations with respect to the training population. Then, we would use the second group of patients to estimate the PLS growth model and to identify anatomical features related to regurgitation severity. This stage would enable to quantify the stability of the selected modes relevant to the clinical parameters. Finally, we would use the third group to test the models. However, to carry out such an experiment we would need at least 100 patients, which were not available during this work.

This leads us to the main clinical limitation: too few patients were involved. As a result, we considered all the available patients, indifferently of the type of initial repair, to maximise the statistical power. Recent studies suggested that the type of initial surgery can lead to different RV remodelling [Samyn et al., 2007; Frigiola et al., 2008]. It would be interesting to investigate the growth patterns in these two populations separately. Nevertheless, this question is extremely difficult to handle as surgical techniques vary from one clinical centre to the other and, above all, improve over time. Finally, the lack of normal data prevented us from studying the differences between ToF patients and healthy subjects. Group-wise analyses would further help in identifying pathological shape patterns.

Other interesting clinical questions related to ToF could be investigated using our approach. First, we could apply the PLS method to investigate how the RV shape evolves when regurgitation volume increases. Preliminary analyses exhibited a very similar trend, which could be explained by cross-effects between growth, dilation and regurgitation severity. Further work however needs to be done in order to confirm these first results. Another possible research direction is to investigate the effect of the genes that regulate myocardium stiffness on the long-term RV remodelling. Patients with stiffer myocardium are known to be more protected against regurgitations. The decision for valve implant may therefore be based on different features for these patients. The LV-RV interaction could also be investigated by creating models of the bi-ventricular myocardium. These models could help in understanding how the septum deforms when the RV dilates. Preliminary results supported this idea. Finally, the approach could also be applied on postoperative data to study the long-term impact of valve implant on RV anatomy.

Part IV

THERAPY: PERSONALISED SIMULATION OF CARDIAC THERAPIES

Personalised Model of Cardiac Electromechanics

Contents

8.1	Motivation	156
8.2	Cardiac Anatomy Model	157
8.2.1	3D Model of Cardiac Geometry for Personalised Simulations	159
8.2.2	Anatomical Labelling for Regional Personalisation	161
8.2.3	Model of Cardiac Fibre Directions	165
8.3	Cardiac Electrophysiology Model	166
8.3.1	Available Models	167
8.3.2	Selected Model: Anisotropic Eikonal Model	168
8.3.3	Personalisation Strategy	169
8.4	Cardiac Biomechanics Model	169
8.4.1	Available Models	169
8.4.2	Selected Model: Simplified Multi-Scale Linear Model	171
8.4.3	Biomechanical Boundary Conditions	173
8.4.4	Personalisation Strategies	176
8.5	Cardiac Hemodynamics Model	178
8.5.1	Simulation of the Cardiac Cycle	179
8.5.2	Simulation of the Arterial Windkessel Effect	181
8.5.3	Simulation of Valve Regurgitations	185
8.6	Discussion	187

In the Parts I and II of this thesis, we presented methods to quantify and model patient-specific cardiac shape and function from medical images that aim to improve diagnosis and prognosis. Yet, these are only the first steps of the clinical framework. Treatment must be planned and applied to a patient. In most of the cases, several therapeutical strategies are available. One needs to figure out which one is optimal for a given patient. So far, that choice is based on evidences drawn from clinical trials on large populations. However, complex pathologies like cardiovascular diseases, congenital diseases or cancers require personalised approaches to accommodate the specificities of each patient.

In this chapter, we propose to use an electromechanical model of the heart compatible with clinical data for personalised predictions of therapy effects. In the

following we describe the proposed methodology. How it is used to predict the post-operative effects of pulmonary valve replacement in patients with repaired ToF is detailed in Chapter 9.

8.1 Motivation

Since the seminal works in cardiac modelling of the early 60's [Fitzhugh, 1961; Noble, 1962], more and more detailed electromechanical (EM) models of the heart have been formulated to simulate the biological phenomena that govern the cardiac activity, from electrophysiology to biomechanics, from the cells to the organ [Glass et al., 1991; McCulloch et al., 1998; Hunter et al., 1998; Nash and Hunter, 2000; Bestel et al., 2001; Hunter et al., 2003b; Sermesant et al., 2006b,a]. These models greatly contributed to the understanding of the heart by allowing the modellers to quantitatively test hypotheses *in-silico*. It is now recognised among the modelling and the clinical communities that such models can have a tremendous impact on the clinical practice, and especially in the personalised management of patients and therapy planing [Ayache, 2004; Crampin et al., 2004; Bassingthwaighe et al., 2009; Ayache et al., 2009]. One could use a virtual heart to test on a computer different therapies and to choose the most suited treatment for a given patient. Nonetheless, translating these tools into precise clinical applications is dauntingly complex, challenged by the discrepancy between the numerous model parameters and the sparse clinical data. Various international initiatives like the Physiome Project [Hunter et al., 2003a] or the Virtual Physiological Human (VPH) framework [Ayache et al., 2006] aim to develop solutions to these scientific challenges.

Despite the difficulties, recent research aims to apply models of the beating heart for patient-specific simulations [Sermesant et al., 2006a; Wong et al., 2007; Tang et al., 2007a; Yang et al., 2008a; Sermesant et al., 2009; Mihalef et al., 2009]. The major difficulty is to find the right compromise between accurate models controlled by large sets of parameters and too simplistic approaches that cannot represent the pathology of the patient. A model is valid only for the question it has been designed for: There is likely no general model that can answer all clinical questions. The strategy employed in these studies, and in this thesis, consists in focusing on a very particular clinical question, on a particular therapy. The model is adapted to the specificities of the pathology and used to predict the effects of the therapy on the cardiac function of the patients. In addition to its possible clinical impact, this approach enables to validate the models from a practical point of view, where their predictive power is preferred to the accuracy of the simulation.

In this thesis, we built upon existing models [Sermesant et al., 2006a] to develop a modular framework for personalised simulations of congenital heart diseases from medical images, with particular emphasis on repaired tetralogy of Fallot (see Chapter 2.3). How is the cardiac function of a patient? Can we simulate the variation of the right ventricular pressure? Can we predict the effects of pulmonary valve replacement or of right ventricle volume reduction on cardiac function? The idea is to use an electromechanical model of the cardiac ventricles to test these therapies on

a computer. The problem is all the more challenging that our patients are children. The quality of the cardiac images is often poor, and an invasive but comprehensive assessment of cardiac function are almost impossible (no pressure measurements, no endocardial mappings). Therefore, the complexity of the electromechanical model must be tailored to the available clinical data for calibration and validation [Serme-sant et al., 2006a; Tang et al., 2007b].

As illustrated in Figure 8.1, the model is constituted of four elements: anatomy, electrophysiology, biomechanics and hemodynamics. The anatomical model is a computational representation of the bi-ventricular myocardium of the patient on which the partial differential equations (PDE) that model the cardiac activity are solved. In this work, the anatomical model is built from medical images, onto which knowledge about cardiac fibres and myocardium lesions are automatically mapped (Section 8.2). A model of cardiac electrophysiology is then used to simulate the electrical wave that commands cardiac contraction (Section 8.3). This model is essential as it enables one to simulate disturbances in the cardiac rhythm, in particular the bundle branch blocks that often occur in repaired tetralogy of Fallot. A biomechanical model of the myocardium is coupled to the cardiac electrophysiology to simulate the cardiac contraction (Section 8.4). The finite element method (FEM) is used to solve the PDE's. Finally, cardiac hemodynamics are modelled as constraints of the electromechanical model (Section 8.5). That model is adapted to the specificities of tetralogy of Fallot by integrating in particular a simple constraint-based model of regurgitations. For each component, we describe the main concepts of the model, the parameters that govern it and strategies for personalisation from clinical data that enabled us to perform personalised simulations of the cardiac function in tetralogy of Fallot (Chapter 9).

8.2 Cardiac Anatomy Model

The first question to ask when one aims to simulate the cardiac function of a patient is how to represent the cardiac anatomy. What is the best computational representation of the patient heart that satisfies the requirements of the study? What anatomical information is needed and what data is available for personalisation? What are the requirements of the electromechanical models in terms of domain discretisation? The answers to these questions will guide the construction of an accurate anatomical model suited for the simulations to be performed.

An anatomical model of the heart is a discrete representation of the patient myocardium on which the constitutive laws of the cardiac electromechanics are solved. The discrete model represents the cardiac geometry and integrates knowledge about myocardial structure such as the orientation of the cardiac fibres and the location of lesions. A 3D object is commonly discretised in a connected mesh when the constitutive equations are solved with finite differences or finite element methods (FEM) [Bathe, 1996]. The mesh is discretised into small elements like hexahedron (six rectangular faces) or tetrahedron (four triangular faces). On the one hand, hexahedral meshes offer a greater numerical accuracy compared to tetrahedral meshes.

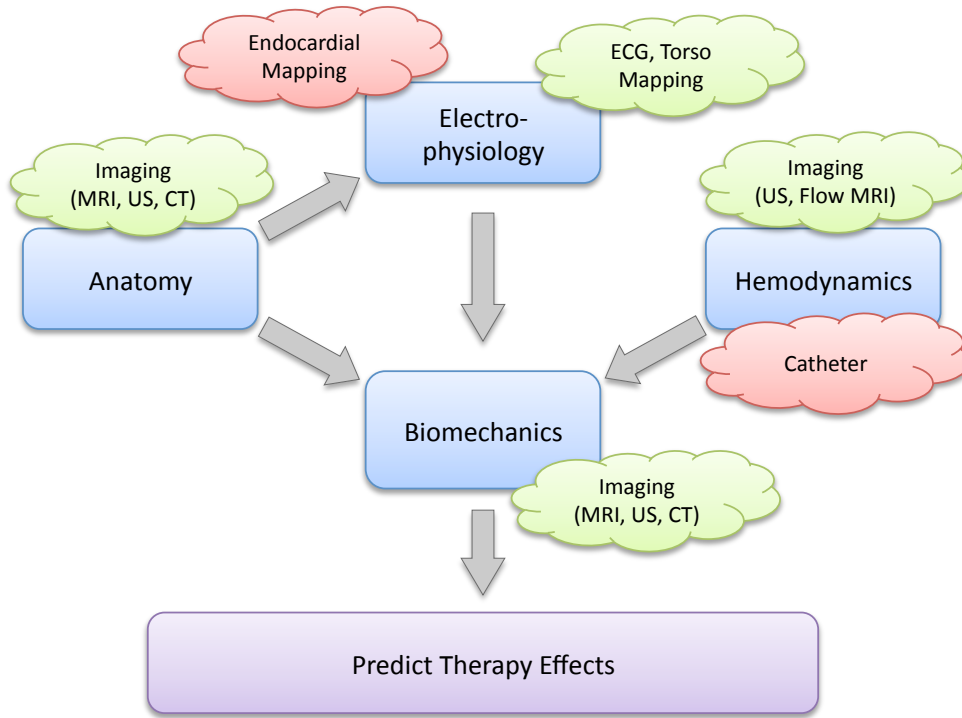


Figure 8.1: Framework for the personalised simulation of cardiac electromechanics in patients. Each part must be personalised from clinical data that can be “patient-friendly” (*in green*) or invasive (*in red*). Arrows point out the dependencies between the elements of the model.

One hexahedron can be enough to discretise a space-varying field according to the basis functions. On the other hand, linear tetrahedral meshes require many more elements to precisely represent the domain since the field inside a tetrahedron is linear, although some techniques are available for high-order polynomial tetrahedral interpolation schemes. Nonetheless, tetrahedral meshes can be generated from surfaces of any shape in contrast with hexahedral meshes, which are more difficult to obtain. Furthermore, because they do not rely on complex interpolation schemes, linear tetrahedral FEM methods can be faster to solve when the number of elements stay reasonable.

Recently, mesh-free approaches have been proposed as an alternative to FEM [Belytschko et al., 1996]. With this framework, the governing equations are solved using radial basis functions on point clouds, which can have any shape. For that matter, mesh-free methods require no remeshing at large deformations and non-linear interpolation can be easily implemented, resulting in accurate simulations even with few nodes [Wong et al., 2010]. However, discontinuous domain boundaries, for instance between the myocardium and the blood pools, are difficult to obtain because of the continuity of the basis function.

Our aim is to simulate the cardiac function of a patient. We thus need to represent any heart shapes. Moreover, efficient discretisation schemes are required for




Figure 8.2: Different discretisation schemes. Hexahedra and tetrahedra are used in FEM. Hexahedra provide more accurate results than tetrahedra but their personalisation to a patient anatomy is more challenging as it is often based on mesh deformation. Mesh-free methods combine the advantages of hexahedra (high accuracy) and tetrahedra (personalisation ability) but domain boundaries are more difficult to represent faithfully.

fast simulations in order to personalise the model. Based on these constraints, we decided to rely on linear tetrahedral meshes despite their limitations. A personalised model of the cardiac anatomy is generated directly from myocardium segmentations (Section 8.2.1) and the resulting mesh is automatically labelled for regional personalisation (Section 8.2.2). A model of fibre direction (but not of fibre sheets as we shall see) is finally mapped onto the mesh to account for myocardium anisotropy (Section 8.2.3).

8.2.1 3D Model of Cardiac Geometry for Personalised Simulations

8.2.1.1 Available Models: Generic Anatomies of the Bi-Ventricular Myocardium

Most of the electromechanical models published in the literature are solved on generic ventricular anatomies. The simplest model is the bi-ellipsoidal model, where the left and right ventricles are represented by truncated ellipsoids [Mercier et al., 1982] (Figure 8.3, left panel). These simplistic models have been widely used as they can be parameterised by closed-form coordinate systems [Nash, 1998] and enable the easy computation of generic fibre models (see Section 8.3) [Guccione and McCulloch, 1991; Arts et al., 2001; Sermesant et al., 2006a; Kroon et al., 2009]. These models can be discretised using either hexahedra or tetrahedra. However, they are too ideal for being representative of the phenomena that occur in real-life geometries.

For this exact reason, highly detailed geometries have been constructed from *ex-vivo* hearts. The cardiac mechanics research group of the University of California, San Diego (UCSD) built a very detailed model of rabbit and pig hearts from histological slices¹ [Vetter and McCulloch, 1998] (Figure 8.3, mid panel). The ven-

¹Models as part of the *Continuity* package, <http://www.continuity.ucsd.edu/Continuity>




Figure 8.3: Generic models of cardiac anatomy.

tricles were fixed mechanically and filled with silicon. The authors then sliced the heart and imaged each slice to create the model. The Auckland Bioengineering Institute created a canine heart² using a similar technique but by slicing the heart from epicardium to endocardium to better preserve the cardiac fibre organisation [Nielsen et al., 1991; LeGrice et al., 1995] (Figure 8.3, right panel). Both models are discretised into hexahedral meshes with bicubic Hermite finite elements expressed in prolate spheroidal coordinates. Recently, anatomical models computed using high-resolution MRI of *ex-vivo* animal hearts have been proposed [Vadakkumpadan et al., 2009]. As they embed very accurate details of the cardiac anatomy, they are widely used in the modelling community. Nonetheless, these models come from non-human specimen. It is not clear to which extent they can be used for patient-specific heart simulations.

8.2.1.2 Selected Model: Patient-Specific Anatomy from MRI Data

The previous generic models may not be suited for simulating cardiac function in patients. The geometry of pathological heart can deviate significantly from the normal shape due to the disease. The ventricle can dilate and the myocardium become thinner. The global shape can even be deteriorated by surgery. One must consider these patient-specific variations to increase the accuracy of the simulation.

A first strategy consists in registering one of the above-mentioned generic mesh to the patient geometry. This approach has been used in [Sermesant et al., 2006a; Niederer et al., 2009] for instance. The main advantage is that the registration maps the mesh discretisation and the anatomical priors that are embedded into the model, like the fibre architecture. This is particularly appealing when the finite elements are hexahedra, which are challenging to build in non-standard geometries. However, the registration may yield badly-shaped elements that may result in numerical instabilities. This limitation can be coped by a subsequent re-meshing.

The models used in this work are solved using tetrahedral finite elements. Efficient tools are therefore available for meshing any geometry with tetrahedra, enabling the construction of patient-specific anatomical models from clinical images directly. We first compute a binary mask of the bi-ventricular myocardium by using

²Model as part of the *CMISS* package, <http://www.cmiss.org/>




Figure 8.4: From mask of patient myocardium to surface mesh to tetrahedral mesh

the pipeline presented in Chapter 3 for instance or any other method [Zheng et al., 2008; Ecabert et al., 2008] (Figure 8.4, left panel). The mesh is taken at the mid-diastole time frame (just before the atrial contraction), when the heart is close to the rest position. As we shall see later in this chapter, valve motion, atrial biomechanics and arterial motions are not simulated directly. This simplifies the generation of the model as only the compact bi-ventricular myocardium is required. Papillary muscle, trabeculae, atria and arteries are excluded from the anatomical model. A triangulated mesh of the myocardium is then computed from a binary mask of the segmentations by means of the INRIA software CGAL³ (Figure 8.4, mid panel). The number of triangles is controlled by defining the average edge length. Finally, we transform the surface mesh into a tetrahedral mesh with GHS3D⁴, another INRIA software program (Figure 8.4, right panel). The surface triangulation is used as reference to create the volume mesh, the more regular the surface discretisation the better the volume mesh.

8.2.2 Anatomical Labelling for Regional Personalisation

The tetrahedral mesh is labelled to identify regions that require special treatments during the simulations. Three types of labels are automatically reported on the mesh: the myocardium interfaces (epicardium, left endocardium and right endocardium), the lesions, and some myocardial zones for regional personalisation.

8.2.2.1 Identification of the Myocardium Interfaces

Myocardial interfaces, namely the endocardium and the epicardium, play an important role in cardiac mechanics as they interact with the structures that neighbour the heart. It is thus essential to define them on the personalised anatomy. An interface is defined as a list of triangles that belong to the surface of the myocardium volume mesh. These triangles can be automatically detected as they belong to only one tetrahedron. In the following, those lists are termed *surface zones*. One can manually pick the surface triangles to define the myocardium interfaces [Sermesant, 2003]. However, this task is time consuming, especially if the mesh has a large

³<http://www.cgal.org>

⁴<http://www-roc.inria.fr/gamma/gamma/ghs3d/ghs.php>

number of elements. We thus developed an automatic algorithm that labels the surface of the volume mesh from the triangulations of the epicardium and endocardia segmentations.

A closest point approach is used. Let \mathcal{T}_{LV} , \mathcal{T}_{RV} and \mathcal{T}_{Epi} be the surface meshes of the left endocardium, the right endocardium and the epicardium respectively. For each surface triangle E of the volume mesh, we compute the point-to-surface distances between the centre of E and the segmentations. E is assigned to the surface zone that corresponds to the closest segmentation \mathcal{T}_{Epi} , \mathcal{T}_{LV} or \mathcal{T}_{RV} . To avoid any ambiguities, we also check that normal of the surface triangle E is oriented approximately in the same direction as the normals of proximal triangles of selected segmentation. Figure 8.5 illustrates the endocardial and epicardial surface zones related to a patient.




Figure 8.5: Automatic labelling of myocardium interfaces from segmented triangulations (*white wireframe*). In red is the left endocardium, in blue the right endocardium and in green the epicardium.

The surface zones enable to automatically identify the left and right ventricles for regional personalisation, labelled as *mesh zones*. Given a tetrahedron T of the mesh, we compute the point to surface distances between the barycentre of T and the surface zones, denoted d_{LV} , d_{RV} and d_{Epi} (Figure 8.6). We assign T to a ventricle according to the following rules:

- $d_{LV} \leq d_{RV}$, independently of d_{Epi} . The tetrahedron T is closer to the left endocardium than to the right endocardium, it *belongs to the left ventricle*.
- $d_{RV} \leq d_{LV} \leq d_{Epi}$. The tetrahedron T is closer to the right endocardium but farther from the epicardium than from the left endocardium. T is in the septum, which *belongs to the left ventricle*.
- In all the other cases, the tetrahedron T *belongs the right ventricle*.

The LV and RV zones computed from the mapped surface zones of Figure 8.5 are showed in Figure 8.6, right panel.




Figure 8.6: Automatic left (*in red*) and right (*in blue*) mesh labelling. In transparent the original delineations. A tetrahedron is automatically assigned to a ventricle according to its proximity to the epicardium and to the left and right endocardia (*see text for details*).

8.2.2.2 Mapping of Scars and Lesions

If the myocardium tissue is damaged by a chronic pathology, a myocardium infarction or surgery sequels, we map the lesions onto the mesh to simulate their impact on the cardiac function. Myocardium scars and other lesions can be assessed with late-enhancement MRI. In these images, the clinicians observe how the scar fix a contrast agent, typically gadolinium. As the scar becomes much brighter than the healthy myocardium, it can be segmented manually or automatically [Setser et al., 2003; Prasad et al., 2004]. The mask of the scar is reported onto the anatomical model using the approach proposed by [Sermesant, 2003] based on Bresenham rasterisation algorithm [Bresenham, 1965]. Let M_{scar} be the binary image of the scar and T be a tetrahedron of the myocardium mesh. Using the Bresenham rasterisation algorithm we pick the image voxels that lay inside the tetrahedron T . From the picked voxels we count those that belong to the scar too. If they represent more than 50% of the picked voxels, the tetrahedron T is assigned to the scar zone. Several scars or lesions can be mapped with different identifiers, thus enabling fine modelling of these abnormal tissues. However, in our implementation a tetrahedron belongs to only one zone at a time. Adaptive meshing can be necessary to accurately represent lesions with complex shapes [Lamecker et al., 2009]. Figure 8.7 shows an example of myocardium scar in a patient with heart failure.

8.2.2.3 Myocardium Subdivision in Standard Anatomical Zones

We finally assign to each tetrahedron of the volume mesh an anatomical region as defined by the American Heart Association [Cerqueira et al., 2002] (Chapter 2, Section 2.2) for regional personalisation. Below the base plane, the left ventricle is divided into four layers: base, mid plane, apical plane and apex. The base and the mid planes are divided into six concentric regions of equal size whereas the apical plane is divided into four zones, resulting in 17 divisions of the left ventricle. The

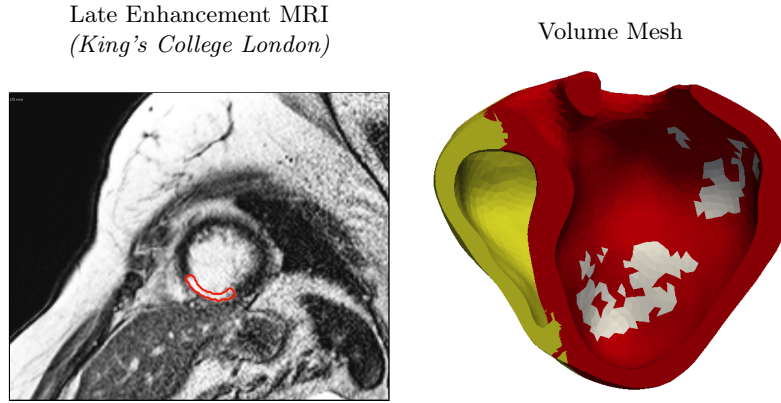


Figure 8.7: *Left panel:* Late enhancement MRI showing a myocardium scar *in red* in a patient with heart failure. The scar is manually delineated. *Right panel:* The scar is automatically mapped onto the patient-specific volume mesh: left ventricle (*in red*), right ventricle (*in yellow*), scar (*in white*).

right ventricle is divided in a similar way (Figure 8.8).

The AHA subdivision is performed automatically from the mapped endocardial and epicardial surface zones. We first define what we call the heart coordinate system Υ of the ventricles (Figure 8.8). The origin of Υ , denoted G , is the LV barycentre computed from the LV endocardium surface. The direction of the LV long axis, represented by the unitary vector \mathbf{e}_l , is given by the line that passes through G and the centre of the mitral valve, namely the barycentre of the edges of the LV endocardium surface. The LV transverse axis is then oriented along the vector $\mathbf{e}_t = (\mathbf{e}_l \wedge \mathbf{l}) / \|\mathbf{e}_l \wedge \mathbf{l}\|$, where \mathbf{l} is the segment that connects G to the RV barycentre. Finally, the LV-RV axis is oriented along the vector $\mathbf{e}_r = (\mathbf{e}_t \wedge \mathbf{e}_l) / \|\mathbf{e}_t \wedge \mathbf{e}_l\|$. It points approximately towards the middle of the septum. The vertices of the volume mesh can now be expressed in the new coordinate system $\Upsilon = (G, \mathbf{e}_l, \mathbf{e}_r, \mathbf{e}_t)$ through the transformation matrix $\mathbf{P} = (\mathbf{e}_l \mathbf{e}_r \mathbf{e}_t)$. Let \mathbf{x} be the spatial position of a mesh vertex in the global coordinate system, the new coordinates are given by:

$$\mathbf{x}_\Upsilon = \mathbf{P}^{-1}\mathbf{x} - \mathbf{P}^{-1}G$$

In the new coordinate system it is straightforward to automatically detect the apex, which is the point of lowest altitude along the LV long axis. The apex zone therefore contains all the tetrahedra whose altitude is bounded by the LV endocardium apex and the epicardium apex. The total height of the LV endocardium is then divided into three evenly spaced regions that correspond to the base, mid-plane and apical planes. We finally use the 2D angle α between the vertex position and the LV-RV axis in the $(\mathbf{e}_l, \mathbf{e}_t)$ plane \mathbf{e}_r (Figure 8.8, right panel) to subdivide each anatomical plane, clockwise, in the appropriate number of regions. The user can rotate the subdivision for final adjustments. Figure 8.8 shows an example of AHA labels.




Figure 8.8: Myocardium subdivision into AHA zones for regional personalisation. The subdivision is computed automatically from the long axis (\mathbf{e}_l), the LV-RV axis (\mathbf{e}_r) and the transverse axis (\mathbf{e}_t)

8.2.3 Model of Cardiac Fibre Directions

8.2.3.1 Available Models

The last information we need to map onto the anatomical model is the cardiac fibre architecture. Cardiac fibres play a pivotal role in the function of the heart (see [Glass et al., 1991] and references therein). The propagation of the electrical wave is faster in the direction of the fibres, the muscle contracts in the fibre direction and it is stiffer in that same direction. It is thus essential to consider the fibre architecture in electromechanical models of the heart.

The distribution of the fibres across the myocardium has been widely studied on *ex-vivo* hearts [Streeter, 1979; Glass et al., 1991; LeGrice et al., 1995; Sanchez-Quintana et al., 1996; Vetter and McCulloch, 1998]. With the advent of diffusion tensor imaging (DTI), a growing literature has been published on the 3D characterisation of the fibre orientation. Contrary to histological studies, DTI preserves the geometry of the heart, thus yielding more accurate measurements [Geerts et al., 2002; Peyrat et al., 2007; Pop et al., 2009]. However, *in-vivo* DTI acquisition of the cardiac fibre orientations remains an open challenge, despite the impressive advances achieved in the last few years in terms of MRI acquisition techniques [Wu et al., 2009]. We must therefore rely on generic models of fibre directions for the simulations.

For patient-specific simulations, a standard approach is to map a generic model of fibre architecture to the patient anatomy. One can for instance map the fibres of the Auckland canine model [LeGrice et al., 1995] or of the UCSD rabbit data [Vetter and McCulloch, 1998] to the geometry of the patient by using mesh registration techniques [Serresant et al., 2006b; Wong et al., 2010]. However, these models were built from animal specimens. Their validity in a patient is therefore questionable. An alternative is to synthesise the fibre architecture from analytical laws formulated from *ex-vivo* studies or post-mortem DTI [Streeter, 1979; Arts et al., 2001]. The

elevation angle of the fibres, i.e. their angle with respect to the short axis plane, varies almost linearly across the myocardium, from -70° on the epicardium to 0° at mid-wall to $+70^\circ$ on the endocardium [Arts et al., 2001]. Nevertheless, such a generic model does not reflect more complex variations in the fibre architecture. It has been observed for instance that the apical fibre architecture is significantly different from the basal architecture [Rijcken et al., 1999]. Recently, [Peyrat et al., 2007] computed an atlas of cardiac fibre architecture from DTI of *ex-vivo* canine hearts⁵. That atlas captures the average fibre architecture and the variability observed among a population of 9 dog hearts. Interestingly, the authors showed that the fibre directions are consistent among the individuals and, more importantly, are consistent with human hearts (comparison with a human heart specimen).

8.2.3.2 Selected Model: Analytical Fibre Model from Atlas Prior

Because *in-vivo* measurements of cardiac fibres are still unreachable in clinical routine, we rely on analytical fibre directions that are adapted to the patient anatomy and whose angles are set according to those observed in the atlas of [Peyrat, 2009]. Fibre elevation is set at the epicardium and endocardium and then linearly interpolated across the myocardium, one tetrahedron receiving one fibre direction. Figure 8.9 shows synthetic fibres in a patient-specific geometry.




Figure 8.9: Synthetic fibre directions mapped on a patient-specific geometry. Colours encode fibre directions.

8.3 Cardiac Electrophysiology Model

Once the cardiac anatomy is represented, we simulate the propagation of the electrical wave, or action potential, across the ventricular myocardium. Simulating cardiac electrophysiology is crucial as it commands the cardiac motion. Sudden deaths and heart failures are often caused by electrophysiological troubles. Less severe, cardiac

⁵Atlas available at <http://www-sop.inria.fr/asclepios/data/heart/index.php>

asynchronous motions are often the consequence of electrical blocks that prevent direct and synchronised excitation of the myocardium. Having a model of cardiac electrophysiology would enable to simulate and predict these adverse events in patients. Numerous models are available in the literature, at different biological scales and theoretical complexity. Which model to choose? Should we prefer complex and detailed models to simplified approaches? The answer to these questions is not trivial and greatly depends on the application. In fact, the answer will certainly be guided by the clinical parameters we want to predict and the availability and the quality of the clinical data used for fitting the model.

8.3.1 Available Models

Since the seminal work of [Hodgkin and Huxley, 1952; Fitzhugh, 1961; Nagumo et al., 1962], a large variety of models have been proposed to reproduce the propagation of the electrical wave across the heart muscle. The reader can find a very nice review of these models in [Clayton and Panfilov, 2008]. Three categories of models can be identified according to the scale at which the problem is approached: biophysical, phenomenological and Eikonal.

Biophysical Models These models work at the cell level to simulate the ionic interactions at the cell membranes [Noble, 1962; Luo and Rudy, 1991; Noble et al., 1998; Ten Tusscher et al., 2004]. They are very detailed models of the underlying biological phenomena, with often more than 50 parameters to control every ion channel. Sophisticated semi-linear reaction-diffusion partial differential equations (PDE) integrate the cell models at the cardiac continuum. Two major categories can be distinguished: mono-domain, which consider the myocardium as a single excitable tissue, or bidomain, which separate the intra and extra cellular domains [Coudière and Pierre, 2006; Bourgault et al., 2009].

Phenomenological Models Historically, these models were the first to be proposed [Fitzhugh, 1961; Nagumo et al., 1962]. Yet, one can see them as simplifications of the biophysical models since they work at a more macroscopic level. In particular, they describe the propagation of the action potential over the cardiac continuum directly, without focusing on the detailed ionic phenomena. Like the phenomenological models, these models can be monodomain [Aliev and Panfilov, 1996] or bidomain [Mitchell and Schaeffer, 2003]. They also rely on semi-linear reaction-diffusion PDEs but they depend on fewer parameters than the biophysical models, typically 2 to 3 parameters.

Eikonal Models These models are the simplest ones. They simulate the propagation of the electrical wave by computing its arrival time at a given point of the myocardium [Colli Franzone et al., 1990; Keener and Sneyd, 1998]. The action potential is not directly simulated. These models are governed by only one parameter and they are solved very efficiently using Fast Marching Methods [Sethian et al.,

1999]. However, because of the extreme simplifications introduced in these models, extreme phenomena like fibrillations are difficult to reproduce, although some recent experiments showed that wave reentry could be reproduced [Pernod et al., 2010].

Biophysical models are extremely complex and rely on too many parameters, too many of them relate to ionic interactions that cannot be measured in clinics. As a result, although such models are of great interest for the understanding of the organ and of complex pathological dysfunctions, they cannot be readily used for personalised simulations in clinics. Phenomenological and Eikonal models are more appropriate and choosing one or the other depends mostly on the electrical abnormalities one wants to simulate. Phenomenological models can capture some pathological patterns, but they are still computationally demanding as they rely on PDE's. Conversely, Eikonal models enable near real-time simulations. One heart beat can be simulated in few seconds in a standard computer using anisotropic multi-front Fast Marching algorithm [Sermesant et al., 2007]. However, complex pathologies like fibrillations are very challenging to simulate. Recent work demonstrated that both types of model can be personalised from clinical data ([Moreau-Villéger et al., 2006; Relan et al., 2010] for the phenomenological models, [Chinchapatnam et al., 2008] for the Eikonal models). Eikonal models require measurements of depolarisation times, which can be acquired using contact or non-contact endocardial mappings; whereas phenomenological models also require measurements of transmembrane potential measurements, which are more difficult to obtain.

8.3.2 Selected Model: Anisotropic Eikonal Model

For our simulations no clinical data about patient cardiac electrophysiology were available. We therefore had to rely on strong assumptions about model parameters. The less parameters, the less assumptions and degrees of freedom. Furthermore, we sought personalised models for intervention simulation, which requires fast models. Finally, the investigated pathologies did not present with extreme abnormalities of the cardiac electrophysiology. Eikonal models were therefore a good compromise.

Eikonal models calculate the depolarisation time T_d at each vertex of the anatomical mesh. We used the dynamic multi-front Eikonal model proposed by [Sermesant et al., 2007]:

$$v^2 (\nabla T_d^T \mathbf{D} \nabla T_d) = 1 \quad (8.1)$$

v is the local conduction velocity and \mathbf{D} is the tensor relating to the conduction anisotropy. \mathbf{D} is defined by $\mathbf{D} = \mathbf{A} \bar{\mathbf{D}} \mathbf{A}^T$, where \mathbf{A} is the matrix defining the fibre orientation in the global coordinate system and $\bar{\mathbf{D}} = \text{diag}(1, \lambda, \lambda)$. λ is the conduction anisotropy ratio between longitudinal and transverse directions. It is about 0.3 to 0.5 in human myocardium [Clayton and Panfilov, 2008; Tomlinson, 2000]. The equation is solved using the Fast Marching Method [Sethian et al., 1999], which considers each tetrahedron of the volume mesh as an automata with incremental discrete states: rest, unknown, excited. The action potential duration (APD) is

modelled by forcing an element to stay in the “excited” state a fixed period of time. The reader is referred to [Sermesant et al., 2007] for further details.

Parameters The Eikonal model is controlled by two main parameters, the anisotropy ratio λ and the conduction velocity v , which can vary locally on the myocardium. The former parameter can be assumed constant over the myocardium. We fixed it in all the experiments as its direct measurement in clinics is nowadays impossible. The latter parameter must be adjusted to capture the cardiac electrophysiology and the resulting motion of the patient heart. In addition to these model parameters, the user must define the regions from where the electrical propagation starts (Figure 8.10, left panel), the trigger time T_{d_0} when the electrical wave starts its propagation, and the action potential duration (APD). All these parameters can be set locally to finely reproduce the cardiac activity.

8.3.3 Personalisation Strategy

Recent studies on adults with heart failure demonstrated that the conduction velocity v of the Eikonal model can be estimated locally from endocardial mappings of depolarisation times [Chinchapatnam et al., 2008; Sermesant et al., 2009]. However, in our studies no electrophysiological data were available. We had to personalise the electrophysiology from the visible motion in MRI. The bundle branches are approximated by large initialisation surface zones on the left and right septum (Figure 8.10, left panel). Purkinje fibres are simulated by setting high conduction velocity on the endocardial surfaces ($v = 2000 \text{ mm.s}^{-1}$, the nominal velocity measured in the Purkinje fibres [Murgatroyd and Krahm, 2002]). The conduction velocity in the myocardium is set to 500 mm.s^{-1} but it can be changed according to the cardiac motion visible in the MRI. Each initialisation zone can be “excited” at different times to simulate bundle branch blocks. APD is as described in [Murgatroyd and Krahm, 2002], 300 ms to the free wall and 400 ms to the septum (Figure 8.10, right panel). APD can also be personalised regionally. Figure 8.11 illustrates simulated isochrones with nominal conductivity velocities and septal electrical initialisations as defined in Figure 8.10, left panel. Table 8.1 summarises the default parameters controlling cardiac electrophysiology.

8.4 Cardiac Biomechanics Model

8.4.1 Available Models

The myocardium is an active, non-linear, anisotropic visco-elastic tissue whose motion is controlled by the cardiac electrophysiology. Its constitutive law is complex, it includes an active element, which simulates the active contraction of the muscle controlled by the action potential, and a passive element, which simulates the mechanical elasticity of the tissue. In practice, the active contraction is viewed as a transient external force that makes the myocardium contract. The passive prop-

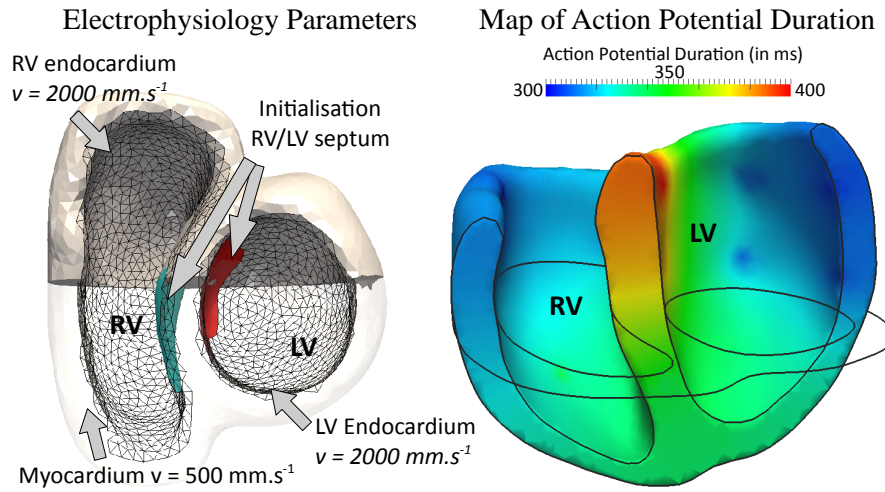


Figure 8.10: *Left panel*: Default configuration of the cardiac electrophysiology model. *Right panel*: Default action potential duration of the myocardium.




Figure 8.11: Simulated isochrones (each colour correspond to a given depolarisation time). The electrical wave is initialised on the septal endocardia, then propagated on the endocardial surfaces and through the myocardium.

erties of the tissue are internal forces that ensures realistic motions [Hunter et al., 2003b]. Three categories of active contraction models can be distinguished:

Biophysical They model the ion interactions and the actin-myosin bindings that generate the cardiac motion [Hunter and Smaill, 1988; Hunter et al., 1996, 1998; Nash and Hunter, 2000; Niederer et al., 2006; Niederer and Smith, 2008]. These very detailed models have been designed from experimental studies on *ex-vivo* animal hearts. They reproduce the biological mechanisms of sarcomere contraction, resulting in complex equations at the tissue level. These models are controlled by a large number of parameters (20 and more) related to ionic phenomena, which cannot be directly measured in patients.

Multi-Scale Phenomenological These models rely on multi-scale integrations to derive from the mechanisms at the actin-myosin bindings phenomenological equations at the organ level [Bestel et al., 2001; Chapelle et al., 2001; Caillerie et al., 2003]. They simplify the biophysical models, resulting in laws controlled by few parameters (usually 4 to 5 parameters).

Lumped These models are analytical models of the fibre contraction that do not consider spatial variability: They do not require meshes to be solved. Lumped models focus on one single myocyte, for which several laws characterise the contraction of the cell [Arts et al., 1979, 1991]. These models can be solved very efficiently but they cannot capture regional abnormalities of the myocardium in patients.

The models of the active contraction are coupled with constitutive laws of the myocardium passive properties. Here again, a large variety of models have been proposed in the last decades. A first technique consists in using transverse isotropic linear elasticity [Serresant et al., 2006b], a simple but fast elastic model. The passive law is controlled by two parameters: the Young modulus E , which controls the stiffness of the tissue, and the Poisson coefficient ν ; which controls the incompressibility of the tissue. The downside of this model however is that it is inaccurate for large deformations, artificial dilations can appear. Non-linear models are therefore preferred by modellers to simulate myocardium biomechanics. One could use non-linear extensions of linear elasticity, termed hyper-elastic models, such as the St Venant-Kirchhoff model [Picinbono et al., 2003; Delingette and Ayache, 2004]. But the standard approach is to derive the elastic stresses from non-linear stress-strain energies established from *in-vitro* experiments where slabs of myocardium tissues are stretched in several directions [Hunter and Smaill, 1988; Humphrey et al., 1990a,b; Häfner et al., 2002]. This approach has produced several well-known constitutive laws, like the pole-zero law [Hunter and Smaill, 1988] or the more recent Costa law [Costa et al., 2001]. Nowadays, the Costa law is the reference model as it has demonstrated the most accurate simulations compared to *ex-vivo* experimental data [Schmid, 2006]. A reason for this success is that it considers the entire fibre architecture: the anisotropy along the fibre direction but also the effect of the fibre sheets on the myocardium motion and in particular on its thickening. However, improving the model accuracy is achieved at the price of complexity, with increasing number of parameters. The Costa law for instance is governed by seven parameters, most of them difficult to estimate *in-vivo*.

8.4.2 Selected Model: Simplified Multi-Scale Linear Model

For the purposes of this work, very complex models are not adapted as they cannot be personalised without strong assumptions on the numerous parameters. Furthermore, the model must be solved in a reasonable amount of time to test different sets of parameters. That is why we decided to use the model developed by [Serresant et al., 2006b], which relies on a linear version of the multi-scale phenomenological

model proposed by [Bestel et al., 2001] and on linear elasticity. Despite its relative simplicity with respect to the more detailed approach, this model is able to simulate the main features of cardiac motion as observed in images of healthy subjects and patients with less severe diseases [Sermesant et al., 2006a]. It is controlled by few clinically-related parameters and is fast enough to allow personalisation from clinical data. This is exactly the level of details that are needed for our purpose. Furthermore, such modelling can constitute a first estimate of the cardiac function that can be used to personalise more complex models.

The equations of the model are solved using the Finite Element Method (FEM) on linear tetrahedrisation. The model is based on a Hill-Maxwell rheological scheme. The muscle is represented by two parallel elements: an active contractile element, which accounts for the active force controlled by the electrical command $u(t)$, and a passive elastic element, which accounts for the passive material properties (Figure 8.12, left panel).

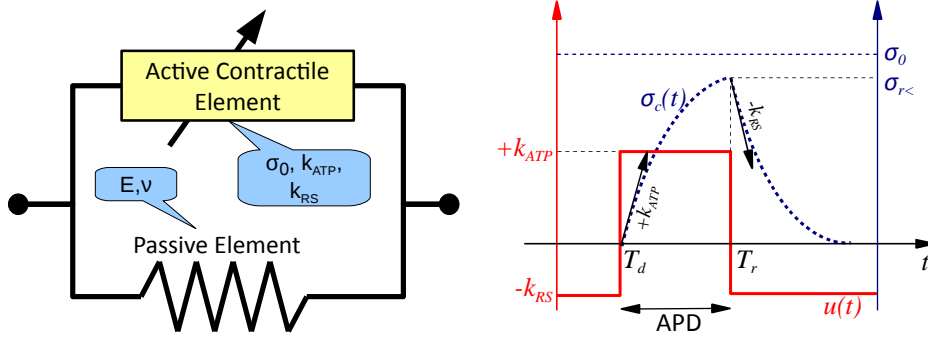


Figure 8.12: *Left panel:* Circuit diagram of the biomechanical model. In the balloons are the related parameters. *Right panel:* Variation of the active contraction stress $\sigma_c(t)$ (in blue) with respect to the electrical command $u(t)$ (in red) governed by the cardiac electrophysiology model.

The passive element is linear, anisotropic, visco-elastic. It is controlled by two parameters: the Young modulus E , i.e. the stiffness of the tissue, and the Poisson ratio ν , namely the incompressibility of the tissue. The active element is controlled by a command $u(t)$ that directly depends on the depolarisation time T_d and the repolarisation time T_r provided by the Eikonal model of the cardiac electrophysiology. $u(t)$ is a switch function that depends on the electrical state of the cell. When the cell depolarises ($T_d < t < T_r$), $u(t)$ is constant and equals $+k_{ATP}$. When the cell repolarises ($T_r < t < T_d + HP$, HP is the heart period), $u(t)$ is switched to $-k_{RS}$ (Figure 8.12, right panel). The command $u(t)$ modifies the active stress $\sigma_c(t)$ of the tetrahedral element according to the linearisation of the multi-scale model proposed by [Bestel et al., 2001]:

$$\frac{d\sigma_c(t)}{dt} + |u(t)|\sigma_c(t) = |u(t)|_+\sigma_0 \quad (8.2)$$

In this equation, σ_0 is the maximum asymptotic contraction and $|u(t)|_+$ is the positive part of the command function $u(t)$. The analytical resolution of this equation

writes:

$$\begin{cases} \text{if } T_d \leq t \leq T_r : & \sigma_c(t) = \sigma_0 [1 - e^{+k_{ATP}(T_d-t)}] \\ \text{if } T_r < t < T_d + HP : & \sigma_c(t) = \sigma_c(T_r) e^{-k_{RS}(T_r-t)} \end{cases} \quad (8.3)$$

The parameters k_{ATP} and k_{RS} are therefore directly related to the myocardium stress: they control the contraction and relaxation rates respectively (Figure 8.12, right panel). For a given tetrahedron with fibre direction \mathbf{f} (\mathbf{f} is a column vector), the active stress $\sigma_c(t)$ results in the 3D anisotropic stress tensor $\Sigma(t) = \sigma_c(t) \mathbf{f} \mathbf{f}^T$, from which we get the contraction force vector:

$$F_c = \frac{1}{4} \int_S \Sigma(t) \mathbf{n} dS \quad (8.4)$$

where S is the surface of the tetrahedron and \mathbf{n} the surface normal.

The dynamic system for all the nodes of the mesh writes:

$$M\ddot{U} + C\dot{U} + KU = F_c + F_p + F_b \quad (8.5)$$

In this equation U is the displacement vector of the mesh nodes, \dot{U} is the velocity of the nodes and \ddot{U} their acceleration. M is the diagonal mass matrix (mass lumping), C is a Rayleigh damping matrix ($C = d \cdot M$, where d is a damping constant) and K is the anisotropic linear elastic stiffness matrix. In linear elasticity, K is constant, it is built and preconditioned only once at the beginning of the simulation. The construction of K is based on the linear Hookean constitutive law between the Cauchy stress tensor and the linear strain tensor (the linear approximation of the Lagrangian strain tensor, see [Sermesant, 2003] for more details). In Equation 8.5, F_p captures the pressures applied to the endocardia during the various cardiac phases (see Section 8.5.1) and F_b accounts for the external boundary conditions as described in the next section.

Parameters The biomechanical model is controlled by 4 free parameters: the maximum contraction σ_0 , the contraction and relaxation rates k_{ATP} and k_{RS} respectively, and the Young modulus E . In addition, one can control the time delay between the cell depolarisation and the beginning of the contraction T_δ , the anisotropy ratio γ of the passive properties, the damping d , the myocardium mass density ρ and incompressibility through the Poisson ratio ν . However, these parameters were fixed in all our experiments to their nominal values: $T_\delta = 0.01s$, $\gamma = 3$ [Costa et al., 1999], $d = 3000$, $\rho = 1.07g/mL$ and $\nu = 0.48$ (near incompressibility).

8.4.3 Biomechanical Boundary Conditions

The dynamic equation (8.5) is well-defined only under boundary conditions that reflect the external conditions of the myocardium. In particular, the virtual heart must be fixed in the 3D space like the heart is fixed in the body. A standard approach is to fix the vertices of the heart base with springs whose stiffness κ is set by the user. However, this approach results in non-realistic cardiac motions: the base is fixed and the apex moves, exactly the reverse of the true cardiac motion.

Furthermore, it has been shown that such a constraint can have a significant impact on the simulations, even on simulated global parameters like end-diastole and end-systole ventricular volumes [Niederer et al., 2009]. One could use image information to get realistic base motion [Billet et al., 2009] but that approach is not suitable for prediction.

To cope with this limitation, we propose a new boundary condition that fix the pericardium of the heart while keeping the base free to contract. In a body, the heart is fixed to the neighbouring organs through the pericardium. The heart lays on top of the diaphragm, just below the ribs and in front of the lungs. It is connected to the great arteries and veins: it cannot translate nor rotate very significantly. It can only contract within the pericardial bag, facilitated by the pericardial fluid between the pericardium and the epicardium.

The idea thus consists in limiting the transverse motions of the heart outside an authorised zone while preserving contractions, slidings and twistings (Figure 8.13). Let Ω be this “free” zone. Ω is approximately the pericardium bag at its maximal extension. A rough estimation of it is obtained by artificially dilating the epicardium mask at end-diastole of about $1 - 2mm$. The heart motion is constrained to Ω by applying a contact force \mathbf{f}_s to the vertices of the epicardium. The intensity and direction of \mathbf{f}_s depends on the position of the vertex. Let P be such a vertex and \mathbf{x} its spatial position. If P is inside the free zone, it can move freely in any direction: the contact force is null $\mathbf{f}_s(\mathbf{x}) = \mathbf{0}$. If P is outside the free zone, it must be brought back inside Ω while preserving the sliding and twisting motions of the heart. One way to achieve it is to push P along the direction normal to the epicardium. The contact force $\mathbf{f}_s(\mathbf{x})$ is therefore aligned with the normal at the closest point of the boundary $\delta\Omega$ of the free zone. This direction is easily approximated at any spatial position by the gradient of a distance map $D_{\delta\Omega}$ computed from $\delta\Omega$. Furthermore, the farther P from $\delta\Omega$, the stronger $\mathbf{f}_s(\mathbf{x})$, to bring the point back into the free zone. The strength of $\mathbf{f}_s(\mathbf{x})$ is therefore controlled by the distance map $D_{\delta\Omega}$. However, to avoid numerical instabilities we upper-bound the amplitude of the contact force $\mathbf{f}_s(\mathbf{x})$ using Geman-McLure M-estimator. A parameter k is added to explicitly control the strength of \mathbf{f}_s . The contact force finally writes:

$$\begin{cases} \mathbf{f}_s(\mathbf{x}) = \mathbf{0}, & \mathbf{x} \in \Omega \\ \mathbf{f}_s(\mathbf{x}) = k \frac{D_{\delta\Omega}(\mathbf{x})^2}{D_{\delta\Omega}(\mathbf{x})^2 + m^2} \nabla D_{\delta\Omega}(\mathbf{x}), & \mathbf{x} \notin \Omega \end{cases} \quad (8.6)$$

where $\nabla D_{\delta\Omega}(\mathbf{x})$ is the gradient direction of the distance map at \mathbf{x} . Figure 8.13, left panel, shows how m and k influence the amplitude of the contact force \mathbf{f}_s . k controls the maximum amplitude of the force far from the free zone boundary. It is called *contact force amplitude* in the following. m controls how fast the contact force reaches its maximum amplitude. It is called *contact force rate*. In practice, m defines a smooth transition between the free zone and the “forbidden” area, which guarantees numerical stability. In our experiment, m and k were fixed and equal to $m = 2mm$ and $k = 10N$. Finally, the contact force \mathbf{f}_s for each node of the mesh are gathered into the boundary force vector F_b of the dynamic system Equation 8.5.

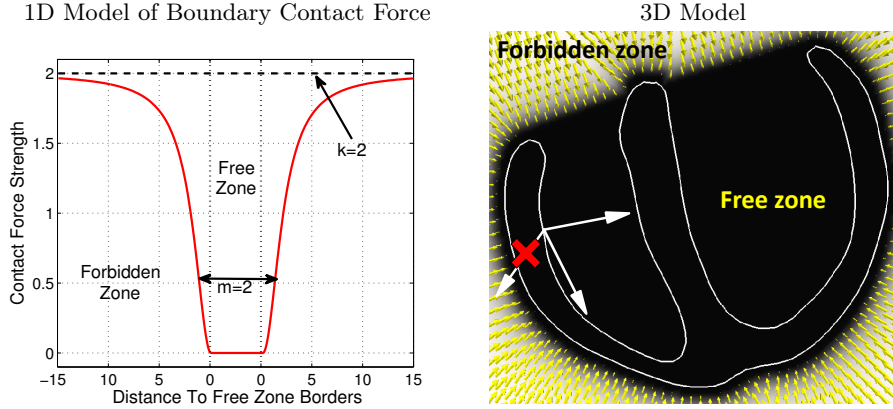


Figure 8.13: *Left panel:* 1D model of the boundary contact force. In the free zone, the force is null. As soon as we quit this zone, the strength of the contact force quickly rises according to the Geman-McLure equation to reach a user-defined maximum. *Right panel:* 3D implementation of the boundary contact force.

Figure 8.14 illustrates the improvements provided by the proposed contact boundary condition on the heart position at end-systole. As one can see, strong basis stiffness ($\kappa = 40 \text{ N.m}^{-1}$) prevents the base plane from moving towards the apex, as observed in real hearts. Furthermore, the blood pressures on the endocardial surfaces make the heart geometry at end-systole look unrealistically elongated. The contact boundary condition alleviates these issues. The base is now free to move towards the apex (a slight basis stiffness was kept to simulate atrial and arterial tensions on the ventricles, $\kappa = 4 \text{ N.m}^{-1}$). Furthermore, the contact force preserves the heart shape by avoiding unrealistic elongations. Finally, myocardium thickening improves as a consequence of the basal contraction.

It is worth stressing on the fact that the proposed boundary contact force is not a model of the pericardium bag but rather a spatial boundary condition that fixes the heart in the 3D space. An appropriate model of the pericardium, seen as a pressure field that applies on the epicardium and whose strength depends on the variation of volume of the pericardial bag is under study [Freeman and LeWinter, 1984]. However, these models introduce new parameters that need to be estimated.

Figure 8.15 illustrates a generic simulation on a patient-specific anatomy but with nominal parameters [Bestel, 2000; Shi and Liu, 2003; Sermesant et al., 2006b] (Table 8.1). The biomechanical model is triggered by the electrical command illustrated in Figure 8.10, with the APD showed in Figure 8.10. Figure 8.16 reports the simulated volume and pressures variations in the blood pool and the main arteries. PV loops can be obtained, which, after personalisation, can be used to extract parameters quantifying the cardiac pump efficacy.

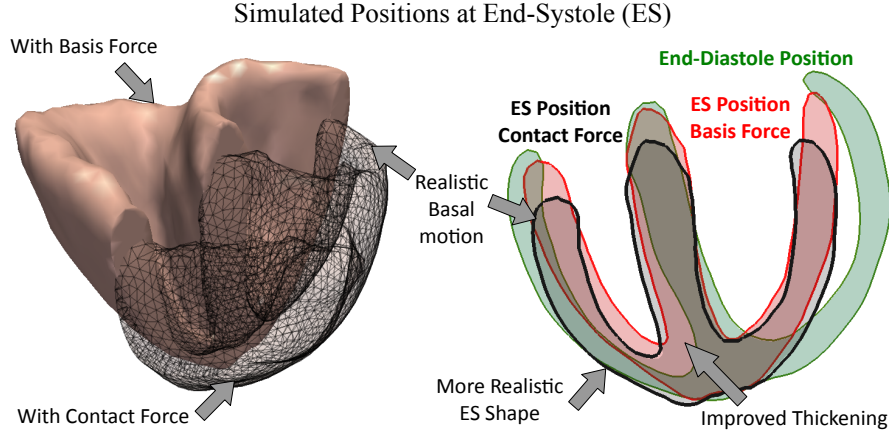


Figure 8.14: Effects of the contact force over the simulated cardiac motion compared to basis boundary conditions only. Strong basis constraints prevent the base plane to contract towards the apex and can lead to “elongated” end-systole (ES) shapes. The contact boundary condition alleviates these limitations: the base contracts downwards, the shape is preserved and myocardium thickening is improved.

8.4.4 Personalisation Strategies

The automatic calibration of the biomechanical parameters remains an open challenge despite the increasing research interest in this direction [Shi and Liu, 2003; Sermesant et al., 2006a; Wong et al., 2007; Billet et al., 2008; Moireau et al., 2008; Billet et al., 2009; Zhang et al., 2010a]. One of the key of success is to use models especially tailored to the sparse clinical data and to the clinical question to investigate. We choose to use simpler model in order to have less but clinically-meaningful parameters.

The personalised simulations performed during this thesis were calibrated manually through trial and errors. This was possible since the simulation of the whole cardiac cycle took about 15-30 minutes only on a standard 2.4GHz Intel Core 2 Duo computer with 4GB of memory, on meshes of about 50000 elements. The manual calibration was performed as follows.

Assume we have dynamic 3D images of the beating heart (3D ultrasound, CT or MRI), on which the heart has been segmented over the entire cardiac cycle, using approaches like the one presented in Chapter 3. The dynamic segmentations quantify the variation of ventricular volumes over time. The dynamic images show the apparent motion of the heart. We also assume that the models of cardiac electrophysiology and hemodynamics have been fitted to the patient physiology (see Section 8.3 and Section 8.5). We first start with nominal parameter values as reported in the literature ($\sigma_0 = 100 \text{ kPa.mm}^{-2}$; $k_{ATP} = 10 \text{ s}^{-1}$; $k_{RS} = -20 \text{ s}^{-1}$; $E = 50 \text{ MPa}$ [Bestel, 2000; Shi and Liu, 2003; Sermesant et al., 2006b], see Table 8.1). The contractile element of the lesions, if any, is decreased or disabled to reproduce their abnormal motion. The maximum contraction σ_0 is then set iteratively such that the simu-

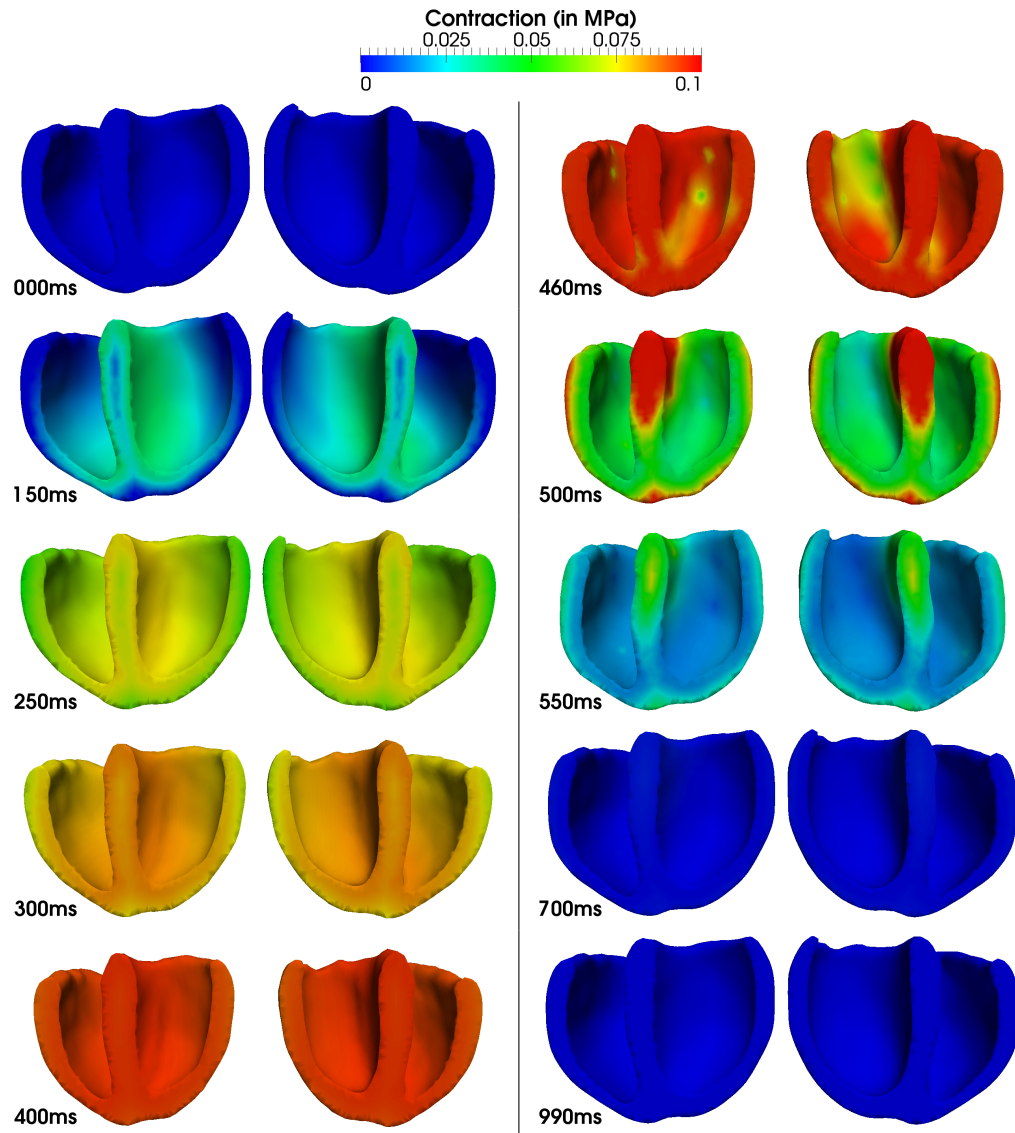


Figure 8.15: Example of biomechanical simulation. Colours encode the active contraction triggered by an electrical command simulated using the Eikonal model. This motion is simulated on a pathological anatomy but with nominal parameters (Table 8.1).

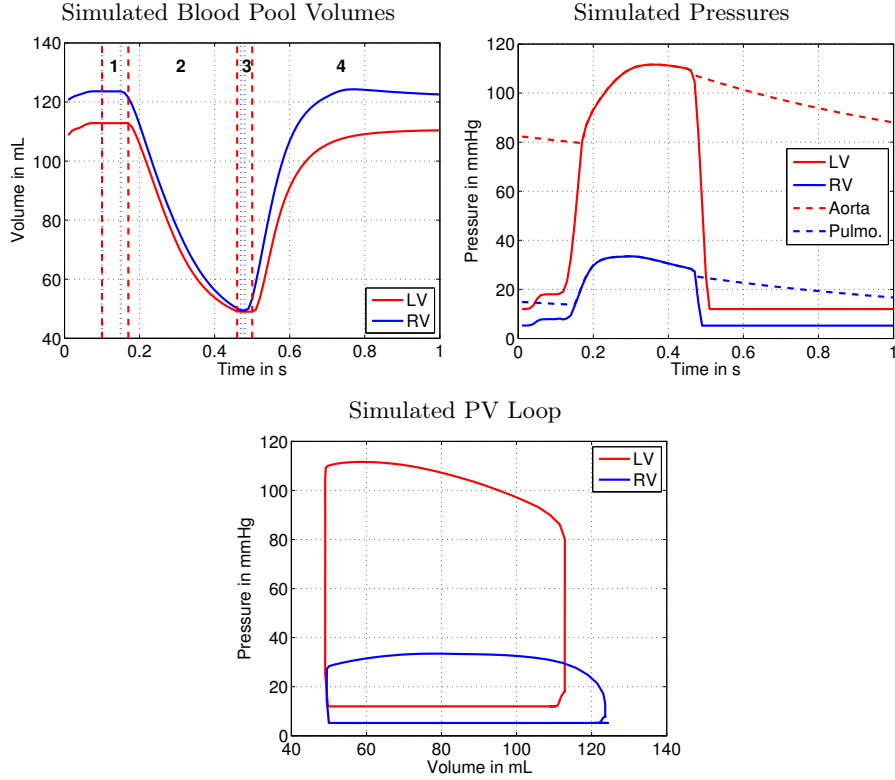


Figure 8.16: Quantitative indices of the left (LV) and right (RV) ventricles corresponding to the simulation of Figure 8.15. On the volume curves, vertical bars denote the cardiac phases: 1) Isovolumetric contraction, 2) Ejection, 3) Isovolumetric relaxation, 4) Ejection. Simulated ejection fractions: 57% (LV) and 60% (RV)

lated end-systole volumes match the measured ones. The Young modulus E is set such that the simulated cardiac relaxation matches the observed one, as the motion of this cardiac phase is mostly due to the relaxation of cardiac myocytes. Finally, the contraction and relaxation rates k_{ATP} and k_{RS} are set such that the slopes of volume variation during systole and diastole respectively match.

8.5 Cardiac Hemodynamics Model

The myocardium function largely depends on external conditions that are mostly determined by the blood flows and the circulatory system. Cardiac hemodynamics must therefore be considered in the model. A standard approach is to simulate the blood flow with fluid-dynamic models (CFD) coupled with the myocardium through fluid-structure interaction (FSI) models. This approach is very detailed but also tremendously complex to solve, with coupled systems controlled by large set of parameters (see [McQueen and Peskin, 2000; Taylor and Draney, 2004; Mihalef et al., 2009] and references therein). An alternative strategy simplifies the problem

by modelling the blood flow as pressure constraints of the biomechanical model. The constraints are computed from lumped models of the circulatory system [Keener and Sneyd, 1998; Arts et al., 2005; Kerckhoffs et al., 2007] or directly input by the user [Sainte-Marie et al., 2006; Sermesant et al., 2006b]. These simplified approaches greatly reduces the computational complexity of the models and decrease the number of parameters to set. However, these models completely ignore the flow patterns in the ventricles, patterns that may have long-term impact in the cardiac function or on the effects of a therapy.

Again, the choice of a strategy depends on the clinical question we want to investigate, the clinical parameters to predict and the available data for the personalisation. If the focus is on the flow patterns through valves or arteries and that nice flow data is available (3D US doppler or more recently 4D flow MRI), then a complete model would be of interest, as in [Mihalef et al., 2009]. Conversely, if only global parameters like ventricular pressure or volume variation must be predicted and if few flow data are available, then simpler model could be more appropriate. In this work, we were mainly interested in predicting the cardiac motion and the global output of the heart, as ejection fractions, volume variations and global pressure changes. We thus decided to use a simple model where hemodynamics is considered as a constraint of the electromechanical model. In particular, we improved the cardiac cycle model proposed by [Sermesant et al., 2006a], briefly described in Section 8.5.1, to simulate the Windkessel effect of arteries (Section 8.5.2) and valve regurgitations (Section 8.5.3).

8.5.1 Simulation of the Cardiac Cycle

8.5.1.1 A Rule-Based Valve Model

We simulate the four cardiac phases – filling, isovolumetric contraction, ejection, isovolumetric relaxation – (Figure 8.17) using a rule-based approach [Sermesant et al., 2006a]. In simple words, the phases are simulated by alternating different boundary conditions of the electromechanical model.

Filling During filling, the ventricular pressure is close to the atrial pressure. A pressure force \mathbf{f}_p whose amplitude equals the atrial pressure is thus applied to the endocardial surface to make the ventricle dilate. The atrial contraction triggered by the P-wave is simulated by artificially increasing the atrial pressure (Figure 8.17A).

Isovolumetric Contraction As soon as the ventricle starts contracting, the cavity volume starts decreasing. The ventricular flow Φ changes sign (positive during dilation, negative during contraction) and the atrial valves in healthy hearts close. We thus switch from filling to isovolumetric contraction by detecting the change of flow sign. During this short phase the cavity volume is constant since all the valves are closed (the ventricular pressure is still lower than the arterial pressure). This constraint is implemented using a Lagrangian approach [Billet, 2010]. At each time step, a penalty pressure \mathbf{f}_p is applied to

the endocardial vertices such that the ventricular volume stays constant. The amplitude of \mathbf{f}_p is exactly related to the ventricular pressure, which raises as the ventricle contracts (the higher the contraction strength, the higher the penalty) (Figure 8.17B).

Ejection Once the ventricular pressure is high enough to open the arterial valves, the blood can flow into the arteries. The ventricular pressure equals the arterial pressure, which is simulated using a Windkessel model (see Section 8.5.2). We apply a pressure force \mathbf{f}_p to the endocardium whose amplitude equals the arterial pressure (Figure 8.17C).

Isovolumetric Relaxation When the cardiac cells repolarises and the ventricle relaxes, the cavity volume increases again: the ventricular flow Φ changes sign (from negative to positive). This would close the arterial valves: we enter in the isovolumetric relaxation. Because the ventricular pressure is still higher than the atrial pressure, all the valves are closed and the ventricular volume stays constant despite myocardium relaxation. The same strategy as the one used during the isovolumetric contraction is applied (Figure 8.17D).




Figure 8.17: The four cardiac phases

The phase model is therefore controlled by both pressures (to exit isovolumetric phases) and blood flows (to enter isovolumetric phases), which are computed as the variation of the cavity volume V per time step Δt : $\Phi(t) = [V(t) - V(t - \Delta t)]/\Delta t$. Denoting $p(t)$ the ventricular pressure, $p_{at}(t)$ the atrial pressure and $p_{ar}(t)$ the arterial pressure, we have the following rules:

- $\Phi(t - \Delta t) > 0$ and $\Phi(t) \leq 0$: switch from filling to isovolumetric contraction.
- $p(t) \geq p_{ar}(t)$: switch from isovolumetric contraction to ejection.
- $\Phi(t - \Delta t) < 0$ and $\Phi(t) \geq 0$: switch from ejection to isovolumetric relaxation.
- $p(t) \leq p_{at}(t)$: switch from isovolumetric relaxation to filling.

During the phases, the pressure force applied to the endocardial surfaces is computed for each vertices, \mathbf{f}_p , and gathered into the global pressure vector F_p of the dynamic systems (Equation 8.5).

In this model, the phases are handled independently for the left and right ventricles to simulate asynchronous cardiac motions. No constraint on the cardiac synchrony is employed. However, the phases are alternated sequentially. We cannot come back to the previous phase. Although this seems quite reasonable for normal cases, extreme pathologies with very impaired myocardium function can lead to transient abnormal valve closure/opening, which cannot be modelled by our algorithm.

Parameters This model of the cardiac cycle is controlled by the following parameters:

- Left and right atrial baseline pressures p_{Lat} and p_{Rat} . These pressures correspond to the pressure in the atria at rest.
- Left and right atrial pressures multiplicative coefficients c_{Lat} and c_{Rat} . When the atria are excited by the electrical wave, they contract and the atrial pressures increase. These coefficients are used to calculate the atrial pressures at highest contraction, when the valves open
- Initial aorta and pulmonary artery pressures at time $t = 0$, $p_{0_{aorta}}$ and $p_{0_{pulmo}}$. These pressures evolve over time according to the Windkessel model, as described in Section 8.5.2.

8.5.1.2 Personalisation Strategy

These parameters can be measured in clinical environment. In routine, ultrasound Doppler imaging can provide a rough estimate of the peak pressures in the atria and arteries when regurgitation are present. This is achieved using Bernoulli's principle between flow speed and pressure [Yock and Popp, 1984]. More accurate measurements can be obtained using catheter pressure probes, but these procedures are invasive and employed only during some interventions. These measurements are directly input into the model as constraints. However, this approach is limited when the aim is to predict the cardiac function after a therapy. What pressures should we use then? How will they change? Immediate postoperative effects could be simulated by assuming no or little changes in atrial and arterial pressures (according to the type of therapy of course). This approximation resulted in promising results [Mansi et al., 2009a; Sermesant et al., 2009], but long-term effects require more sophisticated models.

8.5.2 Simulation of the Arterial Windkessel Effect

8.5.2.1 Available Models

As mentioned in Chapter 2, the arterial pressures vary during ejection to accommodate the cardiac pulse. The arterial walls dilate, gathering the blood that cannot be ejected to the circulatory system, which has finite capacity. When the valves close, the arteries relax and the excess of blood is released to the system. This is

the Windkessel effect (WK) [Frank, 1899]. Several WK models have been developed, from the most simplified lumped models of the arteries to the more complex travelling-wave theory [Crépeau and Sorine, 2007] (see [Westerhof et al., 2009] and reference therein). In this work, we improved the current electromechanical model to reproduce the WK effect using the popular 3-element model [Westerhof et al., 1971; Stergiopoulos et al., 1999; Serresant et al., 2006b], as this model has shown good fits to real curves.

8.5.2.2 Selected Model: 3-Element Windkessel Model

The 3-element WK model is derived from electrical circuit analogies where the blood flow is the current and the arterial pressure is the voltage. First introduced by [Westerhof et al., 1971], this model was designed as an extension of the original 2-element model [Frank, 1899] to faithfully recover the variations of the arterial pressures during the cardiac cycle. More advanced models have been proposed since then, in particular the 4-element WK model, to increase the accuracy of the fitted parameters but the 3-element model remains widely used thanks to its simplicity, its small number of parameters to set and its ability to reproduce realistic pressure patterns.

Figure 8.18 shows the circuit diagram of the 3-element WK model. The first element of the model is a peripheral resistance R_p , which accounts for the distal resistance of the circulatory system mainly due to the small vessels. The compliance C accounts for the elasticity of the arterial walls whereas the characteristic resistance R_c accounts for the blood mass and for the compliance of the artery proximal to the valves. Let $\Phi_{ar}(t)$ be the arterial flow at time t , defined as the opposite of the ventricular flow $\Phi_{ar}(t) = -\Phi(t)$, $p_{ar}(t)$ be the arterial pressure at time t and p_r be a constant low pressure of reference (typically the pressure of the remote venous system). When the blood flows into the arteries ($\Phi_{ar}(t) > 0$), during ejection, the 3-element model writes:

$$\frac{dp_{ar}(t)}{dt} = R_c \frac{d\Phi_{ar}(t)}{dt} + \left(1 + \frac{R_c}{R_p}\right) \frac{\Phi_{ar}(t)}{C} - \frac{p_{ar}(t) - p_r}{R_p C} \quad (8.7)$$

When the arterial valves are closed, the blood flow is stopped ($\Phi_{ar}(t) = 0$) and the model writes:

$$\frac{dp_{ar}(t)}{dt} = -\frac{p_{ar}(t) - p_r}{R_p C} \quad (8.8)$$

These equations are integrated using first-order implicit schemes. Two independent WK models are used for the aorta and the pulmonary artery as we are mainly interested in phenomena that occur at the heart-beat scale. Long-term interactions between the pulmonary and systemic systems are not considered. Figure 8.19 shows simulated aortic pressures given an input flow and standard WK parameters ($p_r = 82 \text{ mmHg}$, $C = 1.6 \text{ mL/mmHg}$, $R_p = 0.5 \text{ mmHg/mL/s}$, $R_c = 0.03 \text{ mmHg/mL/s}$).




Figure 8.18: Circuit analogy of the 3-element Windkessel model




Figure 8.19: Example of simulated pressure curves in the aorta according to input flows (simulated using the electromechanical model).

Parameters The pressure of each artery (aorta and pulmonary artery) is controlled by four parameters:

- Artery compliance, C : elasticity of the arterial wall.
- Peripheral resistance, R_p : distal resistance of the circulatory system mainly due to small vessels.
- Characteristic resistance, R_c : proximal compliance and local inertia.
- Remote pressure, p_r : constant distal pressure in the remote venous system.

These parameters must be set independently for the aorta and the pulmonary artery.

8.5.2.3 Personalisation Strategies

Several studies reported in the literature have shown that calibrated 3-element WK models can faithfully reproduce and predict clinical measurements ([[Stergiopoulos et al., 1995](#); [Segers et al., 2008](#); [Westerhof et al., 2009](#)] and references therein), although the compliance C may be overestimated and the characteristic resistance

R_c underestimated [Stergiopoulos et al., 1999]. This model is therefore well-suited if one seeks to reproduce the curves without interpreting the values of the parameters. Despite this limitation, the 3-element WK has received a great interest from the community as it can be easily personalised from pressure and volume curves. When such data are available, a large variety of methods can be employed, based on exponential fitting or parameter estimation methods [Lankhaar et al., 2006; Westerhof et al., 2009]. However, when no pressure curves are available (as during this thesis), the calibration becomes more complex.

Studies have shown that the time constant $\tau = R_p C$, which characterises the pressure decay when no blood is flowing in the artery, is relatively constant among individuals ($\tau \approx 0.80s$ [Lankhaar et al., 2006]). We can thus fix τ , and start from nominal values reported in the literature to adjust the WK parameters. Figure 8.20 shows the effect of simultaneously increasing R_p and decreasing C while keeping the time variable τ constant. Intuitively, this consists in increasing the pressure in the small vessels (hypertension) while making the artery more elastic. In this case, the peak systolic pressure also increases but the whole curve is shifted upwards, making the calibration difficult when only the pressure at peak systole is available. We then fixed R_p and C and made the characteristic resistance R_c vary. We can see from Figure 8.21 that in this case, only the peak systole pressure is increased and the overall shape of the pressure curve is preserved. Calibrating the model with this parameter only is thus much easier. This is the strategy employed during this thesis.




Figure 8.20: Variations of the simulated aorta pressure with respect to the peripheral resistance R_p and artery compliance C . The time constant $\tau = R_p C = 0.80 ms$, R_c and p_r are kept constant. Increasing R_p and decreasing C shift the curves towards high pressures (C in $mL/mmHg$ and R_p in $mmHg/mL/s$).




Figure 8.21: Variations of the simulated aorta pressure with respect to the characteristic resistance R_c . Remaining WK parameters are kept constant. Increasing R_c increases the pressure at peak systole but does not shift the curve. R_c can be used to adjust the WK model when no pressure curves are available (R_c in $mmHg/mL/s$).

8.5.3 Simulation of Valve Regurgitations

8.5.3.1 Available Models

During this thesis, we had to simulate the cardiac function of patients with repaired tetralogy of Fallot. This disease is characterised by chronic pulmonary valve regurgitations, in addition to extreme dilation of the right ventricle and impaired cardiac function. In simple words, the blood can flow backwards from the pulmonary artery to the ventricles, which is normally prevented by the pulmonary valves. We thus needed a model of valve regurgitations.

The vast majority of regurgitation models rely on fluid-dynamics models and fluid-structure interactions to simulate valves, their kinematics and the blood flow across them [Formaggia et al., 2001; Taylor and Draney, 2004; Gerbeau et al., 2005; Korakianitis and Shi, 2006; Wenk et al., 2010]. These approaches are ideal for studying the flow patterns and the causes and effects of regurgitations. However, linking these models with electromechanical models of the heart is very complex as the two models are tightly coupled [McQueen and Peskin, 2000; Tang et al., 2007b; Yang et al., 2008a]. Furthermore, adding more models into the simulation adds complexity and parameters, which then need to be personalised. We therefore choose to use a simplified model of regurgitations compatible with our electromechanical model and with the available data. Measured regurgitation flows are directly used as constraints.

8.5.3.2 Selected Model: Regurgitations as Hemodynamics Constraints

Regurgitations may be translated in our model by isovolumetric phases that are no longer isovolumetric. Indeed, during the ejection (*resp.* filling) phase, the variation of ventricle volumes is the direct consequence of the myocardium contraction (*resp.* relaxation). We do not have access to the amount of blood that passes through the valves. However, during the isovolumetric phases, we know that the ventricle volumes should be constant. When valves regurgitate, this volume changes, according to the regurgitation flow that depends on the level of valve leakage. This flow is measured in clinics using flow imaging (ultrasound Doppler imaging or flow MRI). We can therefore include this data directly in our model by relieving the isovolumetric constraints.

The isovolumetric phases are modified as follows. Let Φ_c and Φ_r be the regurgitation flows measured at isovolumetric contraction and relaxation respectively. At each instant t , we first estimate the infinitesimal volume variation ΔV that we would have without isovolumetric constraint. Then:

- If $|\Delta V| > |\Phi_{\{c,r\}} \Delta t|$, a penalty constraint is applied to each vertex of the endocardium such that the resulting ventricle volume varies of $\Delta V = \Phi_{\{c,r\}} \Delta t$. The effect of myocardium contraction or relaxation is therefore partially compensated.
- Otherwise, no penalty constraint is applied. The myocardium can contract/relax freely, all the variation of volume is authorised.

With this simple model, the blood pool volume can change during the isovolumetric phases according to the measured regurgitation flows, as illustrated in Figure 8.22.




Figure 8.22: Simulated pressure-volume loops of the right ventricle without (*left panel*) and with (*right panel*) regurgitations. Regurgitations are implemented by allowing the cavity volume to vary during the isovolumetric phases.

Parameters Regurgitations are controlled by two parameters per ventricle:

- The regurgitation flow during isovolumetric contraction Φ_c
- The regurgitation flow during isovolumetric relaxation Φ_r

8.5.3.3 Personalisation Strategies

An approximate estimation of the regurgitation flows can be obtained from ultrasound Doppler imaging or flow MRI. As the isovolumetric phases are very brief, we assume that the regurgitation flows during these phases are constant. We can thus estimate Φ_c and Φ_r by dividing the average regurgitant volume V_{reg} measured during the shortest period of time t_{iso} that includes the isovolumetric phases by t_{iso} : $\Phi_{\{c,r\}} = V_{reg}/t_{iso}$

8.6 Discussion

This chapter described a framework for personalised simulations of the cardiac function whose complexity is compatible with the available clinical data. Relying on simplified models, the framework is computationally efficient and controlled by few parameters, all having clinical meaning (Table 8.1). The proposed framework has been especially designed for simulating the cardiac function of patient with repaired tetralogy of Fallot, as described in Chapter 9. However, each element can be adapted individually to simulate other pathologies, like in [Sermesant et al., 2009].

Modelling the cardiac function is a complex task and the pipeline presented in this chapter is the result of incremental improvements on existing algorithms. However, contributions were brought to the original model [Sermesant et al., 2006b] to adapt it to our specific problem. In particular:

1. A pipeline for the automatic creation of anatomical models of the heart from clinical images was developed.
2. Implementation of a boundary contact force to constrain the heart in the 3D space while keeping the base free to contract as in the real cardiac motion.
3. Implementation of the Windkessel model to simulate arterial pressures
4. Implementation of a regurgitation model that, although simplified, provided realistic simulations of the cardiac function in tetralogy of Fallot (Chapter 9)
5. Design of personalisation strategies for the manual calibration of the biomechanical parameters, yielding first personalised simulations of the cardiac function in tetralogy of Fallot (Chapter 9)

The simplifications introduced in the model to make it “personalisable” are its strengths but also its weaknesses in contrast with more detailed models. These limitations open several perspectives of research for finding computationally efficient solutions still controlled by few clinically-relevant parameters.

For the anatomical model, the most important limitation is the definition of the fibres, which is still generic in our framework. It is indeed not clear how well a generic model of fibre architecture can fit the anatomy of a patient. For instance, fibres surrounding myocardial scars reorient to limit the loss in cardiac output [Geerts-Ossevoort, 2002; Wu et al., 2007]. To cope with this uncertainty, one could model the fibre reorientation [Kroon et al., 2009]. However, such models consider only one remodelling stimulus at a time (homogeneous myocardium stress-strain in [Rijcken et al., 1999; Kroon et al., 2009]) and introduce new parameters to set. The solution will most probably come from the recent advances in *in-vivo* DTI acquisition and post-processing [Wu et al., 2009; Toussaint et al., 2010], which promise to provide patient-specific data that we could use in the model. In parallel, an increasing literature reports that fibre orientations have little influence on global parameters such as ventricular volume or pressures (see [Niederer et al., 2009] and references therein), although their effect on local indices such as strain may be significant. The question of which parameter to predict is therefore crucial to evaluate the reliability of the simulation given a set of fibre orientations. At last, we do not consider the cardiac fibre sheets, although it has been shown that the spatial arrangement of these sheets plays an important role in the myocardium thickening [Costa et al., 1999]. However, integrating this additional constraint is achieved at the price of increased computational complexity, with more complex constitutive laws. Furthermore, fibre sheets are even more difficult to acquire *in-vivo* and statistical analyses on a population of dogs showed that their variations among individual could not be neglected [Peyrat et al., 2007]. Because of these difficulties we decided not to consider them.

Simulating the cardiac electrophysiology using Eikonal models is very appealing as they are fast and controlled by very few parameters. However, complex pathological patterns cannot be simulated with these models. The patients studied during this thesis did not report any major abnormality in their cardiac electrophysiology, which justified this choice. The most severe cases were affected by right bundle branch block, which can be simulated by the Eikonal model. However, in pathologies where tachycardia or fibrillation happens, more sophisticated models are necessary. Mono-domain approaches are good candidates [Aliev and Panfilov, 1996; Mitchell and Schaeffer, 2003] and can be integrated easily in the model with the price of additional computational complexity.

The biomechanical model we employed is transverse isotropic linear elastic with linear electromechanical coupling. This is probably the coarsest simplification of the framework. Rigid body rotation of linear elastic models results in wrong positive strain. Moreover, the myocardium is a non-linear elastic tissue, with non-linear electromechanical coupling. Yet, despite that simplification, realistic cardiac motion was obtained (see Chapter 9, [Serresant et al., 2006b, 2009; Billet et al., 2009]). More advanced models could be used, in particular the reference Costa model of the passive myocardium [Costa et al., 2001] coupled with the complete Bestel-Clement-Sorine multi-scale model of cardiac contraction [Bestel et al., 2001]. Solving these models is more demanding but recent approaches pave the way to efficient implementations of non-linear tissue models [Delingette, 2008b,a]. Benchmark studies

between linear and non-linear model, in the framework of personalised models for clinical prediction, would be of great interest, effectively quantifying the loss in accuracy of the linear models in contrast with the increase in computational complexity. This would enable to implement hybrid models mixing non-linear and linear passive models of the myocardium tissue depending on the length of the deformations.

The biomechanical model was personalised manually, by trial and errors. This approach was possible on few cases but is clearly not a definite solution to the personalisation problem, in particular when the number of cases to treat is large. Automatic parameter estimation is a very intense research topic, challenged by the very large number of degrees of freedom to estimate. Yet, the small number of parameters enable to use variational methods [Billet, 2010], filtering methods [Moireau et al., 2008, 2009; Wong et al., 2010], direct methods such as Powel NEWUOA [Powell, 2008] or more exhaustive approaches based on genetic algorithms. Experiments in these directions are in progress.

Finally, the model of cardiac hemodynamics can also be improved. The current implementation of the cardiac cycle does not allow the simulation of complex regurgitation patterns for instance nor severe valve abnormalities due to extreme cardiac asynchrony. The opening and closure of the valves should be done based on the ventricular pressures only as in [Sainte-Marie et al., 2006; Niederer et al., 2009]. Current work aims to implement a pressure-based model of the cardiac phases. Besides, the current model had also a direct consequence on the modelling of the valve regurgitations, which can be done only during the isovolumetric phases. We chose to use a constraint approach for easy personalisation. This was appropriate for our problem since the simulated therapies consisted in stopping the regurgitations. However, this technique is not suitable for predicting the regurgitations after cardiac remodelling for instance. Which regurgitation flow should we use? A solution would consist in using a 0D lumped model of regurgitations, where the valves are modelled as a resistance to the blood flow. Valve closure would be modelled by an infinite resistance. However, such an approach is possible only if the cardiac cycle is pressure controlled rather than flow controlled. Other more sophisticated improvements would consist in plugging a lumped model of the cardiovascular system [Keener and Sneyd, 1998] to our electromechanical model of the heart, similar to [Arts et al., 2005; Kerckhoffs et al., 2007]. Or, more interestingly, we can couple the heart model with a personalised fluid-dynamics model of the blood flow across the valves [McQueen and Peskin, 2000; Gerbeau et al., 2005; Nordsletten et al., 2009].

Table 8.1: Parameter list of the electromechanical model (1/2). All the parameters can be set for specific regions of the myocardium for regional personalisation. Anisotropy ratios are for fibre/cross-fibre ratios. *DTI*: Diffusion tensor imaging, *ECG*: Electrocardiogram, *tMRI*: tagged MRI.

Notation	Parameters	Clinical Data for Personalisation	Default Values	References
Anatomical Model				
\times e	Heart geometry Myocyte elevation angle	Anatomical images <i>in-vivo</i> DTI	\times $e_{endo} = +70^\circ$, $e_{epi} = -70^\circ$	\times [Arts et al., 2001]
Electrophysiology Model				
T_{d_0}	Trigger time	ECG	$T_{d_0} = 100 \text{ ms}$	[Murgatroyd and Krahn, 2002]
v	Electrical conductivity	Endocardial mapping	$v_{endo} = 2000 \text{ mm.s}^{-1}$, $v_{myo} = 500 \text{ mm.s}^{-1}$	[Murgatroyd and Krahn, 2002]
λ	Conductivity anisotropy	\times	3	[Keener and Sneyd, 1998]
APD	Action potential duration	Endocardial mapping, ECG	$APD = 300 \text{ ms}$	[Murgatroyd and Krahn, 2002]
Δt_{el}	Electrical time step	\times	$\Delta t_{el} = 1 \text{ ms}$	\times
Biomechanical Model				
ρ	Muscle mass density	\times	1070 g/L	[Glass et al., 1991]
E	Young modulus	Dynamic images, tMRI, etc.	$E = 50 \text{ MPa}$	[Shi and Liu, 2003]
γ	Elasticity anisotropy	\times	3	[Costa et al., 1999]
ν	Poisson coefficient	\times	0.48	[Glass et al., 1991]
σ_0	Maximum contraction	Dynamic images, tMRI, etc.	$\sigma_0 = 100 \text{ kPa}$	[Bestel, 2000]
k_{ATP}	Contraction rate	Dynamic images, tMRI, etc.	$k_{ATP} = 10 \text{ s}^{-1}$	[Bestel, 2000]
k_{RS}	Relaxation rate	Dynamic images, tMRI, etc.	$k_{RS} = -20 \text{ s}^{-1}$	[Bestel, 2000]
d	Damping	\times	$d = 3000 \text{ N.m.g}^{-1}$	\times
Δt_{mec}	Mechanical time step	\times	$\Delta t_{mec} = 1 \text{ ms}$	\times
Boundary Conditions				
κ	Base stiffness	Dynamic images	4000 N.m^{-1}	\times
k	Contact force amplitude	\times	10 N	\times
m	Contact force rate	\times	2	\times

Table 8.2: Parameter list of the electromechanical model (2/2). *Art*: arterial, *US*: ultrasound.

Notation	Parameters	Clinical Data for Personalisation	Default Values	References
Hemodynamics Model				
p_{at}	Atrial pressures	Pressure probes, echocardiography	$p_{Lat} = 12 \text{ mmHg}$, $p_{Rat} = 5 \text{ mmHg}$	[Schäffler and Schmidt, 1999]
c_{at}	Atrial contraction coefficient	Pressure probes	$c_{Lat} = c_{Rat} = 1.5$	×
p_0	Initial art. pressures	Pressure probes, echocardiography	$p_{0aorta} = 80 \text{ mmHg}$, $p_{0pulmo} = 15 \text{ mmHg}$	[Schäffler and Schmidt, 1999]
C	Art. compliance	Pressure probes + ventricle volumes	$C_{aorta} = 1.6 \text{ mL/mmHg}$, $C_{pulmo} = 2.1 \text{ mL/mmHg}$	[Westerhof et al., 2009]
R_p	Art. peripheral resistance	Pressure probes + ventricle volumes	$R_{paorta} = 0.5 \text{ mmHg/mL/s}$, $R_{ppulmo} = 0.4 \text{ mmHg/mL/s}$	[Westerhof et al., 2009]
R_c	Art. characteristic resistance	Pressure probes + ventricle volumes	$R_{caorta} = 0.03 \text{ mmHg/mL/s}$, $R_{cpulmo} = 0.015 \text{ mmHg/mL/s}$	[Westerhof et al., 2009]
p_r	Remote art. pressure	Pressure probes + ventricle volumes	$p_{raorta} = 60 \text{ mmHg}$, $p_{rpulmo} = 4 \text{ mmHg}$	×
Φ_c	Systolic regurgitation flows	Flow imaging (MRI, US)	$\Phi_c = 0 \text{ mL.s}^{-1}$	×
Φ_r	Diastolic regurgitation flows	Flow imaging (MRI, US)	$\Phi_r = 0 \text{ mL.s}^{-1}$	×

In-Silico Pulmonary Valve Replacement in ToF

Contents

9.1	Motivation	193
9.2	Real-Time Simulation of Cardiac Volume Reduction	194
9.2.1	Tissue Resection	195
9.2.2	Tissue Attachment	196
9.2.3	Myocardium Fibre Recovery	196
9.3	Personalised Simulation of Pulmonary Valve Replacement	197
9.3.1	Patient Selection and Data Preparation	197
9.3.2	Personalised Simulation of the Preoperative Cardiac Function	198
9.3.3	Personalised Simulation of PVR in ToF Patients	201
9.4	Discussion	203

We described in Chapter 8 the tools to realise personalised simulations of the beating heart in patients with repaired tetralogy of Fallot (ToF). The idea is to use the virtual heart to test the postoperative effects of different pulmonary valve replacement (PVR) strategies on the cardiac function of the patients. We investigate in the following the feasibility of that approach on two young adults.

9.1 Motivation: Can we Predict the Postoperative Effects of Pulmonary Valve Replacement?

In this chapter we are in the situation where the cardiologist decided to replace the pulmonary valves to a patient. This decision was made based on the current cardiac function, estimated using among other indices the iLogDemons for instance, and on a bad prognosis of the heart condition, estimated from common clinical indices like blood pool volume or based on a statistical model of the heart growth like the one we presented in the second part of this thesis. The problem the cardiologist now has to tackle is to evaluate the type of valve replacement strategy he can apply to that patient. Should he replace the valves only, and let the heart remodel itself to its new condition? Or should he directly reduce the right ventricle, taking advantage of the surgery to remove scars and fibrosis? Nowadays, there is no clear-cut consensus on which strategy one should employ for specific patients.

To guide the cardiologist, we propose in this chapter to personalise the simplified electromechanical model (EM) of the heart presented in the previous chapter to reproduce the cardiac function of the patient *in-silico*. With a virtual heart, the cardiologist can test the two PVR strategies and predict their postoperative outcomes on the patient's cardiac function.

A first simulation of PVR for ToF patients was presented by [Tang et al., 2007b]. In that work, the authors made use of a fluid-structure interaction model coupled with a passive isotropic model of the myocardium to simulate patient-specific PVR and right ventricle (RV) volume reduction surgery. Obtained results were promising and supported the surgical RV reduction. However, the authors did not simulate the preoperative regurgitations and, most importantly, the active contraction of the myocardium, which may be affected by the surgery. We propose instead to use our active anisotropic electromechanical (EM) model of the heart. Figure 9.1 shows the different steps of the experiment. The compact bi-ventricular myocardium is segmented from clinical 4D cine MRI. The mid-diastole myocardium mesh is used as 3D anatomical model to simulate patient cardiac function and PVR therapies. The variation of the blood pool volumes throughout the cardiac cycle and the apparent motion on MRI are used to manually calibrate the EM model. After calibration, the resulting virtual heart reproduces the cardiac function of the patient on the computer. We can then test the effects of pulmonary valve replacement, by disabling the regurgitations in the model, and RV volume reduction, by virtually resecting tissues of the RV outflow tract (RVOT) using SOFA¹, an open source soft-tissue intervention platform [Allard et al., 2007]. In the following, we briefly describe the virtual surgery framework and present the results on two young adults with repaired ToF.

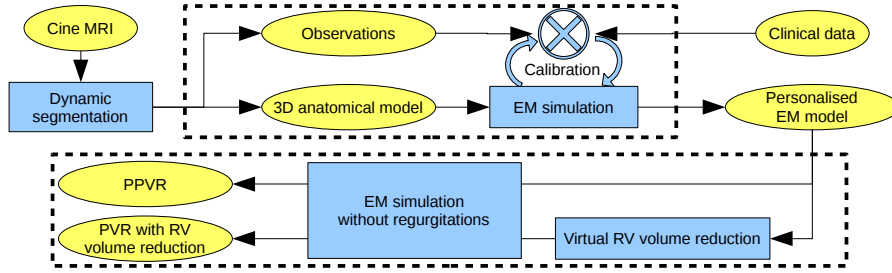


Figure 9.1: Pipeline for personalised model-based simulation of PVR therapies (*see details in text*).

9.2 Real-Time Simulation of Cardiac Volume Reduction

RV volume reduction is performed in real-time using SOFA [Allard et al., 2007], a soft-tissue intervention platform (Figure 9.2). The user reduces the RV volume by resecting tissues that have been identified as pathological. After resection, the RV

¹Available at www.sofa-framework.org

cavity is reconstructed in real-time by sewing the free wall to the septum. Large displacements and rotations of the elements are coped with co-rotational finite element models [Nesme et al., 2005]. A GPU-based implicit solver is used to update the mesh position [Comas et al., 2008].

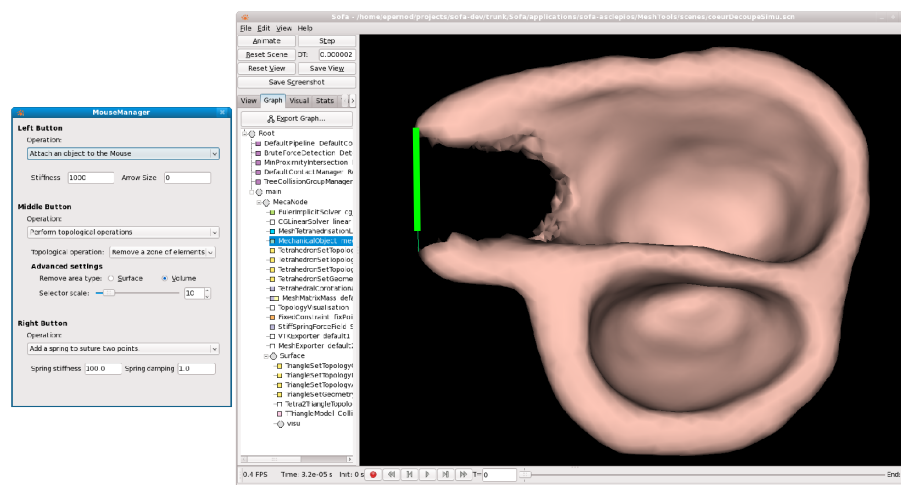


Figure 9.2: Screenshot of the SOFA platform. The anatomical model of a heart is being remodelled by the user. After resection, the user is closing the free-wall by pulling it close to the septum (*green line*). Grabbing, cutting and closing interactions are controlled in real-time during the virtual intervention.

9.2.1 Tissue Resection

Tissue resection is performed interactively by removing all the connected tetrahedra of the mesh that are inside a user-defined sphere region. The centre of the sphere follows the mouse pointer, enabling the user to define interactively the zone to remove. The radius of the sphere is defined through a pop-up menu. The resection of myocardium tissue is achieved successively by varying the centre and the radius of the sphere region (Figure 9.3, left panel). Internally, tetrahedra are removed using the method proposed by [Andre and Delingette, 2008]. Intuitively, the indices of the mesh elements and the meta data attached to them are stored into contiguous array, ensuring fast and direct access to the elements. In that way, when the properties of the mesh are locally modified, the time to update the data structure does not depend on the total number of mesh elements but only on the number of modified elements. However, this approach implies element renumbering when elements are removed from the data structure, in particular during interactive tissue resection. Nonetheless, this procedure is transparent to the user in our experiment as the number of mesh elements was usually fairly low (less than 60 000 elements).




Figure 9.3: Virtual suture of the right ventricle. Each side of the resected area (*left panel*) is brought close to each other (*mid panel*) by adding a “spring” between two user-defined vertices \mathbf{p}_w and \mathbf{p}_s . Several springs can be added (*right panel*) to reconstruct the RV free wall. All the steps are performed in real-time.

9.2.2 Tissue Attachment

After resection, the RV cavity is reconstructed interactively by progressively drawing the RV free wall close to the inter-ventricular septum (Figure 9.3, mid panel). To that end, the user “grabs” the RV free wall and pulls it close to the septum². Then, two vertices are picked, one on the free wall and the other on the septum. A spring between the two triangles is created, with a stiffness of $k = 100N\,m^{-1}$, which progressively brings the two walls close to each other. During the “sewing” interaction, the user can place several springs as a surgeon would sew the myocardium at several points. Once all the connections are established, the user exits the “sewing” interaction: the stiffness of the springs becomes tenfold higher ($k = 1000N\,m$) to seal the free wall to the septum (Figure 9.3, right panel). During the process, the computational efficiency of the sewing is improved by fixing the vertices of the left ventricle (LV) in the 3D space and disabling the biomechanical model on the LV, which does not move to preserve its shape during the virtual surgery. The anatomical model is re-meshed after rasterisation as a binary image to remove holes at the surgical junction. Finally, a scar zone is mapped onto the postoperative anatomy to simulate the effects of the surgical scar on the cardiac function.

9.2.3 Myocardium Fibre Recovery

During the virtual surgery, the RV is resected and deformed but the orientation of the cardiac fibre must be preserved locally as they are intrinsic to the myocardium tissue. This constraint is ensured by encoding the fibre directions in the local barycentric coordinate system of each tetrahedron. In that way, their relative orientation within a given tetrahedron is preserved during element deformations. During the final re-meshing, the fibres are transferred to the final postoperative anatomical model by using an intermediate rasterisation of the fibre directions as a vector image.

²See video at www-sop.inria.fr/members/Erik.Pernod/Movies/heart_suture_HeC.php

9.3 Personalised Simulation of Pulmonary Valve Replacement

9.3.1 Patient Selection and Data Preparation

Virtual pulmonary valve replacement (PVR) was tested on two randomly selected patients with repaired tetralogy of Fallot. Patient evaluation, echocardiography and magnetic resonance imaging (MRI) was performed for all patients. MRI was performed using 1.5T MR scanner (Avanto, Siemens Medical Systems, Erlangen, Germany) in all the centres. Retrospective gated steady-state free precession (SSFP) cine MRI of the heart were acquired in the short-axis view covering the entirety of both ventricles (Figure 9.4A). To date, these patients have undergone no PVR therapy to date. This means that no ground truth was available to validate the simulated post-PVR cardiac function.

9.3.1.1 Patient Selection

Patient 1 The first patient was a 16-year old boy recruited at Hôpitaux de Paris, Necker-Enfants Malade, Paris, France. Echocardiography showed moderate pulmonary and tricuspid regurgitations with moderate RV dilation (third level out of four regurgitation grades). RV pressure at end-systole was about 50 mmHg , value estimated from the regurgitation flow using Bernoulli's principle. Peak regurgitation flows Φ_c and Φ_r were estimated at 50 mL.s^{-1} . Visual inspection of the SSFP cine MRI (10 slices; 1.33 mm isotropic in-plane resolution; 8 mm slice thickness; 25 temporal frames) showed a dilated RVOT with a large aneurysm and severe dyskinetic motion (the RVOT dilates when the heart contracts). However, despite these abnormalities, LV and RV ejection fractions were just below normal and electrophysiology was almost normal.

Patient 2 The second patient was a 21-year old boy recruited at Great Ormond Street Hospital, London, U.K. Echocardiography showed moderate pulmonary regurgitations, mild tricuspid regurgitations and mild RV dilation. The analysis of the SSFP cine MRI (10 slices; 1.77 mm isotropic in-plane resolution; 10 mm slice thickness; 40 temporal frames) confirmed the mild RV dilation but with a relatively preserved RV shape. A large translation of the left ventricle was visible at end-systole. Ejection fractions were low but electrophysiology was almost normal.

9.3.1.2 Image and Mesh Preparation

Cine MRI images underwent slice misalignment correction, contrast enhancement and tri-linear resampling for isotropic voxel size as described in Section 3.2.2. After image preparation, the compact bi-ventricular myocardium was segmented as follows.

For patient 1, RV endocardium at the end-diastole time frame was segmented using the method proposed by [Zheng et al., 2008], as in Section 7.2.1, page 134. LV

endocardium and bi-ventricular epicardium were segmented interactively using the 3D surface modeller presented in Section 3.2.3, page 40. The myocardium was then reconstructed from the endocardial and epicardial surfaces and tracked throughout the cardiac sequence as described in Section 3.2.4, page 45 (Figure 9.4).

For patient 2, the RV endocardium, the LV endocardium and the LV epicardium were segmented using the method proposed by [Zheng et al., 2008] and propagated over the cardiac sequence by using an optical flow method with one-step prediction [Yang et al., 2008b]. The propagated RV epicardium was artificially reconstructed by dilating the RV endocardium by 5 mm, which is slightly larger than the normal RV myocardium thickness to account for RV atrophy in ToF and avoid possible numerical instabilities during the electromechanical simulation. Finally, the compact bi-ventricular myocardium was reconstructed from the propagated surfaces.

For all the patients, visual assessment of the dynamic segmentations showed good agreement despite the dilated RV and the abnormal motions. From the dynamic meshes, we computed the volume curves, reported in Figure 9.6 (dashed green curves) and the ejection fractions (EF) of each ventricle, reported in Table 9.2. EF is defined by the formula:

$$EF = 100. \frac{EDV - ESV}{EDV} \%$$

where, EDV is the end-diastole volume and ESV is the end-systole volume. Ejection fractions, volume curves and MRI were used as reference for adjusting the EM model to the cardiac function of the patient.




Figure 9.4: Segmentation of the compact bi-ventricular myocardium of patient 1 from cine MRI. **A**- Cine MRI. **B**- Contours at end-diastole. **C**- Tracked myocardium contour at end-systole.

9.3.2 Personalised Simulation of the Preoperative Cardiac Function

9.3.2.1 Model Personalisation

Anatomy The anatomical models of the bi-ventricular myocardium were generated from the mid-diastole mesh. The resulting tetrahedral models comprised 59768 (patient 1) and 43549 (patient 2) elements. Left ventricle, right ventricle and AHA

myocardium zones were automatically mapped onto the mesh for regional adjustment (Section 8.2.2). Synthetic fibre orientations were generated as described in Section 8.2.3. Finally, abnormal regions like the dyskinetic RVOT of patient 1 were manually reported on the anatomical models for their simulation. Figure 9.5 illustrates the personalised anatomical model of the two patients.

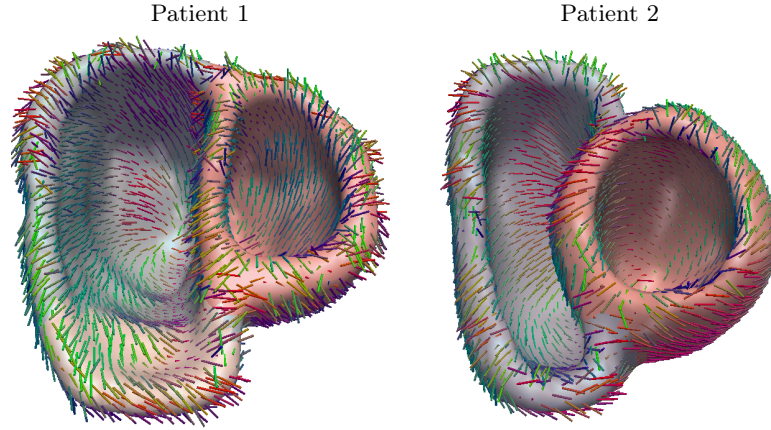


Figure 9.5: Personalised anatomical models. *In red LV. In blue RV. In brown dyskinetic area. Colour lines Myocardium fibres .*

Electrophysiology Since there was no visible anomaly in electrophysiology of these patients, electrophysiology parameters were adjusted as in a healthy heart (Table 8.1). Time synchronisation was performed using the beginning of systole. Without biomechanics, electrophysiology simulation took about 1 minute to compute thanks to the efficient multi-front fast marching scheme.

Biomechanics Passive biomechanical properties were kept to their nominal values (see Table 8.1) while the parameters related to the active element were manually adjusted. This was made possible by the relatively fast computation: 1 full cycle took between 15 to 30 minutes on a MacPro $2 \times 3.2\text{GHz}$ Quad-Core Intel Xeon, 16GB of RAM, depending on the number of mesh elements. We started with nominal values (Table 8.1). The contractile element of the lesions were disabled to reproduce their abnormal motion. The contact boundary force was used (Section 8.4.3) and some stiffness along the long-axis ($k = 2 \cdot 10^4 \text{ N.m}^{-1}$) was added to the basis to simulate the effect of atria and arteries on the bi-ventricular motion. Then, through trial and error, we adjusted the parameters for both ventricles to simulate the observed cardiac function. The simulation was compared to the MRI and the estimated volume curves and ejection fractions. Table 9.1 reports the final parameters for each patient. It should be noted that the parameters did not vary significantly from one patient to another. The personalisation of successive patients was much easier than the first one, only a few trials were necessary to adjust the model. In all the cases,

RV contractility was lower than normal, probably because of its dilated morphology and possible fibrosis.

Hemodynamics Aorta Windkessel parameters were fixed to their nominal values. When the RV end-systolic pressure was available, we personalised the Windkessel parameters of the pulmonary artery by modifying only the pulmonary characteristic resistance, as suggested by the results on synthetic data (Section 8.5.2). Pulmonary peripheral resistance and compliance were kept to their nominal values. Pulmonary regurgitation flows Φ_c and Φ_r were estimated directly from the clinical observations when available, standard realistic values were used otherwise (Table 9.1).

Table 9.1: Adjusted parameters of the cardiac model. Non-reported parameters were kept at their nominal values (see Table 8.1).

Parameters	Patient 1	Patient 2
Maximum contraction σ_0 (in $kPa.mm^{-2}$)	$\sigma_{0LV} = 100$ $\sigma_{0RV} = 70$ $\sigma_{0dysk} = 0$	$\sigma_{0LV} = 70$ $\sigma_{0RV} = 70$
Contraction rate k_{ATP} (in s^{-1})	$k_{ATP} = 10$	$k_{ATPLV} = 15$ $k_{ATPRV} = 5$
Relaxation rate k_{RS} (in s^{-1})	$k_{RS} = -10$	$k_{RS} = -10$
Pulmo. regurgitations flows Φ_c, Φ_r (in $ml.s^{-1}$)	$\Phi_c = \Phi_r \simeq 50$	$\Phi_c = \Phi_r \simeq 30$
RV end-systole pressure p_{RVES} (in $mmHg$)	$P_{RVES} = 50$	$P_{RVES} = 33$
Pulmo. characteristic resistance R_c (in $mmHg mL^{-1}$)	$R_c = 0.03$	$R_c = 0.04$

9.3.2.2 Preoperative Simulation

After model adjustment, realistic ejection fractions (Table 9.2) and volume variations (Figure 9.6, solid curves) were obtained. The discrepancy observed for patient 2 comes from the different meshes used to compute the volumes, the anatomical model was different from the segmented meshes that also captured the arterial valves and trunks. Simulated radial displacements computed from the mid-diastole position were locally consistent with those computed from the segmentation. In particular, the simulation of the dyskinetic RVOT observed in patient 1 was consistent with the observations (Figure 9.7). The cardiac motion of patient 2 was also satisfyingly recovered, as displayed by the images in Figure 9.8. In particular, the abnormal LV translation was captured by the model, suggesting that this motion is a consequence of the RV dilation and weak contractility. The EM model managed to provide, for these patients, realistic contraction patterns. Capturing the relaxation

motion was more difficult, in particular for the second patient. Probable reasons are a coarsely fitted electrical activation pattern (no clinical data was available to finely adjust electrophysiology) and the limitations of the model, which does not consider the early active relaxation of the myocardium [Sengupta et al., 2008]. Despite our simple regurgitation model, the simulated pressure-volume (PV) loops were consistent with measurements in ToF reported in the literature [Redington et al., 1990] (unfortunately, these data were not available for these patients for comparisons).

Table 9.2: Ejection fractions (EF, in percentage) computed from the dynamic segmentation and the simulations. One can see that the adjusted model managed to capture patients EF. If pulmonary valve replacement alone did not exhibit changes in EF, RV volume reduction improved RV and LV EF, suggesting a relationship between RV and LV functions.

Patient 1		Patient 2	
<i>LVEF</i>	<i>RVEF</i>	<i>LVEF</i>	<i>RVEF</i>
Segmentation			
61%	41%	42%	40%
Simulation: Preoperative			
59%	40%	41%	37%
Simulation: PVR			
59%	40%	41%	36%
Simulation: PVR and RV reduction			
63%	51%	54%	46%

9.3.3 Personalised Simulation of PVR in ToF Patients

Valve replacement was simulated by stopping the pulmonary regurgitations of the adjusted heart model. Virtual RV volume reduction was performed as illustrated in Figure 9.9. The RVOT was resected, manually remodelled and reconstructed in real-time. The dyskinetic area of patient 1 was entirely removed (Figure 9.5, left panel). Finally, a postoperative scar was simulated by setting the local conduction velocity v and maximum contraction σ_0 near the surgical junction to $0 m.s^{-1}$ and $0 MPa$ respectively. The frame-rate of the virtual surgery was about $20 - 25 fps$ on a 64b-Linux Core2Duo 2GHz machine with 4GB of RAM and GeForce 8800 Ultra GPU.

Replacing only the valves reduced the duration of the isovolumetric phases and slightly improved the end-systolic pressure as showed by the simulated volume curves and pressure-volume diagrams (Figure 9.6). Yet, no global improvements of the pump function were obtained, ejection fraction stayed unchanged (Table 9.2). This relatively surprising result may be due to our simple regurgitation model that alters the isovolumetric phases only and discard the impact of regurgitations on the other

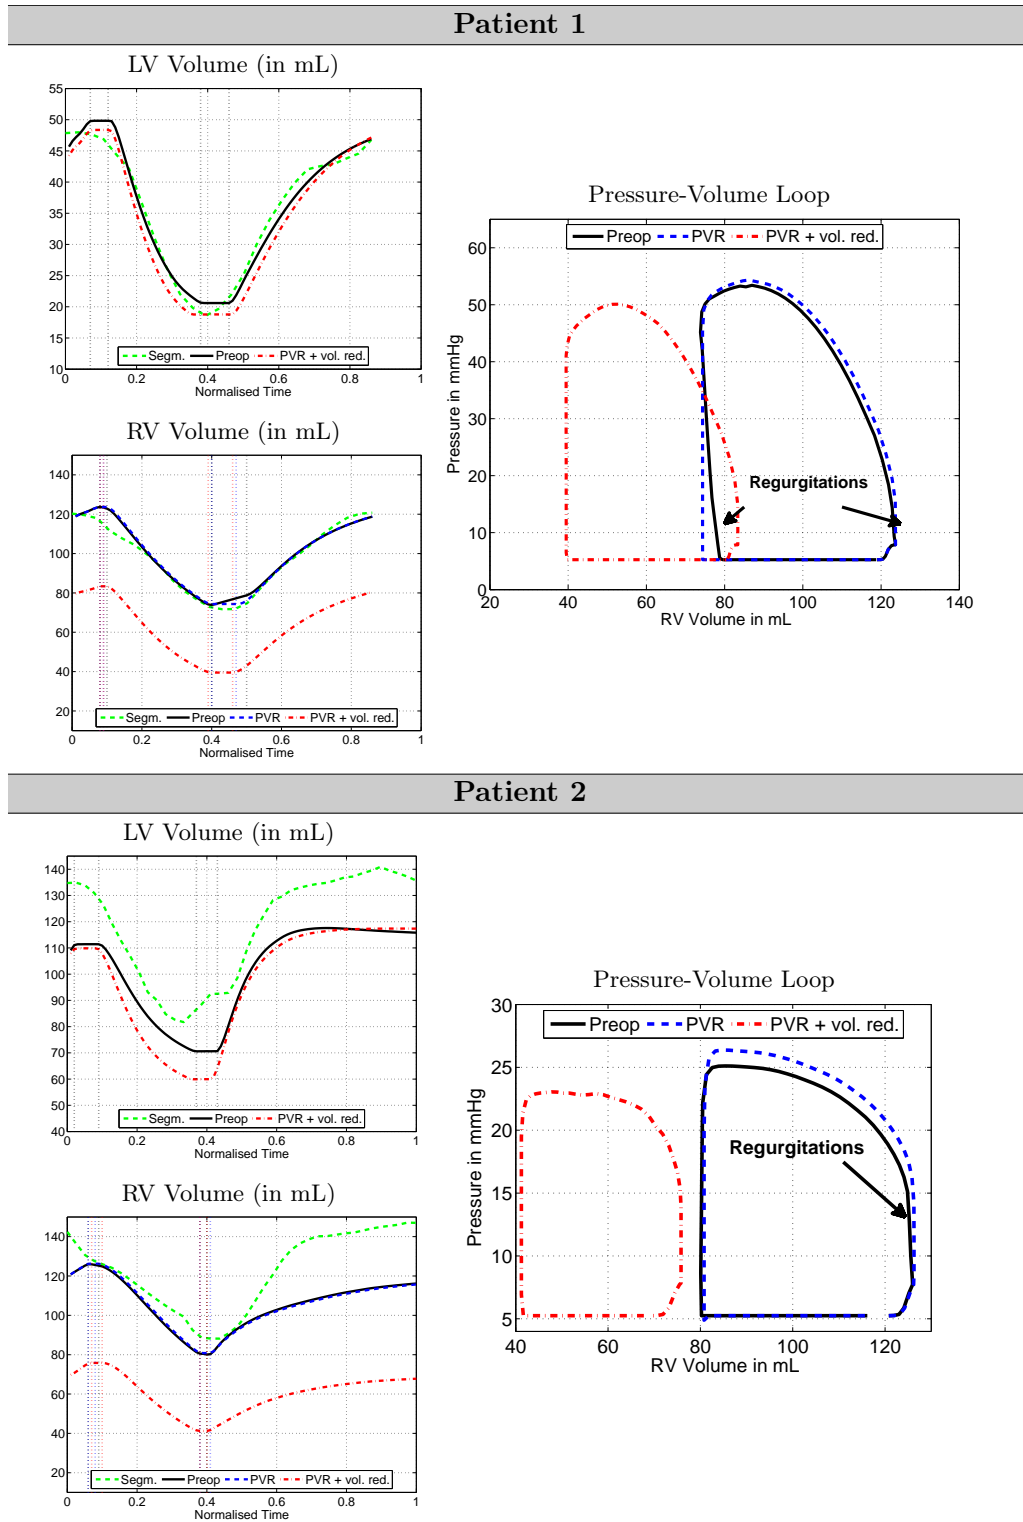


Figure 9.6: RV volume curves and pressure-volume loops. *Left panels*: cavity volume curves computed from segmentation (green curve) and simulations (black, blue and red curves). Vertical bars delineate the simulated cardiac phases. The volume discrepancy between segmentation and simulation in Patient 2 was due to differences between the segmentation meshes (whole ventricles + arterial roots) and anatomical meshes (ventricles cut at the base plane). *Right panel*: simulated pressure-volume loops of the right ventricle. (See text for details)




Figure 9.7: Radial displacements (in mm) at end-systole of patient 1, computed with respect to the end-diastole position. Positive values denote outward motion. Similar colour patterns between the simulated motion and segmentation confirm that the simulated model was able to exhibit realistic motion, in particular the dyskinetic right ventricle outflow tract (RVOT).

cardiac phases. However, personal discussions with cardiologists confirmed that the cardiac function is not significantly improved by PVR just after valve implant: the most important changes in cardiac function come late after the intervention, after natural remodelling of the heart, which we do not consider here.

PVR with direct RV volume reduction yielded significant improvements in both RV and LV function for all the patients. RV volume effectively decreased (Figure 9.6, red curve) and RV postoperative EF improved significantly (Table 9.2). What was more surprising, is that the simulated LV function also improved although we did not modified its shape nor its electromechanical parameters. This experiment confirmed the tight relationship between the two cavities through the inter-ventricular septum.

9.4 Discussion

In this chapter, we have tested the potential of image processing techniques, EM models and virtual soft-tissue intervention platforms to perform virtual and personalised assessment of PVR therapies on ToF patients. The results were promising and suggested that such tools might be used, after comprehensive validation, by clinicians to test different PVR therapies.

For this experiment, we used the modular cardiac simulation framework presented in Chapter 8 as few clinical data were available to personalise more complex models. Yet, despite the simplifications, the model was able to capture the cardiac function of all the considered patients, in particular the systolic motion. For these patients, we found that PVR with direct RV volume reduction would have better results just after intervention than PVR alone. By removing lesions and scars and by reducing the RV volume manually, the cardiologist improved the RV function




Figure 9.8: Comparison of personalised simulation (*yellow*) with MRI and segmentation (*red*: LV, *blue*: RV). The model is not guided by the image. One can see that the model managed to capture the abnormal leftward translation of the LV due to the dilated and impaired RV.




Figure 9.9: Virtual RV volume reduction surgery. From left to right: original mesh, after resection, sewing and final mesh. *Colour lines*: fibre orientations. *Black area*: postoperative scar.




Figure 9.10: Postoperative anatomies with the surgical scar in black.

and, simultaneously, improved the LV function by minimising abnormal septal motions. However, this procedure is invasive (open-heart surgery), with high risks for the patient. Furthermore, the long term advantages of this procedure over the PVR approach is not clear [Therrien et al., 2000]. On the one hand, the effects of PVR alone are visible few months after implant, when the heart has remodelled to its new loading condition. On the other hand, heart remodelling can reduce the effects of the surgical RV volume reduction to accommodate the surgical scar. It is therefore crucial to model postoperative cardiac remodelling to provide the cardiologist with a full vision of the effects of PVR treatment on a patient. A possible direction would consist in using the statistical models presented in this thesis for instance.

This work constitutes one of the first attempts towards personalised simulation of PVR in ToF. The lack of postoperative data prevented us from validating the simulated PVR effects and, as a consequence, our cardiac model. Yet, obtained results were found qualitatively reasonable by cardiologists, which encouraged further work in that direction.

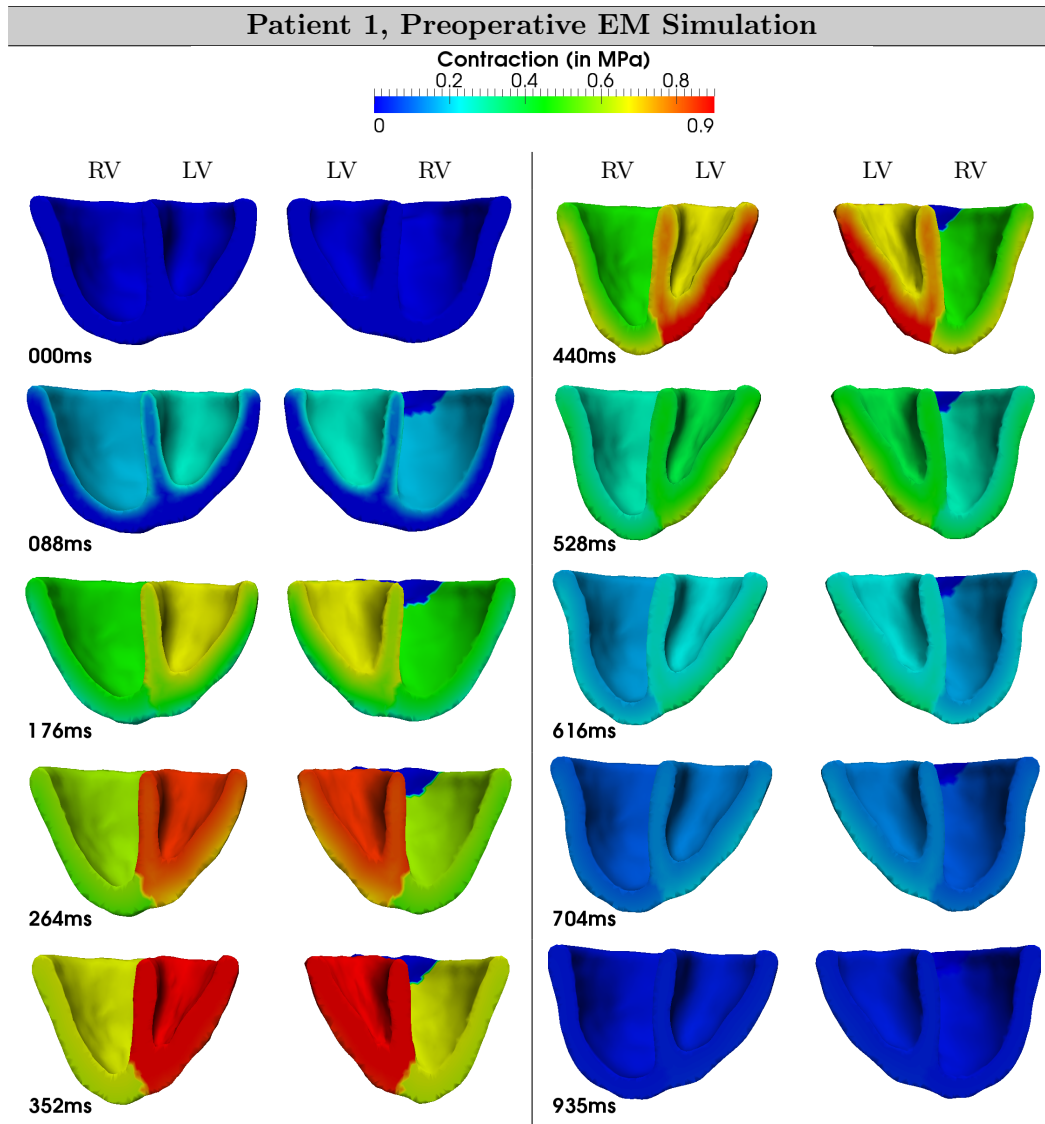


Figure 9.11: Simulated cardiac function of the Patient 1. Colours encode the active contraction in MPa ($0 MPa$ in the dyskinetic area). RVOT dyskinetic motion was recovered by disabling its active contraction. RV contractility was decreased to capture the pathological RV motion.

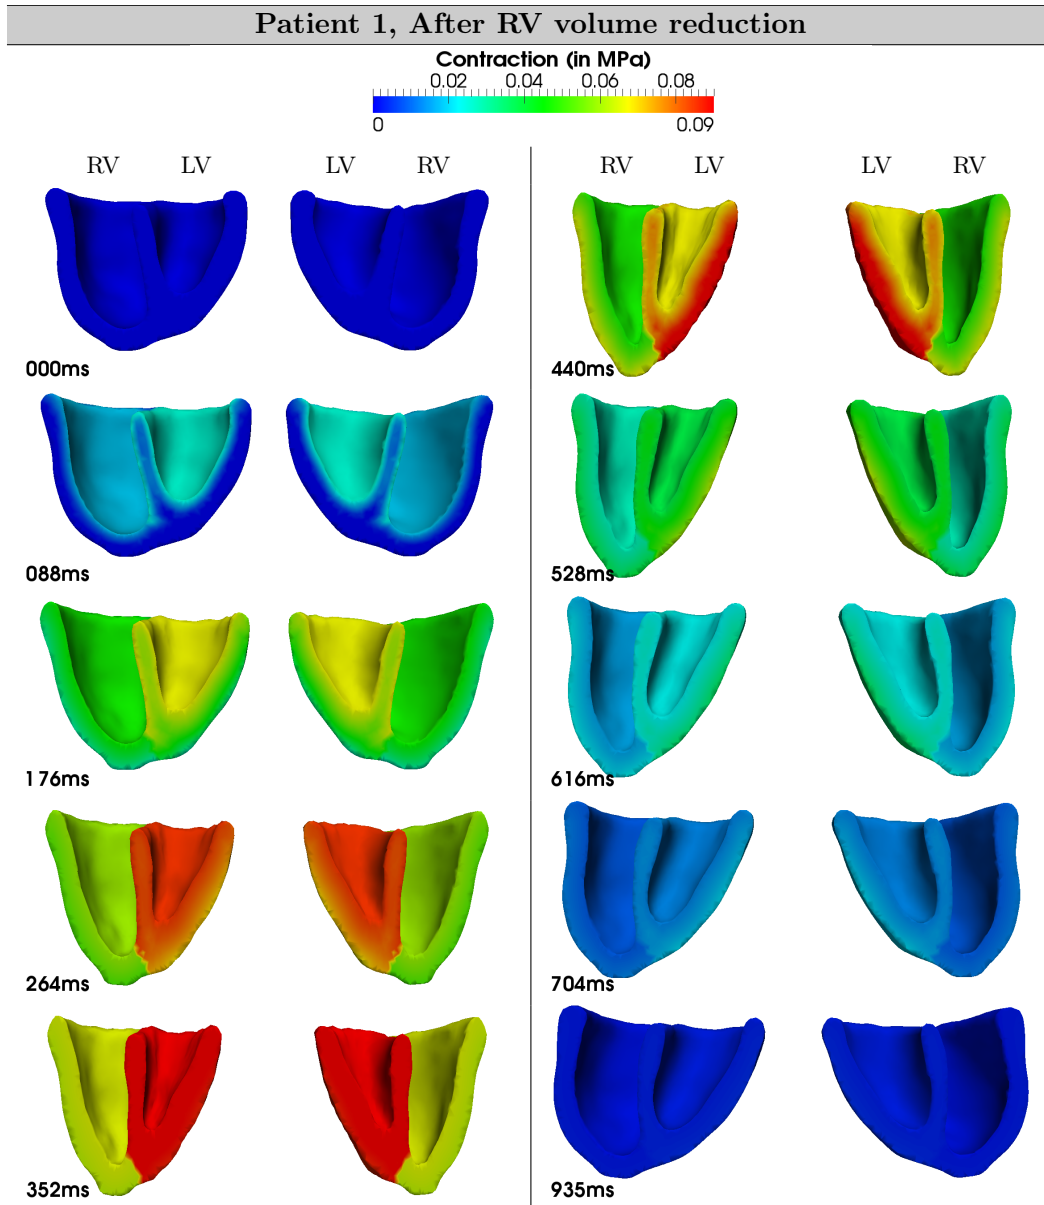


Figure 9.12: Simulated cardiac function of Patient 1 after virtual RV volume reduction. Colours encode the active contraction in MPa. The dyskinetic area was removed and a surgical scar was simulated. The scar did not contract (no electrical conductivity and no contractility). RV volume reduction significantly improved the function of both ventricles (Table 9.2).




Figure 9.13: Simulated cardiac function of Patient 2. Colours encode the active contraction in *MPa*. Observe the abnormal leftward translation of the LV due to the dilated and impaired RV, in particular the apical septal region.

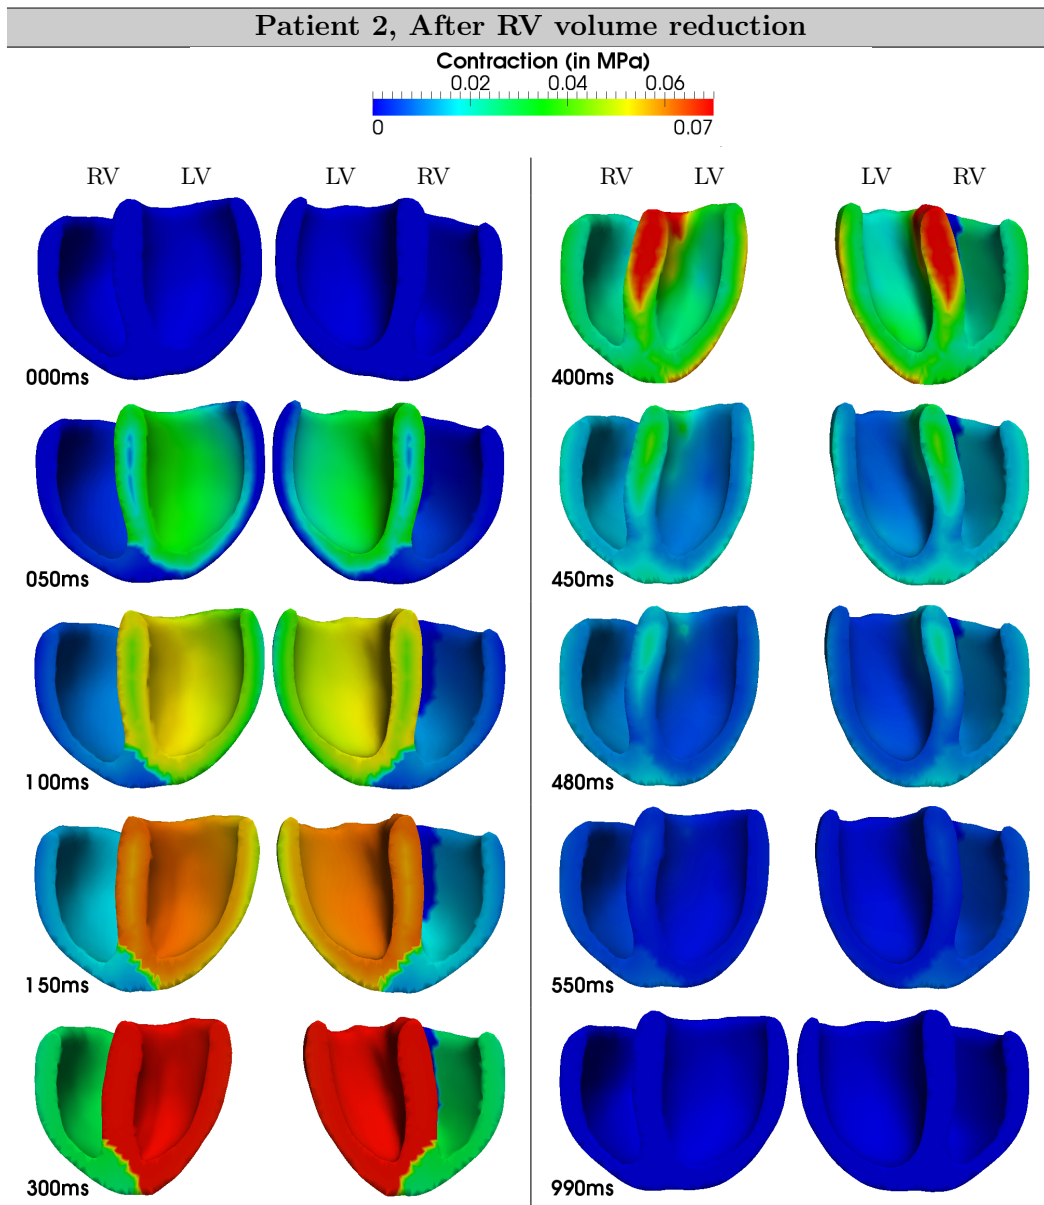


Figure 9.14: Simulated cardiac function of Patient 2 after virtual RV volume reduction. Colours encode the active contraction in MPa. The RV was surgically remodelled and a surgical scar was simulated. RV volume reduction significantly improved the function of both ventricles (Table 9.2).

Several research directions can be explored to improve the framework. First of all, the approach must be validated *a posteriori* on postoperative data or in animal experiments. This is a mandatory step before applying these tools in the clinical workflow. It has to be noted however that similar experiments on adults with heart failure showed promising predictive powers of the model [Serresant et al., 2009]. Second, the regurgitation model can be improved to take into account the other compartments and the other cardiac phases. In particular, models like those presented in [Keener and Sneyd, 1998] could be implemented if the cardiac cycle is pressure-controlled [Sainte-Marie et al., 2006]. Another alternative is to couple fluid-dynamics models of the pulmonary flow with our electromechanical model of the ventricles [Gerbeau et al., 2005; Mihalef et al., 2009; Nordsletten et al., 2009]. Cardiac biomechanics can also be enhanced by using more sophisticated models of the passive myocardium, like the orthotropic non-linear Costa law for instance [Costa et al., 2001]. It would be interesting however to quantify the added value of more detailed models with respect to the increased computational time and number of parameters. For Costa law for instance, it would be important to evaluate how sensitive is the simulation to the orientation of the fibre sheets as this data is difficult to attain in clinics. Another improvement would be to model the active relaxation to improve the simulation of diastole. This could be done by modifying the mechano-electrical coupling, like in the original model of [Bestel et al., 2001]. The mechano-electrical feedback could also be studied as the extreme dilation of the right ventricle may lead to arrhythmias. Finally, automated parameter estimation methods would greatly benefit these applications. Identified parameters would then be used as quantitative features of the cardiac condition for diagnostic support.

Part V

CONCLUSIONS

Conclusions and Perspectives

Contents

10.1 Summary of the Contributions	213
10.1.1 Methodological Contributions	214
10.1.2 Clinical Impacts	217
10.1.3 Software Development	218
10.2 Perspectives	219
10.2.1 Publications	225

Linking advanced physiological and statistical models of the heart with medical images can have tremendous repercussions in the clinical management of patients. In particular, it can provide a new set of tools for a more patient-specific medicine [Hunter et al., 2003a; Ayache et al., 2006; Bassingthwaighe et al., 2009]. Guided by the clinical context of repaired tetralogy of Fallot (ToF), we investigated in this thesis this precept for three specific clinical questions: How can physiological priors of the cardiac tissue improve an image registration algorithm to estimate the cardiac deformation from standard anatomical images? How can we model the cardiac growth observed in a population and identify anatomical patterns related to a pathology? How can we predict the postoperative outcomes on the cardiac function of a patient?

Of course, this thesis does not provide the final answers to these questions but rather constitutes small steps towards the global scope of model-based cardiology. In the following sections the main contributions of this work are summarised, followed by some general research perspectives.

10.1 Summary of the Contributions

This thesis was realised in a unique multidisciplinary setting. The work was performed in the Asclepios research group, INRIA, with a strong involvement within the Health-e-Child European Project¹ and motivated by close collaborations with cardiologists from different countries. As a result, the methodological developments of this work also had some clinical impacts. Moreover, some of the results were integrated into the Health-e-Child platform, thus contributing to the global achievements of the project.

¹<http://www.health-e-child.org/>

10.1.1 Methodological Contributions

Cardiac Segmentation Pipeline

Although not the main topic of that thesis, we had to develop a multi-purpose segmentation pipeline to delineate the cardiac geometry from medical images as this step was mandatory for most of our investigations (Chapter 3, page 3). The most important contribution in that matter is probably the interactive 3D surface modeller that enables to delineate the 3D myocardium boundaries. The method relies on variational implicit functions that are generated on-the-fly through user interactions. Publicly available as part of the CardioViz3D software, this tool was presented among other CardioViz functionalities at the VCBM workshop, 2008 [Toussaint et al., 2008]. Since then, it has been used as initial stage of several studies, in particular for myocardium tissue tracking [Billet et al., 2008, 2009; Mansi et al., 2009d, 2010b] and personalised simulations of cardiac electromechanics [Mansi et al., 2009a; Sermesant et al., 2009; Chabiniok et al., 2009].

It has to be noted that the investigations performed in the framework of Health-e-Child, like the statistical analyses of the right ventricle shape (Chapter 7, page 133) and the personalised simulations of the heart in ToF (Chapter 9, page 193), relied on segmentations computed by Siemens Corporate Research using the methods presented in [Zheng et al., 2008; Yang et al., 2008b].

Estimation of Myocardium Deformation from Clinical Images

To estimate the three-dimensional myocardium strain from dynamic anatomical images, we proposed to constrain the logDemons registration algorithm [Vercauteren et al., 2008] with physiological priors about the cardiac tissue (Chapter 4, page 61, and Chapter 5, page 89). Since the myocardium is an elastic incompressible tissue, we constrained the demons algorithm to provide elastic and volume-preserving deformations to cope with the lack of textures and the low image quality of clinical routine MRI that hinder image registration. The proposed algorithm, named iLogDemons, is the result of the following methodological contributions:

- **New insights into the *ad-hoc* Gaussian regularisation of the demons algorithm were provided.** Contrary to [Cahill et al., 2009], who justifies demons as an approximation of the stationary solution of a diffusion equation, we explained demons Gaussian regularisation as a Tikhonov regularisation thanks to the intermediate correspondence field and its coupling to the transformation to estimate. Our interpretation is well-posed and enables to seamlessly integrate more enhanced regularisation schemes.
- **We integrated an elastic-like regulariser based on isotropic differential quadratic forms (IDQF) whose solution is exactly the separable vector filter proposed by [Cachier and Ayache, 2004].** This was achieved by demonstrating the link between multi-order IDQF and separable Gaussian vector filters. As a result, no linear system must be solved, the filter is separable and calculable through Gaussian smoothing, yielding fast elastic-like

regularisation. In our experiments, the estimated elastic deformations were smooth at any order.

- **The logDemons was constrained to be incompressible by parameterising the diffeomorphic deformations with stationary divergence-free velocities.** Since deformations parameterised by divergence-free velocities are incompressible, the volume-preserving condition is ensured by constraining the velocities to be divergence-free using the Lagrangian multiplier method. We demonstrated that this constraint amounts to projecting the stationary velocity fields onto the space of divergence-free vector fields through Helmholtz decomposition. As a result, the proposed constraint is strong, it does not require *ad-hoc* numerical schemes and it can be imposed in subregions of the image domain only.
- On synthetic data, our method provided quantitatively more accurate estimations of incompressible deformations. On real data, **iLogDemons significantly enhanced the accuracy of the cardiac deformations estimated on standard clinical images when compared to tagged MRI and 2D-strain echocardiography.** Elastic regularisation and incompressibility improved the registration in regions with low textures by reorienting the displacements such that incompressibility is ensured. iLogDemons could constitute a complementary tool to get hints about the myocardium integrity when no myocardium displacement imaging modalities are available.

The justification of demons Gaussian regularisation and the incompressibility constraint presented in this thesis have been accepted for publication at the MICCAI 2010 conference [Mansi et al., 2010b]. The complete details of the elastic incompressible algorithm have been submitted to an international journal [Mansi et al., 2010a]. Note that a previous version of that work was presented at the FIMH 2009 conference [Mansi et al., 2009d].

Statistical Analysis of the Heart Remodelling

Having in mind the long-term myocardium damages observed in repaired ToF, we investigated a statistical approach to quantify and to model the remodelling of the heart (Chapter 6, page 107, and Chapter 7, page 133). To that end, the non-parametric framework based on currents proposed by [Durrleman et al., 2009a] was used. In that framework, the observed shapes were encoded by the deformations that map an ideal atlas to the patients.

Instead of directly modelling the biological phenomena that are involved in cardiac remodelling, we observed how the heart shapes varied in a population of patients with different ages. The idea, commonly used in neuroscience and socio-economics, consists in relating the shape variations to clinical indices to identify shape features relevant to the pathology and to study the anatomical evolutions. Despite the change of paradigm, from an explicative approach to an explorative one, obtained

results were surprisingly plausible. This approach yielded the following methodological contributions:

- **We proposed a statistical framework to identify shape patterns relevant to a pathology.** The heart shape of a patient was compactly represented by a shape vector defined on a PCA subspace of deformations. Then, standard statistical analyses were performed between the PCA shape vectors and clinical variables to identify the deformation modes, which encoded the shape variabilities, that were relevant to the pathology.
- **We developed an algorithm based on partial least squares (PLS) and canonical correlation analysis (CCA) to estimate a generative model of heart growth.** The model computes an average heart shape for a given clinical index, the body surface area (BSA) in this thesis. Contrary to most studies that use PLS to predict new variables [McIntosh and Lobaugh, 2004; Rao et al., 2008; Yang et al., 2008c], in this thesis PLS was employed to estimate an optimal subspace of deformations that was also relevant to BSA. The heart shapes were then represented by a shape vector defined on the PLS subspace. CCA between BSA and the PLS shape vectors provided the generative model of heart growth. This strategy enabled us to **model how the shape evolves during growth.**
- For the specific case of the statistical modelling of heart growth, **we demonstrated that PLS space decomposition was much more efficient than the traditional PCA approach, with better generalisation and increased plausibility of the generative heart growth model.** Contrary to PCA, PLS automatically extracts the deformation modes that are also relevant to the external clinical parameter under study. This may include deformation modes that could be considered as noise by PCA (i.e. with low variance).

A first version of that work, without the PLS method, was presented at the MICCAI 2009 conference [Mansi et al., 2009c]. A methodological paper based on Chapter 6 and Chapter 7 is in preparation, with a planned submission in August-September 2010. It has to be noted that the proposed statistical framework could be applied to other organs or clinical questions, which opens new methodological and clinical perspectives.

Electromechanical Modelling of the Heart for Personalised Therapy Simulations

We finally tackled the question whether direct electromechanical models of the heart could help in planning therapies by predicting their postoperative effects (Chapter 8, page 155, and Chapter 9, page 193). As modelling the cardiac function is a complex task, we started from an existing model [Serresant et al., 2006b], which we improved to simulate the specificities of repaired ToF. In particular:

- We developed tools for the **automatic generation of anatomical models of the heart from clinical images**. These tools paved the way for patient-specific simulations as they enabled to build the geometrical models on which electromechanical PDE's were solved. These tools are now used in the modelling experiments performed in Asclepios [Billet et al., 2009; Sermesant et al., 2009].
- We proposed a **new boundary contact force to constrain the heart position in the 3D space**, rather than imposing unrealistically strong stiffness at the heart base. In our experiments, **the resulting motion was more realistic, the base could contract almost freely while the apex did not move**. The stiffness of the base constraint was used exclusively to simulate the constraints related to the arteries and atria.
- We improved the model of cardiac hemodynamics to adapt it to the specificities of ToF. In particular, **we proposed a simple regurgitation model where measured regurgitation flows were used as constraints of the model**. Despite the simplicity of the approach, simulated hemodynamics were fairly plausible. **A 3-element Windkessel model was also implemented** to simulate the arterial pressures [Stergiopoulos et al., 1999]. To avoid numerical instabilities, an implicit integration scheme was used.
- **We proposed a method based on SOFA framework² for the personalised simulation of pulmonary valve replacement in repaired ToF**. In that framework, the user simulate some steps of a virtual surgery, in real-time. The virtual surgery in SOFA was performed with Barbara André and Erik Pernod from Asclepios.
- These contributions, together with simple personalisation strategies, enabled us to simulate the cardiac function of two patients with repaired tetralogy of Fallot before and after pulmonary valve replacement.

The simulation pipeline and the results on the first patient with ToF were published in [Mansi et al., 2009a] and presented at the 3DPH workshop, Zermatt, 2008. The tools for the automatic generation of anatomical models have also been used in other studies, in particular [Chinchapatnam et al., 2009; Sermesant et al., 2009]. A journal paper describing the complete methodology with the new simulation results is in preparation. The submission is planned in August-September 2010.

10.1.2 Clinical Impacts

The developments presented in this thesis were motivated by clinical questions related to repaired ToF. As a result, the methodological contributions that were proposed naturally yielded the following clinical use:

²<http://www.sofa-framework.org>

- **Morphological changes of the right ventricle shape related to the regurgitations in repaired ToF were identified** (Section 7.3, page 137). The exhibited 3D pathological shape patterns could constitute quantitative metrics of ToF severity. Our findings were consistent with observations reported in the literature [Geva, 2006; Sheehan et al., 2007; Bodhey et al., 2008]. To the best of our knowledge, this study is the first that relates the complete 3D right ventricle shape to functional abnormalities in repaired ToF.
- **A first model of right ventricle growth in repaired ToF estimated from a population of 32 patients was proposed** (Section 7.4, page 141). The model provided information on the 3D evolution of the right ventricle in these patients. These results received a great interest among the clinicians, which are now willing to use these techniques to investigate other questions related to repaired ToF and other cardiac pathologies. New projects are being set up to further investigate cardiac remodelling. To the best of our knowledge, this work shows for the first time the evolution of the 3D right ventricle shape in repaired ToF.
- Using personalised models of the cardiac function in two ToF patients, **we observed that valve replacement with right ventricle reconstruction may improve both right and left ventricular functions** (Chapter 9, page 193). Contrary to previous studies [Yang et al., 2008a], we employed here an active model of cardiac biomechanics with a simple model of regurgitations in order to investigate the consequences of the surgery on the systolic motion. Despite the lack of quantitative validation, this result was found clinically promising by cardiologists, leading to new projects to validate the approach and the findings.

The personalised simulations of pulmonary valve replacement were presented at the Congress of the French Society of Paediatrics [Mansi et al., 2008, 2009b]. In 2010, a special session for *in-silico* models in Paediatrics was organised by this congress, at which we were invited to present our recent results about the cardiac modelling and the statistical analyses of the right ventricle in ToF [Mansi et al., 2010c]. Finally, the great interest of the clinical community raised by the statistical analyses of ToF hearts motivates us to submit a clinical paper about these findings. This matter is currently in discussion with cardiologists.

10.1.3 Software Development

Software Contributions

The interactive 3D surface modeller presented in Chapter 3, page 33 is freely available as part of the CardioViz3D package³. Other contributions to CardioViz3D were also provided, among them the automatic estimation of kinematic metrics about cardiac simulations (strains, displacements, volume variation, etc.).

³<http://www-sop.inria.fr/asclepios/software/CardioViz3D/>

Similarly, the model of the right ventricle growth has been released to the Virtual Physiological Human community and freely accessible⁴ along with the source codes of the “currents” framework developed by S. Durrleman and J. Glaunès.

Finally, the iLogDemons was developed from the publicly available ITK source code of the logDemons [Dru and Vercauteren, 2009]. It is available in the Asclepios image processing library, MIPS. A public release is planned for Autumn 2010.

Software Integration for the Health-e-Child Project

Most of the works performed during this thesis were integrated in the Health-e-Child European project. In particular:

- Our results were included in numerous demonstrations of the Health-e-Child technical achievements. The most relevant are:
 - The personalised simulations of pulmonary valve replacement in ToF patients were part of the Health-e-Child demonstration that won **the Exhibit Grand Prize at the European ICT 2008 conference**.
 - Two presentations were given at the final public conference organised by the project at Sestri Levante, April 23-24, 2010⁵. In that conference, the personalised electromechanical models and the statistical analyses in ToF were presented, with positive and fruitful feedbacks.
 - Our work was presented in collaboration with clinicians at the annual reviews of the project at the European Community.
- CardioViz3D and SOFA⁶, a soft-tissue intervention platform, were integrated into the Health-e-Child platform with the help of Nicolas Toussaint and Erik Pernod, from Asclepios. The tool was released to the clinicians for visualising personalised simulations and performing virtual surgeries.
- We published a multimedia website⁷ that describes the various results achieved by Health-e-Child in terms of disease modelling [Pennec et al., 2007-2009-2010]. Interactive contents (movies, 3D models, etc.) were used to explain the principles of advanced modelling of the heart. The website was regularly updated to disseminate our findings to the community.

10.2 Perspectives

Many short-term research directions have been reported in the conclusion sections of this manuscript. Those perspectives are mainly focused on the specific questions tackled by each chapter. In the following, we summarise the most important ones

⁴<http://www-sop.inria.fr/asclepios/projects/Health-e-Child/ShapeAnalysis/index.php>

⁵<http://conference.health-e-child.org/>

⁶<http://www.sofa-framework.org>

⁷<http://www-sop.inria.fr/asclepios/projects/Health-e-Child/DiseaseModels/index.html>

and provide some more general perspectives in the context of the Virtual Physiological Human [Hunter et al., 2003a; Ayache et al., 2006; Bassingthwaite et al., 2009].

An Electromechanical Model of the Heart for Patient-Specific Simulations

The development of an electromechanical model of the heart for personalised simulation and therapy prediction is a dauntingly complex task. In this thesis, we studied the feasibility of such an approach in patients with repaired ToF (Chapter 9, page 193). Nonetheless, a lot of work still needs to be done before these methods are used routinely in the clinics.

A first crucial task to perform is validation. The results reported in this thesis were found qualitatively plausible by the cardiologists. However, because of the lack of postoperative data, they could not be validated. Ideally, one would first personalise the model on preoperative data and then simulate the therapy *in-silico*, comparing the simulated outcomes with the real ones *a posteriori*. This approach, already employed by [Serresant et al., 2009] for instance, enables to validate the model for the specific clinical question that is tackled. Indeed, here the objective is not to develop the most accurate cardiac model but rather to design a model that predicts the effects of a given therapy with enough accuracy. This paradigm simplifies the problem as only a subset of phenomena needs to be represented. Of course, these simplifications must be controlled in order to ascertain the generalisation of the simulations. “Models should be made as simple as possible, but not simpler”, as stated by Einstein about scientific theories [Einstein, 1934].

The limitations of our model open new research directions. In the following we mention few of them.

Myocardium Biomechanics It is now accepted that linear elasticity is not enough to reliably simulate the cardiac function [Hunter and Smaill, 1988; Costa et al., 2001; Schmid, 2006]. This limitation was noticeable in our simulations during diastole, whose calibration was challenging. Ideally one would use non-linear models [Hunter and Smaill, 1988; Costa et al., 2001]. It would be interesting though to benchmark the different models in order to quantify their domain of validity. In some cases, linear elasticity may be enough and thus preferred for its computational efficiency, and *vice versa*. Such a benchmarking would enable to implement model-adaptive simulators of the beating heart as in [Picinbono et al., 2003]. Regions with small displacements would be modelled using linear transverse isotropic elasticity, whereas regions with large displacements would be modelled using the non-linear Costa model.

Regurgitation Model Ideally, regurgitations should be simulated and not constrained. As regurgitations are pressure-dependent, we should control the different phases of the cardiac cycle through the ventricular pressures [Sainte-Marie et al., 2006; Niederer et al., 2009]. We are currently working on this.

More advanced approaches could be investigated. For instance, fluid-dynamics models could be coupled with the electromechanical model to simulate the regurgitations and the flow patterns [McQueen and Peskin, 2000; Gerbeau et al., 2005; Nordsletten et al., 2009]. This would enable to simulate the impact of the regurgitant blood flow on the myocardium wall stress, suspected to be a source of pathological myocardium stiffening and remodelling.

Automatic Parameter Estimation In this work the electromechanical model was personalised manually, through trials and errors. Despite the promising results, automatic methods would greatly facilitate this task and reduce user variability. Nonetheless, automatic parameter estimation is extremely challenging due to the large number of degrees of freedom. Scientists are striving to develop methods to reach this goal. Preliminary works are showing promising results in that sense [Moireau et al., 2009; Wong et al., 2010; Billet, 2010]. Several possibilities are available, from variational approaches [Billet, 2010] to filtering methods [Moireau et al., 2009; Wong et al., 2010] or trust-region techniques [Chinchapatnam et al., 2008] or even more exhaustive techniques like genetic algorithms [Khalil et al., 2006]. Yet, all these methods require that the model is compatible with the available clinical data. To cope with this issue, coarse-to-fine approaches could be investigated [Relan et al., 2010], where simplified models would be used as input to more complex models.

Improving the Quantification and the Simulation of the Cardiac Function Using 3D Myocardium Strain Maps

The initial rationale of the iLogDemons algorithm was to estimate the 3D myocardium strain from standard clinical images to quantify the cardiac motion when no tagged MRI (tMRI) or related imaging modalities are available. The long-term objective was to integrate the estimated strain maps into the electromechanical model of the heart. Of course, we do not claim that iLogDemons method should substitute advanced medical imaging technologies like tMRI. It rather constitutes a tool to evaluate the cardiac deformation when such data are not available. The cardiologist would always have a way to estimate the myocardium strain, which would open new diagnostic opportunities [Moore et al., 2000]. Regions with abnormal strains could highlight localised cardiac asynchrony for instance [Helm et al., 2005b] due to myocardium lesions. In repaired ToF, myocardium strain could constitute a quantitative feature for pulmonary valve replacement [Eyskens et al., 2010].

In addition to the short-term research perspectives listed in Section 4.6, page 80 and Section 5.4, page 102 (validation of the algorithm, comparison with existing tools, etc.), we can identify more long-term research directions towards the scope of model-based medicine. In particular:

Myocardium Strain Atlases After validation, iLogDemons could be used to estimate myocardium strain atlases from large databases of cardiac images by using the method proposed by [Peyrat et al., 2007] on strain tensors. Tech-

niques based on parallel transport [DoCarmo, 1992; Rao et al., 2002; Qiu et al., 2008], inter-subject 4D image registration [Peyrat et al., 2009] or template comparison [Aljabar et al., 2008] could be investigated to compare the cardiac function of different patients. Thanks to well-posed frameworks for statistics on tensor fields [Arsigny et al., 2006b], techniques like those presented in Chapter 6 could be used to identify changes in myocardium strain relevant to pathologies and to estimate a model of cardiac function remodelling. From a clinical point of view, such atlases would provide new insights into the pathologies.

Electromechanical Model Personalisation from Strain Estimated 3D strains could be used to automatically personalise the electromechanical model of the heart. Nowadays, this task is achieved using the apparent myocardium velocity [Moireau et al., 2008] or the position of its boundaries [Moireau et al., 2009; Billet, 2010; Wong et al., 2010]. Tissue motion inside the myocardium is discarded. Using 3D myocardium strain would further constrain the inverse problem, thus reducing the uncertainty on the model parameters as the full 3D motion would be considered. To that end, approaches based on the minimisation of a cost function could be employed. Another alternative would be to integrate the electromechanical model directly into the iLogDemons algorithm, similar to what [Sundar et al., 2009b] proposed. The cardiac model would be implemented as a regularisation term, which is now possible thanks to our justification of demons regularisation.

Strain-Structure Interaction Strain-driven personalisation of electromechanical models of the heart would enable to investigate the interactions between strain and structure. One could use for instance the estimated myocardium strain to study how the fibre architecture remodels [Rijcken et al., 1999; Ubbink et al., 2006; Kroon et al., 2009]. Investigations in that direction could greatly benefit from the recent advances in *in-vivo* cardiac DTI [Wu et al., 2007; Toussaint et al., 2010]. Similarly, estimated strains could also drive a model of myocardial growth [Rodriguez et al., 1994; Kroon et al., 2007].

Non-Invasive Electrophysiology Quantification Finally, iLogDemons could be applied on 3D ultrasound images with high frame-rate to recover the propagation of the electrical wave from the observed cardiac biomechanics [Sanchez-Ortiz et al., 2005]. This work would result in image-based assessment of patient electrophysiology, which could be employed to further personalise the electromechanical model of the heart. We could also study the mechano-electrical coupling in patients to whom invasive measurements cannot be performed. Investigated by Adityo Prakosa, this strategy already provided promising results.

Towards a “Growing” Model of the Heart

The identification and modelling of the right ventricle remodelling in repaired ToF received a great enthusiasm among the clinicians. Our approach constitutes another way to study the 3D alterations of the cardiac anatomy due to a pathology. Of course, a more thorough validation is required to confirm these results, using for instance leave-one-out cross-validation. We refer the reader to Section 7.5, page 150 for specific comments on that matter. Still, this work opens several research directions, both from methodological and clinical point of views.

Growth Model of the Beating Heart From a methodological point of view, it would be interesting to extend the analysis to the 4D cardiac motion to build a growth model of the beating heart. The approach would be different from the method proposed by [Durrleman et al., 2009b] since, contrary to follow-up data, the cardiac motion is periodical, with well-defined temporal landmarks and driven by controlled biomechanical phenomena. A first direction would be to analyse the heart shape and the myocardium deformation jointly, using methods like in [Aljabar et al., 2008]. A more challenging approach would be to build a complete 4D template of the heart. The idea is to create an atlas of the heart shape from the observations. Then, the cardiac deformations of each patient would be mapped to the atlas space using parallel transport in order to create an average beating heart [DoCarmo, 1992; Rao et al., 2002]. Cross-sectional analyses would finally be applied to get a 5D model of the long-term remodelling of the beating heart.

Predicting Patient Heart Shape How can we apply the average statistical model of growth to the patient anatomy? In this thesis, we only studied the average growth of the right ventricle observed in a population. Yet, the true clinical question that motivated our work is to predict the heart shape of a specific patient. The answer to this question is far from being straightforward. First, one needs to map the growth model to the patient geometry. Again, this could be achieved using parallel transport but this would probably not be enough. Indeed, the pathological time course can vary tremendously from one patient to another, variability that may not be captured by cross-sectional analyses which tend to “smooth” the temporal observations. A solution would consist in using longitudinal data [Thompson et al., 2000; Aljabar et al., 2008; Durrleman et al., 2009b]. On the one hand, longitudinal data would make the model of heart remodelling more accurate. On the other hand, we could use the previous exams of a patient to “deform” the average growth model such that it matches the pathology evolution of that patient. In that way, acceleration or deceleration in the evolution of the pathology would be better captured.

Left-Right Ventricle Interaction in ToF Other key clinical questions related to tetralogy of Fallot and other pathologies could be investigated using the same approach. For instance, left and right ventricles could be analysed simultaneously to study the inter-ventricular interactions [Zervan et al., 2009]. Pre-

liminary results not reported in this manuscript showed that the left ventricular shape also varies over time, with visible flattening of the inter-ventricular septum as the right ventricle dilates. More thorough work is required to confirm these results but full-heart analyses would without any doubt provide more complete insights into the pathology.

Long-Term Effects of a Therapy after Remodelling A second interesting study would be to apply the proposed strategy to investigate the long-term effects of therapies on the cardiac anatomy. For instance, one could study the effects of the initial repair of ToF on the right ventricle shape and function. This would help in identifying the main factors of right ventricle degeneration, which could then be minimised. Is the ventricular patch related to the RVOT aneurysm? Similarly, the method could be applied to study the postoperative remodelling due to pulmonary valve replacement or cardiac resynchronisation.

A Statistical Physiological Model of the Heart

Finally, an appealing research direction is to investigate how statistical models of the heart growth could guide direct models of cardiac electromechanics.

A first idea would be to **use the statistical model to identify the changes in the heart shape related to external parameters of interest and to derive constitutive laws about the cardiac remodelling due to those parameters**. Those laws would reflect the visible changes of the cardiac shape and function identified by the statistics. Such an approach could help in enhancing the models of cardiac remodelling already proposed in the literature [Rodriguez et al., 1994; Kroon et al., 2007]. An alternative direction would consist in using a statistical model of heart remodelling to validate direct biomechanical models of growth as they represent the average remodelling observed in a population. At a longer term, models of cardiac remodelling could be used to predict the cardiac remodelling in patients due to a pathology or after therapy.

In parallel, **statistical models of the heart shape could be directly integrated in the electromechanical model to study the variations of the simulated cardiac functions when the anatomy varies**. Stochastic partial differential equations and polynomial chaos theory would constitute a suitable framework to mix the two approaches (see [Jakeman and Roberts, 2009] and references therein). Stochastic PDE are PDE with additional random variables. In that framework, the statistical anatomical model could be seen as the random Gaussian variable of the stochastic model. Simulations would automatically consider the random geometry and provide a set of solutions that depends on the stochastic anatomy. Such an analysis would provide precious insights into the sensitivity of the model parameters but also help in personalising the simulations to the patient cardiac function by choosing, among the results, the simulation that best matches the patient heart function.

Future

We conclude this manuscript with a long-term, blue-sky scenario. In the future, electromechanical and statistical models of the heart should play an important role in the clinical management of patients. As illustrated in Figure 10.1, electromechanical models of the heart would be able to integrate heterogeneous clinical data in a common framework, providing a global view of patient pathology. One would query them to get additional information for a more personalised cardiology. For instance, a cardiologist would input all the clinical data he acquired for a given patient into the model to get a virtual, *in-silico*, representation of the cardiac function of that patient. The *diagnosis* would be improved as the personalised virtual heart would be able to return quantitative features of the cardiac integrity that cannot be measured in clinics. The virtual heart would also provide hints on patient *prognosis* thanks to statistical or direct models of cardiac remodelling. The cardiologist would then have quantitative predictors for therapy planning. Finally, all the possible therapeutical strategies would be tested *in-silico* to choose the optimal option for the patient. The combination of medical imaging, statistical analysis and biophysical modelling promise powerful tools for a personalised computer-aided medicine.

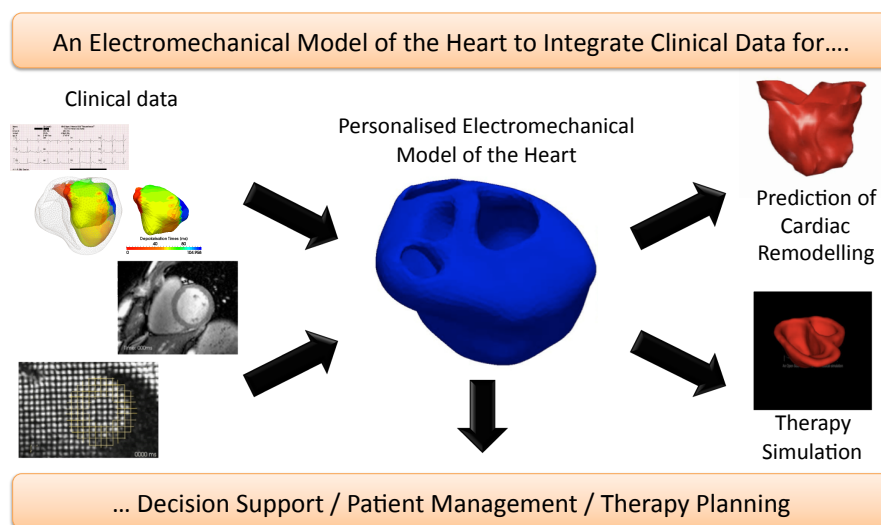


Figure 10.1: Towards a virtual physiological heart for patient-specific simulations.

10.2.1 Publications

The work realised during this thesis led to the following publications.

Article in Journals or Book Chapters

Published articles and book chapters:

1. [Mansi et al., 2009a] **T. Mansi**, B. André, M. Lynch, M. Sermesant, H. Delingette, Y. Boudjemline and N. Ayache. *Virtual Pulmonary Valve Replacement In-*

terventions with a Personalised Cardiac Electromechanical Model. In Recent Advances in the 3D Physiological Human, volume 5528 of *Lecture Notes in Computer Science*, pages 201–210. Springer, 2009.

2. [Sermesant et al., 2008] M. Sermesant, J.M. Peyrat, P. Chinchapatnam, F. Billel, **T. Mansi**, K. Rhode, H. Delingette, R. Razavi and N. Ayache. *Toward patient-specific myocardial models of the heart.* Heart Failure Clinics, vol. 4, no. 3, pages 289–301, 2008.

Submitted articles:

1. [Mansi et al., 2010a] **T. Mansi**, X. Pennec, M. Sermesant, H. Delingette et N. Ayache. *iLogDemons: A Demons-Based Registration Algorithm for Tracking Incompressible Elastic Biological Tissues.* International Journal of Computer Vision. *Submitted.*

Articles in preparation:

1. **T. Mansi et al.** *A Generative Statistical Analysis of the Heart Shape for Prediction of Remodelling and Therapy Planning.* Planned submission: August/September 2010.
2. **T. Mansi et al.** *An Image-Based Framework for Personalised Simulations of Cardiac Therapies: Application to Pulmonary Valve Replacement.* Planned submission: August/September 2010.

A clinical paper about the statistical analyses of the right ventricle shape in repaired ToF is in discussion with cardiologists.

International Conference Articles (Full, peer-reviewed and archived articles)

1. [Mansi et al., 2010b] **T. Mansi**, X. Pennec, M. Sermesant, H. Delingette et N. Ayache. *LogDemons Revisited: Consistent Regularisation and Incompressibility Constraint for Soft Tissue Tracking in Medical Images.* In Medical Image Computing and Computer Assisted Intervention (MICCAI), Lecture Notes in Computer Science. Springer, 2010. *In press.*
2. [Mansi et al., 2009c] **T. Mansi**, S. Durrleman, B. Bernhardt, M. Sermesant, H. Delingette, I. Voigt, P. Lurz, A. M Taylor, J. Blanc, Y. Boudjemline, X. Pennec and N. Ayache. *A Statistical Model of Right Ventricle in Tetralogy of Fallot for Prediction of Remodelling and Therapy Planning.* In Medical Image Computing and Computer Assisted Intervention (MICCAI), volume 5761 of *Lecture Notes in Computer Science*, pages 214–221, London, UK, September 2009. Springer.
3. [Mansi et al., 2009d] **T. Mansi**, J.M. Peyrat, M. Sermesant, H. Delingette, J. Blanc, Y. Boudjemline and N. Ayache. *Physically-Constrained Diffeomorphic*

Demons for the Estimation of 3D Myocardium Strain from Cine-MRI. In Proceedings of Functional Imaging and Modeling of the Heart 2009 (FIMH'09), volume 5528 of *Lecture Notes in Computer Science*, pages 201–210, 3-5 June 2009. Springer.

4. [Pop et al., 2009] M. Pop, M. Sermesant, **T. Mansi**, E. Crystal, J. Det-sky, Y. Yang, P. Fefer, E.R. McVeigh, A. Dick, N. Ayache and G. A. Wright *Characterization of Post-infarct Scars in a Porcine Model – A Combined Experimental and Theoretical Study.* In Proceedings of Functional Imaging and Modeling of the Heart 2009 (FIMH'09), volume 5528 of *Lecture Notes in Computer Science*, page 10, 3-5 June 2009. Springer.
5. [Sermesant et al., 2009] M. Sermesant, F. Billet, R. Chabiniok, **T. Mansi**, P. Chinchapatnam, P. Moireau, J.M. Peyrat, K. Rhode, M. Ginks, P. Lambiase, S. Arridge, H. Delingette, M. Sorine, A. Rinaldi, D. Chapelle, R. Razavi and N. Ayache. *Personalised Electromechanical Model of the Heart for the Prediction of the Acute Effects of Cardiac Resynchronisation Therapy.* In Proceedings of Functional Imaging and Modeling of the Heart 2009 (FIMH'09), volume 5528 of *Lecture Notes in Computer Science*, pages 239–248, 3-5 June 2009. Springer.
6. [Chinchapatnam et al., 2009] P. Chinchapatnam, K. Rhode, M. Ginks, **T. Mansi**, J.M. Peyrat, P. Lambiase, A. Rinaldi, R. Razavi, S. Arridge et M. Sermesant. *Estimation of Volumetric Myocardial Apparent Conductivity from Endocardial Electro-Anatomical Mapping.* In Proceedings of 31st Annual International IEEE on Engineering in Medicine and Biology Society Conference (EMBS), page 2907, 2009.

International Conference Articles (Full and peer-reviewed articles)

1. [Lamecker et al., 2009] H. Lamecker, **T. Mansi**, J. Relan, F. Billet, M. Sermesant, N. Ayache and H. Delingette. *Adaptive Tetrahedral Meshing for Personalized Cardiac Simulations.* In MICCAI Workshop on Cardiovascular Interventional Imaging and Biophysical Modelling (CI2BM09), pages 149–158, London United Kingdom, 2009.
2. [Toussaint et al., 2008] N. Toussaint, **T. Mansi**, H. Delingette, N. Ayache and M. Sermesant. *An Integrated Platform for Dynamic Cardiac Simulation and Image Processing: Application to Personalised Tetralogy of Fallot Simulation.* In Proc. Eurographics Workshop on Visual Computing for Biomedicine (VCBM), Delft, The Netherlands, 2008.

Articles and Abstracts in French Clinical Journals

1. [Mansi et al., 2010c] **T. Mansi**, M. Sermesant, H. Delingette, X. Pennec, N. Ayache et Boudjemline Y. *In-Silico Models for the Simulation and Prediction*

- of the Cardiac Function*. Archives de Pédiatrie, vol. 17, no. 6, pages 611-612, Juin 2010.
2. [Mansi et al., 2009b] **T. Mansi**, B. André, M. Sermesant, H. Delingette, N. Ayache and Y. Boudjemline. *Simulation personnalisée de remplacements valvulaires pulmonaires grâce à l'utilisation d'un modèle mathématique du coeur*. Archives de Pédiatrie - Congrès des Sociétés Françaises Médico-chirurgicales Pédiatriques, June 2009.
 3. [Mansi et al., 2008] **T. Mansi**, M. Sermesant, M. Huber, A. Taylor, G. Pongiglione, X. Pennec and Y. Boudjemline. *Modélisation électromécanique du coeur et analyse d'images*. Archives de Pédiatrie - Congrès des Sociétés Françaises Médico-chirurgicales Pédiatriques, February 2008. Archives de Pédiatrie 15(5):1032 – Special award for best presentation.

Others

1. [Pennec et al., 2007-2009-2010] X. Pennec, C. Basso, Y. Boudjemline, S. Durleman, E. Konukoglu, **T. Mansi**, G. Pongiglione, M. Santoro, M. Sermesant, B. Stos, N. Toussaint and G. Trocchio. *Third generation disease models: Image analysis tools, pediatric heart diseases, inflammatory diseases and brain tumors*. Deliverables D11.2, D11.3 and D11.4⁸, European project Health-e-Child (IST-2004-027749), June 2007-2009-2010.

Invited Talks

1. *In-Silico Models for the Simulation and Prediction of the Cardiac Function in Patients*, Congress of French Paediatrics Society, Round-table conference about “*In-Silico Models in Pediatrics*”, Paris, France, 16 June, 2010.
2. *Personalised Simulation of Pulmonary Valve Replacement in Repaired Tetralogy of Fallot Patients*, Necker-Enfants Malades, Hôpitaux de Paris, Paris, France, 15 October 2008. Hosted by: Dr. Younes Boudjemline, M.D.

⁸<http://www-sop.inria.fr/asclepios/projects/Health-e-Child/DiseaseModels>

Part VI

Appendices

On the Image Noise Parameter of the LogDemons Algorithm

Contents

A.1 Motivation	231
A.2 Methods	231
A.2.1 Current Formulation of Demons Optimisation	231
A.2.2 On the Importance of the Noise Parameter σ_i^2	232
A.2.3 Levenberg Demons Optimisation	233
A.3 Preliminary Experiments and Discussions	234
A.4 Conclusions	235

A.1 Motivation

In the logDemons algorithm (Chapter 4), the level of noise in the images, σ_i^2 , is estimated locally, voxel by voxel, to control the amplitude of the update velocity. Although this definition implicitly ensures the numerical stability of the algorithm (it constraints the gradient to be finite), it is mathematically inconsistent as σ_i^2 is spatially-dependent and varies with the demons iterations. Region-wise estimation of the image noise should be preferred but this may yield numerical instabilities as the amplitude of the update velocity and its gradient are not bounded anymore. In this appendix we propose a formulation that enables one to estimate the noise in the images in a region-wise manner to have a theoretically grounded optimisation method, without hampering the stability and the performances of the algorithm. The idea is to modify the Gauss-Newton scheme proposed by [Vercauteren et al., 2008, 2009] (Chapter 4), by a Levenberg strategy that always ensures bounded update velocity.

A.2 Methods

A.2.1 Current Formulation of Demons Optimisation

The logDemons algorithm alternately minimises the energy functional:

$$\mathcal{E}(\mathbf{v}_c, \mathbf{v}) = \frac{1}{\sigma_i^2} \|R - T \circ \exp(\mathbf{v}_c)\|_{L_2}^2 + \frac{1}{\sigma_x^2} \|\log(\exp(-\mathbf{v}) \circ \exp(\mathbf{v}_c))\|_{L_2}^2 + \frac{1}{\sigma_d^2} \mathcal{R}(\mathbf{v})$$

In the previous equation, R is the fixed reference image, T is the moving template image, \mathbf{v} and \mathbf{v}_c are the stationary velocity fields that parameterise the deformation field $\phi = \exp(\mathbf{v})$ and the correspondence field $\phi_c = \exp(\mathbf{v}_c)$ respectively. $\mathcal{R}(\mathbf{v})$ is the regulariser. σ_i^2 accounts for the noise in the images. The question we address is how to estimate σ_i^2 from the images R and $T \circ \phi$ in a consistent way.

We focus on the optimisation step, which minimises $\mathcal{E}(\mathbf{v}_c, \mathbf{v})$ with respect to \mathbf{v}_c to calculate the optimal correspondence field ϕ_c . This is achieved using the diffeomorphic update rule $\phi_c = \phi \circ \exp(\delta\mathbf{v})$, where ϕ is the current estimate of the transformation and $\delta\mathbf{v}$ is the optimal update velocity to find (see Chapter 4). With this rule, the correspondence energy writes:

$$\mathcal{E}_{corr}(\delta\mathbf{v}) = \frac{1}{\sigma_i^2} \|R - T \circ \phi \circ \exp(\delta\mathbf{v})\|_{L_2}^2 + \frac{1}{\sigma_x^2} \|\delta\mathbf{v}\|_{L_2}^2 \quad (\text{A.1})$$

A closed form minimiser of this energy is found with a Gauss-Newton approach. The first term is linearised using the ESM scheme proposed in [Vercauteren et al., 2009]. Let $J(\mathbf{x})$ be the symmetric gradient at the spatial position \mathbf{x} defined by $J(\mathbf{x}) = (\nabla R(\mathbf{x}) + \nabla(T \circ \phi)(\mathbf{x}))/2$. The optimal condition $\partial\mathcal{E}_{corr} = 0$ writes:

$$\left(JJ^T + \frac{\sigma_i^2}{\sigma_x^2} \text{Id} \right) \delta\mathbf{v} = -(R - T \circ \phi)J \quad (\text{A.2})$$

which gives the optimal Gauss-Newton update:

$$\delta\mathbf{v}(\mathbf{x}) = -\frac{R(\mathbf{x}) - T \circ \phi(\mathbf{x})}{\|J(\mathbf{x})\|^2 + \sigma_i^2/\sigma_x^2} J(\mathbf{x}) \quad (\text{A.3})$$

In [Cachier et al., 1999; Vercauteren et al., 2009], the authors estimate σ_i^2 locally, at every voxel of the images:

$$\sigma_i^2(\mathbf{x}) = |R(\mathbf{x}) - T \circ \phi(\mathbf{x})|^2 \quad (\text{A.4})$$

A.2.2 On the Importance of the Noise Parameter σ_i^2

With the local noise estimator (Eq. A.4), it can be demonstrated that the norm of the Gauss-Newton update is always upper bounded by $\sigma_x/2$ [Cachier et al., 1999; Vercauteren et al., 2009]. This guarantees the stability of the algorithm. However, σ_i^2 depends on the spatial position \mathbf{x} and on the current estimate of the transformation ϕ . Therefore, it must be considered when calculating the functional derivative of the correspondence energy (Eq. A.1). Moreover, the estimated noise is sensitive to local outliers (for instance spikes) as it does not consider neighbouring information. A more regional estimator of the image noise is thus needed.

Lets consider the following region-wise noise estimator:

$$\sigma_i^2(\mathbf{x}) = [(G_{\sigma_n} \star |R - T \circ \phi|)(\mathbf{x})]^2 \quad (\text{A.5})$$

In the previous equation, G_{σ_n} is a Gaussian kernel with standard deviation σ_n , which defines the region around the spatial position \mathbf{x} where the noise is estimated.

This definition is more grounded from an estimation point of view. When σ_n is sufficiently high, σ_i^2 does not depend on \mathbf{x} and ϕ at the first order, locally. The Gauss-Newton update (Eq. A.3) is valid at the first order. However, its amplitude is not bounded anymore. Lets $\alpha = |R - T \circ \phi|$. We have:

$$\|\delta \mathbf{v}(\mathbf{x})\| = \frac{\alpha(\mathbf{x}) \|J(\mathbf{x})\|}{\|J(\mathbf{x})\|^2 + (G_{\sigma_n} \star \alpha)^2 / \sigma_x^2}$$

which is upper bounded by (according to the inequality $a^2 + b^2 > 2ab$)

$$\|\delta \mathbf{v}(\mathbf{x})\| \leq \frac{\sigma_x}{2} \frac{\alpha(\mathbf{x})}{(G_{\sigma_n} \star \alpha)(\mathbf{x})}$$

In this equation we can see that if the noise is estimated locally using Eq. A.4, the update velocity is upper bounded by $\|\delta \mathbf{v}(\mathbf{x})\| < \sigma_x/2$. This is not the case anymore with the proposed regional noise. For instance, when the algorithm is close to convergence, R and $T \circ \phi$ are similar and $\|J(\mathbf{x})\| \rightarrow 0$. If the images present white noise (independent at each voxel), we can have $\alpha(\mathbf{x}) \neq 0$ at \mathbf{x} but $(G_{\sigma_n} \star \alpha)(\mathbf{x})$ very close to 0. $\alpha(\mathbf{x})/(G_{\sigma_n} \star \alpha)(\mathbf{x})$ thus tends to $+\infty$. This behaviour seriously hampers the stability of the algorithm. Similar problems occur when the image similarity metric is changed from the Sum of Squared Differences (SSD) to the Sum of Absolute Differences (SAD) or multimodal ones like mutual information.

A.2.3 Levenberg Demons Optimisation

As a solution to the above-mentioned stability issue we propose a Levenberg approach. The idea is to damp the optimal condition (Eq. A.2) such that the numerical stability of the algorithm is ensured everywhere in the images, at any iteration. A similar approach was already employed in [Yeo et al., 2009] to ensure bounded Gauss-Newton updates calculated on a non-regular grid. The Levenberg modification of the optimal condition writes:

$$\left(JJ^T + \frac{\sigma_i^2}{\sigma_x^2} \text{Id} + \lambda \text{Id} \right) \delta \mathbf{v} = -(R - T \circ \phi)J \quad (\text{A.6})$$

where λ is a non-negative damping factor. We choose λ such that the update velocity $\delta \mathbf{v}$ is always bounded. We have:

$$\delta \mathbf{v}(\mathbf{x}) = -\frac{R(\mathbf{x}) - T \circ \phi(\mathbf{x})}{\|J(\mathbf{x})\|^2 + \sigma_i^2/\sigma_x^2 + \lambda} J(\mathbf{x}) \quad (\text{A.7})$$

whose norm is $\|\delta \mathbf{v}(\mathbf{x})\| = \alpha(\mathbf{x}) \|J(\mathbf{x})\| / (\|J(\mathbf{x})\|^2 + \sigma_i^2/\sigma_x^2 + \lambda)$. By applying the same technique as in the previous section, an upper bound of $\|\delta \mathbf{v}(\mathbf{x})\|$ is found:

$$\|\delta \mathbf{v}(\mathbf{x})\| \leq \frac{\alpha}{2\sqrt{\sigma_i^2/\sigma_x^2 + \lambda}} \quad (\text{A.8})$$

The norm of the Gauss-Newton update is upper bounded by $\sigma_x/2$ if the damping factor verifies:

$$\frac{\alpha^2 - \sigma_i^2}{\sigma_x^2} < \lambda \quad \text{and} \quad 0 < \lambda \quad (\text{A.9})$$

This condition is valid for any estimator of the image noise σ_i^2 . As λ must be positive, when $\alpha^2 \leq \sigma_i^2$, we set $\lambda = 0$.

Optimal convergence speed is achieved when the norm of the Gauss-Newton update is maximal, under the constraint of numerical stability. We thus define λ equal to the lower bound given by Eq. A.9 and the stabilised Gauss-Newton update becomes:

$$\delta \mathbf{v}(\mathbf{x}) = -\frac{R(\mathbf{x}) - T \circ \phi(\mathbf{x})}{\|J(\mathbf{x})\|^2 + \max(\sigma_i^2, \alpha^2)/\sigma_x^2} J(\mathbf{x}) \quad (\text{A.10})$$

When the local noise estimator is used (Eq. A.4), $\sigma_i^2 = \alpha^2$, we retrieve the original update velocity.

A.3 Preliminary Experiments and Discussions

The proposed Levenberg stabilisation was tested on synthetic data with known ground truth. A 3D isotropic Steady-State Free Precession (SSFP) MR image of the heart ($53 \times 60 \times 60$ slices, 1 mm^3 isotropic voxel spacing), henceforth called test image, was warped by a random diffeomorphic deformation field computed from a random velocity field. The test image and the warped image were then altered with a slight Gaussian noise (Figure A.1).




Figure A.1: Synthetic 3D image warped with a random diffeomorphic deformation field (represented by a warped grid).

The test image was registered to the warped image using the original local noise estimator, $\sigma_i(\mathbf{x}) = |R(\mathbf{x}) - T \circ \phi(\mathbf{x})|^2$, and a global estimator of the image noise defined as the average over the image domain of the squared intensity differences, $\sigma_i(\mathbf{x}) = \text{mean}(|R(\mathbf{x}) - T \circ \phi(\mathbf{x})|^2)$, which is thus constant. The global estimator is tested with and without the proposed Levenberg stabilisation.

Registration results are reported in Figure A.2. As one can see, the non-stabilised global noise estimator diverges after some iterations, when the images are almost registered (low MSE). This is exactly the situation described in Section A.2.2. As the registration evolves, the two images become similar. The global noise and the

amplitude of the gradient J thus tend to 0, which may yield unbounded update velocity at the grey level spikes. Figure A.3, mid panel, illustrates the resulting deformation. In this example, strong, non diffeomorphic deformations appeared at the image boundaries. The proposed stabilisation approach solves this issue. Numerical stability is guaranteed. Obtained registration is very similar to the reference result obtained with the local noise estimator. Non-reported experiments with a regional image noise estimator resulted in similar conclusions.

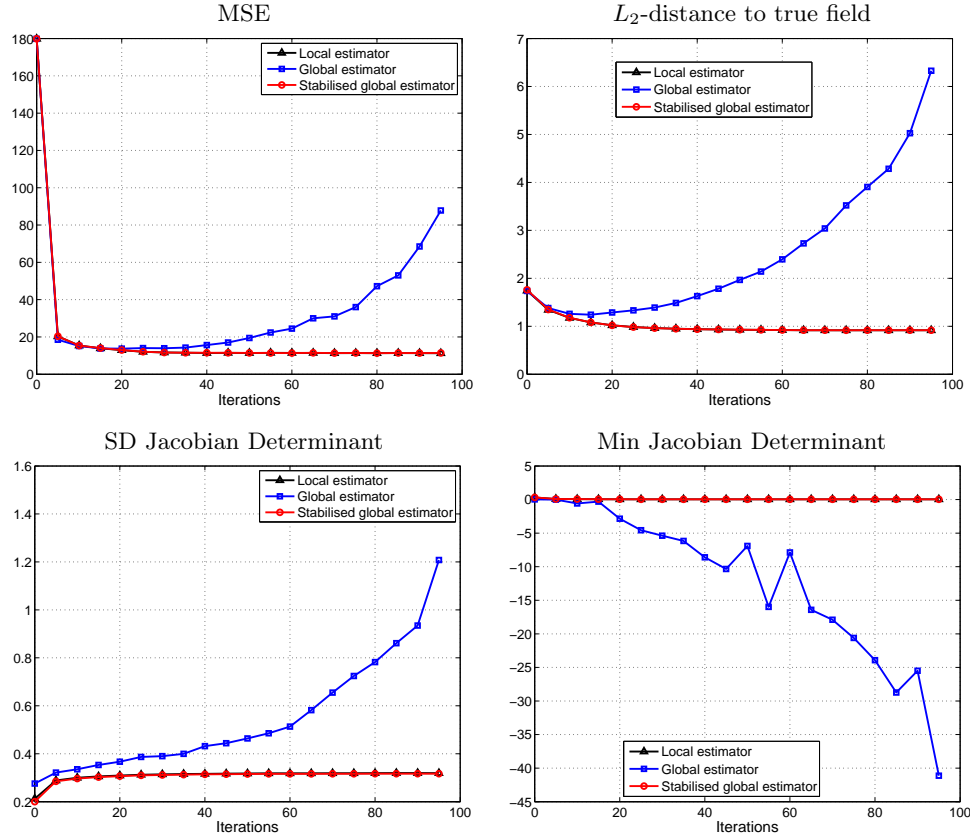


Figure A.2: Without Levenberg-like stabilisation, using a global estimator of image noise yields numerical instabilities after few iterations (deviation from the optimal solution, negative Jacobian determinants). The proposed stabilisation prevents the instabilities and ensures similar performances as the local noise estimator.

A.4 Conclusions

To conclude, by decoupling the noise estimation problem from the algorithm stabilisation with bounded updates, the Levenberg method that we propose is providing a generic and theoretically fully grounded optimisation method without any loss of performances. This opens new possibilities for the generalisation of the demons. Other more advanced noise estimators could be used to cope with intensity non-

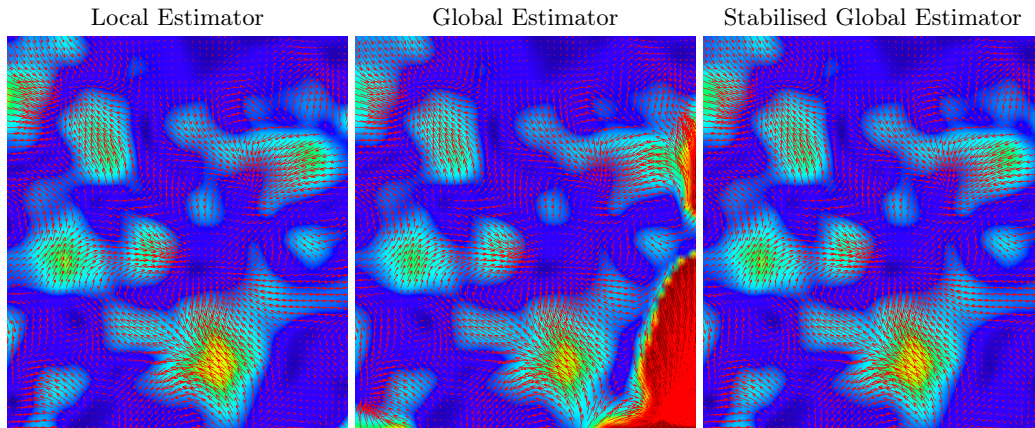


Figure A.3: Estimated deformation fields. The proposed stabilisation method prevents the strong deformations that appear, in this example, at the image boundaries.

uniformity artefacts. Similarly, different image similarity metrics can now be used as we can bound the updates whatever criterion the forces are coming from, without having to tweak with a theoretically unjustified point-wise noise estimation.

Bibliography

- N. Ablitt, J. Gao, J. Keegan, L. Stegger, D. Firmin, and G. Yang. Predictive cardiac motion modeling and correction with partial least squares regression. *IEEE transactions on medical imaging*, 23(10):1315–1324, 2004. [126](#)
- H. Akaike. A new look at the statistical model identification. *IEEE transactions on automatic control*, 19(6):716–723, 1974. [124](#)
- A. Aletras, S. Ding, R. Balaban, and H. Wen. Dense: displacement encoding with stimulated echoes in cardiac functional mri. *Journal of Magnetic Resonance*, 137(1):247–252, 1999. [24](#)
- R. R. Aliev and A. V. Panfilov. A simple two-variable model of cardiac excitation. *Chaos, Solitons & Fractals*, 7(3):293–301, March 1996. [167](#), [188](#)
- P. Aljabar, K. Bhatia, M. Murgasova, J. Hajnal, J. Boardman, L. Srinivasan, M. Rutherford, L. Dyet, A. Edwards, and D. Rueckert. Assessment of brain growth in early childhood using deformation-based morphometry. *Neuroimage*, 39(1):348–358, 2008. [125](#), [222](#), [223](#)
- J. Allard, S. Cotin, F. Faure, P.-J. Bensoussan, F. Poyer, C. Duriez, H. Delingette, and L. Grisoni. SOFA – an Open Source Framework for Medical Simulation. In *Medicine Meets Virtual Reality (MMVR’15)*, 2007. [194](#)
- S. Allasonniere, Y. Amit, and A. Troune. Towards a coherent statistical framework for dense deformable template estimation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69(1):3–29, 2007. [109](#), [111](#)
- B. Andre and H. Delingette. Versatile design of changing mesh topologies for surgery simulation. *Lecture Notes in Computer Science*, 5104:147–156, 2008. [195](#)
- A. Andreopoulos and J. Tsotsos. Efficient and generalizable statistical models of shape and appearance for analysis of cardiac MRI. *Medical Image Analysis*, 12(3):335–357, 2008. [34](#)
- V. Arnold. *Mathematical methods of classical mechanics*. Springer-Verlag, 1989. [62](#)
- V. Arsigny, O. Commowick, X. Pennec, and N. Ayache. A log-euclidean framework for statistics on diffeomorphisms. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, volume 4190 of *Lecture Notes in Computer Science*, page 924. Springer, 2006a. [62](#), [65](#), [83](#)
- V. Arsigny, P. Fillard, X. Pennec, and N. Ayache. Log-Euclidean metrics for fast and simple calculus on diffusion tensors. *Magnetic Resonance in Medicine*, 56(2):411–421, 2006b. [87](#), [222](#)

- T. Arts, R. Reneman, and P. Veenstra. A model of the mechanics of the left ventricle. *Annals of biomedical engineering*, 7(3):299–318, 1979. 171
- T. Arts, P. Bovendeerd, F. Prinzen, and R. Reneman. Relation between left ventricular cavity pressure and volume and systolic fiber stress and strain in the wall. *Biophysical journal*, 59(1):93–102, 1991. 171
- T. Arts, K. D. Costa, J. W. Covell, and A. D. McCulloch. Relating myocardial laminar architecture to shear strain and muscle fiber orientation. *Am J Physiol Heart Circ Physiol*, 280(5):H2222–2229, 2001. 159, 165, 166, 190
- T. Arts, T. Delhaas, P. Bovendeerd, X. Verbeek, and F. Prinzen. Adaptation to mechanical load determines shape and properties of heart and circulation: the circadapt model. *American Journal of Physiology- Heart and Circulatory Physiology*, 288(4):H1943, 2005. 179, 189
- T. Arts, F. Prinzen, T. Delhaas, J. Milles, A. Rossi, and P. Clarysse. Mapping displacement and deformation of the heart with local sine-wave modeling. *IEEE Transactions on Medical Imaging*, 2010. 103
- K. Arun, T. Huang, and S. Blostein. Least-squares fitting of two 3-D point sets. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(5):698–700, September 1987. 112, 136
- J. Ashburner, C. Hutton, R. Frackowiak, I. Johnsrude, C. Price, and K. Friston. Identifying global anatomical differences: deformation-based morphometry. *Human Brain Mapping*, 6(5-6):348–357, 1998. 62, 108
- N. Ayache, editor. *Computational Models for the Human Body*. Handbook of Numerical Analysis (Ph. Ciarlet series editor). Elsevier, 2004. 670 pages. 156
- N. Ayache, J.-P. Boissel, S. Brunak, G. Clapworthy, G. Lonsdale, J. Fingberg, A. Frangi, G. Deco, P. Hunter, P. Nielsen, M. Halstead, R. Hose, I. Magnin, F. Martin-Sanchez, P. Sloot, J. Kaandorp, A. Hoekstra, S. Van Sint Jan, and M. Viceconti. 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, 2006. 156, 213, 220
- N. Ayache, O. Clatz, H. Delingette, G. Malandain, X. Pennec, and M. Sermesant. Asclepios: a Research Project-Team at INRIA for the Analysis and Simulation of Biomedical Images. In Y. Bertot, G. Huet, J.-J. Lévy, and G. Plotkin, editors, *From semantics to computer science: essays in honor of Gilles Kahn*, pages 415–436. Cambridge University Press, 2009. 156
- P. Basser, J. Mattiello, and D. LeBihan. Mr diffusion tensor spectroscopy and imaging. *Biophysical journal*, 66(1):259–267, 1994. 25
- J. Bassingthwaite, P. Hunter, and D. Noble. The cardiac physiome: perspectives for the future. *Experimental Physiology*, 94(5):597, 2009. 156, 213, 220

- K. Bathe. *Finite Element Procedures*. Prentice-Hall, 1996. 157
- M. Beg, P. Helm, E. McVeigh, M. Miller, and R. Winslow. Computational cardiac anatomy using MRI. *Magnetic resonance in medicine: official journal of the Society of Magnetic Resonance in Medicine/Society of Magnetic Resonance in Medicine*, 52(5):1167, 2004. 125, 151
- M. Beg, M. Miller, A. Trounev , and L. Younes. Computing Large Deformation Metric Mappings via Geodesic Flows of Diffeomorphisms. *International Journal of Computer Vision*, 61(2):139–157, 2005. 62, 83, 114
- T. Belytschko, Y. Krongauz, D. Organ, M. Fleming, and P. Krysl. Meshless methods: an overview and recent developments. *Computer methods in applied mechanics and engineering*, 139(1-4):3–47, 1996. 158
- J. Bestel. *Mod le diff rentiel de la contraction musculaire contr l e. Application au syst me cardiovasculaire*. PhD thesis, Universit  Paris IX Dauphine, 2000. 175, 176, 190
- J. Bestel, F. Clement, and M. Sorine. A biomechanical model of muscle contraction. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, Lecture Notes in Computer Science, pages 1159–1161. Springer, 2001. 156, 171, 172, 188, 210
- F. Billet. *Assimilation de donn es images pour la personnalisation d’un mod le  lectrom canique du c ur*. PhD thesis, Universit  de Nice-Sophia Antipolis, 2010. 179, 189, 221, 222
- F. Billet, M. Sermesant, H. Delingette, and N. Ayache. Cardiac motion recovery by coupling an electromechanical model and cine-MRI data: First steps. In K. Miller and P. Nielsen, editors, *Proc. of the Workshop on Computational Biomechanics for Medicine III. (Workshop MICCAI-2008)*, September 2008. 55, 56, 176, 214
- F. Billet, M. Sermesant, H. Delingette, and N. Ayache. Cardiac motion recovery and boundary conditions estimation by coupling an electromechanical model and cine-MRI data. In *Proceedings of Functional Imaging and Modeling of the Heart 2009 (FIMH’09)*, volume 5528 of *Lecture Notes in Computer Science*, pages 376–385, 3-5 June 2009. 55, 56, 174, 176, 188, 214, 217
- A. Bistoquet, J. Oshinski, and O. Skrinjar. Left ventricular deformation recovery from cine MRI using an incompressible model. *IEEE Transactions on Medical Imaging*, 26(9):1136–1153, 2007. 34, 45
- A. Bistoquet, J. Oshinski, and O. Skrinjar. Myocardial deformation recovery from cine MRI using a nearly incompressible biventricular model. *Medical Image Analysis*, 12(1):69–85, 2008. 62, 63, 82, 90

- R. Blanc, M. Reyes, C. Seiler, and G. Székely. Conditional variability of statistical shape models based on surrogate variables. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, Lecture Notes in Computer Science, pages 84–91. Springer, 2009. 119
- J. Bloomenthal and C. Bajaj. *Introduction to implicit surfaces*. Morgan Kaufmann, 1997. 40
- N. Bodhey, P. Beerbaum, S. Sarikouch, S. Kropf, P. Lange, F. Berger, R. Anderson, and T. Kuehne. Functional Analysis of the Components of the Right Ventricle in the Setting of Tetralogy of Fallot. *Circulation: Cardiovascular Imaging*, 1(2):141, 2008. 134, 140, 144, 148, 150, 218
- H. Bogren and M. Buonocore. 4D magnetic resonance velocity mapping of blood flow patterns in the aorta in young vs. elderly normal subjects. *Journal of Magnetic Resonance Imaging*, 10(5):861–869, 1999. 24
- J. Boisvert, F. Cheriet, X. Pennec, H. Labelle, and N. Ayache. Geometric variability of the scoliotic spine using statistics on articulated shape models. *IEEE Transactions on Medical Imaging*, 27(4):557–568, 2008. 119
- F. Bookstein. Size and shape spaces for landmark data in two dimensions. *Statistical Science*, 1(2):181–222, 1986. 108
- G. Borgefors. On digital distance transforms in three dimensions. *Computer Vision and Image Understanding*, 64(3):368–376, 1996. 50
- M. Bossa, M. Hernandez, and S. Olmos. Contributions to 3D diffeomorphic atlas estimation: application to brain images. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, volume 10 of *Lecture Notes in Computer Science*, page 667. Springer, 2007. 62, 66
- Y. Bourgault, Y. Coudière, and C. Pierre. Existence and uniqueness of the solution for the bidomain model used in cardiac electrophysiology. *Nonlinear analysis: Real world applications*, 10(1):458–482, 2009. 167
- J. Bresenham. Algorithm for computer control of a digital plotter. *IBM Systems journal*, 4(1):25–30, 1965. 163
- C. Broit. *Optimal registration of deformed images*. PhD thesis, University of Pennsylvania Philadelphia, PA, USA, 1981. 63
- P. Cachier and N. Ayache. Isotropic energies, filters and splines for vectorial regularization. *J. of Math. Imaging and Vision*, 20(3):251–265, 2004. 63, 64, 68, 69, 70, 80, 214
- P. Cachier, X. Pennec, and N. Ayache. Fast Non Rigid Matching by Gradient Descent: Study and Improvements of the Demons Algorithm. Tech. Report RR-3706, INRIA, 06 1999. 66, 232

- P. Cachier, E. Bardinet, D. Dormont, X. Pennec, and N. Ayache. Iconic feature based nonrigid registration: the PASHA algorithm. *Computer Vision and Image Understanding*, 89(2-3):272–298, 2003. [62](#), [65](#), [67](#)
- N. D. Cahill, J. A. Noble, and D. J. Hawkes. A demons algorithm for image registration with locally adaptive regularization. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, volume 5761 of *Lecture Notes in Computer Science*, pages 574–581. Springer, 2009. ISBN 978-3-642-04267-6. [63](#), [68](#), [83](#), [214](#)
- D. Caillerie, A. Mourad, and A. Raoult. Cell-to-muscle homogenization. application to a constitutive law for the myocardium. *Mathematical Modelling and Numerical Analysis*, 37(4):681–698, 2003. [171](#)
- V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *International journal of computer vision*, 22(1):61–79, 1997. [47](#)
- J. Cates, P. Fletcher, and R. Whitaker. A hypothesis testing framework for high-dimensional shape models. In *MICCAI Workshop on Mathematical Foundations of Computational Anatomy*, page 170, 2008. [118](#)
- M. Cerqueira, N. Weissman, V. Dilsizian, A. Jacobs, S. Kaul, W. Laskey, D. Pennell, J. Rumberger, T. Ryan, and M. Verani. Standardized myocardial segmentation and nomenclature for tomographic imaging of the heart a statement for health-care professionals from the cardiac imaging committee of the council on clinical cardiology of the american heart association. *Circulation*, 105(4):539–542, 2002. [20](#), [21](#), [93](#), [163](#)
- R. Chabiniok, D. Chapelle, P. Lesault, A. Rahmouni, and J. Deux. Validation of a biomechanical heart model using animal data with acute myocardial infarction. In *CI2BM09 - MICCAI Workshop on Cardiovascular Interventional Imaging and Biophysical Modelling*, London United Kingdom, 2009. [55](#), [56](#), [214](#)
- J. Chambers. *Echocardiography in clinical practice*. Informa HealthCare, 2001. [20](#)
- A. Chandler, R. Pinder, T. Netsch, J. Schnabel, D. Hawkes, D. Hill, and R. Razavi. Correction of misaligned slices in multi-slice cardiovascular magnetic resonance using slice-to-volume registration. *Journal of Cardiovascular Magnetic Resonance*, 10(1):13, 2008. [36](#)
- R. Chandrashekara, R. Mohiaddin, and D. Rueckert. Analysis of 3-D myocardial motion in tagged MR images using nonrigid image registration. *IEEE Transactions on Medical Imaging*, 23(10):1245–1250, 2004. [103](#)
- D. Chapelle, F. Clément, F. Génot, P. Tallec, M. Sorine, and J. Urquiza. A physiologically-based model for the active cardiac muscle contraction. *Functional Imaging and Modeling of the Heart*, pages 128–133, 2001. [171](#)

- J. Chen, W. Liu, H. Zhang, L. Lacy, X. Yang, S. Song, S. Wickline, and X. Yu. Regional ventricular wall thickening reflects changes in cardiac fiber and sheet structure during contraction: quantification with diffusion tensor mri. *American Journal of Physiology- Heart and Circulatory Physiology*, 289(5):H1898, 2005. 26
- P. Chinchapatnam, K. Rhode, M. Ginks, C. Rinaldi, P. Lambiase, R. Razavi, S. Arridge, and M. Sermesant. Model-based imaging of cardiac apparent conductivity and local conduction velocity for diagnosis and planning of therapy. *IEEE Transactions on Medical Imaging*, 27(11):1631–1642, 2008. 168, 169, 221
- P. Chinchapatnam, K. Rhode, M. Ginks, T. Mansi, J. Peyrat, P. Lambiase, A. Rinaldi, R. Razavi, S. Arridge, and M. Sermesant. Estimation of volumetric myocardial apparent conductivity from endocardial electro-anatomical mapping. In *Proceedings of 31st Annual International IEEE on Engineering in Medicine and Biology Society Conference (EMBS)*, page 2907, 2009. 217, 227
- E. Chung, A. Leon, L. Tavazzi, J. Sun, P. Nihoyannopoulos, J. Merlino, W. Abraham, S. Ghio, C. Leclercq, J. Bax, et al. Results of the predictors of response to crt (prospect) trial. *Circulation*, 117(20):2608, 2008. 22, 23, 100
- O. Clatz, H. Delingette, I. Talos, A. Golby, R. Kikinis, F. Jolesz, N. Ayache, and S. Warfield. Robust nonrigid registration to capture brain shift from intraoperative MRI. *IEEE Transactions on Medical Imaging*, 24(11):1417–1427, 2005. 64
- R. Clayton and A. Panfilov. A guide to modelling cardiac electrical activity in anatomically detailed ventricles. *Progress in biophysics and molecular biology*, 96(1-3):19–43, 2008. 167, 168
- P. Colli Franzone, L. Guerri, and S. Rovida. Wavefront propagation in an activation model of the anisotropic cardiac tissue: asymptotic analysis and numerical simulations. *Journal of mathematical biology*, 28(2):121–176, 1990. 167
- O. Comas, Z. Taylor, J. Allard, S. Ourselin, S. Cotin, and J. Passenger. Efficient nonlinear fem for soft tissue modelling and its gpu implementation within the open source framework sofa. In *International Symposium on Computational Models for Biomedical Simulation (ISBMS)*, volume 5104 of *Lecture Notes in Computer Science*, pages 28–39. Springer, 2008. 195
- T. Cootes, C. Taylor, D. Cooper, J. Graham, et al. Active shape models-their training and application. *Computer vision and image understanding*, 61(1):38–59, 1995. 34, 108
- T. Cootes, G. Edwards, and C. Taylor. Active appearance models. *Computer Vision—ECCV’98*, page 484, 1998. 34
- K. Costa, Y. Takayama, A. McCulloch, and J. Covell. Laminar fiber architecture and three-dimensional systolic mechanics in canine ventricular myocardium. *American Journal of Physiology- Heart and Circulatory Physiology*, 276(2):H595, 1999. 173, 188, 190

- K. Costa, J. Holmes, and D. McCulloch. Modelling cardiac mechanical properties in three dimensions. *Philosophical Transactions A*, 359(1783):1233, 2001. 171, 188, 210, 220
- Y. Coudière and C. Pierre. Stability and convergence of a finite volume method for two systems of reaction-diffusion equations in electro-cardiology. *Nonlinear analysis: real world applications*, 7(4):916–935, 2006. 167
- J. Cousty, L. Najman, M. Couprie, S. Clément-Guinaudeau, T. Goissen, and J. Garot. Segmentation of 4d cardiac mri: Automated method based on spatio-temporal watershed cuts. *Image and Vision Computing*, 2010. 34
- E. Crampin, M. Halstead, P. Hunter, P. Nielsen, D. Noble, N. Smith, and M. Tawhai. Computational physiology and the physiome project. *Experimental Physiology*, 89(1):1, 2004. 156
- E. Crépeau and M. Sorine. A reduced model of pulsatile flow in an arterial compartment. *Chaos, Solitons & Fractals*, 34(2):594–605, 2007. 182
- A. Cuzol, P. Hellier, and E. Mémin. A low dimensional fluid motion estimator. *International Journal of Computer Vision*, 75(3):329–349, 2007. 63
- R. Davies, T. Cootes, and C. Taylor. A minimum description length approach to statistical shape modelling. In *Information Processing in Medical Imaging*, Lecture Notes in Computer Science, pages 50–63. Springer, 2001. 108
- J. Declerck, J. Feldmar, and N. Ayache. Definition of a 4d continuous polar transformation for the tracking and the analysis of lv motion. In *CVRMed-MRCAS’97*, pages 33–42. Springer, 1997. 57
- M. DeCraene, O. Camara, B. Bijmens, and A. Frangi. Large diffeomorphic ffd registration for motion and strain quantification from 3d-us sequences. In *Functional Imaging and Modeling of the Heart: 5th International Conference, Fimh 2009 Nice, France, June 3-5, 2009 Proceedings*, page 437. Springer-Verlag New York Inc, 2009. 63, 103
- P. del Nido. Surgical management of right ventricular dysfunction late after repair of tetralogy of fallot: right ventricular remodeling surgery. In *Seminars in Thoracic and Cardiovascular Surgery: Pediatric Cardiac Surgery Annual*, volume 9, pages 29–34. Elsevier, 2006. 30, 31
- B. Delhay, P. Clarysse, and I. Magnin. Locally Adapted Spatio-temporal Deformation Model for Dense Motion Estimation in Periodic Cardiac Image Sequences. In *FIMH*, volume 4466, page 393. Springer, 2007. 57, 90, 103
- H. Delingette. Biquadratic and quadratic springs for modeling st venant kirchhoff materials. In *Fourth International Symposium on BioMedical Simulation (ISBMS’08)*, volume 5104 of *Lecture Notes in Computer Science*, pages 40–48, London, UK, July 2008a. Springer. 188

- H. Delingette. General object reconstruction based on simplex meshes. *International Journal of Computer Vision*, 32(2):111–146, 1999. [36](#), [45](#), [48](#), [49](#)
- H. Delingette. Triangular springs for modeling nonlinear membranes. *IEEE Transactions on Visualisation and Computer Graphics*, 14(2), March/April 2008b. [188](#)
- H. Delingette and N. Ayache. Soft tissue modeling for surgery simulation. In N. Ayache, editor, *Computational Models for the Human Body*, Handbook of Numerical Analysis, pages 453–550. Elsevier, ph. ciarlet edition, 2004. [171](#)
- H. Delingette, E. Bardinet, D. Rey, J.-D. Lemarechal, J. Montagnat, S. Ourselin, A. Roche, D. Dormont, J. Yelnik, and N. Ayache. Yav++: a software platform for medical image processing and visualization. In *Workshop on Interactive Medical Image Visualization and Analysis satellite symposia of MICCAI, IMIVA'01*, Utrecht, The Netherlands, October 2001. [56](#)
- R. Deriche. Recursively implementing the gaussian and its derivatives. *Rapports de recherche- INRIA*, 1993. [71](#), [82](#)
- H. Dinh, G. Turk, and G. Slabaugh. Reconstructing surfaces using anisotropic basis functions. In *Eighth IEEE International Conference on Computer Vision, 2001. ICCV 2001. Proceedings*, volume 2, 2001. [41](#)
- M. DoCarmo. Riemannian geometry (transl. by francis flaherty), math. theory appl, 1992. [83](#), [148](#), [151](#), [222](#), [223](#)
- F. Dru and T. Vercauteren. An ITK implementation of the symmetric log-domain diffeomorphic demons algorithm. *Insight Journal – 2009 January - June*, May 2009. [72](#), [219](#)
- D. DuBois and E. DuBois. The measurement of the surface area of man. *Archives of Internal Medicine*, 15(5_2):868, 1915. [126](#), [135](#)
- P. Dupuis, U. Grenander, and M. Miller. Variational problems on flows of diffeomorphisms for image matching. *Quarterly of Applied Mathematics*, 56(3):587–600, 1998. [114](#)
- S. Durrleman. *Statistical models of currents for measuring the variability of anatomical curves, surfaces and their evolution*. Thèse de sciences (phd thesis), Université de Nice-Sophia Antipolis, March 2010. [112](#)
- S. Durrleman, X. Pennec, A. Trouvé, and N. Ayache. A forward model to build unbiased atlases from curves and surfaces. In X. Pennec and S. Joshi, editors, *Proc. of the International Workshop on the Mathematical Foundations of Computational Anatomy (MFCA-2008)*, September 2008. [109](#), [111](#)
- S. Durrleman, X. Pennec, A. Trouvé, and N. Ayache. Statistical models on sets of curves and surfaces based on currents. *Medical Image Analysis*, 13(5):793–808, Oct. 2009a. [8](#), [111](#), [112](#), [113](#), [115](#), [116](#), [215](#)

- S. Durrleman, X. Pennec, A. Trouvé, G. Gerig, and N. Ayache. Spatiotemporal atlas estimation for developmental delay detection in longitudinal datasets. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, Lecture Notes in Computer Science, pages 297–304. Springer, 2009b. [125](#), [151](#), [223](#)
- O. Ecabert, J. Peters, H. Schramm, C. Lorenz, J. von Berg, M. J. Walker, M. Vembar, M. E. Olszewski, K. Subramanyan, G. Lavi, and J. Weese. Automatic model-based segmentation of the heart in ct images. *IEEE Transactions on Medical Imaging*, 27(9):1189–1201, 2008. [34](#), [35](#), [161](#)
- A. Einstein. On the method of theoretical physics. *Philosophy of science*, 1(2):163–169, 1934. [220](#)
- A. Elen, J. Hermans, J. Ganame, D. Loeckx, J. Bogaert, F. Maes, and P. Suetens. Automatic 3-d breath-hold related motion correction of dynamic multislice mri. *IEEE Transactions on Medical Imaging*, 29(3):868, 2010. [37](#)
- M. Endo. Calcium release from the sarcoplasmic reticulum. *Physiological Reviews*, 57(1):71, 1977. [15](#)
- L. C. Evans. *Partial Differential Equations*. American Mathematical Society, June 1998. ISBN 0821807722. [64](#), [71](#), [84](#)
- B. Eyskens, S. Brown, P. Claus, S. Dymarkowski, M. Gewillig, J. Bogaert, and L. Mertens. The influence of pulmonary regurgitation on regional right ventricular function in children after surgical repair of tetralogy of fallot. *European Journal of Echocardiography*, 2010. [221](#)
- L. Feng, R. Donnino, J. Babb, L. Axel, and D. Kim. Numerical and in vivo validation of fast cine displacement-encoded with stimulated echoes (dense) mri for quantification of regional cardiac function. *Magnetic Resonance in Medicine*, 62(3):682–690, 2009. ISSN 1522-2594. [24](#)
- W. Feng, T. Denney, S. Lloyd, L. Dell’Italia, and H. Gupta. Contour regularized left ventricular strain analysis from cine MRI. In *Biomedical Imaging: From Nano to Macro, 2008. ISBI 2008. 5th IEEE International Symposium on*, pages 520–523, 2008. [90](#), [92](#)
- M. Ferrant, A. Nabavi, B. Macq, F. Jolesz, R. Kikinis, and S. Warfield. Registration of 3D intraoperative MR images of the brain using a finite element biomechanical model. *IEEE Transactions on Medical Imaging*, 20(12):1384–1397, 2001. [64](#)
- R. Fitzhugh. Impulses and physiological states in theoretical models of nerve membrane. *Biophysical Journal*, 1(6):445–466, 1961. [156](#), [167](#)
- L. Formaggia, J. Gerbeau, F. Nobile, and A. Quarteroni. On the coupling of 3d and 1d navier-stokes equations for flow problems in compliant vessels* 1. *Computer Methods in Applied Mechanics and Engineering*, 191(6-7):561–582, 2001. [185](#)

- A. Frangi, W. Niessen, and M. Viergever. Three-dimensional modeling for functional analysis of cardiac images: A review. *IEEE Transactions on Medical Imaging*, 20(1):2–5, 2001. 34
- O. Frank. Die grundform des arteriellen pulses. *Z Biol*, 37:483–526, 1899. 19, 182
- G. Freeman and M. LeWinter. Pericardial adaptations during chronic cardiac dilation in dogs. *Circulation research*, 54(3):294, 1984. 175
- J. Freund, D. Comaniciu, Y. Ioannis, P. Liu, R. McClatchey, E. Morley-Fletcher, X. Pennec, G. Pongiglione, and X. Zhou. Health-e-child: an integrated biomedical platform for grid-based paediatric applications. *Studies in health technology and informatics*, 120:259, 2006. 4
- A. Frigiola, A. Redington, S. Cullen, and M. Vogel. Pulmonary regurgitation is an important determinant of right ventricular contractile dysfunction in patients with surgically repaired tetralogy of fallot. *Circulation*, 110(11_suppl_1), 2004. 134
- A. Frigiola, V. Tsang, J. Nordmeyer, P. Lurz, C. van Doorn, A. Taylor, P. Bonhoeffer, and M. de Leval. Current approaches to pulmonary regurgitation. *European Journal of Cardio-Thoracic Surgery*, 34(3):576, 2008. 134, 152
- Y. Fujikoshi and L. Veitch. Estimation of dimensionality in canonical correlation analysis. *Biometrika*, 66(2):345–351, 1979. 130
- M. A. Gatzoulis, J. A. Till, J. Somerville, and A. N. Redington. Mechano-electrical interaction in tetralogy of Fallot : QRS prolongation relates to right ventricular size and predicts malignant ventricular arrhythmias and sudden death. *Circulation*, 92(2):231–237, 1995. 30
- L. Geerts, P. Bovendeerd, K. Nicolay, and T. Arts. Characterization of the normal cardiac myofiber field in goat measured with mr-diffusion tensor imaging. *American Journal of Physiology- Heart and Circulatory Physiology*, 283(1):H139, 2002. 25, 165
- L. Geerts-Ossevoort. *Cardiac Myofiber Reorientation: A Mechanism for Adaptation*. PhD thesis, Technische Universiteit Eindhoven, 2002. 108, 188
- J. Gerbeau, M. Vidrascu, and P. Frey. Fluid-structure interaction in blood flows on geometries based on medical imaging. *Computers & Structures*, 83(2-3):155–165, 2005. 185, 189, 210, 221
- T. Geva. Indications and timing of pulmonary valve replacement after tetralogy of Fallot repair. In *Seminars in Thoracic and Cardiovascular Surgery: Pediatric Cardiac Surgery Annual*, volume 9, pages 11–22. Elsevier, 2006. 30, 134, 144, 218
- L. Glass, P. Hunter, and A. McCulloch. *Theory of Heart: Biomechanics, Biophysics, and Nonlinear Dynamics of Cardiac Function*. Springer-Verlag, 1991. 61, 89, 90, 156, 165, 190

- J. Glaunes. *Modeles déformables en appariement de formes. Transport par difféomorphismes de points, de mesures et de courants pour la comparaison de formes et l'anatomie numérique*. PhD thesis, Phd Thesis (in french), Université Paris 13, 2005. [114](#), [115](#)
- H. Glockner. Fundamental problems in the theory of infinite-dimensional lie groups. *Arxiv preprint math/0602078*, 2006. [66](#)
- J. Gorce, D. Friboulet, and I. Magnin. Estimation of three-dimensional cardiac velocity fields: assessment of a differential method and application to three-dimensional CT data. *Medical Image Analysis*, 1(3):245–261, 1997. [63](#)
- U. Grenander and M. Miller. Computational anatomy: An emerging discipline. *Quarterly of applied mathematics*, 56(4):617–694, 1998. [108](#), [109](#)
- L. Grosse-Wortmann and A. Redington. Doing the right thing at the right time: is there more to pulmonary valve replacement than meets the eye? *European heart journal*, 30(17):2076, 2009. [141](#)
- J. Guccione and A. McCulloch. Theory of heart: biomechanics, biophysics, and nonlinear dynamics of cardiac function. *Finite element modeling of ventricular mechanics*. New York: Springer-Verlag, pages 121–44, 1991. [159](#)
- A. Guimond, J. Meunier, and J.-P. Thirion. Average brain models: A convergence study. *Computer Vision and Image Understanding*, 77(2):192–210, 2000. [109](#), [110](#), [111](#)
- E. Haber and J. Modersitzki. Numerical methods for volume preserving image registration. *Inverse Problems*, 20(5):1621–1638, 2004. [63](#), [82](#), [102](#)
- J. Häfner, F. Sachse, C. Sansour, G. Seemann, and O. Dössel. Hyperelastic description of elastomechanic properties of the heart: a new material law and its application. *Biomedizinische Technik. Biomedical engineering*, 47:770, 2002. [171](#)
- J. Hamm, D. Ye, R. Verma, and C. Davatzikos. Gram: A framework for geodesic registration on anatomical manifolds. *Medical Image Analysis*, 2010. [151](#)
- D. Heitz, E. Mémin, and C. Schnörr. Variational fluid flow measurements from image sequences: synopsis and perspectives. *Experiments in Fluids*, pages 1–25, 2009. [63](#)
- W. Helbing and A. De Roos. Clinical applications of cardiac magnetic resonance imaging after repair of tetralogy of fallot. *Pediatric Cardiology*, 21(1):70–79, 2000. [23](#)
- P. Helm, H. Tseng, L. Younes, E. McVeigh, and R. Winslow. Ex vivo 3d diffusion tensor imaging and quantification of cardiac laminar structure. *Magnetic resonance in medicine*, 54(4):850, 2005a. [25](#)

- R. Helm, C. Leclercq, O. Faris, C. Ozturk, E. McVeigh, A. Lardo, and D. Kass. Cardiac dyssynchrony analysis using circumferential versus longitudinal strain: implications for assessing cardiac resynchronization. *Circulation*, 111(21):2760, 2005b. [221](#)
- M. Hernandez, M. Bossa, and S. Olmos. Registration of anatomical images using paths of diffeomorphisms parameterized with stationary vector field flows. *International Journal of Computer Vision*, 85(3):291–306, 2009. [62](#), [83](#)
- J. Hinkle, P. Fletcher, B. Wang, B. Salter, and S. Joshi. 4D MAP Image Reconstruction Incorporating Organ Motion. In *Proceedings of the 21st International Conference on Information Processing in Medical Imaging*, page 687. Springer, 2009. [64](#), [73](#), [83](#), [102](#)
- A. Hodgkin and A. Huxley. A quantitative description of ion currents and its applications to conduction and excitation in nerve membranes. *Journal of Physiology*, 117(4):500–544, 1952. [167](#)
- J. Hoffman and S. Kaplan. The incidence of congenital heart disease. *Journal of the American College of Cardiology*, 39(12):1890–1900, 2002. [28](#)
- B. Horn and B. Schunck. Determining optical flow. *Artificial intelligence*, 17(1-3):185–203, 1981. [63](#)
- A. Höskuldsson. Partial least square regression methods. *Journal of Chemometrics*, 2:211–228, 1988. [127](#), [128](#)
- H. Hotelling. Relations between two sets of variates. *Biometrika*, 28(3-4):321, 1936. [129](#)
- H. Hufnagel, X. Pennec, J. Ehrhardt, H. Handels, and N. Ayache. Point-based statistical shape models with probabilistic correspondences and affine em-icp. In Springer, editor, *Bildverarbeitung fuer die Medizin 2007*, Informatik aktuell, pages 434–438. Springer, March 2007. [112](#)
- J. Humphrey, R. Strumpf, and F. Yin. Determination of a constitutive relation for passive myocardium: I. a new functional form. *Journal of biomechanical engineering*, 112:333, 1990a. [171](#)
- J. Humphrey, R. Strumpf, and F. Yin. Determination of a constitutive relation for passive myocardium: II.—parameter estimation. *Journal of Biomechanical Engineering*, 112:340, 1990b. [171](#)
- P. Hunter and B. Smaill. The analysis of cardiac function: a continuum approach. *Prog Biophys Mol Biol*, 52(2):101–164, 1988. [170](#), [171](#), [220](#)
- P. Hunter, M. Nash, and G. Sands. Computational electromechanics of the heart. *Computational biology of the heart*, pages 347–409, 1996. [170](#)

- P. Hunter, A. McCulloch, and H. Ter Keurs. Modelling the mechanical properties of cardiac muscle. *Progress in biophysics and molecular biology*, 69(2-3):289–331, 1998. 156, 170
- P. Hunter, T. Borg, and P. Box. Integration from proteins to organs: the physiome project. *Nature Reviews Molecular Cell Biology*, 4:237–243, 2003a. 156, 213, 220
- P. J. Hunter, A. J. Pullan, and B. H. Smaill. Modeling total heart function. *Annu Rev Biomed Eng*, 5:147–177, 2003b. 156, 170
- R. Ionasec, Y. Wang, B. Georgescu, I. Voigt, N. Navab, and D. Comaniciu. Robust motion estimation using trajectory spectrum learning: Application to aortic and mitral valve modeling from 4d tee. In *IEEE Int’l Conf. Computer Vision and Pattern Recognition (ICCV)*, 2009. 57
- R. I. Ionasec, I. Voigt, B. Georgescu, Y. Wang, H. Houle, F. Vega-Higuera, N. Navab, and D. Comaniciu. Patient-specific modeling and quantification of the aortic and mitral valves from 4d cardiac ct and tee. *IEEE Transactions on Medical Imaging*, pages 1–1, 2010. 57
- J. Jakeman and S. Roberts. Stochastic galerkin and collocation methods for quantifying uncertainty in differential equations: a review. *ANZIAM*, 50:C815–C830, 2009. 224
- Y. Ji, P. Hervé, U. Aickelin, and A. Pitiot. Parcellation of fmri datasets with ica and pls-a data driven approach. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, LNCS, pages 984–991. Springer, 2009. 126
- B. Jian and B. Vemuri. A robust algorithm for point set registration using mixture of Gaussians. In *Proc. ICCV 2005*, volume 2, pages 1246–1251, 2005. 112
- M. Jolly. Automatic segmentation of the left ventricle in cardiac mr and ct images. *International Journal of Computer Vision*, 70(2):151–163, 2006. 34
- S. Joshi, B. Davis, M. Jomier, and G. Gerig. Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImage*, 23:151–160, 2004. 109, 110
- M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International journal of computer vision*, 1(4):321–331, 1988. 47
- M. R. Kaus, J. von Berg, J. Weese, W. Niessen, and V. Pekar. Automated segmentation of the left ventricle in cardiac MRI. *Medical Image Analysis*, 8(3):245–254, 2004. 35
- J. Keener and J. Sneyd. *Mathematical physiology*. Springer-Verlag, 1998. 14, 15, 167, 179, 189, 190, 210
- D. Kendall. A survey of the statistical theory of shape. *Statistical Science*, 4(2): 87–99, 1989. 108

- R. C. P. Kerckhoffs, L. N. Maxwell, G. Quan, B. J. Bassingthwaite, J. H. Omens, and A. D. McCulloch. Coupling of a 3D finite element model of cardiac ventricular mechanics to lumped systems models of the systemic and pulmonic circulation. *Annals of Biomedical Engineering*, 35(1):1–18, January 2007. 179, 189
- A. Khalil, B. Bouma, and M. Kaazempur Mofrad. A combined fem/genetic algorithm for vascular soft tissue elasticity estimation. *Cardiovascular Engineering*, 6(3):93–102, 2006. 221
- S. Khambadkone, L. Coats, A. Taylor, Y. Boudjemline, G. Derrick, V. Tsang, J. Cooper, V. Muthurangu, S. R. Hegde, R. S. Razavi, D. Pellerin, J. Deanfield, and P. Bonhoeffer. Percutaneous pulmonary valve implantation in humans: Results in 59 consecutive patients. *Circulation*, 112(8):1189–1197, 2005. 30
- D. Kim, W. Gilson, C. Kramer, and F. Epstein. Myocardial tissue tracking with two-dimensional cine displacement-encoded mr imaging: Development and initial evaluation1. *Radiology*, 230(3):862, 2004. 24
- B. Kimia, A. Tannenbaum, and S. Zucker. On the evolution of curves via a function of curvature. i. the classical case. *Journal of Mathematical Analysis and Applications*, 163(2):438–458, 1992. 48
- T. Korakianitis and Y. Shi. Numerical simulation of cardiovascular dynamics with healthy and diseased heart valves. *Journal of Biomechanics*, 39(11):1964 – 1982, 2006. ISSN 0021-9290. 185
- R. Krishnamurthy. Pediatric cardiac mri: anatomy and function. *Pediatric radiology*, 38:192–199, 2008. 23
- W. Kroon, T. Delhaas, T. Arts, and P. Bovendeerd. Constitutive modeling of cardiac tissue growth. In *Functional Imaging and Modeling of the Heart*, volume 4466 of *Lecture Notes in Computer Science*, page 340. Springer, 2007. 108, 222, 224
- W. Kroon, T. Delhaas, P. Bovendeerd, and T. Arts. Computational analysis of the myocardial structure: Adaptation of cardiac myofiber orientations through deformation. *Medical Image Analysis*, 13(2):346–353, 2009. 108, 159, 188, 222
- H. Lamecker, T. Mansi, J. Relan, F. Billet, M. Sermesant, N. Ayache, and H. Delingette. Adaptive tetrahedral meshing for personalized cardiac simulations. In *MICCAI Workshop on Cardiovascular Interventional Imaging and Biophysical Modelling (CI2BM09)*, pages 149–158, London United Kingdom, 2009. 163, 227
- R. Lang, M. Bierig, R. Devereux, F. Flachskampf, E. Foster, P. Pellikka, M. Picard, M. Roman, J. Seward, J. Shanewise, et al. Recommendations for chamber quantification: a report from the american society of echocardiography’s guidelines and standards committee and the chamber quantification writing group, developed in conjunction with the european association of echocardiography, a branch of the european society of cardiology. *Journal of the American Society of Echocardiography*, 18(12):1440–1463, 2005. 21

- J. Lankhaar, N. Westerhof, T. Faes, K. Marques, J. Marcus, P. Postmus, and A. Vonk-Noordegraaf. Quantification of right ventricular afterload in patients with and without pulmonary hypertension. *American Journal of Physiology- Heart and Circulatory Physiology*, 291(4):H1731, 2006. [184](#)
- D. Le Bihan, E. Breton, D. Lallemand, P. Grenier, E. Cabanis, and M. Laval-Jeantet. Mr imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders. *Radiology*, 161(2):401, 1986. [25](#)
- M. Ledesma-Carbayo, J. Kybic, M. Sühling, P. Hunziker, M. Desco, A. Santos, and M. Unser. Cardiac ultrasound motion detection by elastic registration exploiting temporal coherence. *Proc. IEEE on Biomedical Imaging*, 2:585–588, 2002. [22](#), [103](#)
- I. LeGrice, B. Smaill, L. Chai, S. Edgar, J. Gavin, and P. Hunter. Laminar structure of the heart: ventricular myocyte arrangement and connective tissue architecture in the dog. *American Journal of Physiology- Heart and Circulatory Physiology*, 269(2):H571, 1995. [160](#), [165](#)
- V. Lempitsky, M. Verhoek, J. Noble, and A. Blake. Random forest classification for automatic delineation of myocardium in real-time 3d echocardiography. In *Functional Imaging and Modeling of the Heart*, Lecture Notes in Computer Science, pages 447–456. Springer, 2009. [34](#)
- Z. Liang and P. Lauterbur. *Principles of magnetic resonance imaging: a signal processing perspective*. SPIE Optical Engineering Press, 1999. [20](#), [23](#), [24](#)
- W. Lorensen and H. E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Proc. SIGGRAPH '87*, volume 21, pages 163–169, 1987. [42](#)
- M. Lorenzo-Valdes, G. I. Sanchez-Ortiz, A. G. Elkington, R. H. Mohiaddin, and D. Rueckert. Segmentation of 4D cardiac MR images using a probabilistic atlas and the EM algorithm. *Medical Image Analysis*, 8(3):255–265, 2004. [35](#)
- X. Lu, B. Georgescu, A. Littmann, E. Mueller, and D. Comaniciu. Discriminative Joint Context for Automatic Landmark Set Detection from a Single Cardiac MR Long Axis Slice. In *Functional Imaging and Modeling of the Heart (FIMH)*, Lecture Notes in Computer Science, pages 457–465. Springer, 2009. [35](#)
- C. Luo and Y. Rudy. A model of the ventricular cardiac action potential. depolarization, repolarization, and their interaction. *Circulation Research*, 68(6):1501, 1991. [167](#)
- M. Lynch, O. Ghita, and P. Whelan. Segmentation of the left ventricle of the heart in 3-d+ t mri data using an optimized nonrigid temporal model. *IEEE Transactions on Medical Imaging*, 27(2):195–203, 2008. [35](#)

- R. Malladi, J. Sethian, and B. Vemuri. Shape modeling with front propagation: A level set approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(2):158–175, 1995. [47](#)
- T. Mansi, M. Sermesant, M. Huber, A. Taylor, G. Pongiglione, X. Pennec, and Y. Boudjemline. Modélisation électromécanique du coeur et analyse d’images. Archives de Pédiatrie - Congrès des Sociétés Françaises Médico-chirurgicales Pédiatriques, February 2008. Archives de Pédiatrie 15(5):1032 – Special award for best presentation. [218](#), [228](#)
- T. Mansi, B. André, M. Lynch, M. Sermesant, H. Delingette, Y. Boudjemline, and N. Ayache. Virtual pulmonary valve replacement interventions with a personalised cardiac electromechanical model. In *Recent Advances in the 3D Physiological Human*, volume 5528 of *Lecture Notes in Computer Science*, pages 201–210. Springer, 2009a. [181](#), [214](#), [217](#), [225](#)
- T. Mansi, B. André, M. Sermesant, H. Delingette, N. Ayache, and Y. Boudjemline. Simulation personnalisée de remplacements valvulaires pulmonaires grâce à l’utilisation d’un modèle mathématique du coeur. Archives de Pédiatrie - Congrès des Sociétés Françaises Médico-chirurgicales Pédiatriques, June 2009b. [218](#), [228](#)
- T. Mansi, S. Durrleman, B. Bernhardt, M. Sermesant, H. Delingette, I. Voigt, P. Lurz, A. M. Taylor, J. Blanc, Y. Boudjemline, X. Pennec, and N. Ayache. A Statistical Model of Right Ventricle in Tetralogy of Fallot for Prediction of Remodelling and Therapy Planning. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, volume 5761 of *Lecture Notes in Computer Science*, pages 214–221, London, UK, Sept. 2009c. Springer. [124](#), [126](#), [144](#), [152](#), [216](#), [226](#)
- T. Mansi, J.-M. Peyrat, M. Sermesant, H. Delingette, J. Blanc, Y. Boudjemline, and N. Ayache. Physically-Constrained Diffeomorphic Demons for the Estimation of 3D Myocardium Strain from Cine-MRI. In *Proceedings of Functional Imaging and Modeling of the Heart 2009 (FIMH’09)*, volume 5528 of *Lecture Notes in Computer Science*, pages 201–210, 3-5 June 2009d. [55](#), [63](#), [214](#), [215](#), [226](#)
- T. Mansi, X. Pennec, M. Sermesant, H. Delingette, and N. Ayache. ilogdemons: A demons-based registration algorithm for tracking incompressible elastic biological tissues. *International Journal of Computer Vision*, 2010a. Submitted. [215](#), [226](#)
- T. Mansi, X. Pennec, M. Sermesant, H. Delingette, and N. Ayache. LogDemons Revisited: Consistent Regularisation and Incompressibility Constraint for Soft Tissue Tracking in Medical Images. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, Lecture Notes in Computer Science. Springer, 2010b. In press. [214](#), [215](#), [226](#)

- T. Mansi, M. Sermesant, H. Delingette, X. Pennec, N. Ayache, and B. Y. In-silico models for the simulation and prediction of the cardiac function. *Archives de Pédiatrie*, 17(6):611–612, Juin 2010c. [218](#), [227](#)
- C. Mao, S. Qin, and D. Wright. A sketch-based approach to human body modelling. *Computers & Graphics*, 33(4):521–541, 2009. [56](#)
- E. Marbán. Cardiac channelopathies. *Nature*, 415(6868):213–218, 2002. [16](#)
- W. Massy. Principal components regression in exploratory statistical research. *Journal of the American Statistical Association*, 60(309):234–256, 1965. [123](#)
- A. McCulloch, J. Bassingthwaite, P. Hunter, and D. Noble. Computational biology of the heart: from structure to function. *Prog Biophys Mol Biol*, 69(2-3), 1998. [156](#)
- T. McInerney and D. Terzopoulos. A dynamic finite element surface model for segmentation and tracking in multidimensional medical images with application to cardiac 4D image analysis. *Computerized Medical Imaging and Graphics*, 19(1):69–83, 1995. [34](#)
- A. McIntosh and N. Lobaugh. Partial least squares analysis of neuroimaging data: applications and advances. *Neuroimage*, 23:S250–S263, 2004. [126](#), [216](#)
- D. McQueen and C. Peskin. A three-dimensional computer model of the human heart for studying cardiac fluid dynamics. *ACM SIGGRAPH Computer Graphics*, 34(1):56–60, 2000. [178](#), [185](#), [189](#), [221](#)
- E. McVeigh. Regional myocardial function. *Cardiology Clinics*, 16(2):189–206, 1998. [24](#), [89](#)
- J. Mercier, T. DiSessa, J. Jarmakani, T. Nakanishi, S. Hiraishi, J. Isabel-Jones, and W. Friedman. Two-dimensional echocardiographic assessment of left ventricular volumes and ejection fraction in children. *Circulation*, 65(5):962, 1982. [159](#)
- V. Mihalef, D. Metaxas, M. Sussman, V. Hurmusiadis, and L. Axel. Atrioventricular blood flow simulation based on patient-specific data. In *Functional Imaging and Modeling of the Heart*, Lecture Notes in Computer Science, pages 386–395. Springer, 2009. [156](#), [178](#), [179](#), [210](#)
- M. Miller, A. Trouvé, and L. Younes. On the metrics and Euler-Lagrange equations of computational anatomy. *Annual review of biomedical engineering*, 4(1):375–405, 2002a. [114](#)
- M. I. Miller, A. Trouvé, and L. Younes. On the metrics and Euler-Lagrange equations of computational anatomy. *Annual Review of Biomed. Eng.*, 4:375–405, 2002b. [62](#)
- C. Mitchell and D. Schaeffer. A two-current model for the dynamics of cardiac membrane. *Bulletin of mathematical biology*, 65(5):767–793, 2003. [167](#), [188](#)

- S. Mitchell, J. Bosch, B. Lelieveldt, R. Van der Geest, J. Reiber, and M. Sonka. 3-D active appearance models: segmentation of cardiac MR and ultrasound images. *IEEE Transactions on Medical Imaging*, 21(9):1167–1178, 2002. 34
- J. Modersitzki. *Numerical Methods for Image Registration*. Oxford University Press, 2004. 63, 68
- P. Moireau, D. Chapelle, and P. Tallec. Joint state and parameter estimation for distributed mechanical systems. *Computer Methods in Applied Mechanics and Engineering*, 197(6-8):659–677, 2008. 176, 189, 222
- P. Moireau, D. Chapelle, and P. Tallec. Filtering for distributed mechanical systems using position measurements. *Inverse Problems*, 25:035010, 2009. 189, 221, 222
- J. Montagnat and H. Delingette. 4D deformable models with temporal constraints: application to 4D cardiac image segmentation. *Medical Image Analysis*, 9(1):87–100, 2005. 35, 57
- J. Montagnat and H. Delingette. Globally constrained deformable models for 3D object reconstruction. *Signal Processing*, 71(2):173–186, 1998. 35
- J. Montagnat, H. Delingette, and N. Ayache. A review of deformable surfaces: topology, geometry and deformation. *Image and Vision Computing*, 19(14):1023–1040, December 2001. 47
- C. Moore, C. Lugo-Olivieri, E. McVeigh, and E. Zerhouni. Three-dimensional systolic strain patterns in the normal human left ventricle: Characterization with tagged MR imaging. *Radiology*, 214(2):453–466, 2000. 24, 61, 93, 96, 221
- V. Moreau-Villéger, H. Delingette, M. Sermesant, H. Ashikaga, E. McVeigh, and N. Ayache. Building maps of local apparent conductivity of the epicardium with a 2d electrophysiological model of the heart. *IEEE Transactions on Biomedical Engineering*, 53(8):1457–1466, August 2006. 168
- B. Morse, T. Yoo, D. Chen, P. Rheingans, and K. Subramanian. Interpolating implicit surfaces from scattered surface data using compactly supported radial basis functions. In *Proceedings of the International Conference on Shape Modeling & Applications*, page 89. IEEE Computer Society, 2001. 41, 43, 56
- G. Mosier. Determining a simple structure when loadings for certain tests are known. *Psychometrika*, pages 149–162, 1939. 112, 136
- F. Murgatroyd and A. Krahn. *Handbook of cardiac electrophysiology: A practical guide to invasive EP studies and catheter Ablation*. Remedica, 2002. 15, 169, 190
- J. Murphy, B. Gersh, D. Mair, V. Fuster, M. McGoon, D. Ilstrup, D. McGoon, J. Kirklin, and G. Danielson. Long-term outcome in patients undergoing surgical repair of tetralogy of fallot. *The New England journal of medicine*, 329(9):593, 1993. 30

- J. Nagumo, S. Arimoto, and S. Yoshizawa. An active pulse transmission line simulating nerve axon. *Proceedings of the IRE*, 50(10):2061–2070, 1962. 167
- J. Narula, G. Buckberg, and B. Khanderia, editors. *Function Follows Form*, volume 4. Elsevier, July 2008. 108
- M. Nash. *Mechanics and material properties of the heart using an anatomically accurate mathematical model*. PhD thesis, University of Auckland, 1998. 159
- M. Nash and P. Hunter. Computational mechanics of the heart. *Journal of elasticity*, 61(1):113–141, 2000. 156, 170
- M. Nesme, Y. Payan, and F. Faure. Efficient, physically plausible finite elements. *Eurographics (short papers)*, pages 77–80, 2005. 195
- S. Niederer and N. Smith. An improved numerical method for strong coupling of excitation and contraction models in the heart. *Progress in biophysics and molecular biology*, 96(1-3):90–111, 2008. 170
- S. Niederer, P. Hunter, and N. Smith. A quantitative analysis of cardiac myocyte relaxation: a simulation study. *Biophysical Journal*, 90(5):1697–1722, 2006. 170
- S. Niederer, K. Rhode, R. Razavi, and N. Smith. The importance of model parameters and boundary conditions in whole organ models of cardiac contraction. *Functional Imaging and Modeling of the Heart*, pages 348–356, 2009. 160, 174, 188, 189, 220
- M. Nielsen, L. Florack, and R. Deriche. Regularization and scale space. Tech. report, INRIA, 1994. 67, 68
- P. Nielsen, I. Le Grice, B. Smaill, and P. Hunter. Mathematical methods of classical mechanical model of geometry and fibrous structure of the heart. *American Journal of Physiology- Heart and Circulatory Physiology*, 260(4):H1365, 1991. 160
- D. Noble. A modification of the hodgkin—huxley equations applicable to purkinje fibre action and pacemaker potentials. *The Journal of Physiology*, 160(2):317, 1962. 156, 167
- D. Noble, A. Varghese, P. Kohl, and P. Noble. Improved guinea-pig ventricular cell model incorporating a diadic space, IKr and IKs, and length- and tension-dependent processes. *Can J Cardiol*, 14(1):123–134, 1998. 167
- J. Noble and D. Boukerroui. Ultrasound image segmentation: A survey. *IEEE Transactions on Medical Imaging*, 25(8):987–1010, 2006. 34
- M. Nollert et al. Long-term survival in patients with repair of tetralogy of fallot: 36-year follow-up of 490 survivors of the first year after surgical repair. *Journal of the American College of Cardiology*, 30(5):1374–1383, 1997. 30

- D. Nordsletten, S. Niederer, M. Nash, P. Hunter, and N. Smith. Coupling multi-physics models to cardiac mechanics. *Progress in Biophysics and Molecular Biology*, 2009. 189, 210, 221
- T. Oosterhof. *Cardiovascular magnetic resonance in patients with corrected Tetralogy of Fallot*. PhD thesis, University of Amsterdam, 2006. 30, 108
- T. Oosterhof, B. Mulder, H. Vliegen, and A. de Roos. Cardiovascular magnetic resonance in the follow-up of patients with corrected tetralogy of fallot: a review. *American Heart Journal*, 151(2):265–272, 2006. 30, 108
- N. Osman, S. Sampath, E. Atalar, and J. Prince. Imaging longitudinal cardiac strain on short-axis images using strain-encoded mri. *Magnetic Resonance in Medicine*, 46(2):324–334, 2001. 24
- S. Ourselin, A. Roche, G. Subsol, X. Pennec, and N. Ayache. Reconstructing a 3d structure from serial histological sections. *Image and Vision Computing*, 19(1-2): 25–31, January 2001. 37, 38
- X. Papademetris, A. Sinusas, D. Dione, R. Constable, and J. Duncan. Estimating 3D strain from 4D cine-MRI and echocardiography: In-vivo validation. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, Lecture Notes in Computer Science, pages 678–686. Springer, 2000. 64, 90, 92
- A. Pednekar, U. Kurkure, R. Muthupillai, S. Flamm, and I. Kakadiaris. Automated left ventricular segmentation in cardiac mri. *IEEE Transactions on Biomedical Engineering*, 53(7):1425, 2006. 35
- X. Pennec. Statistical computing on manifolds: from Riemannian geometry to computational anatomy. In *Emerging Trends in Visual Computing*, volume 5416 of *Lecture Notes in Computer Science*, pages 347–386. Springer, 2008. 108
- X. Pennec, P. Cachier, and N. Ayache. Understanding the demon’s algorithm: 3D non-rigid registration by gradient descent. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, volume 1679 of *Lecture Notes in Computer Science*, pages 597–605, Cambridge, UK, September 1999. Springer. 63, 68
- X. Pennec, C. Basso, Y. Boudjemline, S. Durrleman, E. Konukoglu, T. Mansi, G. Pongiglione, M. Santoro, M. Sermesant, B. Stos, N. Toussaint, and G. Trocchio. Third generation disease models: Image analysis tools, pediatric heart diseases, inflammatory diseases and brain tumors. Deliverables D11.2-D11.3-D11.4, European project Health-e-Child (IST-2004-027749), June 2007-2009-2010. 219, 228
- E. Pernod, M. Sermesant, J. Relan, and H. Delingette. Interactive real time simulation of cardiac radio-frequency ablation. In *Proc. of Eurographics Workshop on Visual Computing for Biology and Medicine (2010) (VCBM’2010)*, Leipzig, August 2010. 168

- D. Perperidis, R. Mohiaddin, and D. Rueckert. Spatio-temporal free-form registration of cardiac MR image sequences. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, Lecture Notes in Computer Science, pages 911–919. Springer, 2004. 57, 103
- D. Perperidis, R. Mohiaddin, and D. Rueckert. Construction of a 4D statistical atlas of the cardiac anatomy and its use in classification. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, Lecture Notes in Computer Science, pages 402–410. Springer, 2005. 45
- J. Peters, O. Ecabert, C. Meyer, R. Kneser, and J. Weese. Optimizing boundary detection via simulated search with applications to multi-modal heart segmentation. *Medical Image Analysis*, 14(1):70–84, 2010. 34
- J.-M. Peyrat. *Comparison of Cardiac Anatomy and Function: Statistics on Fibre Architecture from DT-MRI and Registration of 4D CT Images*. Phd thesis, Nice Sophia Antipolis University, November 2009. 166
- J.-M. Peyrat, M. Sermesant, X. Pennec, H. Delingette, C. Xu, E. R. McVeigh, and N. Ayache. A computational framework for the statistical analysis of cardiac diffusion tensors: Application to a small database of canine hearts. *IEEE TMI*, 26(11):1500–1514, 2007. 165, 166, 188, 221
- J.-M. Peyrat, H. Delingette, M. Sermesant, C. Xu, and N. Ayache. Registration of 4d cardiac ct sequences under trajectory constraints with multichannel diffeomorphic demons. *IEEE Transactions on Medical Imaging*, In Press 2009. 55, 56, 222
- N. Phatak, S. Maas, A. Veress, N. Pack, E. Di Bella, and J. Weiss. Strain measurement in the left ventricle during systole with deformable image registration. *Medical Image Analysis*, 13:354–361, 2009. 64
- G. Picinbono, H. Delingette, and N. Ayache. Non-linear anisotropic elasticity for real-time surgery simulation. *Graphical Models*, 65(5):305–321, September 2003. 171, 220
- G. Pongiglione and G. Trocchio. Report on diagnostic coding system and data collection protocols. Public Deliverable D9.1, Health-e-Child, 2006. 135
- M. Pop, M. Sermesant, T. Mansi, E. Crystal, J. Detsky, Y. Yang, P. Fefer, E. McVeigh, A. Dick, N. Ayache, and G. A. Wright. Characterization of post-infarct scars in a porcine model—a combined experimental and theoretical study. In *Proceedings of the 5th International Conference on Functional Imaging and Modeling of the Heart*, page 10. Springer, 2009. 165, 227
- M. Powell. Developments of newuoa for minimization without derivatives. *IMA Journal of Numerical Analysis*, 2008. 189

- S. Prasad, P. Kotwinski, and R. Assomul. The role of cardiovascular magnetic resonance in the evaluation of patients with heart failure. *Expert Review of Cardiovascular Therapy*, 2(1):53–59, 2004. 163
- A. Qiu, L. Younes, M. Miller, and J. Csernansky. Parallel transport in diffeomorphisms distinguishes the time-dependent pattern of hippocampal surface deformation due to healthy aging and the dementia of the alzheimer’s type. *NeuroImage*, 40(1):68–76, 2008. 125, 151, 222
- L. Querol, P. Büchler, D. Rueckert, L. Nolte, and M. Ballester. Statistical finite element model for bone shape and biomechanical properties. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, Lecture Notes in Computer Science, pages 405–411. Springer, 2006. 119
- A. Rao, G. Sanchez-Ortiz, R. Chandrashekar, M. Lorenzo-Valdes, R. Mohiaddin, and D. Rueckert. Comparison of cardiac motion across subjects using non-rigid registration. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, Lecture Notes in Computer Science, pages 722–729. Springer, 2002. 151, 222, 223
- A. Rao, P. Aljabar, and D. Rueckert. Hierarchical statistical shape analysis and prediction of sub-cortical brain structures. *Medical Image Analysis*, 12(1):55–68, 2008. 126, 216
- A. N. Redington, M. L. Rigby, E. A. Shinebourne, and P. J. Oldershaw. Changes in the pressure-volume relation of the right ventricle when its loading conditions are modified. *Br. Heart J.*, 63(1):45–49, 1990. 201
- J. Relan, P. Chinchapatnam, H. Delingette, K. Rhode, R. Razavi, N. Ayache, , and M. Sermesant. Coupled personalisation of electrophysiology models for simulation of induced ischemic ventricular tachycardia. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, LNCS. Springer, 2010. To appear. 168, 221
- W. Remme and K. Swedberg. Guidelines for the diagnosis and treatment of chronic heart failure. *European Heart Journal*, 22(17):1527, 2001. 108
- J. Rijcken, P. Bovendeerd, A. Schoofs, D. Van Campen, and T. Arts. Optimization of cardiac fiber orientation for homogeneous fiber strain during ejection. *Annals of biomedical engineering*, 27(3):289–297, 1999. 166, 188, 222
- E. Rodriguez, A. Hoger, and A. McCulloch. Stress-dependent finite growth in soft elastic tissues. *Journal of biomechanics*, 27(4):455–467, 1994. 108, 222, 224
- T. Rohlfing, C. Maurer Jr, D. Bluemke, and M. Jacobs. Volume-preserving nonrigid registration of MR breast images using free-form deformation with an incompressibility constraint. *IEEE Transactions on Medical Imaging*, 22(6):730–741, 2003. 63, 82

- R. Rosipal and N. Krämer. Overview and recent advances in partial least squares. *Subspace, Latent Structure and Feature Selection*, pages 34–51, 2006. [123](#), [126](#)
- R. Rosipal and L. Trejo. Kernel partial least squares regression in reproducing kernel hilbert space. *The Journal of Machine Learning Research*, 2:123, 2002. [151](#)
- D. Rueckert, L. Sonoda, C. Hayes, D. Hill, M. Leach, and D. Hawkes. Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Transactions on Medical Imaging*, 18(8):712–721, 1999. [63](#), [90](#)
- D. 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)*, Lecture Notes in Computer Science, pages 77–84. Springer, 2001. [109](#)
- D. Rueckert, P. Aljabar, R. Heckemann, J. Hajnal, and A. Hammers. Diffeomorphic registration using b-splines. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, volume 9 of *Lecture Notes in Computer Science*, page 702. Springer, 2006a. [63](#)
- D. Rueckert, R. Chandrashekara, P. Aljabar, K. Bhatia, J. Boardman, L. Srinivasan, M. Rutherford, L. Dyet, A. Edwards, J. Hajnal, et al. Quantification of growth and motion using non-rigid registration. *Computer Vision Approaches to Medical Image Analysis*, pages 49–60, 2006b. [62](#), [103](#)
- Y. Saad. *Iterative methods for sparse linear systems*. Society for Industrial Mathematics, 2003. [73](#)
- K. A. Saddi, C. Chéfd’hotel, and F. Cheriet. Large deformation registration of contrast-enhanced images with volume-preserving constraint. In *Proceedings of SPIE Medical Imaging*, volume 6512. SPIE, 2007. [63](#), [77](#), [102](#)
- J. Sainte-Marie, D. Chapelle, R. Cimrman, and M. Sorine. Modeling and estimation of the cardiac electromechanical activity. *Computers & Structures*, 84(28):1743–1759, 2006. [179](#), [189](#), [210](#), [220](#)
- M. Samyn, A. Powell, R. Garg, L. Sena, and T. Geva. Range of ventricular dimensions and function by steady-state free precession cine MRI in repaired tetralogy of Fallot: right ventricular outflow tract patch vs. conduit repair. *Journal of Magnetic Resonance Imaging*, 26(4):934–940, 2007. [134](#), [152](#)
- G. Sanchez-Ortiz, M. Sermesant, K. Rhode, R. Chandrashekara, R. Razavi, D. Hill, and D. Rueckert. Localization of abnormal conduction pathways for tachyarrhythmia treatment using tagged mri. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, Lecture Notes in Computer Science, pages 425–433. Springer, 2005. [222](#)
- D. Sanchez-Quintana, R. Anderson, and S. Ho. Ventricular myoarchitecture in tetralogy of fallot. *British Medical Journal*, 76(3):280, 1996. [165](#)

- G. Saporta. *Probabilités, analyse des données et statistique*. Editions Technip, 2006. 123, 124
- R. Schaback. Creating surfaces from scattered data using radial basis functions. *Mathematical methods for curves and surfaces*, pages 477–496, 1995. 41, 56
- J. Schaerer, C. Casta, J. Pousin, and P. Clarysse. A dynamic elastic model for segmentation and tracking of the heart in mr image sequences. *Medical Image Analysis*, 2010. 90
- A. Schäffler and S. Schmidt. *Anatomie Physiologie Biologie*. Maloine, Paris, 1999. 191
- S. Schievano, L. Coats, F. Migliavacca, W. Norman, A. Frigiola, J. Deanfield, P. Bonhoeffer, and A. Taylor. Variations in right ventricular outflow tract morphology following repair of congenital heart disease: implications for percutaneous pulmonary valve implantation. *Journal of Cardiovascular Magnetic Resonance*, 9(4): 687–695, 2007. 30
- H. Schmid. *Passive myocardial mechanics: constitutive laws and materials parameter estimation*. PhD thesis, University of Auckland, 2006. 171, 220
- P. Schoenhagen, A. Stillman, S. Halliburton, and R. White. CT of the heart: principles, advances, clinical uses. *Cleveland Clinic journal of medicine*, 72(2):127, 2005. 27
- G. Schwartz. Estimating the dimension of a model. *Annals of Statistics*, 6(2): 461–464, 1978. 124
- P. Segers, E. Rietzschel, M. De Buyzere, N. Stergiopulos, N. Westerhof, L. Van Bortel, T. Gillebert, and P. Verdonck. Three-and four-element windkessel models: assessment of their fitting performance in a large cohort of healthy middle-aged individuals. *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine*, 222(4):417–428, 2008. 183
- P. Sengupta, B. Khandheria, and J. Narula. Twist and untwist mechanics of the left ventricle. *Heart Failure Clinics*, 4(3):315–324, 2008. 201
- M. Sermesant. *Modèle électromécanique du coeur pour l’analyse d’image et la simulation (Electromechanical Model of the Heart for Image Analysis and Simulation)*. PhD thesis, Université de Nice Sophia Antipolis, May 2003. 26, 161, 163, 173
- M. Sermesant, C. Forest, X. Pennec, H. Delingette, and N. Ayache. Deformable biomechanical models: Application to 4D cardiac image analysis. *Medical Image Analysis*, 7(4):475–488, 2003. 34
- M. Sermesant, H. Delingette, and N. Ayache. An electromechanical model of the heart for image analysis and simulation. *IEEE TMI*, 25(5):612–625, 2006a. 156, 157, 159, 160, 172, 176, 179

- M. Sermesant, P. Moireau, O. Camara, J. Sainte-Marie, R. Andriantsimiavona, R. Cimirman, D. L. Hill, D. Chapelle, and R. Razavi. Cardiac function estimation from MRI using a heart model and data assimilation: Advances and difficulties. *Medical Image Analysis*, 10(4):642–656, 2006b. 9, 156, 165, 171, 175, 176, 179, 182, 187, 188, 216
- M. Sermesant, E. Konukoglu, H. Delingette, Y. Coudiere, P. Chinchapatnam, K. Rhode, R. Razavi, and N. Ayache. An anisotropic multi-front fast marching method for real-time simulation of cardiac electrophysiology. *Proc. FIMH 2007*, 9:160–169, 2007. 168, 169
- M. Sermesant, J. Peyrat, P. Chinchapatnam, F. Billet, T. Mansi, K. Rhode, H. Delingette, R. Razavi, and N. Ayache. Toward patient-specific myocardial models of the heart. *Heart Failure Clinics*, 4(3):289–301, 2008. 226
- M. Sermesant, F. Billet, R. Chabiniok, T. Mansi, P. Chinchapatnam, P. Moireau, J.-M. Peyrat, K. Rhode, M. Ginks, P. Lambiase, S. Arridge, H. Delingette, M. Sorine, A. Rinaldi, D. Chapelle, R. Razavi, and N. Ayache. Personalised electromechanical model of the heart for the prediction of the acute effects of cardiac resynchronisation therapy. In *Proceedings of Functional Imaging and Modeling of the Heart 2009 (FIMH'09)*, volume 5528 of *Lecture Notes in Computer Science*, pages 239–248, 3-5 June 2009. 55, 56, 156, 169, 181, 187, 188, 210, 214, 217, 220, 227
- J. Sethian et al. *Level set methods and fast marching methods*. Cambridge university press Cambridge, 1999. 47, 167, 168
- R. Setser, D. Bexell, T. O'Donnell, A. Stillman, M. Lieber, P. Schoenhagen, and R. White. Quantitative assessment of myocardial scar in delayed enhancement magnetic resonance imaging. *Journal of Magnetic Resonance Imaging*, 18(4):434–441, 2003. 163
- D. Shah, R. Judd, and R. Kim. Technology insight: Mri of the myocardium. *Nature Clinical Practice Cardiovascular Medicine*, 2(11):597–605, 2005. 20, 24
- F. Sheehan, S. Ge, G. Vick III, K. Urnes, W. Kerwin, E. Bolson, T. Chung, J. Kovalchin, D. Sahn, M. Jerosch-Herold, et al. Three-Dimensional Shape Analysis of Right Ventricular Remodeling in Repaired Tetralogy of Fallot. *The American Journal of Cardiology*, 2007. 134, 136, 140, 144, 148, 150, 218
- P. Shi and H. Liu. Stochastic finite element framework for simultaneous estimation of cardiac kinematic functions and material parameters. *Medical Image Analysis*, 7(4):445–464, 2003. 175, 176, 190
- P. Y. Simard and G. E. Mailloux. A projection operator for the restoration of divergence-free vector fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(2):248–256, 1988. 70, 72

- A. Sinusas, X. Papademetris, R. Constable, D. Dione, M. Slade, P. Shi, and J. Duncan. Quantification of 3-D regional myocardial deformation: shape-based analysis of magnetic resonance images. *American Journal of Physiology- Heart and Circulatory Physiology*, 281(2):698–714, 2001. [64](#)
- S. Song and R. Leahy. Computation of 3-D velocity fields from 3-D cine CT images of a human heart. *IEEE Transactions on Medical Imaging*, 10(3):295–306, 1991. [63](#)
- M. Sonka and M. Fitzpatrick. *Handbook of medical imaging*. SPIE, 2000. [34](#)
- C. Sorzano, P. Thévenaz, and M. Unser. Elastic registration of biological images using vector-spline regularization. *IEEE Transactions on Biomedical Engineering*, 52(4):652–663, 2005. [63](#), [82](#)
- R. Stefanescu, X. Pennec, and N. Ayache. Grid powered nonlinear image registration with locally adaptive regularization. *Medical Image Analysis*, 8(3):325–342, September 2004. [83](#)
- N. Stergiopoulos, J. Meister, and N. Westerhof. Evaluation of methods for estimation of total arterial compliance. *American Journal of Physiology- Heart and Circulatory Physiology*, 268(4):H1540, 1995. [183](#)
- N. Stergiopoulos, B. Westerhof, and N. Westerhof. Total arterial inertance as the fourth element of the windkessel model. *American Journal of Physiology- Heart and Circulatory Physiology*, 276(1):81–88, 1999. [182](#), [184](#), [217](#)
- D. Streeter. Gross morphology and fiber geometry of the heart. *Handbook of physiology*, 1:61–112, 1979. [165](#)
- M. Styner, J. Lieberman, D. Pantazis, and G. Gerig. Boundary and medial shape analysis of the hippocampus in schizophrenia. *Medical Image Analysis*, 8(3):197–203, 2004. [109](#)
- H. Sundar, C. Davatzikos, and G. Biros. Biomechanically-Constrained 4D Estimation of Myocardial Motion. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, volume 5762 of *Lecture Notes in Computer Science*, pages 257–265. Springer, 2009a. [64](#)
- H. Sundar, H. Littb, and D. Shena. Estimating myocardial motion by 4d image warping. *Pattern Recognition*, 42:2514–2526, 2009b. [103](#), [222](#)
- D. Tang, C. Yang, T. Geva, and P. del Nido. Two-layer passive/active anisotropic fsi models with fiber orientation: Mri-based patient-specific modeling of right ventricular response to pulmonary valve insertion surgery. *Molecular & cellular biomechanics: MCB*, 4(3):159, 2007a. [156](#)
- D. Tang, C. Yang, T. Geva, and P. J. del Nido. Patient-specific virtual surgery for right ventricle volume reduction and patch design using MRI-based 3D FSI RV/LV/patch models. In *Proc. CME 2007*, pages 157–162, 2007b. [157](#), [185](#), [194](#)

- C. Tanner, J. Schnabel, A. Degenhard, A. Castellano-Smith, C. Hayes, M. Leach, D. Hose, D. Hill, and D. Hawkes. Validation of volume-preserving non-rigid registration: application to contrast-enhanced mr-mammography. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, Lecture Notes in Computer Science, pages 307–314. Springer, 2002. [63](#), [82](#), [102](#)
- C. Taylor and M. Draney. Experimental and computational methods in cardiovascular fluid mechanics. *Annual Review of Fluid Mechanics*, 2004. [178](#), [185](#)
- K. Ten Tusscher, D. Noble, P. Noble, and A. Panfilov. A model for human ventricular tissue. *American Journal of Physiology- Heart and Circulatory Physiology*, 286(4):H1573, 2004. [167](#)
- D. Terzopoulos. Multiresolution computation of visible-surface representations. *Ph. D. thesis, Department of Electrical Engineering and Computer Science, MIT*, 1984. [47](#)
- A. Teske, B. De Boeck, P. Melman, G. Sieswerda, P. Doevendans, and M. Cramer. Echocardiographic quantification of myocardial function using tissue deformation imaging, a guide to image acquisition and analysis using tissue doppler and speckle tracking. *Cardiovascular Ultrasound*, 5:27, 2007. [20](#), [21](#), [22](#), [89](#), [98](#)
- J. Therrien, S. Siu, P. McLaughlin, P. Liu, W. Williams, and G. Webb. Pulmonary valve replacement in adults late after repair of tetralogy of Fallot: are we operating too late? *Journal of the American College of Cardiology*, 36(5):1670–1675, 2000. [30](#), [205](#)
- J. Therrien, Y. Provost, N. Merchant, W. Williams, J. Colman, and G. Webb. Optimal timing for pulmonary valve replacement in adults after tetralogy of Fallot repair. *The American journal of cardiology*, 95(6):779–782, 2005. [30](#), [108](#), [134](#)
- J. P. Thirion. Image matching as a diffusion process: an analogy with Maxwell’s demons. *Medical Image Analysis*, pages 243–260, 1998. [62](#), [65](#)
- D. Thompson. *On Form and Growth*. Cambridge University Press, 1917. [109](#)
- P. Thompson, J. Giedd, R. Woods, D. MacDonald, A. Evans, and A. Toga. Growth patterns in the developing brain detected by using continuum mechanical tensor maps. *Nature*, 404(6774):190–193, 2000. [108](#), [125](#), [223](#)
- K. Tomlinson. *Finite element solution of an eikonal equation for excitation wavefront propagation in ventricular mycodium*. PhD thesis, University of Auckland, 2000. [168](#)
- E. Topol and R. Califf. *Textbook of cardiovascular medicine*. Lippincott Williams & Wilkins, 2007. [20](#), [21](#), [27](#)

- N. Toussaint, T. Mansi, H. Delingette, N. Ayache, and M. Sermesant. An Integrated Platform for Dynamic Cardiac Simulation and Image Processing: Application to Personalised Tetralogy of Fallot Simulation. In *Proc. Eurographics Workshop on Visual Computing for Biomedicine (VCBM)*, Delft, The Netherlands, 2008. 56, 214, 227
- N. Toussaint, C. Stoeck, S. Kozerke, M. Sermesant, and P. G. Batchelor. In-vivo Human 3D Cardiac Fibre Architecture: Reconstruction Using Curvilinear Interpolation of Diffusion Tensor Images. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, LNCS. Springer, 2010. To appear. 26, 188, 222
- W. Tseng, T. Reese, R. Weisskoff, T. Brady, and V. Wedeen. Myocardial fiber shortening in humans: Initial results of mr imaging1. *Radiology*, 216(1):128, 2000. 26
- Z. Tu. Probabilistic boosting-tree: Learning discriminative models for classification, recognition, and clustering. In *IEEE International Conference on Computer Vision, 2005. ICCV 2005*, volume 2, 2005. 35
- G. Turk and J. O'Brien. Variational implicit surfaces. Technical report, Georgia Institute of Technology, 1999a. 41
- G. Turk and J. O'Brien. Shape transformation using variational implicit functions. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, page 342. ACM Press/Addison-Wesley Publishing Co., 1999b. 41
- G. Turk and J. O'brien. Modelling with implicit surfaces that interpolate. *ACM Transactions on Graphics (TOG)*, 21(4):873, 2002. 40, 41, 43, 44
- S. Ubbink, P. Bovendeerd, T. Delhaas, T. Arts, and F. van de Vosse. Towards model-based analysis of cardiac mr tagging data: relation between left ventricular shear strain and myofiber orientation. *Medical Image Analysis*, 10(4):632–641, 2006. 222
- F. Vadakkumpadan, L. Rantner, B. Tice, P. Boyle, A. Prassl, E. Vigmond, G. Plank, and N. Trayanova. Image-based models of cardiac structure with applications in arrhythmia and defibrillation studies. *Journal of electrocardiology*, 42(2):157, 2009. 160
- M. Vaillant and J. Glaunes. Surface matching via currents. In *Proc. IPMI 2005*, page 381. Springer, 2005. 114, 115
- H. C. van Assen, M. G. Danilouchkine, A. F. Frangi, S. Ordas, J. J. Westenberg, J. H. Reiber, and B. P. Lelieveldt. SPASM: A 3D-ASM for segmentation of sparse and arbitrarily oriented cardiac MRI data. *Medical Image Analysis*, 10(2):286–303, April 2006. 34
- P. Van Dijk. Direct cardiac nmr imaging of heart wall and blood flow velocity. *Journal of computer assisted tomography*, 8(3):429, 1984. 24

- T. Vercauteren, X. Pennec, A. Perchant, and N. Ayache. Symmetric log-domain diffeomorphic registration: A demons-based approach. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, volume 5241 of *Lecture Notes in Computer Science*, pages 754–761, New York, USA, September 2008. Springer-Verlag. 7, 62, 65, 66, 67, 80, 214, 231
- T. Vercauteren, X. Pennec, A. Perchant, and N. Ayache. Diffeomorphic demons: Efficient non-parametric image registration. *NeuroImage*, 45(1S1):61–72, 2009. 35, 45, 47, 62, 66, 86, 231, 232
- A. Veress, G. Gullberg, and J. Weiss. Measurement of Strain in the Left Ventricle during Diastole with cine-MRI and Deformable Image Registration. *Journal of Biomechanical Engineering*, 127:1195, 2005. 64, 90
- F. Vetter and A. McCulloch. Three-dimensional analysis of regional cardiac function: a model of rabbit ventricular anatomy. *Progress in biophysics and molecular biology*, 69(2-3):157–183, 1998. 159, 165
- H. Wendland. *Scattered data approximation*. Cambridge University Press, 2005. 41, 42
- J. Wenk, Z. Zhang, G. Cheng, D. Malhotra, G. Acevedo-Bolton, M. Burger, T. Suzuki, D. Saloner, A. Wallace, and J. Guccione. First finite element model of the left ventricle with mitral valve: Insights into ischemic mitral regurgitation. *The Annals of Thoracic Surgery*, 89(5):1546, 2010. 185
- N. Westerhof, G. Elzinga, and P. Sipkema. An artificial arterial system for pumping hearts. *Journal of Applied Physiology*, 31(5):776, 1971. 182
- N. Westerhof, J. Lankhaar, and B. Westerhof. The arterial windkessel. *Medical and Biological Engineering and Computing*, 47(2):131–141, 2009. 182, 183, 184, 191
- H. Wold. Estimation of principal components and related models by iterative least squares. *Multivariate analysis*, pages 391–420, 1966. 126
- S. Wold, M. Sjöström, and L. Eriksson. Pls-regression: a basic tool of chemometrics. *Chemometrics and intelligent laboratory systems*, 58(2):109–130, 2001. 126
- K. Wong, L. Wang, H. Zhang, H. Liu, and P. Shi. Integrating functional and structural images for simultaneous cardiac segmentation and deformation recovery. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, volume 4791 of *Lecture Notes in Computer Science*, page 270. Springer, 2007. 156, 176
- K. Wong, L. Wang, H. Zhang, H. Liu, and P. Shi. Meshfree implementation of individualized active cardiac dynamics. *Computerized Medical Imaging and Graphics*, 34(1):91–103, 2010. 158, 159, 165, 189, 221, 222

- K. Worsley. An overview and some new developments in the statistical analysis of PET and fMRI data. *Human Brain Mapping*, 5(4):254–258, 1997. [126](#)
- E. Wu, Y. Wu, J. Nicholls, J. Wang, S. Liao, S. Zhu, C. Lau, and H. Tse. Mr diffusion tensor imaging study of postinfarct myocardium structural remodeling in a porcine model. *Magnetic Resonance in Medicine*, 58(4):687–695, 2007. [188](#), [222](#)
- M. Wu, M. Su, Y. Huang, K. Chiou, P. Yang, H. Pan, T. Reese, V. Wedeen, and W. Tseng. Sequential changes of myocardial microstructure in patients postmyocardial infarction by diffusion-tensor cardiac mr: correlation with left ventricular structure and function. *Circulation: Cardiovascular Imaging*, 2(1):32, 2009. [26](#), [165](#), [188](#)
- C. Yang, D. Tang, T. Geva, and P. del Nido. Mri-based patient-specific computational modeling of right ventricular response to pulmonary valve insertion surgery: a passive anisotropic fsi model with fiber orientationspecific computational modeling of right ventricular response to pulmonary valve insertion surgery: a passive anisotropic fsi model with fiber orientation. In *International Conference on BioMedical Engineering and Informatics*, 2008a. [156](#), [185](#), [218](#)
- L. Yang, B. Georgescu, Y. Zheng, D. Foran, and D. Comaniciu. A fast and accurate tracking algorithm of left ventricles in 3D echocardiography. In *Proceedings/IEEE International Symposium on Biomedical Imaging: from nano to macro. IEEE International Symposium on Biomedical Imaging*, volume 5, page 221. NIH Public Access, 2008b. [57](#), [136](#), [198](#), [214](#)
- Y. Yang, D. Rueckert, and A. Bull. Predicting the shapes of bones at a joint: application to the shoulder. *Computer Methods in Biomechanics and Biomedical Engineering*, 11(1):19–30, 2008c. [216](#)
- B. Yeo, M. Sabuncu, T. Vercauteren, N. Ayache, B. Fischl, and P. Golland. Spherical demons: Fast diffeomorphic landmark-free surface registration. *IEEE Transactions on Medical Imaging*, 22:55, 2009. [233](#)
- P. Yock and R. Popp. Noninvasive estimation of right ventricular systolic pressure by doppler ultrasound in patients with tricuspid regurgitation. *Circulation*, 70(4):657, 1984. [181](#)
- T. Yoo, B. Morse, K. Subramanian, P. Rheingans, and M. Ackerman. Anatomic modeling from unstructured samples using variational implicit surfaces. *Studies in Health Technology and Informatics*, pages 594–600, 2001. [56](#)
- A. Young, D. Kraitchman, L. Dougherty, and L. Axel. Tracking and finite element analysis of stripe deformation in magnetic resonance tagging. *IEEE Transactions on Medical Imaging*, 14(3):413–421, 1995. [103](#)

- A. Youssef, E. Ibrahim, G. Korosoglou, M. Abraham, R. Weiss, and N. Osman. Strain-encoding cardiovascular magnetic resonance for assessment of right-ventricular regional function. *Journal of Cardiovascular Magnetic Resonance*, 10(1):33, 2008. [24](#)
- J. L. Zamorano. Development in 3D echocardiography. *European Cardiology*, 5:3–7, 2009. [21](#), [22](#)
- E. Zerhouni, D. Parish, W. Rogers, A. Yang, and E. Shapiro. Human heart: tagging with MR imaging—a method for noninvasive assessment of myocardial motion. *Radiology*, 169(1):59–63, 1988. [24](#), [89](#)
- K. Zervan, C. Male, T. Benesch, and U. Salzer-Muhar. Ventricular interaction in children after repair of tetralogy of fallot: a longitudinal echocardiographic study. *European Journal of Echocardiography*, 2009. [223](#)
- H. Zhang, B. Li, A. Young, and P. Hunter. Recovery of myocardial kinematic function without the time history of external loads. *EURASIP Journal on Advances in Signal Processing*, pages 1–9, 2010a. [176](#)
- H. Zhang, A. Wahle, R. Johnson, T. Scholz, and M. Sonka. 4D Cardiac MR Image Analysis: Left and Right Ventricular Morphology and Function. *IEEE Transactions on Medical Imaging*, 29(2):350–364, February 2010b. [34](#), [35](#), [134](#)
- Y. Zheng, A. Barbu, B. Georgescu, M. Scheuering, and D. 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*, 27(11):1668–1681, 2008. [34](#), [35](#), [136](#), [161](#), [197](#), [198](#), [214](#)
- Y. Zheng, B. Georgescu, and D. Comaniciu. Marginal Space Learning for Efficient Detection of 2D/3D Anatomical Structures in Medical Images. In *Information Processing in Medical Imaging (IPMI)*, Lecture Notes in Computer Science, pages 411–422. Springer, 2009. [35](#)
- L. Zhong, R. Tan, Y. Su, S. Yeo, T. Chua, T. Koh, D. Ghista, and G. Kassab. *Computational Cardiovascular Mechanics*, chapter Noninvasive Assessment of Left Ventricular Remodeling: Geometry, Wall Stress, and Function, pages 179–196. Springer, 2010. [134](#)
- Y. Zhu, X. Papademetris, A. Sinusas, and J. Duncan. A coupled deformable model for tracking myocardial borders from real-time echocardiography using an incompressibility constraint. *Medical Image Analysis*, 2010. [34](#), [35](#), [56](#)