



HAL
open science

Reduced-order statistical models of cardiac growth, motion and blood flow : application to the tetralogy of Fallot heart

Kristin Mcleod

► **To cite this version:**

Kristin Mcleod. Reduced-order statistical models of cardiac growth, motion and blood flow : application to the tetralogy of Fallot heart. Other. Université Nice Sophia Antipolis, 2013. English. NNT : 2013NICE4091 . tel-00942556

HAL Id: tel-00942556

<https://theses.hal.science/tel-00942556>

Submitted on 6 Feb 2014

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

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



UNIVERSITÉ DE NICE - SOPHIA ANTIPOLIS
ÉCOLE DOCTORALE STIC
SCIENCES ET TECHNOLOGIES DE L'INFORMATION
ET DE LA COMMUNICATION

THÈSE

pour obtenir le titre de

Docteur en Sciences

de l'Université de Nice - Sophia Antipolis

Mention : AUTOMATIQUE, TRAITEMENT DU SIGNAL ET DES IMAGES

Présentée et soutenue par

Kristin MCLEOD

Modèles Statistiques Réduits de la Croissance Cardiaque, du Mouvement et de la Circulation Sanguine

- Application à la Tétralogie de Fallot -

Thèse dirigée par Maxime SERMESANT et Xavier PENNEC

préparée à l'INRIA Sophia Antipolis, Projet ASCLEPIOS

soutenue le 8 Novembre 2013

Jury :

<i>Rapporteurs :</i>	Alistair YOUNG	- The University of Auckland
	Daniel RUECKERT	- Imperial College London
<i>Directeurs :</i>	Xavier PENNEC	- INRIA Asclepios
	Maxime SERMESANT	- INRIA Asclepios
<i>Président :</i>	Reza RAZAVI	- King's College London
<i>Examineurs :</i>	Patrick CLARYSSE	- CNRS Lyon
<i>Invités :</i>	Mathieu DE CRAENE	- Philips
	Tommaso MANSI	- Siemens Corporate Technology



UNIVERSITY OF NICE - SOPHIA ANTIPOLIS
DOCTORAL SCHOOL STIC
SCIENCES, INFORMATION TECHNOLOGIES
AND COMMUNICATION

PHD THESIS

to obtain the title of

PhD of Science

from the University of Nice - Sophia Antipolis

Speciality : CONTROL SYSTEMS, SIGNAL AND IMAGE PROCESSING

Defended by

Kristin MCLEOD

Reduced-Order Statistical Models of Cardiac Growth, Motion and Blood Flow - Application to the Tetralogy of Fallot Heart -

Thesis Advisors: Maxime SERMESANT and Xavier PENNEC

prepared at INRIA Sophia Antipolis, ASCLEPIOS research group

defended on November 8, 2013

Jury :

<i>Reviewers :</i>	Alistair YOUNG	-	The University of Auckland
	Daniel RUECKERT	-	Imperial College London
<i>Advisors :</i>	Xavier PENNEC	-	INRIA Asclepios
	Maxime SERMESANT	-	INRIA Asclepios
<i>President :</i>	Reza RAZAVI	-	King's College London
<i>Examiners :</i>	Patrick CLARYSSE	-	CNRS Lyon
<i>Invited :</i>	Mathieu DE CRAENE	-	Philips
	Tommaso MANSI	-	Siemens Corporate Technology

Modèles statistiques réduits de la croissance cardiaque, du mouvement et de la circulation sanguine - Application à la tétralogie de Fallot

Résumé: Cette thèse présente les travaux réalisés en vue de l'élaboration d'un modèle cardiaque associant croissance, mouvement et circulation sanguine pour permettre ensuite la construction d'un modèle prédictif par patient à partir d'un modèle standard de l'ensemble d'une population.

Le premier axe de ce travail est la simulation de la croissance bi-ventriculaire due au vieillissement. Un modèle existant de surface unique, calculé à l'aide de méthodes statistiques, a été généralisé à un modèle bi-ventriculaire puis appliqué à des patients souffrant de la tétralogie de Fallot afin de modéliser l'évolution complexe des ventricules atteints de la pathologie.

Le deuxième axe concerne la modélisation du mouvement cardiaque au niveau de la population. Un modèle d'ordre réduit basé sur un modèle polyaffine et LogDemon a été proposé. Il simule spécifiquement la dynamique cardiaque avec un petit nombre de paramètres. Après calcul des transformations, les paramètres sont analysés par des méthodes statistiques afin d'obtenir des mesures moyennes pour la population. Un modèle de mouvement moyen a été calculé pour représenter le mouvement standard de la population.

Le troisième s'intéresse à la simulation de l'écoulement sanguin à l'échelle de la population. La complexité des simulations spécifiques à un patient a été réduite grâce à l'utilisation de méthodes relatives à l'analyse d'image et à la dynamique des fluides numérique ainsi qu'à des techniques de réduction d'ordre de modèle. La simulation du flux sanguin dans l'artère pulmonaire pour des patients atteints de tétralogie de Fallot a permis de mieux comprendre l'impact du sang régurgité sur la pression et la vitesse.

Étant donné nos contributions sur ces trois axes, nous sommes maintenant en bonne position pour élaborer le modèle couplé des contributions interdépendantes de la croissance, du mouvement et de l'écoulement sanguin. Ce modèle pourrait être utilisé afin d'aider la planification de la thérapie et la prise de décision chez les patients souffrant de maladies cardiaques.

Mots-clés: Analyse d'images médicales; Recalage; Modèles statistiques réduits; Tétralogie de Fallot

Reduced-Order Statistical Models of Cardiac Growth, Motion and Blood Flow - Application to the Tetralogy of Fallot Heart

Abstract: This thesis presents work towards a coupled model of cardiac growth, motion, and blood flow to enable predictive patient-specific models to be built from a population-based model.

The first axis of this work is to simulate bi-ventricular growth through aging. A previously proposed single surface model computed using statistical methods was extended to a bi-ventricular model and applied to Tetralogy of Fallot patients to model the complex evolution of the ventricles due to the pathology.

The second axis concerns the development of a model to simulate cardiac motion at a population level. A reduced-order cardiac-specific motion model was proposed to simulate the motion dynamics with a small number of parameters using a Polyaffine and LogDemons based model. From the computed transformations, the parameters were analysed using statistical methods to obtain population-based measures of normality. A mean motion model was derived to represent the normal motion for a given population.

The third axis is to develop a reduced-order framework for blood flow simulation to obtain population-based flow dynamics. The complexity of patient-specific simulations was reduced by combining image analysis, computational fluid dynamics and model order reduction techniques. Blood flow through the pulmonary artery in Tetralogy of Fallot patients was simulated to better understand the impact of regurgitated blood on pressure and velocity.

Given our contributions on these three axes, we are now in a good position to couple the models in order to capture the interrelated contributions of growth, motion and flow. Such a model could be used to aid in therapy planning and decision making for patients with heart disease.

Keywords: Medical image analysis; Non-rigid registration; Statistical model reduction; Tetralogy of Fallot

Acknowledgments

I would like to firstly thank my PhD supervisors Xavier Pennec and Maxime Sermesant, for your support, encouragement, and open minds. Working together has been an extremely valuable scientific experience. It has been a great pleasure working with people that are not only knowledgeable, but also passionate and creative. It's truly an honour to work with people that you look up to and respect both on a professional and personal level. I am extremely grateful to you both for taking the time whenever needed to discuss new ideas, for being patient, for pushing me into challenging situations that led to great results, and for being able to have a laugh. The last three years haven't felt like work. I will miss the intense lengthy meetings where I come with one question to answer and leave with 5 more interesting unanswered questions!

A big thank you to the thesis reviewers Daniel Rueckert and Alistair Young for taking the time to read this thesis and for your useful insights and feedback. I feel very fortunate to have had the opportunity to interact with you both at conferences and during this process. Thank you also to Reza Razavi, Patrick Clarysse and Mathieu De Craene for participating as jury members for the defense of this thesis. A special thank you to Reza Razavi for giving me the opportunity to do a clinical rotation at Evelyne Children's Hospital. It was an unforgettable and eye-opening experience.

During this PhD I was fortunate to work with a number of very intelligent and interesting people. I would like to take the time to thank these people now. Tommaso Mansi - thank you for inspiring me at the beginning of this thesis, I will always remember our deep, insightful and funny discussions. Christof Seiler and Nicolas Toussaint, with whom I worked with extensively, it was truly a pleasure both inside and outside of the lab. Alfonso Caiazzo - for introducing me to new ideas and projects. Adityo Prakosa - thank you for putting up with all my questions and always welcoming me with a smile. Mathieu De Craene and Catalina Tobon Gomez - for putting up with the millions of emails and for always being a pleasure to work with.

A key element of the success of this thesis was the amazing work environment. Thank you to Nicholas Ayache for providing such a great work place and for encouraging and supporting interaction between the team members and with partners outside of the team. I've enjoyed meeting and working together with everyone from the team. Stephanie Marchesseau - the time at ASCLEPIOS wouldn't have been the same without you. Thank you for all the support, help, and most of all for all the fun times. Thank you to my office mates Marine Breuilly, Ken Wong, Brina Goyette, Rocio Cabrera Lozoya, Sonia Durand and Nina Mio for all the fun times, tea breaks and girly chats. Thanks to Jan Margeta for being a great work husband. Thank you so much to Isabelle Strobant for the help, support and patience starting from before I even arrived to France.

I finally would like to thank my family for their love, support, encouragement and kindness. I truly would never have succeeded without the support of my mum Janice McLeod. Thank you for always pushing me to do what makes me happy, no matter the obstacles or challenges, and for showing me that everything is possible. This thesis is dedicated to you.

Thank you all!

Table of Contents

1	Clinical Context	3
1.1	Introduction	3
1.2	Key Steps to Completing the Jigsaw Puzzle	6
1.3	Thesis Organisation and Main Contributions	7
1.4	List of Publications and Awards	10
2	The Role of Imaging and Analysis in Cardiology	13
2.1	The Human Heart	13
2.2	Medical Imaging in Cardiology	17
2.3	Inter and Intra-Patient Analysis using Image and Mesh Registration	23
2.4	Conclusion	27
I	STATISTICAL MODELLING OF CARDIAC GROWTH	29
3	Bi-ventricular Growth Modeling in the Tetralogy of Fallot Population	31
3.1	Chapter Overview	31
3.2	Introduction	32
3.3	Statistical Shape Analysis	34
3.4	Application on ToF Data	41
3.5	Conclusion	50
II	STATISTICAL MODELLING OF CARDIAC MOTION	51
4	Prerequisite: Quantitative Validation of the iLogDemons Method	53
4.1	Chapter Overview	53
4.2	Application to Real Sequences from Healthy Volunteers	54
4.3	Application to Synthetic Echocardiography Sequences	65
5	Anatomically Constrained Regional Motion Model for Cardiac Strain Analysis	73
5.1	Chapter Overview	74
5.2	Introduction	74
5.3	Background: Polyaffine LogDemons	79
5.4	Cardiac-Specific Polyaffine LogDemons	82
5.5	Experiments	89
5.6	Discussion	96
5.7	Conclusion	99
5.8	Appendix: Numerical Implementation	99

6	Polyaffine Parameter Analysis	101
6.1	Chapter Overview	101
6.2	Regional Analysis of Left Ventricle Function using a Cardiac-Specific Polyaffine Motion Model	102
6.3	Spatio-Temporal Dimension Reduction of Cardiac Motion for Group-Wise Analysis and Statistical Testing	105
7	Regional Mean Trajectory Model of the Left Ventricle for Classification and Population-Based Motion Analysis	111
7.1	Chapter Overview	112
7.2	Introduction	113
7.3	Transformation Parameter Estimation using a Polyaffine Model	116
7.4	Spatio-Temporal Parameter Alignment	117
7.5	Mean Motion Model Construction	121
7.6	Spatial and Temporal Mode Construction using Tensor Decomposition	123
7.7	Healthy vs. Tetralogy of Fallot Motion	126
7.8	Discussion and Perspectives	131
7.9	Conclusion	134
7.10	Appendix	134
III	STATISTICAL MODELLING OF CARDIAC BLOOD FLOW	135
8	Group-wise Construction of Reduced Models for Understanding and Characterization of Pulmonary Blood Flows from Medical Images	137
8.1	Chapter Overview	138
8.2	Introduction	138
8.3	Computation of an Average Geometry	141
8.4	An Atlas-based Reduced Order Model of Blood Flow	145
8.5	Numerical Experiments on the Pulmonary Artery of Repaired Tetralogy of Fallot Patients	150
8.6	Discussion	162
8.7	Conclusion	165
9	A Preliminary Approach for a 5D Cardiac Model	169
9.1	Chapter Overview	169
9.2	Introduction	169
9.3	Static Growth Modelling using Cross-Sectional Statistics	171
9.4	Dynamic Growth Modelling using Regression	171
9.5	Coupled Model of Cardiac Growth, Motion and Blood Flow	173
9.6	Perspectives	174
10	Conclusions and Perspectives	177
10.1	Scope of this Thesis	177
10.2	Main Contributions	178

Table of Contents

ix

10.3 Collaborative Work 180
10.4 Perspectives and Future Applications 181
10.5 Conclusion 183

Bibliography

185

INTRODUCTION

Clinical Context

Contents

1.1 Introduction	3
1.1.1 The Care4Me European Project	4
1.1.2 Global Objective of This Work	4
1.2 Key Steps to Completing the Jigsaw Puzzle	6
1.3 Thesis Organisation and Main Contributions	7
1.4 List of Publications and Awards	10

1.1 Introduction

Medical imaging is becoming more and more widely used in clinical practice to detect abnormalities within tissue, bones, and organs. Imaging can be performed over several time periods to acquire up-to-date images to keep track of disease progression, as well as to thoroughly analyse organ, bone and tissue structure to aid in decision making and therapy planning. Medical imaging provides a vital tool in everyday clinical work-flow to increase the quality of patient care.

Medical image analysis aims to go even further to aid clinicians by modelling what isn't measured from the images directly to obtain quantitative measures of structure and function. In the case of cardiovascular modelling, medical image analysis techniques can be used to model the different structures, to obtain quantitative measures such as size and volume of the ventricles and atria, or to model function such as motion, blood flow, electrophysiology, mechanics etc. Naturally, these measures can be useful in determining and measuring how well the heart is functioning; whether the heart is beating efficiently, pumping enough blood, beating synchronously, and so on. Modelling what isn't observed directly can provide complementary information to aid clinicians with diagnosis, prognosis, and treatment planning.

Typical treatment planning for cardiovascular disease involves a meeting between the cardiologists, physiologists and surgeons to discuss and come to a consensus on the best method of treatment to pursue for a given patient. In this situation, the choice is guided by the medical images, which give an indicator of the severity of the disease, and by the knowledge and experience of the cardiologists and surgeons regarding potential treatment options. Obviously, this decision making process may be flawed and biased by the experience of the clinicians, the willingness to perform potentially risky procedures, the interpretation of the data available, and the ability of a clinician to convince the others of their

opinion on what is the best choice of treatment. Potentially, medical image analysis techniques could be used to improve this process by giving clinicians more indicators to firstly better understand the current status of the disease, and secondly to give a better idea of how a patient will react to a given therapy.

1.1.1 The Care4Me European Project

A large proportion of this thesis was conducted in the context of the Cooperative Advanced REsearch for Medical Efficiency (Care4Me) ITEA2 project, a European project aimed towards developing clinical indicators of disease within three key areas: cancerous, cardiovascular, and neuro-degenerative diseases. The overall aim of the Care4Me project was to offer better patient care by developing innovative tools towards improved diagnosis, prognosis and treatment planning. The project was centered around advancing tools at the software level from advanced structural imaging, towards enhanced functional imaging, as shown in Fig. 1.1.




Figure 1.1: The overall aim of the Care4Me project towards improved imaging by developing methods at the structural imaging level.

The Care4Me project aimed to reach this goal by developing new methods for medical image analysis to aid clinicians by providing more quantitative measures, extracted from potentially large data sets. Using the hospital information systems and known medical knowledge, patient data can potentially be used to drive data analysis and image-based models, leading towards decision support for patient care. Our role in this process is in developing image-based models of anatomy and pathology, as circled in Fig. 1.2, in the area of cardiovascular imaging in particular.

1.1.2 Global Objective of This Work

Given that an average heart beats over a billion times during a lifetime, maintaining healthy motion is important for the heart to be able to sustain such a workload. Since abnormal




Figure 1.2: The key steps in the pipeline towards improved systems for patient care, by incorporating hospital systems and medical knowledge with patient data sets to drive new models to aid in diagnosis and treatment planning. This thesis describes work towards image-based modelling of cardiovascular disease (circled).

motion is commonly observed in patients with cardiovascular disease, clinicians are interested in knowing whether the heart is beating sufficiently well or not, to determine what, if any, course of action is needed. Modelling the heart motion can thus be useful in better determining and understanding cardiac function. Predicting the outcome of a given treatment option can also provide useful information for clinicians since decision making can rely largely on previous experience with a given treatment method. For example, a clinician may be largely inclined to perform a given therapy based on the fact that the same method worked in a majority of cases in the past. However, this may not be the optimal method for a specific patient, and may not result in the best possible outcome for the long-term. Therefore, as an added predictor of the outcome of a therapy, models could be used to give more of an indication of how the patient will develop over time. For example, knowing the typical growth dynamics, the growth of the heart of a specific patient can be modelled over time to analyse the development of the heart by first simulating therapy (such as surgery) on the patient, then propagating over time. Modelling the blood flow dynamics can give even further information since abnormal blood flow can result in inadequate pressure of blood flowing through the heart, increased stress on the tissue wall, and pressure overload. Furthermore, the vast amount of data available could be used in a more productive way to discover new treatment methods and to determine optimal treatment options for a given patient based on the success of such a choice of treatment on similar patients from around the world.

This thesis is aimed towards these exact clinical challenges to model cardiac growth, motion and blood flow. Moreover, by incorporating a range of data, potentially from different time points, different patients, different scanners, different imaging modalities, from different centers and even from different countries, we aim to make use of this data to ob-

tain population-based evaluations of cardiac shape and function. Hence, this work leads towards a 5D+flow population-based model of the heart incorporating the 3-dimensional dynamics of shape, the time dimension of the heart beat, the time dimension as the patient ages, as well as the flow dynamics. A key theme throughout this thesis is reduced-order modelling, where we attempt to simulate the different phenomena with models that are described with a small number of parameters that can be meaningfully analysed.

1.2 Key Steps to Completing the Jigsaw Puzzle

The global objective of this work is to develop a generative model of a beating and growing heart with blood flowing through it. In order to construct such a model, first a low-dimensional model of each individual component is required. The key questions addressed in this thesis are described in the following three sections.

1) Generative Model of a Growing Heart

Understanding how the heart develops as a patient grows can give useful insight for not only determining the prognosis of a disease, but also to aid in therapy planning. By modelling the growth, it may be possible to better understand how the heart will develop and evolve over time given a certain choice of therapy. Thus, it can be of interest to model the growth for a given population to understand the shape evolution over time. Furthermore, given the coupled shape dynamics between the left and the right ventricle, it may be of interest to model the bi-ventricular growth over time. Patient data is becoming more and more widely available for research purposes, however, at the present time, few longitudinal data-sets are available for the same patients to be able to model the development over time for a given patient. Furthermore, we are interested in modelling in a group-wise manner rather than patient-specific modelling, in order to better understand the typical growth patterns for a given population. Thus, the key question is:

- *How can we model bi-ventricular growth over time for a chosen population with few longitudinal data-sets available, that can be generalised to a given population?*

2) Generative Model of a Beating Heart

The motion dynamics observed during the cardiac cycle can be useful indicators for the purposes of diagnosis, prognosis, as well as therapy planning. The way in which the ventricles move during the cardiac cycle can be very different for healthy and pathological cases. The aim of modelling these different dynamics is to attempt to provide better patient care through better understanding of what is normal motion and how much a patient differs from the norm. Given that motion sequences are becoming more and more available for both normal and diseased patients, this source of information could potentially be used to draw meaningful conclusions for a given population. Therefore, the key question here is:

- *How can we model cardiac motion over time at a population level to be able to capture the typical motion dynamics observed for a given population, by making use of a data-set of several subjects drawn from the same population?*

3) Generative Model of Blood Flow

Understanding the blood flow dynamics through the cardiovascular system is important for understanding different pathologies and for predicting the strain on the heart under current conditions. Given the complex flow dynamics and the intricate coupling between organ shape and the flow through that organ, modelling the blood flow at a patient-specific level is already challenging, and modelling at a population-level can be even more challenging. This leads to the key question:

- *How can we simulate blood flow through a given organ at a population level to be able to develop a generative blood flow model for a given population?*

1.3 Thesis Organisation and Main Contributions

This thesis is organised with the aim of reading like a jigsaw puzzle, with each piece joining together to complete the final picture, which is the development of a 5D+flow model of the beating, growing, and flowing heart. Following a brief introduction of the heart, imaging techniques and state-of-the-art (in Chapter 2), the thesis is divided according to the key questions described in the previous section as follows:

Part - I

Statistical Modelling of Cardiac Growth

Chapter 3 describes a method towards answering the first key question. Namely, a methodological framework for constructing a bi-ventricular growth model from a data-set of subjects from the same population at different ages to generate a statistical model of growth is proposed. The key contribution of this chapter is the extension of the previously proposed method for modelling the growth of the right ventricle [Mansi et al., 2011b], to include both ventricles. In the case of the left ventricle, both the endocardial and epicardial surfaces can be modelled since the muscle is thicker than that of the right ventricle. Thus, the described method models the dynamics of three surfaces (LV endocardium and epicardium, RV endocardium) with aging. Obviously, since the ventricles are attached, they have an effect on each other, and so by modelling the growth of both ventricles jointly, we hope to capture this coupling. The method was applied to model the growth dynamics in patients suffering from Tetralogy of Fallot. The computed modes of variation were interpreted in relation to the pathology to obtain meaningful conclusions about the shape variability in the population. This chapter was published in Modeling in Computational Biology and Biomedicine, Lectures Notes in Mathematical and Computational Biology [McLeod et al., 2013a].

Part - II

Statistical Modelling of Cardiac Motion

Chapter 4 describes work towards testing and benchmarking state-of-the-art methods for cardiac motion tracking by applying previous methods to data-sets with given ground truth. Given that a large proportion of the PhD work was devoted to the development of a robust motion model, analysing and testing state-of-the-art methods was an important first step. The incompressible LogDemons (iLogDemons) method was applied first to a data-set of image sequences of healthy volunteers, with manually annotated 4D landmarks used to compute the error of the tracking. The same method was also applied to a data-set of synthetically generated sequences with known ground truth displacement. This method was quantitatively compared to other motion tracking methods. These applications were published in the STACOM workshop proceedings of 2011 [McLeod et al., 2012a] and 2012 ([Prakosa et al., 2012a]).

Chapter 5 presents a reduced-order motion model for tracking cardiac motion along the cycle from medical images. Based on the findings from the analysis of motion tracking algorithms described in Chapter 4, the key missing component was robustness and reproducibility of the current method. The key contribution of this chapter is the proposal of a new method for modelling cardiac motion with a greatly reduced number of parameters using image registration techniques. The method builds on the previous Polyaffine LogDemons method to incorporate cardiac-specific model constraints (namely an incompressibility penalisation, and neighbouring region regularisation), as well as cardiac specific priors (namely, curved weight functions fitting the shape of the ventricle, and clinically relevant Polyaffine region definition). Image-driven spatial priors are added to give varying confidence of the input velocity vectors depending on areas of high gradient in the image. Furthermore, the motion in the epicardium was restricted to prevent the epicardial motion from being dragged by the much larger endocardial motion. Since few parameters are needed to describe the motion, the method is robust, and since the definition of the Polyaffine regions are consistent from subject-to-subject, the method can be applied to study differences between populations. The method was applied to the STACOM 2011 cardiac motion tracking data-set, from which we obtained physiologically reasonable deformations at a level of accuracy consistent with state-of-the-art methods, but with a greatly reduced number of parameters. This chapter is in preparation for submission to a journal and builds on work published in the STACOM workshop proceedings of 2012 ([McLeod et al., 2012b]).

Chapter 6 describes the application of the reduced-order motion model presented in Chapter 5 to analyse the transformation parameters. Given that there are few parameters in the motion model that are consistently represented in different subjects, meaningful comparison of the parameters between subjects can be performed. The parameters were first analysed at a regional level for comparison between healthy and pathological subjects. By analysing the parameters directly, we observed interesting differences in terms of

synchrony of motion between asymptomatic subjects and patients with heart failure. The parameters were further analysed by performing model reduction directly on the Polyaffine matrices for all subjects within a population, to determine group-wise temporal and spatial patterns. A patient-specific reduced model is also described, that allows for further reduction of the parameters needed to describe the motion. We were able to obtain clinically meaningful temporal bases for the population. The first section of this chapter was published in the proceedings of the FIMH conference of 2013 ([McLeod et al., 2013b]), and the second section was published in the proceedings of the MICCAI conference of 2013 ([McLeod et al., 2013a]).

Chapter 7 describes the extension of the statistical analysis of transformation parameters in Chapter 6, to propose a population-based motion model by transporting all the motion parameters computed using the method described in Chapter 5, to a common space using spatio-temporal alignment. Statistical analyses are performed on the aligned parameters to firstly create a mean motion model of a given population. The typical motion patterns observed in a given population are analysed by performing statistical model reduction on the stacked matrices of parameters of all subjects (using SVD and PCA), then on the full tensor of all parameters for all subjects (using Tucker tensor decomposition). The methods were applied to create mean motion models for healthy subjects and for Tetralogy of Fallot patients, to compare the different motion dynamics observed between the two populations. The computed mean motion model gives realistic motion, and leads towards population-based comparison of motion. This chapter will be submitted to Transactions in Medical Imaging.

Part - III

Statistical Modelling of Cardiac Blood Flow

Chapter 8 describes a method for performing reduced-order blood flow simulations. Patient-specific blood flow solutions are transported to an atlas geometry to obtain a mean blood flow solution for the population. Model-order reduction is performed on the mean blood flow using Proper Orthogonal Decomposition, to obtain reduced-order bases of pressure and velocity. These bases can be transported to a new subject by a non-rigid deformation computed using the LDDMM method. These bases can then be re-solved for the subject-specific geometry by solving the Navier Stokes equations on the reduced bases. In this chapter, we describe the application of this method for performing reduced-order blood flow simulations on new patients from the learned group-wise blood flow solutions. The method was applied to blood flow simulation in the pulmonary artery for Tetralogy of Fallot patients. We were able to obtain similar flow patterns from the reduced model as the full flow simulation indicating that such a model can produce reasonable flow simulations with a reduced number of parameters. This chapter was accepted to Medical Image Analysis ([McLeod et al., 2013]) as an extension of the work published in the STACOM workshop proceedings of 2010 ([McLeod et al., 2010a]).

Conclusions

Chapter 9 introduces an approach for constructing a coupled model of growth and motion, to lead towards the global objective of this work for developing a 5D+flow model. We propose to use cross-sectional statistical design applied to transformations that describe cardiac motion, to regress a data-set of transformations from different subjects. We suggest using a similar method as the one proposed in Chapter 3 to regress the transformation parameters (rather than velocities).

Chapter 10 concludes this thesis by summarising the key contributions of this work, discussing the collaborative aspect of the PhD, and suggestions for future work to extend on the methods described in this thesis.

1.4 List of Publications and Awards

Journal articles and Book Chapters

- [McLeod et al., 2013a] **K. McLeod**, T. Mansi, M. Sermesant, G. Pongiglione, X. Pennec. Statistical Shape Analysis of Surfaces in Medical Images Applied to the Tetralogy of Fallot Heart. *In: Modeling in Computational Biology and Biomedicine, Lectures Notes in Mathematical and Computational Biology*. Chapter 3.
- [Tobon-Gomez et al., 2013] C. Tobon-Gomez, M. De-Craene, **K. McLeod**, T. Lautz, W. Shi, A. Hennemuth, A. Prakosa, H. Wang, G. Carr-White, S. Kapetanakis, A. Lutz, V. Rasche, T. Schaeffter, C. Butakoff, O. Friman, T. Mansi, M. Sermesant, X. Zhuang, S. Ourselin, H-O. Peitgen, X. Pennec, R. Razavi, D. Reuckert, A.F. Frangi, K.S. Rhode Benchmarking Framework for Myocardial Tracking and Deformation Algorithms: An Open Access Database. *Published in Medical Image Analysis, 2013*. This article is in relation to the work described in Chapter 4.
- [McLeod et al., 2013b] **K. McLeod**, C. Seiler, L. Le Folgoc, N. Toussaint, M. Sermesant, X. Pennec. Physiologically Constrained Regional Motion Model for Cardiac Strain Analysis. *In preparation for submission to a journal*. Chapter 5.
- [McLeod et al., 2013] **K. McLeod**, R. Guibert, A. Caizzo, T. Mansi, M.A. Fernández, M. Sermesant, X. Pennec, I.E. Vignon-Clementel, Y. Boudjemline, J-F. Gerbeau. Group-wise Construction of Reduced Models for Understanding and Characterization of Pulmonary Blood Flows from Medical Images. *Submitted to Medical Image Analysis*. Chapter 8.

Selective Peer-Reviewed Conference Papers

- [McLeod et al., 2013a] **K. McLeod**, C. Seiler, M. Sermesant, X. Pennec. Spatio-Temporal Dimension Reduction of Cardiac Motion for Group-Wise Analysis and Statistical Testing *Proceedings in Medical Image Computing and Computer Assisted Intervention - MICCAI, 2013*. Partly in Chapter 6.

- [McLeod et al., 2013b] **K. McLeod**, C. Seiler, N. Toussaint, M. Sermesant, X. Pennec. Regional Analysis of Left Ventricle Function using a Cardiac-Specific Polyaffine Motion Model *Proceedings in Functional Imaging and Modeling of the Heart - FIMH*, 2013. Oral podium presentation. Partly in Chapter 6.

Workshop Papers

- [McLeod et al., 2010a] **K. McLeod**, A. Caiazzo, M.A. Fernandez, T.Mansi, I.E. Vignon-Clementel, M. Sermesant, X. Pennec, Y. Boudjemline, J.F. Gerbeau Atlas-Based Reduced Models of Blood Flows for Fast Patient-Specific Simulations. *STACOM Workshop - MICCAI*, 2010. Oral podium presentation. Preliminary version of the method described in Chapter 8.
- [McLeod et al., 2012a] **K. McLeod**, A. Prakosa, T. Mansi, M. Sermesant, X. Pennec An Incompressible Log-Domain Demons Algorithm for Tracking Heart Tissue. *STACOM Workshop - MICCAI*, 2011. Oral podium presentation. Partly in Chapter 4.
- [Prakosa et al., 2012a] A. Prakosa, **K. McLeod**, M. Sermesant, X. Pennec Evaluation of iLogDemons Algorithm for Cardiac Motion Tracking in Synthetic Ultrasound Sequence. *STACOM Workshop - MICCAI*, 2012. Partly in Chapter 4.
- [McLeod et al., 2012b] **K. McLeod**, C. Seiler, M. Sermesant, X. Pennec A Near-Incompressible Poly-Affine Motion Model for Cardiac Function Analysis. *STACOM Workshop - MICCAI*, 2012. Oral podium presentation. Preliminary version of the method described in Chapter 5.

Conference Abstracts

- [McLeod et al., 2010b] **K. McLeod**, T. Mansi, S. Durrleman, M. Sermesant, X. Pennec. Statistical Analysis of the Anatomy: From Digital Patient to Digital Population. *Virtual Physiological Human (VPH) Conference*, Brussels, 2010.

Awards

- The paper "An Incompressible Log-Domain Demons Algorithm for Tracking Heart Tissue" was awarded the Best Paper Award for the Cardiac Motion Tracking Challenge (2011).

The Role of Imaging and Analysis in Cardiology

Contents

2.1 The Human Heart	13
2.1.1 Tetralogy of Fallot	15
2.2 Medical Imaging in Cardiology	17
2.2.1 Cardiac MRI	17
2.3 Inter and Intra-Patient Analysis using Image and Mesh Registration . .	23
2.3.1 Transformation Types	25
2.3.2 Similarity Measures	26
2.3.3 Optimisation Process	26
2.4 Conclusion	27

The purpose of this chapter is to briefly introduce concepts needed to understand the following chapters of this thesis. Since the main application of this work is modelling the Tetralogy of Fallot heart, a brief description of the heart structures is first provided (Sec. 2.1) and of the abnormalities observed in ToF patients (Sec. 2.3). A brief introduction to medical imaging techniques used in cardiology is provided, with further details on cardiac MRI (Sec. 2.2), since this is the main modality studied in this PhD. The key methods proposed in this work use mesh and image registration, thus a brief overview and state-of-the-art of registration methods is provided in Sec. 2.3, in order to position the proposed method with respect to other registration methods.

2.1 The Human Heart

The heart is made up of four main structures consisting of two *ventricles* (left and right) and two *atria* (the left and right atria). The left and right sides of the heart are separated by a thin wall made up of biological tissue called the *septum*. The section of the septum which divides the left and right atrium is called the *inter-atrial septum*, and the section which divides the ventricles is called the *inter-ventricular septum*. The atria and ventricles for each side are divided by valves which control the flow of blood from the atria to the ventricles. The atria collect blood, and once full this blood is pumped into the respective ventricles. The right ventricle is fed by de-oxygenated blood from the body through the *inferior* and *superior vena cava*, which is then pumped to the lungs to gather oxygen to

feed back to the left ventricle through the *pulmonary veins*. A key component of heart function is the valves which ensure that blood flows in only one direction from one section to another (from the right atrium to the right ventricle, for example). The septum ensures that de-oxygenated blood in the right ventricle does not mix with the oxygenated blood in the left ventricle (see Figure 2.1).




Figure 2.1: An illustration of a healthy human heart. Image obtained from <http://www.nhlbi.nih.gov>

The flow of blood through the chambers and arteries of the heart is a complex process. When the pressure in the right (resp. left) atrium is higher than in the right (resp. left) ventricle, the atrium contracts and the *tricuspid* (resp. *mitral*) valve dividing the right (resp. left) atrium and the right (resp. left) ventricle open to allow blood to flow to the ventricle. When the pressure in the ventricle increases, the tricuspid (resp. mitral) valve closes to restrict blood from flowing back into the atrium. Similarly, when there is more pressure in the right (resp. left) ventricle than in the *pulmonary artery* (resp. *aorta*), the *pulmonary valve* (resp. *aortic valve*) opens to allow blood to flow through the artery to the lungs (resp. body). The moment when ventricles pump blood (contract) is known as systole, and when the ventricles relax (while blood is being pumped into the ventricles from the atria) is known as diastole. For a normal heart, the two sides are synchronised so that each side contracts and relaxes in unison.

This activity is controlled by electrical waves that depolarise cardiac cells which causes the cells to contract. Each heart beat starts at the *sinoatrial nodes* in the atria with an electrical wave that then spreads over the atria causing the atrial myocytes to contract, which results in blood being pumped to the ventricles. This is marked by the P wave on an

electrocardiogram (ECG). When the signal arrives at the *atrioventricular* nodes it slows to allow the ventricles to fill with blood. The signal then runs down each side of the septum along the right and left bundle branches (marked by the Q wave on an ECG), and on to the cells in the walls of the heart through the *Purkinje fibers*. The signal then spreads across the cells causing the ventricles to contract. This contraction pushes blood from the ventricles to the respective arteries. Following this, the signal passes the ventricles so they begin to relax. The left ventricle contracts slightly before the right ventricle. On an ECG, the R wave marks the contraction of the left ventricle, and the S wave marks the contraction of the right ventricle (see Fig. 2.2).




Figure 2.2: Left: The key areas of the electrical conduction system of the heart (image from <http://www.nlm.nih.gov>). Right: The activation points at different nodes (image from <http://www.bem.fi/book/06/fi/0607.gif>).

Good heart function depends on both the electrical system and the pumping action. Abnormal electrical function can have severe effects on heart function, such as causing the heart to beat too fast or too slow, so that the body is not fed enough oxygen or so that the blood is not pumped to the body fast enough.

2.1.1 Tetralogy of Fallot

In this work the clinical application we will focus on is Tetralogy of Fallot (ToF). ToF is a severe congenital heart defect that affects approximately four out of every 10,000 babies [Hoffman, 1995]. This condition, which is generally diagnosed prenatally, has four associated defects: a ventricular septal defect, pulmonary stenosis, right ventricular hypertrophy and an overriding aorta.

The *ventricular septal defect* is an opening in the ventricular septum (that separates the two ventricles), which thus allows de-oxygenated blood in the right ventricle to mix freely with oxygenated blood in the left ventricle. This results in blood volume overload in the left ventricle from the blood re-circuiting to the right ventricle and on to the lungs, and poorly oxygenated blood is pumped to the body.




Figure 2.3: Comparison of a normal heart with a heart with Tetralogy of Fallot. Image obtained from <http://www.nhlbi.nih.gov>

As the name indicates, *pulmonary stenosis* is used to describe the obstruction at the right ventricular outflow tract. This consequently causes the right ventricle to pump harder (causing stress to the right ventricle) and restricts blood flow to the lungs. The long term effect of this condition is extreme right ventricular dilation due to severe regurgitation (*right ventricular hypertrophy*). Ventricular hypertrophy can result in irregular heart rhythm and pressure change in the ventricle.

An *overriding aorta* is also present in individuals with ToF, which means that the aorta is misplaced and sits over the ventricular septal defect, rather than over the left ventricle. Consequently, the aorta can receive de-oxygenated blood directly from the right ventricle.

Open heart surgical repair is required within the first year of life to repair the ventricular septal defect, as well as to remove the obstruction to the right ventricular outflow tract, and repairing the other defects depending on the severity of the condition. During this surgery, the pulmonary valves can be removed or damaged, causing long-term pulmonary regurgitation. Therefore, patients may require repeat follow up operations to reduce post-operative sequelae. In this case, the valves need to be replaced with an artificial device to ensure one way flow of blood from the right ventricle to the lungs. The timing of follow-up operations is crucial to ensure that the treatment is instigated before the heart becomes too deformed for intervention, while also maximising the time between interventions because of the limited life span of the replacement valves. The clinical challenge for this condition is to understand how the heart evolves, with the main goal to be able to predict the optimal time for intervention. However, due to the large variability of shape, the optimal time, placement and size of the artificial device can be difficult to determine.

Therefore, in view of assisting cardiologists, understanding the shape of the ventricles is crucial for treatment planning. Previous studies have analysed the shape of the

heart in order to determine the optimal intervention time for pulmonary valve replacement [Geva, 2006]. Recent work has been done to simulate the growth of the right ventricle under these pathological conditions [Mansi et al., 2011b]. In [Schievano et al., 2007], the shape of the pulmonary artery of 12 ToF patients was determined using MR imaging data, and prototype models were constructed using a 3D printer to establish the suitability of the patients for non-invasive valve replacement intervention. However, creating prototype models in this way is time consuming, and expensive if models are created for individual patients. Rather, we aim to use statistical models of the heart by extending on previous work to give a better understanding of the deformation of the shape of the ventricles, under different states and for different pathologies.

2.2 Medical Imaging in Cardiology

Several different techniques for acquiring medical images have been designed to visualise different structures. In cardiology these include structural and functional magnetic resonance imaging (MRI), echocardiography (ECHO) and computed tomography (CT). ECHO is the most widely used imaging technique given that it is quick and easy to use, inexpensive, and non-invasive. ECHO is generally recommended as the first imaging modality used for the purpose of diagnosis and for when there is a change in the patient's condition. In general clinical practice, only 2D sequences are acquired, though 3D imaging is also now possible. The disadvantage of ECHO is that the images are usually subject to high noise levels and can be obtained only from certain angles. Tissue contrast is also generally low, making it difficult to distinguish the boundaries of different structures. Cardiac MRI and CT imaging are often used in cases where more accurate images are needed to better distinguish different structures and for therapy planning to have a better view of the heart. Using CT, high resolution images with good image contrast can be captured and can be used to visual the motion dynamics through the cardiac cycle. However, ionising radiation is emitted during the scan that can be harmful to the patient. Cardiac MRI, on the other hand, also has the advantage of capturing high resolution images, but without the harmful radiation. Therefore, cardiac MR imaging is increasingly being used in clinical practice. For this reason, the focus of this work is to perform image analysis mainly on cardiac MR images, though the techniques themselves are not fixed to this particular modality.

Since this work is focused on image analysis tools applied to cardiac MR images, a brief description of this modality is provided in the following section. For further details on the image acquisition and machines for cardiac MR as well as ECHO and CT, see [Higgins and Albert de Roos, 2006], [Bogaert et al., 2012], [Leeson et al., 2012], [Dewey, 2011].

2.2.1 Cardiac MRI

MRI is becoming increasingly used in clinical practice to image the motion of the heart as it beats. Cardiac MRI is in principal the same as standard MRI, where radio waves, magnets and a computer are used to construct images of different structures and tissue. MRI is commonly used to image the heart due to the good contrast obtained between soft tissue,

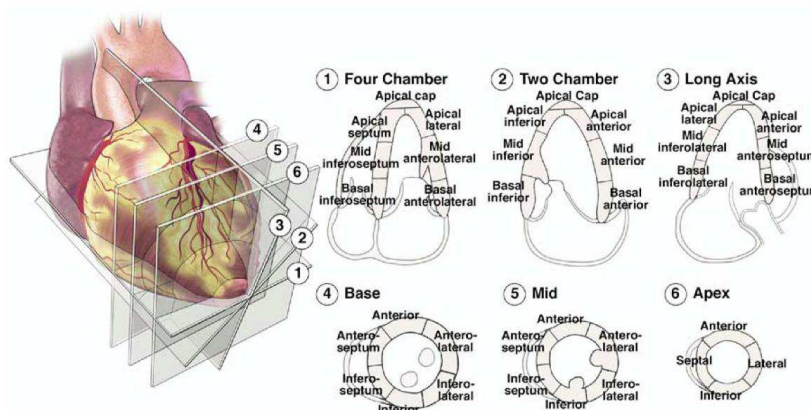


Figure 2.4: The different standard imaging angles of the heart including the four chamber view, two chamber view and long axis view, shown with the American Heart Association subregions of the heart. Image obtained from [Lang et al., 2005]

blood and bone (see Fig. 2.6), and can be acquired from any angle (see Fig. 2.4). Cardiac MRI uses ECG signals to detect the R peak in the sequence, from which 4D sequences can be built (see Fig. fig:ECG).

Cardiac MRI can be used to diagnose and assess a number of cardiovascular conditions. For example, damage caused by heart failure or a heart attack can be assessed through cardiac MR imaging. Cardiac MRI can provide complementary information to other imaging such as ECHO, X-ray and CT. In particular, cardiac MRI imaging can be performed to:

- Evaluate the anatomy and function of the heart chambers, valves, major vessels, and surrounding structures such as the pericardium (the fluid filled sac that surrounds the heart).
- Diagnose a variety of cardiovascular (heart and/or blood vessel) disorders such as tumors, infections, and inflammatory conditions.
- Detect and evaluate the effects of coronary artery disease such as limited blood flow to the heart muscle and scarring within the heart muscle after heart attack.
- Plan a patient's treatment for cardiovascular disorders.
- Monitor a patient's progression over time.
- Evaluate the anatomy of the heart and blood vessels in children with congenital heart disease.
- Examine the size of the heart chambers and the thickness of the heart wall.
- Determine the extent of myocardial (heart muscle) damage and effect on pump function caused by a heart attack or progressive heart disease.
- Detect the buildup of plaque and blockages in blood vessels.
- Assess a patient's recovery following treatment.
- Evaluate the heart anatomy, pump function, heart valve function and vascular blood flow both before and after surgical repair of congenital cardiovascular disease in children and adults [Knauth et al., 2008].

(List courtesy of <http://www.radiologyinfo.org/>).

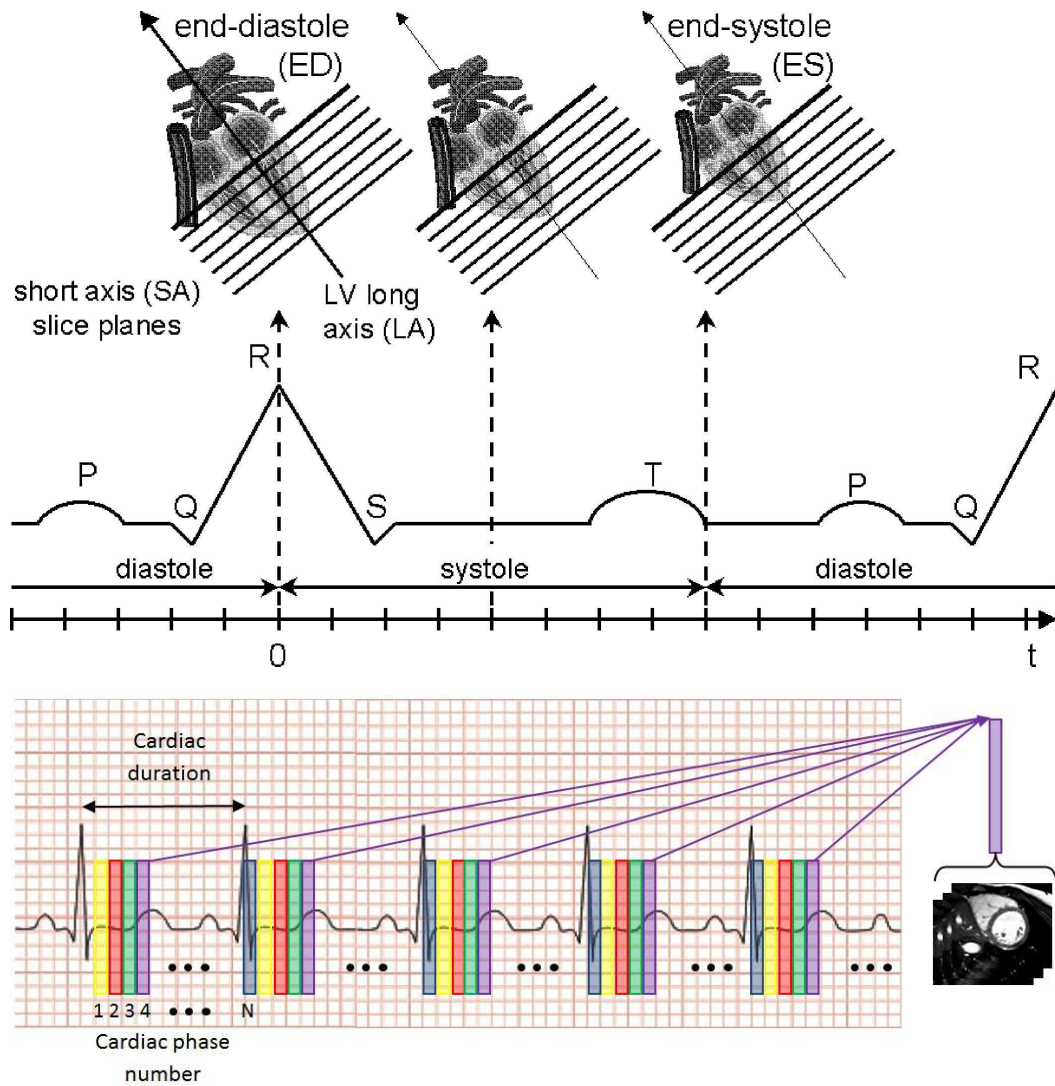


Figure 2.5: Top: Illustration of ECG gating of cardiac MR images for detecting the R peak of the cardiac cycle. Image obtained from [Makela et al., 2002]. Bottom: Cine MR sequences can be constructed by dividing the cycle into N frames and averaging over images acquired for several heart beats at the same frame (where the frame is detected from the ECG signal).




Figure 2.6: Cardiac MRI. Localisation of the heart in the body is shown in (A), with a given MR slice angle. An example of a 3D MR angiogram is shown in (B). (C) shows the left ventricle image acquired from the slice shown in (A) when healthy (left), damaged from heart attack (centre) and when the left ventricle is not receiving enough blood (right). Image obtained from www.daviddarling.info/encyclopedia/C/cardiacMRI.html

2.2.1.1 Benefits vs. Risks

Some of the benefits and risks of using cardiac MR imaging instead of (or as well as) other imaging modalities are:

Benefits:

- + Non exposure to ionising radiation such as those emitted in CT and X-ray imaging.
- + The ability to diagnose a broad range of conditions related to anatomical abnormalities (such as congenital heart defects), functional abnormalities (such as valve failure), tumors and conditions related to coronary artery disease and cardiomyopathy.
- + MR imaging can be used during certain interventional procedures to give interventional cardiologists and surgeons a better view of the structures to both save time and improve the quality of care. For this exact purpose, the use of hybrid operating theaters is becoming increasingly popular.
- + Since MR images can be acquired from any arbitrary angle, structures that would be otherwise hidden behind bone or other material can be visualised.
- + Allergic reaction due to the use of contrast agent is less common than the reaction due to iodine-based agents used in CT and X-ray imaging.
- + Both structure and function can be assessed without the need for ionising radiation.

Risks:

- Some patients need to be sedated for cardiac MR such as children or patients suffering from claustrophobia.
- Due to the high magnetic field within the MR machine, there may be some interaction with metal devices inside the body.
- Contrast agents can induce an allergic reaction in some patients. Some patients with very poor kidney function can also be at risk of Nephrogenic systemic fibrosis, though this is rare.
- The contrast agent may be dangerous for breast-feeding mothers through transfer to infants via breast milk.

2.2.1.2 Different Uses of Cardiac MRI

There is a wide range of uses of this type of imaging, a few of which are briefly described here:

Cardiac Anatomy: Cardiac MRI can be used for defining *cardiac anatomy* using Steady-State Free Precision (SSFP) sequences [Scheffler and Lehnhardt, 2003] in heart conditions such as constrictive pericarditis [Francone et al., 2006], cardiac neoplasm or thrombus [O'Donnell et al., 2009], congenital heart disease [Sørensen et al., 2004], or for demonstrating the presence of a patent foramen ovale [Nusser et al., 2006] (see Fig. 2.6). SSFP is commonly used in clinical practice to acquire detailed images of the structures through the heart beat.

Myocardial Perfusion: Cardiac MRI can also be used to measure *myocardial perfusion* with cardiac MR perfusion imaging for patients suspected to have ischemic heart disease (e.g. angina) to examine potential perfusion defects in the myocardium [Schwitter et al., 2001]. By inducing stress on the heart with either adenosine or inotrope Dobutamine [Paetsch et al., 2004] the patient can be imaged under the stress condition, and again at rest, to visually compare the differences. Myocardial perfusion MR is becoming more commonly used in clinical practice due to the fact that it is well tolerated for most patients and gives good indications for examining stenosis or detecting ischemia.




Figure 2.7: Left: Cardiac MR perfusion imaging with stress testing on a patient with subtotal occlusion of the left anterior descending artery. White arrows indicate stress perfusion defects [Motwani et al., 2012]. Centre: Left ventricle muscle fibers extracted using image analysis on DTMRI images. Right: Tagged MRI. Motion of the heart can be tracked from MR image sequences through tagged MR imaging. <http://www.nature.com/nrcardio/journal/v6/n12/thumbs/nrcardio.2009.189-f1.jpg>

Cardiac Fiber Architecture: Using cardiac diffusion tensor MR imaging (DTMRI), the *cardiac fibers* can be extracted to better understand how the fibers are arranged [Helm et al., 2005]. Since electrical currents propagate along the fiber directions, this can give useful insight for patients suffering from conditions related to abnormal electrophysiology such as arrhythmias (see Fig. 2.7 centre). DTMR images, however, are generally only acquired for research purposes, therefore little data is currently available.

Cardiac Motion: By analysing the *4D dynamics* of the heart through the cardiac cycle, cardiac MRI can be used to quantify left and right ventricular function for patients suffering from cardiomyopathy [Hunold et al., 2005], heart failure [Bellenger et al., 2000], arrhythmogenic right ventricular dysplasia [van der Wall et al., 2000] and pulmonary hypertension [Roeleveld et al., 2005]. Tagged MRI can be used to track the motion of the heart through the cardiac cycle [Park et al., 1996, Clarysse et al., 2002, Clarysse et al., 2011] (see Fig. 2.7 right), however it is not currently a part of standard clinical routine as it increases the imaging acquisition time (and therefore also the cost of the acquisition). Therefore, we aim to analyse the motion from SSFP sequences alone since these are routinely acquired.

Cardiac Blood Flow: Using cardiac MR imaging, it is possible to extract quantitative 1D measures of *blood flow* [McVeigh, 1996] that can be used to diagnose and assess valvular disease (such as regurgitation or stenosis) [Glockner et al., 2003], or to analyse shunts (such as atrial or ventricular septal defects) [Hundley et al., 1995]. (see Fig. 2.8). The measures are often acquired as part of the imaging set, however such measures are limited by the fact that the measures are taken only for one slice of the image. This limitation motivates the use of the method described in Chapter 8.




Figure 2.8: Quantitative 1D measures of blood flow from cardiac MRI.

2.3 Inter and Intra-Patient Analysis using Image and Mesh Registration

As described in the previous section, imaging alone can provide qualitative and quantitative measures of structure and function. However, current imaging techniques are somewhat limited in the ability to provide quantitative measures of motion, 3D blood flow, changes over time etc. For this reason, image analysis can enable deeper analysis of the observed phenomena. Structures can, for example, be measured with respect to size, volume, muscle thickness, and ventricular volume using image segmentation. Blood flow through the ventricles, atria and arteries can be simulated using computation fluid dynamics simulation to estimate the flow dynamics under different conditions and the impact of stress, irregular flow patterns, and strain. By quantitatively comparing two images, differences can be highlighted between the two time points to establish the evolution of an organ over time, or the growth of a tumour, for example.

A common method for quantifying differences between two images (either from the same, or different patients), is to use image (or mesh) registration to estimate the deformation from one image (or mesh) to another [Hajnal and Hill, 2010, Goshtasby, 2012]. In practice, registration amounts to approximating the transformation that optimises the similarity between the fixed image (or mesh) and the transformed moving image (or mesh) (see Fig. 2.9). Several methods to do this in the context of medical imaging have been proposed and successfully applied to a number of different applications to model changes in organs (such as the brain [Klein et al., 2009, Rueckert et al., 2003, Crum et al., 2013], lungs [Gorbunova et al., 2012, Risser et al., 2012, Ehrhardt et al., 2011], heart

[Chandrashekar et al., 2004b, Perperidis et al., 2005, Makela et al., 2002], liver [Velec et al., 2012], breast [Denton et al., 1999, Juh and Suh, 2010]), or fiber structures (such as cardiac muscle fibers [Lombaert et al., 2012, Toussaint et al., 2013], fiber bundles in the brain [O'Donnell et al., 2012]), or vein structures [Osorio et al., 2012] (in the cardiovascular or pulmonary systems for example). Pair-wise image (or mesh) registration is essentially the process of bringing the images (meshes) to a common coordinate system. In doing so, the image (mesh) differences are parameterised by the transformations needed to bring the images (meshes) to the same coordinate space. For example, if a square was transformed to a diamond with sides of equal length, the corresponding transformation would represent a rotation, which would indicate that the difference between the two objects is simply a rotation in space. This illustration is an example of a rigid transformation. In the case of cardiac motion tracking, in general we are rather interested in non-rigid registration that can capture more complex differences such as local stretching and scaling.

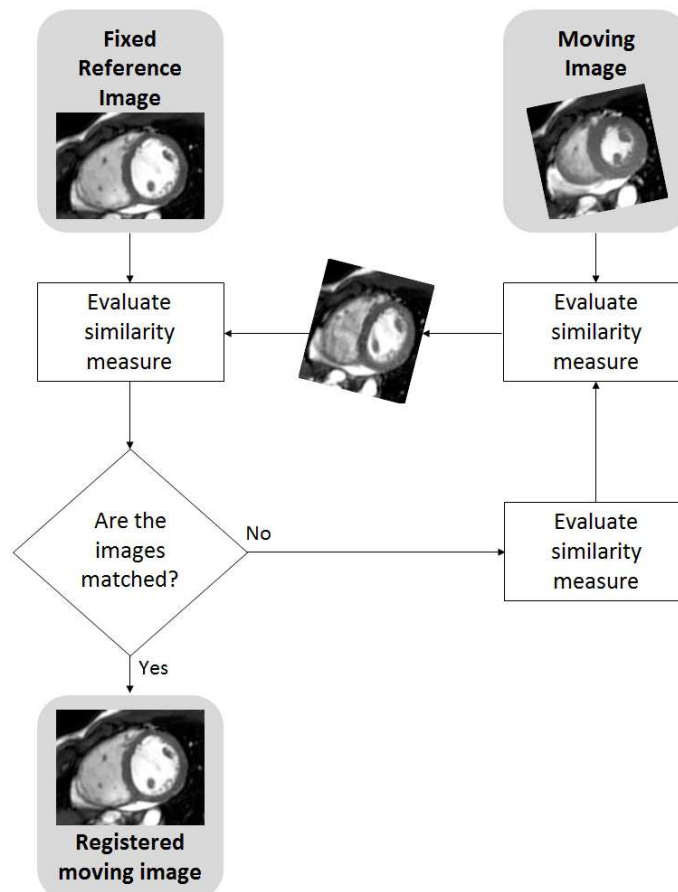


Figure 2.9: The process of registering one image (the moving, or floating image) to another image (the fixed, or reference image). An iterative *optimisation process* can be used to estimate the *transformation* parameters, given a chosen *similarity measure*. Image modified from [Hutton and Braun, 2003].

In the following sections we provide a brief overview of the key components of registration algorithms; namely the types of transformations that we wish to capture with a registration algorithm, ways to measure the similarity between two images, and optimisation methods for finding the best transformations. This is not meant as a full review of registration methods, but rather as an overview of different methods in order to position the methods proposed in this thesis with respect to other types of registration algorithms.

2.3.1 Transformation Types

Possible transformation types for medical registration problems include:

- Rigid
- Affine
- Piecewise affine or Polyaffine
- Non-rigid

Rigid (or linear) transformations include rotations, translations and reflections. Rigid registration assumes that there is no distortion between the images. A rigid transformation model is composed of three rotations and three translations (6 degrees of freedom), and can be represented by a 4×4 homogeneous matrix. Affine transformations include rigid transformations, while also accounting for stretch and shear (12 degrees of freedom) and can also be represented by a 4×4 homogeneous matrix. Piecewise affine transformations [Pitiot et al., 2003] go a bit further to define an affine transformation at a regional level, therefore allowing different regions in an image to transform in different ways. Polyaffine transformations are essentially an extension of piecewise affine transformations to include a smooth transition between regions so that globally the transformation is smooth, rather than piecewise continuous. Polyaffine transformations [Arsigny et al., 2009] include a weight function that controls the fusion between the piecewise affine transformations, therefore the transformation is represented not only by the affine matrices, but also the weight functions. Non-rigid transformations are non-linear and thus cannot be represented by matrices. The transformation can be parameterised in a number of different ways through fluid models [D'Agostino et al., 2003] or spline-based models [Rueckert et al., 2006] for example. The choice of transformation type depends on the level of detail we want to capture. A method for measuring the similarity between the two objects is then required. A brief overview of such methods is provided in the next section.




Figure 2.10: Example of the transformation types needed to get from the source image (left) to the target image (right). Image adapted from [Crum et al., 2004].

2.3.2 Similarity Measures

A similarity measure is needed to determine how similar two images (or meshes) are (or equivalently how much they differ) after alignment. This can be through a measure of the similarity of intensities of two images, or between geometric differences in two images (or meshes).

Geometric approaches, that compute the correspondence between geometric objects, are rather sparsely distributed, therefore the match needs to be interpolated between the data points. However, geometric approaches have the advantage of incorporating anatomical information to the matching. Landmarks can be used to drive the registration algorithms provided that landmarks can be defined consistently in both images (or meshes) [Johnson and Christensen, 2002]. Intensity-based similarity measures imply a cost function that minimises the differences in intensity between two images. Intensity-based approaches generally match intensities over the whole image. The choice of similarity measure when measuring intensity differences is largely dependent on the image modality and whether the registration is mono- or multi-modal.

Commonly used similarity measures are the sum of squared differences (SSD), correlation coefficients (CC), mutual information (MI) and normalised mutual information (NMI). The SSD measure, as the name implies, minimises the total sum of squared differences between image intensities (or between points) and assumes that images differ only up to Gaussian noise. Since the SSD measure is sensitive to intensity differences, histogram normalisation can improve the accuracy of the measure, where large outliers are excluded. CC is an area-based method that searches the intensity correspondences within a given window. CC assumes an affine relationship between image intensities, hence is only suitable for mono-modality registration. MI, a method originating in information theory, is a statistical approach that essentially measures the statistical dependency between two data-sets by using the joint intensity probability distribution. MI, in contrast to SSD and CC, assumes a probabilistic relationship between intensities, and therefore can be used for multi-modality registration. However, MI relies on the computation of the joint histogram, which can be computationally expensive (and requires complex optimisation schemes).

2.3.3 Optimisation Process

Once the type of transformations we are wanting to capture is defined, as well as the measure of similarity that determines the goodness of fit, an optimisation process is required to find the best possible transformation. The key role of the optimisation process is to approximate a set of parameters to maximise (or minimise) a given function. Common methods to do this generally use numerical methods such as gradient descent or conjugate gradient descent for example. Steepest gradient descent approaches can be used to refine the parameters without searching over all the possible parameters by using the past search history. The refinement process continues until the similarity criterion is reached. Regularisation terms may also be included to constrain the search space to obtain smoother or more physiologically plausible deformations.

2.4 Conclusion

This chapter is meant purely as a basic overview of the heart structures (highlighting the differences between healthy and Tetralogy of Fallot hearts), cardiac imaging methods, and a brief introduction to the methods used in the following chapters. This work is focused on image analysis of cardiac MR given the broad use of this modality in clinical practice. A more in depth state-of-the-art in registration methods for specific applications is given in each chapter. Overall, in the proposed work we are dealing with mono-modal registration (MRI-to-MRI), therefore the SSD similarity measure is preferred due to its simplicity and since we consider the assumption of linearity between image intensities (after intensity normalisation) as being valid for the applications described in this thesis. In Chapters 3 and 8, the preferred transformation method is the Large Deformation Diffeomorphic Metric Mapping (LDDMM) algorithm on surfaces (with currents representing the surfaces) since we require a method that is able to capture large deformations (as for inter-subject registration). In Chapters 4-7, we rather prefer stationary velocity fields (SVF's) parameterised either locally (using a Demons based method) or regionally (using a Polyaffine based method) to perform non-rigid registration of images for intra-patient registration to define the transformation voxel-wise over the whole image (Demons) characterised by a few parameters (Polyaffine).

Part I

**STATISTICAL MODELLING OF
CARDIAC GROWTH**

Bi-ventricular Growth Modeling in the Tetralogy of Fallot Population

Contents

3.1 Chapter Overview	31
3.2 Introduction	32
3.2.1 Repaired Tetralogy of Fallot	33
3.2.2 Chapter Organisation	34
3.3 Statistical Shape Analysis	34
3.3.1 Shapes, Forms and Deformations	34
3.3.2 From Points to Surfaces: the Formalism of Currents	37
3.3.3 An Algorithm for Surface Registration using Currents	39
3.3.4 Building an Unbiased Atlas	40
3.4 Application on ToF Data	41
3.4.1 The Analysis Pipeline	41
3.4.2 Diagnosis Parameters	42
3.4.3 Building an Evolution Model	45
3.5 Conclusion	50

This chapter was published as a chapter titled "Statistical Shape Analysis of Surfaces in Medical Images Applied to the Tetralogy of Fallot Heart", in *Modeling in Computational Biology and Biomedicine, Lectures Notes in Mathematical and Computational Biology* [McLeod et al., 2013a]. This chapter addresses the first key question of this work on how to model bi-ventricular growth over time for a chosen population with few available longitudinal data-sets, that can be generalised to a given population. A bi-ventricular growth model was developed to simulate the growth of the left and right ventricles in a Tetralogy of Fallot population to obtain clinically relevant modes of variation related to the observed dynamics for this pathology.

3.1 Chapter Overview

There is an increasing need for shape statistics in medical imaging to provide quantitative measures to aid in diagnosis, prognosis and therapy planning. In view of this, we describe methods for computing such statistics by utilising a well-posed framework for representing

32 Chapter 3. Bi-ventricular Growth Modeling in the Tetralogy of Fallot Population

the shape of surfaces as currents. Given this representation we can compute an atlas as a mean representation of the population and the main modes of variation around this mean. The modes are computed using principal component analysis (PCA) and applying standard correlation analysis to these allows to correlate shape features with clinical indices. Beyond this, we can compute a generative model of growth using partial least squares regression (PLS) and canonical correlation analysis (CCA). In this chapter, we investigate a clinical application of these statistical techniques on the shape of the heart for patients with repaired Tetralogy of Fallot (rToF), a severe congenital heart defect that requires surgical repair early in infancy. We relate the shape to the severity of the pathology and we build a bi-ventricular growth model of the rToF heart from cross-sectional data which gives insights about the evolution of the disease.

3.2 Introduction

During the past ten years, the biophysical modeling of the human body has been a topic of increasing interest in the field of biomedical image analysis. New treatments, therapy plans and surgical techniques are continually being developed and enhanced to improve the outcome for patients, The aim of such modeling is to formulate personalised medicine, where a digital model of an organ can be adjusted to a patient using personal clinical data as input to the given model. This virtual organ would enable to estimate parameters which are difficult to quantify in clinical routine and to test therapies in-silico as well as to predict the evolution of the organ over time and with therapy. These methods can be interesting for example in predicting the long-term outcome of a newly developed treatment in which longitudinal patient data is not yet available, or for testing the outcome of a number of different treatment methods virtually to predict the best plan for a given patient. This notion is largely driven by the fact that testing individual treatments to determine the best choice for a given patient is neither ethical, nor in fact possible and can bridge the gap between clinical population testing by combining this with personalised patient-specific models.

When the biological mechanisms involved are too complex, robust statistical approaches could be used to produce generative models from data. Such statistical analyses can both provide a predictive model and guide the biophysical approach. However, computing statistics on dynamic 3D shapes is very challenging. Traditional methods rely on point based parametrisations of the shapes, where the point-to-point correspondences can be an important limiting factor for the usability of the method since the correspondences must be correct and consistent over all the shapes. New approaches were recently developed to compute such statistics without this limitation [Durrleman et al., 2009b, Durrleman et al., 2008b, Durrleman et al., 2009a, Durrleman, 2010, Mansi et al., 2009, Mansi, 2010, Mansi et al., 2011b, Hufnagel, 2010]. Here, the statistical shape analysis tools are based on currents, a non-parametric representation of shapes (here, the term ‘shape’ is used to refer to geometric data such as curves and surfaces). These tools have a wide range of applications, and provide a well-posed framework for statistical shape analysis of groups. Due to the fact that the methods do

not assume point correspondences between structures (and in fact assigning landmarks to structures such as the heart are arbitrary), a wider range of data can be used. For example, one can use surfaces to model organs such as the heart, brain, and lungs, curves to model sulcal lines on the brain cortex, and sets of curves to represent fiber bundles from diffusion MRI in the brain. The goal of this chapter is to give an overview of this methodology in the context of a specific clinical problem: the prediction of the cardiac shape remodeling in repaired tetralogy of Fallot due to chronic regurgitation.

3.2.1 Repaired Tetralogy of Fallot

Tetralogy of Fallot (ToF) is a congenital heart defect that affects approximately four out of every 10,000 babies [Hoffman and Kaplan, 2002]. The primary defect associated with this condition is a ventricular-septal opening which allows blood to flow freely between the ventricles. Secondary is stenosis in the pulmonary artery which restricts the blood flow from the right ventricle to the lungs. Due to a misalignment of the aorta over the ventricular-septal defect, the aorta is fed by both the left and right ventricle rather than just the left ventricle. Patients may also have hypertrophy in the right ventricle that causes a boot-like shape of the ventricle which is characteristic of this condition.

These abnormalities require open heart surgical repair early in infancy. As part of this surgery, the stenosis in the pulmonary artery is cleared to allow blood to flow more freely through the artery. As a consequence, the pulmonary valves which control one way blood flow from the right ventricle to the pulmonary artery can be destroyed completely or damaged, causing blood to leak back to the right ventricle. The long-term outcome of this condition is right ventricular re-modeling caused by the regurgitated blood (see Fig. 3.1). Nowadays, patients undergo follow up operations to reduce this post-operative sequelae. In particular, the valves are replaced with an artificial device to reduce the regurgitation. A clinical challenge for these patients is determining the optimal time for intervention. On the one hand, it is preferable to wait as long as possible before performing such operations since the artificial devices have a limited life-span resulting in the need for repeat follow-up operations. Either open heart surgery or percutaneous valve replacement may be required, depending on the diameter of the pulmonary annulus. These therapies carry some danger, so minimising the frequency and total number of surgeries is crucial for these patients. On the other hand, the heart undergoes re-modeling due to the regurgitated blood producing volume overload in the right ventricle and if left long enough this damage can be irreversible. Determining the optimal time for intervention is a trade-off between maximising the time between operations and minimising damage to the heart.

Therefore, the clinical application for this condition is to better understand how the heart evolves over a large period of time, often decades, with and without surgery, with the main goal to be able to predict the optimal time for intervention. However, due to the large variability of shape, the optimal time, placement and size of the artificial device can be difficult to determine. Furthermore, the complexity of the biological mechanisms involved in heart growth hinders the development of direct model of cardiac re-modeling. Statistical shape analysis can therefore be employed to aid in further understanding the pathology to assist cardiologists with diagnosis, therapy planning, and long term prognosis.

34 Chapter 3. Bi-ventricular Growth Modeling in the Tetralogy of Fallot Population

More particularly, we are interested in determining quantitative measures of the shape that correlate with the cardiac function and the severity of the disease in these patients (bio-markers of the disease). Such bio-markers could be used to assess the severity of the disease when regurgitation cannot be assessed or to provide structural information that may appear before deterioration of the blood dynamics.




Figure 3.1: Comparison of a normal heart with a Tetralogy of Fallot heart.

3.2.2 Chapter Organisation

In the following sections, we will first review in Section 3.3 statistical shape analysis methods with a particular focus on the statistical analysis of surfaces. We will introduce the formalism of deformations which is at the center of most of the current works in computational anatomy [Thompson et al., 2004]. Then, we will detail the framework of currents to represent surfaces, and show how this can be turned into an effective shape analysis technique. Section 3.4 will apply this methodology to the shape of the heart of 13 patients with repaired tetralogy of Fallot. Correlating the shape with clinical variables will illustrate how we can extract some insight about the relationship between morphology and physiology. Last but not least, we will exemplify how the lack of longitudinal measurements can be bypassed by building a statistical generative growth model from cross-sectional data which summarises the heart shape remodeling at the population level.

3.3 Statistical Shape Analysis

3.3.1 Shapes, Forms and Deformations

There is generally no physical model that can faithfully relate the shape of organs in different patients. Thus, to analyse their variability in a population, one usually extracts some anatomically representative landmarks (or more generally geometric features), and models their statistical distribution across the population, via a mean shape and covariance structure analysis after a group-wise matching for instance. One of the earliest methods

[Bookstein, 1978, Bookstein, 1986] consists in studying the variability of anatomical landmark positions among a population: after a global pose (position and orientation) normalisation and the consistent identification of landmarks within all patients (see Fig. 3.2), a principle component analysis (PCA) is performed to extract the main modes of variation of the shape.




Figure 3.2: Using landmark-based methods to register one surface mesh to another surface mesh, for each point x in patient A , we require the corresponding point $y(x)$ in patient B . This requires a full parameterisation of the patient surface meshes. This is illustrated on the right ventricle of two patients with rToF.

3.3.1.1 Shapes

In such a process, the global pose of our objects is often considered as a nuisance factor related to the arbitrary coordinate system of the acquisition device: generally speaking the shape of an object is understood as the geometric information that is invariant under translation, rotations and rescaling. Thus, one is often interested only in the shape, i.e. what remains if we quotient out the object space by a group action, usually rigid body, similarity or affine transformations. For instance, when we consider the equivalent classes of sets of N (labeled or unlabeled) points under the action of a global similarity (resp. rigid body) transformation, we obtain the celebrated shape (resp. size and shape) spaces of Kendall et al [Kendall, 1989, Dryden and Mardia, 1991, Le and Kendall, 1993, Small, 1996]. One can define similarly the shape spaces of curves by removing the effect of re-parameterisations and the global pose of the curve [Joshi et al., 2006, Mio et al., 2007, Joshi et al., 2007b, Joshi et al., 2007a]. However, in medical image analysis, the natural coordinate system of an organ is its position and orientation with respect to surrounding tissues and organs in the image. Thus, once images are normalised by aligning them to a reference anatomy, the pose of organs is a variable of interest and should be kept for analysis.

3.3.1.2 Deformations

An alternative modeling of shapes was proposed by D’Arcy Thompson in 1917 [Thompson, 1917]. The idea is to assume that there is an atlas object, which represents the reference shape (this can be seen as an atlas or an equivalent of the mean shape). Then, the variability of the shape is analysed through the deformations of this reference object towards the actual observations: a shape difference is encoded by the transformation that deforms one onto the other. As the deformation of a smooth object should be a smooth object, we have to work with diffeomorphisms (invertible, one-to-one mappings with smooth inverses). This formalism was promoted to a generic shape analysis tool by Grenander and Miller [Grenander, 1993, Miller and Younes, 2001] based on advanced mathematical tools to compute on infinite dimensional groups of diffeomorphisms [Trouvé, 1998]. One key feature of this lift of the shape characteristics from the object space to the transformation space is that it allows to apply the typical deformations to other objects than the ones analysed in the first place (provided that they also live in the coordinate system of the atlas). For instance, [Durrleman et al., 2008b] analysed the variability of sulcal curves on the surface of the brain cortex to extract the main deformation modes. These modes could then be used to deform accordingly the surface of the cortex or the full 3D volume of the brain.

However, the problem is even more complex than previously as we want here to perform statistics on large deformations, which are known to belong to a non-linear and infinite dimensional manifold. The Riemannian setting is one of the most powerful structures to generalise simple statistics to non-linear spaces: it provides a definition of the distance between points of our manifold and a notion of shortest path / straight lines using geodesics [Pennec, 2006, Pennec, 2008]. The main difficulty is that the mean value cannot be defined through an integral or a sum as in Euclidean spaces. Instead, one must look for points in the manifold that minimise the dispersion of the other points around it, which is commonly measured by the variance (the mean squared distance). This is what is called the Fréchet or Karcher mean. Then, one can compute the covariance matrix (the directional dispersion around the mean) by developing the manifold onto its tangent space at the mean point. Basically, this amounts to representing each data point by the momentum (initial speed vector) needed to shoot¹ a geodesic to it from the mean.

This generalisation of statistics has to be slightly modified to fit the atlas deformation model: the ‘distance’ between the atlas and a shape is given here by the length of the shortest path in the space of deformations. The atlas is called unbiased if it is centered, i.e. if it minimises (among all possible atlases) the sum of squared distances to the shapes [Joshi et al., 2004]. In this sense, this is the Fréchet mean of the shapes. Then, each shape is represented by the momentum of the deformation that allows to regenerate the shape by deforming the atlas. These momentum all belong to the tangent space² of the diffeomorphisms at the identity, which is a vector space in which we can perform a PCA [Vaillant et al., 2004]: shooting in the space of deformations along the first eigenvectors of the covariance matrix of the momentum give deformation modes that represent the main

¹We use the term "geodesic shooting" to define the integration of the Euler-Lagrange equations, which plays the role of the exponential map in Riemannian geometry

²More precisely to the cotangent space

shape variability when applied to the atlas.

3.3.2 From Points to Surfaces: the Formalism of Currents

Landmarks can be encoded by the probability of their location in space (typically a Gaussian around their expected value). When the noise is going to zero, then this pdf becomes singular. However, we can continue to deal with these types of singularities by considering distributions (generalised functions) instead of functions, which include Diracs. Mathematically, a Dirac is not a function, but an object which can be characterised by the result of its integration against any function of a sufficiently smooth functional space: $\forall f \in W, \int \delta_x(y) \cdot f(y) \cdot dy = f(x)$. This is actually an element of the linear functionals over the space W . In that framework, a set of N points x_k can be represented by the ‘pdf’ $p(x) = 1/N \sum_i \delta_{x_k}(x)$, and its evaluation on a function $f(x)$ results in the mean of the values of f at the points x_k .

When we move to curves and surfaces, we are dealing with the locally singular locus of points in certain directions only: curves and surfaces are continuous along their tangent. The extension of the notion of distributions that allows to take that information into account is the geometric integration theory. The basic idea is to integrate differential forms, but for objects like smooth surfaces in 3D, we shall simply define currents by their action on vector fields, similarly to the way distributions are defined by their action on scalar functions.

3.3.2.1 Currents for Surfaces

Let W be a Hilbert space of vector fields (a possibly infinite complete vector space provided with a scalar product). For a given surface S , we can measure the flux of any vector field $\omega \in W$ through this surface:

$$\mathcal{S}(\omega) = \int_{x \in S} \langle \omega(x) | n(x) \rangle d\sigma(x),$$

where $n(x)$ is the normal to the smooth surface at point $x \in S$ and $d\sigma(x)$ the surface element around point x . The shape of the surface S is characterised by the variation of the flux as the test vector field ω varies in W . Thus, the surface actually defines a continuous linear form on W which can be identified to an element of the dual space W^* , which is the space of linear functionals from W to \mathbb{R} (currents). The nice property of currents is that it is a vector space: we can add or subtract current from each others, or multiply them by a scalar. However, we should keep in mind that the space of currents is larger than the space of smooth surfaces: one can for instance add many pieces of surfaces together in a non continuous way to create a non-continuous object.

Now that we have identified surfaces to currents, we need to define more carefully what is the space W that we consider. The core element proposed by [Glaunès, 2005] is to consider a Kernel metric (typically the Gaussian kernel $K_W(x, y) = \exp(-\|x - y\|^2 / \lambda_W^2)$) in order to turn W into a Reproducible Kernel Hilbert Space (RKHS). The reproducibility property implies that W is the dense span of basis vector fields of the form $\omega_x^\alpha(y) = K_W(x, y)\alpha$. This means that any vector field of W can be written as an infinite linear combination of such basis vectors. The kernel induces a scalar product which is easily computed on two

38 Chapter 3. Bi-ventricular Growth Modeling in the Tetralogy of Fallot Population

basis vectors: $\langle \omega_x^\alpha \mid \omega_y^\beta \rangle_W = \alpha^T K_W(x, y) \beta$. The kernel can be seen as a low pass filter which weights more heavily the high spatial frequencies of the vector field ω than the low frequencies (essentially penalising high frequencies).

Since W is a Hilbert space, by the Riesz representation theorem, there exists an isometric linear mapping \mathcal{L}_W from W to W^* which maps to each $\omega \in W$ the linear form that reproduces the scalar product: $(\mathcal{L}_W(\omega))(\omega') = \langle \omega \mid \omega' \rangle_W$. This canonical isomorphism allows to define the dual of an element, or conversely to map to each surface the vector field that optimises the flux. For instance, the dual of the above basis vectors are the Dirac delta currents $\delta_x^\alpha = \mathcal{L}_W(\omega_x^\alpha)$ as we have $\langle \omega_x^\alpha \mid \omega \rangle_W = \alpha^T \omega(x)$. These can be seen as vector fields whose spatial support is concentrated at one point only. The space of currents W^* is the dense span of these basis elements. For instance, the surface S is represented by the current $\mathcal{S} = \int_{x \in S} \delta_x^{n(x)} d\sigma(x)$. The usual way to define a norm on the dual space is to take the operator norm: $\|\mathcal{S}\|_{W^*} = \sup_{\omega \in W, \|\omega\|=1} |\mathcal{S}(\omega)|$. The distance induced by this norm might seem difficult to use, but thanks to the RKHS properties, we have a closed form for basis vectors:

$$\langle \delta_x^\alpha \mid \delta_y^\beta \rangle_{W^*} = \alpha^T K(x, y) \beta = \langle \delta_x^\alpha \mid \delta_y^\beta \rangle_W \alpha^T \beta$$

It is interesting to notice that this distance can be approximated by

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

when the points x and y are within a fraction of λ_W , while the distance is essentially constant (and equal to $\|\alpha\|_{L_2}^2 + \|\beta\|_{L_2}^2$) when the points x and y are more than a few λ_W apart. This behavior is typical of a robust distance in statistics, meaning that outliers (over a few λ_W in distance) will have (almost) no effect on the optimisation of the distance between the surfaces as it is an almost constant penalty.

In practice, surfaces are often represented by discrete triangulated meshes. Assuming that each face has a support which is smaller than ε times the scale λ_W , we can approximate the surface by the current $\mathcal{S} = \sum_k \delta_{x_k}^{\alpha_k}$, where x_k is the barycenter and α_k the normal weighted by the area of the face. From the above Taylor expansion of the distance, we can see that the approximation error is less than ε^2 for each face in the W^* norm.

The scalar product between two discrete currents is obtained by linearity:

$$\left\langle \sum_k \delta_{x_k}^{\alpha_k} \mid \sum_j \delta_{y_j}^{\beta_j} \right\rangle_{W^*} = \sum_{k,j} \alpha_k^T K_W(x_k, y_j) \beta_j$$

and the distance between two discrete surfaces is simply $\text{dist}^2(S, S') = \|\mathcal{S} - \mathcal{S}'\|_{W^*}^2$.

3.3.2.2 Sparse Representation of Currents using Matching Pursuit Algorithm

Representing shapes using currents gives a nice theoretical framework that allows us to compute simple statistics such as mean and principal modes. However, the complexity of the distance computation is quadratic in the number of Dirac delta currents used to represent the shape. This number can be quite high if we take highly detailed surface meshes, even though it adds nothing for the comparison of surfaces at a given scale λ_W . Therefore,

in order to remain computationally as efficient as possible, we require a sparse representation of currents that retains the information needed for the scale of the analysis. This is the essence of the matching pursuit algorithm as described in [Durrleman et al., 2009a].

In brief, the matching pursuit algorithm is a greedy approach for finding an approximation of the current T that solves $\mathcal{L}_W^{-1}(T) = \gamma$, for a given vector field $\gamma \in W$. This amounts to finding N points (x_k) and vectors (α_k) such that the current $Pk(T) = \sum_{k=1}^N \delta_{x_k}^{\alpha_k}$ is as close as possible to T . If the points are known, then $\Pi(T)$ is the orthogonal projection of T onto $Span(\delta_{x_k}^{\varepsilon_q}; q = 1, 2, 3, k = 1 \dots N)$, where ε_q is the canonical basis of \mathbb{R}^3 . The orthogonality condition: $\langle T, \delta_{x_k}^{\varepsilon_q} \rangle_{W^*} = \langle \Pi(T), \delta_{x_k}^{\varepsilon_q} \rangle_{W^*}$ leads to a linear set of $3N$ equations

$$\sum_{p=1}^N (K(x_i, x_p) \alpha_p)_k = \gamma(x_i)_k \quad (3.1)$$

which can be solved iteratively over the continuous space of Dirac delta currents to find the point positions (x_k) as well as their associated momenta and the residual vector field. Additionally, we can also sample the vector field on a grid Λ . Applying the matching pursuit in the discrete case constrains the estimated momenta to lie on the nodes of the grid which forces the estimated current to belong to a discrete set of currents W_Λ^* . This simplifies the problem in both complexity as well as computation time. Figure 3.3 shows a sparse representation of the right ventricle.




Figure 3.3: The Dirac delta currents of a triangulated mesh are the normal vectors of every face, centered at the face barycenters. A greedy algorithm reduces the amount of delta currents needed to represent the shape while preserving the accuracy of the representation.

3.3.3 An Algorithm for Surface Registration using Currents

Having defined a non-parametric representation of surfaces, we now require an algorithm to register one surface to another. Here, the term registration is meant to define the action of transforming two objects into the same co-ordinate system. Since we are comparing topologically similar shapes, the transformations are restricted to those which preserve the topology of the object and give a smooth one-to-one (invertible) transformation (i.e. a diffeomorphism). The space of diffeomorphisms give non-linear deformations

that allow local smooth variations to be captured in the registration. However, diffeomorphic transformations have an infinite degree of freedom, therefore optimising over the whole group of diffeomorphisms may not be possible. So, we use a smaller infinite group of diffeomorphisms to allow computations with discrete parameterisations using the Large Deformation Diffeomorphic Metric Mapping (LDDMM) method described in [Beg et al., 2005, Miller et al., 2003] for images.

The LDDMM framework uses a group of diffeomorphisms constructed through integration of time-varying vector fields that belong to a RKHS. This gives a geodesic flow of diffeomorphisms ϕ_t for a continuous parameter t within the interval $[0, 1]$. At time $t = 0$, we have the identity mapping ϕ_0 . The mapping at time 1 gives the desired transformation ϕ_1 which is required for mapping one image to the other. The path of any point x is defined by $\phi_t(x)$ and leads to the final position $\phi(x) = \phi_1(x)$. By following the reverse path, we can compute the inverse deformation.

The basic idea of the LDDMM framework is to minimise the distance between objects after transformation (this is the similarity energy or data attachment term) with a penalisation for the length of the deformation trajectory (which is the regularisation energy). A time t , the speed at point $y = \phi_t(x)$ is $v_t(y) = d\phi_t(y)/dt$. This suggests to define the energy of a velocity field at deformation ϕ_t using a right-invariant metric $\|v_t\|_{\phi_t}^2 = \|v_t \circ \phi_t^{(-1)}\|_W^2$, where $\|\cdot\|_W$ is the norm of the RKHS of the velocity fields at the identity transformation. The energy of the deformation trajectory is thus $E(\phi_t) = \int_0^1 \|v_t \circ \phi_t^{(-1)}\|_W^2 dt$, and the optimal curves joining $\phi_0 = Id$ to ϕ_1 are geodesics in the space of diffeomorphisms. Since optimal transformations are geodesics, we know that they are completely determined by their initial value, here the velocity field $v_0(x)$ (or more specially the initial momentum $K_W^{-1}v_0$). Denoting ϕS the action of the transformation ϕ on the source object S , the registration criterion to the target object T is thus

$$C(\phi) = \text{dist}(\phi_1 S, T)^2 + \int_0^1 \|v_t \circ \phi_t^{(-1)}\|_W^2 dt$$

For our surface registration problem, the objects are surfaces represented by their associated currents \mathcal{S} and \mathcal{T} and the distance is taken in the space of currents W^* . Since our currents have a finite point-wise support, we can show that the optimal initial momentum of the deformation has the same point-wise support (if $\mathcal{S} = \sum_i \delta_{x_k}^{\alpha_k}$, then $v_0(x) = \sum_i K_W(x, x_k) \beta_k$ for some set of vectors β_k), which means that we are left with a finite dimensional optimisation problem [Vaillant and Glaunes, 2005, Glaunès, 2005].

3.3.4 Building an Unbiased Atlas

For building the optimally centered atlas, we took the ‘forward’ strategy that models the set of surfaces as the deformation of an unknown ‘ideal’ atlas plus some residuals [Allasonnière et al., 2007, Durrleman et al., 2008a, Durrleman et al., 2009a]. This can be expressed as:

$$T_i = \phi_i \hat{T} + \varepsilon_i. \quad (3.2)$$

for i patients, where \hat{T} is the atlas we are estimating, ϕ is the deformation that maps the atlas to the surfaces T , and ε represents the residuals (shape features not captured by the

atlas such as changes in topology etc.). The mean shape information is described in the atlas \hat{T} while the shape variability is encoded in the transformation ϕ_k .

The atlas is first initialised by taking the mean of the patient meshes. This initial atlas is then registered to each of the patients individually. A new atlas \hat{T} that minimises the error

$$\Lambda(\hat{T}) = \|T_i - \phi_i(\hat{T})\|_{W^*}^2 \quad (3.3)$$

is computed. We then register the updated atlas to the individuals, recompute the atlas and loop until convergence (see Algorithm 1).

Algorithm 1 Atlas Estimation

Require: N segmented patient images (surface meshes).

- 1: Rigidly align meshes to a reference patient using rotations and translations.
 - 2: Create initial atlas \hat{T}_0 as the mean of the patient meshes.
 - 3: **loop** {over N until convergence}
 - 4: Estimate the transformations ϕ_k that register the atlas \hat{T}_{N-1} to the individual T_k .
 - 5: Update the atlas by minimising the error in 8.7 using the estimated transformations ϕ_k and the atlas \hat{T}_{N-1}
 - 6: **return** Final atlas \hat{T}_N and the related transformations ϕ_k^N .
-

3.4 Application on ToF Data

To demonstrate the usefulness of our statistical shape analysis in a clinical context, we consider a population of patients with repaired tetralogy of Fallot.

As mentioned in the introduction, the clinical problem in rToF patients is in better understanding the shape re-modeling over time to aid in determining the optimal time for surgical intervention. In view of this problem, we are first interested in identifying the clinically relevant shape patterns by investigating the relationship between shape and given clinical indices. The assumption here is that changes in heart morphology may reveal structural and functional dysfunctions due to the chronic regurgitation. In particular, we aim to establish which shape patterns are related to the pathology in order to give further insights into the condition. We can then go beyond the identification of pathological shape patterns, to estimate a predictive growth model of the heart to give an indication of how the heart will grow and re-model in time for this pathology.

Making use of currents to represent shapes, and an atlas to give an average representation of the population, we can apply statistical methods to address key questions for diagnosis, prognosis and prediction.

3.4.1 The Analysis Pipeline

Patient Data We consider a data-set of thirteen patients (10 males, mean age \pm SD = 19 ± 9) with repaired tetralogy of Fallot. Steady-state Free Precision cine MRI were acquired using a 1.5T scanner (Avanto, Siemens AG, Erlangen) in the short axis view

42 Chapter 3. Bi-ventricular Growth Modeling in the Tetralogy of Fallot Population

covering entirely both ventricles (10-15 slices; isotropic in-plane resolution: 1.1x1.1mm to 1.7x1.7mm; slice thickness: 5mm to 10mm; 25-40 phases).

Surface Mesh Delineation In the case of ToF patients, we are interested in studying the shape of the left and right ventricles of the heart. Using the statistical shape analysis methods described in the previous section, we can consider the ventricles as surfaces represented by triangulated meshes. The surface meshes of the right and left ventricle endocardium (inner layer of heart tissue), and left ventricle epicardium (outer layer of heart tissue) are defined using image segmentation by delineating the boundaries of each ventricle at end-diastole using the methods proposed by Zheng et al. [Zheng et al., 2008] (see Fig. 3.4). This method defines the delineated boundaries using an anatomical model and thus establishes a point correspondence between meshes. Therefore, in this instance we used a standard least-squares method to first rigidly align the surface meshes to reduce the effect of patient positioning.




Figure 3.4: Delineated boundaries of the right ventricle (pink), left ventricle endocardium (green) and left ventricle epicardium (yellow) shown on one image slice (left) and the 3D reconstruction (center and right).

Mean Atlas Construction Using the algorithm for computing a mean shape with the currents method described in Section 3.3, we estimated an atlas for the ToF data, see Fig. 3.5. Eleven iterations of the alternate minimisation for the shape atlas (Algorithm 1) were needed to reach convergence. The resulting atlas \hat{T} was well centered (mean over standard deviation of the deformations was 0.36946). Atlas-to-patient registration is implemented in parallel on a cluster of computers, which means that the atlas creation time is minimally dependent of the number of patients used to create it.

3.4.2 Diagnosis Parameters

A key topic in computational shape analysis is identifying pathologically specific shape features in populations of diseased patients compared to controls (see [Cates et al., 2008] and references therein for instance). Beyond simply identifying the pathological shape features, we would like to also quantify the degree to which the shape is altered due to the pathology. One way in which this can be done is by correlating the shapes with clinical features to determine the severity of the disease, which requires a consistent representation




Figure 3.5: Three views of the mean atlas estimated from 13 patients with repaired ToF with the right ventricle in blue, left ventricle endocardium in white and left ventricle epicardium in wire-frame.

of the patient shapes. For this we use firstly principal component analysis (PCA) to extract the main modes of shape variation followed by standard statistical design to exhibit those that are correlated to the pathology (see Fig. 3.6 top row).

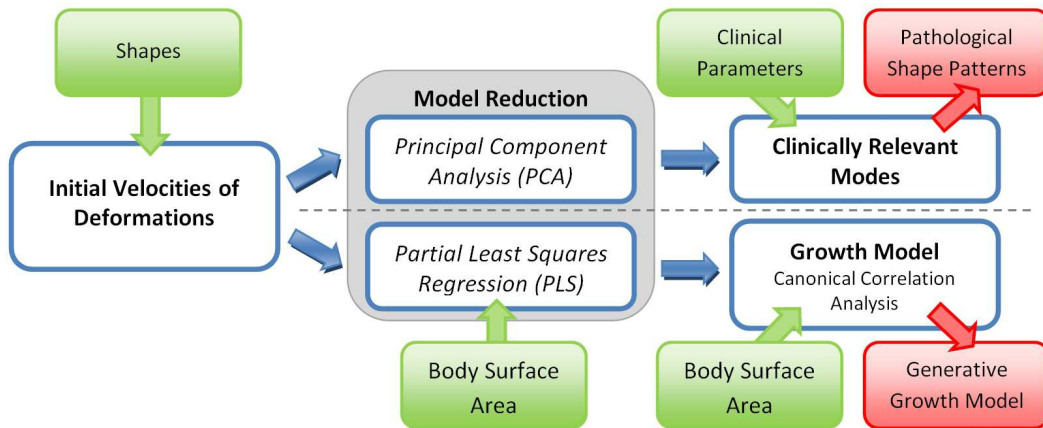


Figure 3.6: Shape analysis pipeline with both clinically relevant model reduction and growth model generation. Using PCA modes and standard correlation analysis, the pathological shape patterns can be identified. Using a combination of PLS regression and CCA a statistical generative growth model can be derived.

3.4.2.1 Model Reduction using Principal Component Analysis

Since statistical shape analysis is a high dimensional problem with a large number of parameters and variables to solve for (despite the matching pursuit reduction, a shape can still be represented by hundreds of moments), we first reduce the dimension of the problem by applying principal component analysis (PCA). This gives the modes of deformation that describe the amount of variation of shape observed in the population.

PCA is applied on the initial velocity fields $v_0^{(i)}$ to extract the main deformation modes

44 Chapter 3. Bi-ventricular Growth Modeling in the Tetralogy of Fallot Population

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 p are calculated by solving the eigenvalue problem $\Sigma p = \mu p$, where the elements σ_{ij} of the covariance matrix Σ are calculated in the kernel space W . Assuming that the deformation from the atlas to patient i is parameterised by the initial vector field $v_0^{(i)}(x) = \sum_k K_W(x, x_k) \beta_k^{(i)}$, where the x_k are the point positions of the delta Dirac currents of the atlas, and $\beta_k^{(i)}$ the moment vector at x_k , then the mean initial vector field is $\bar{v}_0(x) = \sum_k K_W(x, x_k) \bar{\beta}_k$ and the covariance is

$$\sigma_{ij} = \langle v_0^i - \bar{v}_0, v_0^j - \bar{v}_0 \rangle_V = \sum_{k,l} (\beta_k^{(i)} - \bar{\beta}_k) K_W(x_k, x_l) (\beta_l^{(j)} - \bar{\beta}_l) \quad (3.4)$$

The principal components are obtained by computing the spectral decomposition $\Sigma = PMP^T$. 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 p^m . The m^{th} loading l^m of the PCA decomposition is given by the formula:

$$l^m = \sum_i p^m[i] \beta^{(i)} \quad (3.5)$$

In this equation, $p^m[i]$ is the i^{th} element of the m^{th} eigenvector of Σ , $\beta^{(i)} = (\beta_1^{(i)} - \bar{\beta}_1, \dots, \beta_n^{(i)} - \bar{\beta}_n)^T$ is the $n \times 3$ matrix that gathers the n centered moment of patient i . As a result, the initial velocity field of the m^{th} mode is $v_{0,m}(x) = \sum_k K_W(x, x_k) l^m$. The variability captured by this m^{th} deformation mode between $[\eta \sigma^m; +\eta \sigma^m]$ is visualised by deforming the atlas T with the deformations ϕ^{-m} and ϕ^{+m} parameterised by the moments $\bar{\beta} - \eta l^m$ and $\bar{\beta} + \eta l^m$ respectively. Selecting the first p modes only among the $N - 1$ possible modes (where N is the number of patients) allows to explain a percentage $\sum_{m=1}^p \sigma^m / \text{trace}(\Sigma)$ of the total variance.

The orthogonal projection of each patient's initial velocity field onto the selected PCA subspace gives a unique shape vector. This simply corresponds to the coordinates of the projection in the basis constituted by the chosen eigenmodes:

$$s^{i,m} = \langle v_0^{(i)} - \bar{v}_0, v_{0,m} \rangle_W = \sum_{j,k} [\beta_j^{(i)} - \bar{\beta}_j] K_W(x_j, x_k) l_k^m \quad (3.6)$$

Using PCA 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 and could be considered as noise in the model. Consequently, we are 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.

3.4.2.2 Identifying Factors Between Shape and Clinical Features

The $s^{i,m}$ quantify the amount of variability along the m^{th} mode present in patient deformation. We can thus investigate the heart shape by relating these shape vectors to clinical

parameters that quantify the pathology. Ordinal clinical parameters are investigated using non-parametric rank-based statistics. Kruskal-Wallis analysis of variance is applied to find effects between the investigated parameters and shape [R Development Core Team, 2009]. If an effect is found, post-hoc two-sample Wilcoxon test is used to determine which levels differ [R Development Core Team, 2009]. Continuous clinical parameters are investigated using linear regression and Akaike Information Criterion (AIC) model reduction [Akaike, 1974] to detect relevant modes and the direction of correlation.

To illustrate this method, in [Mansi et al., 2011b] we showed that these methods enable one to identify shape features related to the severity of the regurgitation for a data-set of 49 repaired ToF patients. The relationship between RV shape and pulmonary regurgitation were investigated by relating the PCA shape vectors with tricuspid regurgitation, transpulmonary valve regurgitation and pulmonary regurgitation volume indices taken from colour Doppler ultrasound and phase-contrast magnetic resonance images (PC-MRI). 90% of the spectral energy was explained by 18 PCA modes.

3.4.3 Building an Evolution Model

As explained beforehand, understanding and quantifying heart remodelling in these patients is crucial for planning pulmonary valve replacement. Given that there is a lack of longitudinal data available for these patients, we make use of the atlas as the mean of the population and cross-sectional statistics to formulate a generative growth model. Such a model could be used as reference, from which the pathology evolution of one patient could be quantitatively compared. In cross-sectional statistical design, one does not propagate the evolution over time for a single patient but rather considers each patient as an instance in the growth evolution. In this way we can model the growth of the population given these instances using regression analysis (see Fig. 3.7). By making use of the initial velocities v_0^i that parametrise the deformations ϕ^i computed using the methods described in the previous sections, we can regress the velocities against an index of patient growth using standard statistical techniques.




Figure 3.7: Cross-sectional regression of shapes. Each patient is associated to a point in time (patient age for instance). A regression model is derived from the temporal data.

In order to obtain statistically significant results, we first need to reduce the dimen-

46 Chapter 3. Bi-ventricular Growth Modeling in the Tetralogy of Fallot Population

sionality of the problem to consider just the factors related to patient growth, while also removing any co-linearity between factors. In the previous section the model reduction was performed using PCA. In this case we chose instead to use partial least squares (PLS) regression since it has the added advantage of computing the components that are most related to a given external parameter (i.e. patient growth). Using PLS allows us to compute the components that best describe the variance of both the matrix of predictors (X) and the matrix of responses (Y), as well as the covariance between X and Y , in a manner such that the regression $Y = f(X)$ is optimal.

In the case of ToF patients, we would ideally like to model the atlas deformation of a patient as a function of growth (i.e. $deformation = f(growth)$), however solving this problem is not possible due to the large number of deformation parameters that would need to be predicted with a single, one-dimensional, parameter (growth). Rather, we revert the problem to be a function of the deformations: $growth = f(deformations)$ which has a much lower number of parameters to predict. The outputted values are then projected onto the reduced PLS subspace and from that we can revert the problem to the desired form as a function of the shape using canonical correlation analysis (CCA) to give a generative growth model of the heart.

For this example we use body surface area (BSA) as the index of growth rather than patient age to better represent the growth given the variable age at which children enter puberty. We use the Dubois formula [DuBois and DuBois, 1915] to compute the BSA for each patient:

$$BSA(m^2) = 0.007184 \times weight(kg)^{0.425} \times height(cm)^{0.725} \quad (3.7)$$

The pipeline we have just described for computing a generative growth model is shown in Fig. 3.6 (bottom row).

3.4.3.1 Model Reduction using Partial Least Squares Regression

As introduced in the previous section, regression using the partial least squares regression (PLS) method is based on finding an optimal basis of the predictor variables X that maximises the variances of X and Y as well as their covariances. The method can be considered as the optimal estimation of two weight vectors r and s that satisfy

$$\max_{|r|=|s|=1} cov(Xr, Ys) = \max_{|r|=|s|=1} var(Xr)corr(Xr, Ys)^2var(Ys) \quad (3.8)$$

under the constraint that the regression between X and Y is optimal. Mathematically, the centered variables $X_c = X - \bar{X}$ and $Y_c = Y - \bar{Y}$ are modeled by

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

where T and U are the matrices of the PLS modes, P and Q are the loading matrices which describe the weight of each variable in X and Y respectively and E and F are the

residual terms which are the same size as X_c and Y_c respectively. Additionally the following regression condition is imposed for the PLS modes:

$$U = TD + G \quad (3.9)$$

where D is a diagonal matrix of weights and G a matrix of residuals. Due to this added condition, the PLS loadings P and Q are not necessarily orthogonal as is the case for PCA modes.

Algorithm 2 Partial Least Squares Regression (PLS)

Require: Variables X and Y , number of components $p \leq N - 1$.

- 1: $X_c^0 = X - \bar{X}$, $Y_c^0 = Y - \bar{Y}$
 - 2: **for** $n = 1$ to p **do**
 - 2: $r^n \leftarrow$ first eigenvector of $X_c^{nT} Y_c^n Y_c^{nT} X_c^n$
 - 2: $t^n \leftarrow X_c^n r^n / \|r^n\|$ n^{th} PLS component of X
 - 2: $s^n \leftarrow Y_c^n t^n / t^{nT} t^n$
 - 2: $u^n \leftarrow Y_c^n s^n / \|s^n\|$ n^{th} PLS component of Y
 - 2: $p^n \leftarrow X_c^n t^n / t^{nT} t^n$ n^{th} loading of X
 - 2: $q^n \leftarrow Y_c^n u^n / u^{nT} u^n$ n^{th} loading of Y
 - 2: $X_c^{n+1} \leftarrow X_c^n - t^n p^{nT}$ deflation of X_c
 - 2: $Y_c^{n+1} \leftarrow Y_c^n - t^n [t^{nT} Y_c^n / t^{nT} t^n]$ deflation of Y_c
 - 3: **return** $T = (t^n)_{n=1\dots p}$, $P = (p^n)_{n=1\dots p}$, $U = (u^n)_{n=1\dots p}$, $Q = (q^n)_{n=1\dots p}$
-

Several algorithms have been proposed to compute the PLS modes. In this work, we use the PLS1 method, an efficient iterative algorithm that does not require matrix inversion as summarised in Algorithm 2. X is the matrix of the initial velocity field moments for all patients and Y is the vector of the BSA values for all patients. The first five PLS modes are shown in Figure 3.8 and the explained variance and correlation of modes is shown in Figure 3.9. These modes account for 99% of the covariance between shape and BSA in the population and 61% of the shape variability and are oriented along increasing BSA. Visually we can see that Modes 1 and 3 display an overall dilation in both the left and right ventricles. The second mode shows a narrowing in the right ventricular outflow tract with a noticeable dilation in the left ventricle which can be seen in the top view of the mode. Modes 4 and 5 show an elongation at the right ventricular outflow tract. Mode 5 also shows an aneurysm in the right ventricle outflow tract which is characteristic of ToF patients, although it is still not clear if this is due to pulmonary regurgitations (although both are correlated) or the initial surgical patch.

3.4.3.2 Generating a Growth Model using Canonical Correlation Analysis

Using PLS as described above allows us to predict BSA given the shape, however what we would like is to estimate the shape given BSA. To reverse the relationship we use canonical correlation analysis (CCA) on the PLS shape vectors (the t_i^j 's of Algorithm 2). CCA computes the vectors r and s that maximise the correlation between the two sets X




Figure 3.8: The first five PLS modes of variation that describe 99% of the observed BSA variability and 61% of the observed shape variability in the population.

and Y :

$$\max_{|r|=|s|=1} \text{corr}(Xr, Ys)^2 \quad (3.10)$$

where Y is the vector of BSA values and $X = [t_1, \dots, t_N]^T$ is the matrix of shape descriptors.

If we define the covariance matrices $V_{UZ} = \frac{1}{N-1} U^T Z$, then 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 cor-

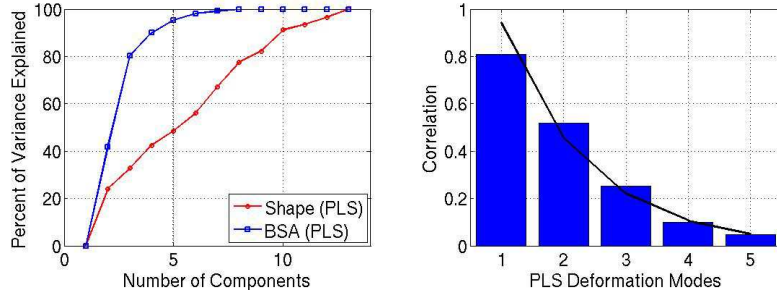


Figure 3.9: Cumulative variance of PLS modes with respect to shape and BSA (left panel) and CCA correlation coefficients between BSA and PLS modes with exponential fitted curve in black (right panel).

relation coefficient $r = \sigma_{XY} / \sqrt{\sigma_{XX} \sigma_{YY}}$. The sought correlations are obtained by SVD decomposition of Γ :

$$\Gamma = ASB^T \quad (3.11)$$

S is the diagonal matrix of the correlation coefficients between correlation vectors and A and B are rotation matrices of correlation vectors, i.e. $A^T A = B^T B = Id$. 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 atlas T with the growth deformation Φ parametrised by the moments $\mu = B.R.\Sigma_k \rho[k] p^k$, where p^k is the k^{th} PLS loading.

3.4.3.3 Interpretation

The growth model computed on the ToF data-set is shown in Fig. 3.10. This model shows an expected overall growth of both ventricles as body surface area increases. We can also see the caving of the septum into the right ventricle as time passes and the elongation of the right ventricular outflow tract which is observed in these patients over time.

Given a larger data-set this model can be improved further by increasing the certainty in the model and by capturing more variability in shape observed in a wider population. As well, given more data we can divide the patients according to an external parameter such as treatment group to formulate a growth model for a given course of treatment. In the case of ToF, this could allow clinicians to have a better idea of how the heart re-models after different types of valve replacement surgery and more importantly, the effect of the initial surgery on the long-term outcome. This is the key question.




Figure 3.10: Mean growth model computed from a population of 13 rToF patients. Both ventricles grow as BSA increases

3.5 Conclusion

In this chapter we presented methods for computing statistics on shapes. The proposed methods rely on currents to represent the anatomical shapes in a consistent way without the need for defining landmarks or point correspondences between shapes. This allows the computation of population statistics such as the mean and variance observed within a population and to correlate shape with clinical indices, for instance to quantify the severity of a disease. We also present a framework for formulating a generative statistical growth model to simulate the growth of an organ over time. This is especially important for our clinical application where understanding the shape remodeling of the heart in congenital diseases like Tetralogy of Fallot is crucial to better understand the evolution of the disease and ultimately to aid for therapy planning.

Part II

STATISTICAL MODELLING OF CARDIAC MOTION

Prerequisite: Quantitative Validation of the iLogDemons Method

Contents

4.1 Chapter Overview	53
4.2 Application to Real Sequences from Healthy Volunteers	54
4.2.1 Methodology	54
4.2.2 Implementation	56
4.2.3 Image Pre-Processing	58
4.2.4 Application to Challenge Data	59
4.2.5 Strain Estimation	61
4.2.6 Discussion	61
4.2.7 Conclusion	65
4.3 Application to Synthetic Echocardiography Sequences	65
4.3.1 Introduction	65
4.3.2 Application to Challenge Data	66
4.3.3 Discussion	71
4.3.4 Conclusion	72

In Part I of this thesis the question of modelling cardiac growth over time was addressed by formulating a generative model of growth using statistical methods. Part II of this thesis is rather aimed towards a group-wise model of cardiac motion. This chapter is a first step towards this goal as a pre-requisite to the following chapters of this section to test state-of-the-art cardiac motion tracking algorithms for benchmarking and comparison. The first section of this chapter (Sec. 4.2) was published in [McLeod et al., 2012a] and contributes to the work published in [Tobon-Gomez et al., 2013]. Sec. 4.3 is based on the work published in [Prakosa et al., 2012a] (note that the description of the iLogDemons method has been removed, to avoid repetition).

4.1 Chapter Overview

Quantifying current methods for cardiac image registration was a key element of this work to analyse the state-of-the-art methods in order to identify the important features needed in future methods, and to benchmark future methods by providing quantitative analyses and

comparisons to previous methods. To this end, the incompressible LogDemons method (iLogDemons for short), was tested on a real data-set of healthy volunteers in the STACOM MICCAI cardiac motion tracking challenge of 2011 (Sec. 4.2) and to a data-set of synthetic sequences in the STACOM MICCAI cardiac motion tracking challenge of 2012 (Sec. 4.3).

4.2 Application to Real Sequences from Healthy Volunteers

The application of the previously proposed iLogDemons algorithm to the STACOM motion-tracking challenge data is described. The iLogDemons algorithm is a consistent and efficient framework for tracking left-ventricle heart tissue using an elastic incompressible non-linear registration algorithm based on the LogDemons algorithm. This method has shown promising results when applied to previous data-sets. Along with having the advantages of the LogDemons algorithm such as computing deformations that are invertible with smooth inverse, the method has the added advantage of allowing physiological constraints to be added to the deformation model. The registration is entirely performed in the log-domain with the incompressibility constraint strongly ensured and applied directly in the demons minimisation space. Strong incompressibility is ensured by constraining the stationary velocity fields that parameterise the transformations to be divergence-free in the myocardium. The method is applied to a data-set of 15 volunteers and one phantom, each with echocardiography, cine-MR and tagged-MR images. We are able to obtain reasonable results for each modality and good results for echocardiography images with respect to quality of the registration and computed strain curves.

4.2.1 Methodology

4.2.1.1 Cardiac Motion Tracking using Physiological Constraints

Tracking cardiac motion from 3D images is a difficult task due to the complex movement of the myocardium through the cardiac cycle. The left ventricular (LV) movement includes a contraction of the ventricle with a longitudinal motion towards the apex as well as a twisting motion from the base of the ventricle in the circumferential direction. Common methods for motion tracking using non-rigid registration are able to capture the dilation of the ventricle, however capturing the twisting motion is a difficult task. The incompressible log-domain demons algorithm described in [Mansi et al., 2011a] (iLogDemons for short) aims to tackle this problem by imposing physiological constraints (such as incompressibility and elasticity in the myocardium) in the previously proposed log-domain demons algorithm (LogDemons) [Vercauteren et al., 2008]. For the purpose of this work we don't provide here a state of the art on cardiac motion tracking algorithms, but rather refer the reader to [Mansi et al., 2011a] and references therein. We apply the iLogDemons method to a 3D data-set of 15 volunteers and one phantom with echocardiography, cine-MR and tagged-MR image sequences. The method is described here in brief, for a more thorough and descriptive analysis see [Mansi et al., 2011a].

4.2.1.2 Review of the Log-Domain Demons Algorithm

The iLogDemons algorithm is an extension of the LogDemons algorithm [Vercauteren et al., 2008]. The LogDemons algorithm estimates a dense non-linear transformation ϕ that best aligns a template image T to a reference image R . The transformation ϕ is parameterised by stationary velocity fields \mathbf{v} through the exponential map $\phi = \exp(\mathbf{v})$ [Arsigny et al., 2006]. The images R and T are registered by minimising in the space of velocities (the log-domain) the energy functional: $\varepsilon(\mathbf{v}, \mathbf{v}_c) = 1/\sigma_i^2 \|R - T \circ \exp(\mathbf{v}_c)\|_{L_2}^2 + 1/\sigma_x^2 \|\log(\exp(-\mathbf{v}) \circ \exp(\mathbf{v}_c))\|_{L_2}^2 + 1/\sigma_d^2 \|\nabla \mathbf{v}\|^2$, where σ_i^2 relates to the noise in the images and σ_d^2 controls the regularisation strength. The velocity field \mathbf{v} parameterises the transformation ϕ , and \mathbf{v}_c parameterises an intermediate transformation $\phi_c = \exp(\mathbf{v}_c)$ that models the *correspondences* between the voxels of the two images. During the *optimisation step*, $\varepsilon(\mathbf{v}, \mathbf{v}_c)$ is minimised with respect to \mathbf{v}_c . Under the diffeomorphic update rule $\phi_c \leftarrow \phi \circ \exp(\delta \mathbf{v})$, the optimal update velocity writes $\delta \mathbf{v}(\mathbf{x}) = (R(\mathbf{x}) - T \circ \phi(\mathbf{x})) / (\|J(\mathbf{x})\|^2 + \sigma_i/\sigma(\mathbf{x}))J(\mathbf{x})$. In this equation, $J(\mathbf{x})$ is the symmetric gradient $J(\mathbf{x}) = (\nabla R(\mathbf{x}) + \nabla(T \circ \phi)(\mathbf{x}))/2$. The correspondence velocity \mathbf{v}_c is then updated using the first order approximation of the Baker-Campbell-Hausdorff (BCH) formula $\mathbf{v}_c = Z(\mathbf{v}, \delta \mathbf{v}) = \mathbf{v} + \delta \mathbf{v} + 1/2[\mathbf{v}, \delta \mathbf{v}] + 1/12[\mathbf{v}, [\mathbf{v}, \delta \mathbf{v}]] + O(\|\delta \mathbf{v}\|^2)$, where the Lie bracket $[\cdot, \cdot]$ is defined by $[\mathbf{v}, \delta \mathbf{v}] = (\nabla \mathbf{v})\delta \mathbf{v} - (\nabla \delta \mathbf{v})\mathbf{v}$. Finally, the *regularisation step* estimates the optimal regularised transformation ϕ by minimising $\varepsilon(\mathbf{v}, \mathbf{v}_c)$ with respect to \mathbf{v} , which is approximated by smoothing the correspondence velocity \mathbf{v}_c with a Gaussian kernel G_σ .

4.2.1.3 Modeling Elasticity in the Myocardium

In order to incorporate an elastic regulariser into the LogDemons framework, a consistent mathematical formulation of the LogDemons regularisation is required. In [Mansi et al., 2011a] a closed-form expression of the demons Gaussian regulariser $\varepsilon_{reg}(\mathbf{v}) = 1/\sigma_x^2 \|\log(\exp(-\mathbf{v}) \circ \exp(\mathbf{v}_c))\|_{L_2}^2 + 1/\sigma_d^2 \|\nabla \mathbf{v}\|^2$ is given by linearising the first term using the BCH formula and replacing the second term with the infinite sum Tikhonov regulariser. We could then replace the Gaussian regulariser by an elastic-like one, in a consistent way. The proposed elastic regulariser amounts to filtering the correspondence velocities by the elastic-like kernel:

$$\mathbf{v} = \left(G_\sigma Id + \frac{\sigma^2 \kappa}{1 + \kappa} HG_\sigma \right) \star \mathbf{v}_c = G_{\sigma, \kappa} \star \mathbf{v}_c \quad (4.1)$$

where $\sigma^2 = 2/\sigma_d^2$, HG_σ is the Hessian of the Gaussian kernel G_σ and $G_{\sigma, \kappa}$ is the elastic-like vector filter. In this formulation, $\kappa > 0$ penalises the global compressibility, and setting $\kappa = 0$ gives the Gaussian filter used in the LogDemons algorithm.

4.2.1.4 Incorporating Strong Incompressibility in the Myocardium

Incorporating incompressibility into the LogDemons consists in constraining the velocity fields \mathbf{v} to be divergence-free. Demons optimisation step is not modified, as it optimises

\mathbf{v}_c only, but demons regularisation energy is now optimised under the divergence-free constraint, which amounts to minimising the Lagrange function:

$$P(\mathbf{v}, p) = \frac{1}{\sigma_x^2} \|\mathbf{v}_c - \mathbf{v}\|_{L_2}^2 + \int_{\Omega} \sum_{k=1}^{+\infty} \frac{Q_{el}^k(\mathbf{v})}{\sigma_x^2 \sigma_d^{2k}} - \frac{2}{\sigma_x^2} \int_{\Omega} p \nabla \cdot \mathbf{v}. \quad (4.2)$$

where Q_{el}^k is the k^{th} order isotropic differential quadratic form (IDQF) of a vector field \mathbf{v} defined by $Q_{el}^k(\mathbf{v}) = \alpha_k \delta_{i_1 \dots i_k} \mathbf{v}_{i_{k+1}} \delta_{i_1 \dots i_k} \mathbf{v}_{i_{k+1}} + \beta_k \delta_{i_1 \dots i_k} \mathbf{v}_{i_{k+1}} \delta_{i_2 \dots i_k} \mathbf{v}_{i_1}$. In this equation, the Lagrange multiplier p is a scalar function of the Sobolev space $H_0^1(\Omega)$ that vanishes at infinity. The second term is the elastic-like regulariser that leads to the filter previously mentioned. We refer the reader to [Mansi et al., 2011a] for details.

Optima of (4.2) are found by solving $\delta_{\mathbf{v}} P(\mathbf{v}, p) = 0$:

$$\mathbf{v} + \sum_{k=1}^{\infty} \frac{(-1)^k}{\sigma_d^{2k}} (\alpha_k \Delta^k \mathbf{v} + \beta_k \Delta^{k-1} \nabla \nabla^T \mathbf{v}) = \mathbf{v}_c - \nabla p \quad (4.3)$$

with $p = 0$ at the domain boundaries $\delta\Omega$. The divergence of (4.3) under the optimal condition $\nabla \cdot \mathbf{v} = 0$ yields the Poisson equation $\Delta p = \nabla \cdot \mathbf{v}_c$ with 0-Dirichlet boundary conditions, which can be solved independently of \mathbf{v} to get p . The right hand side of (4.3) is thus the L_2 projection of \mathbf{v}_c onto the space of divergence-free vector fields. Computationally, the divergence-free constraint on the velocity fields is enforced by smoothing the velocity field then projecting onto the space of divergence-free velocity fields. This is theoretically the same as projecting onto the space of divergence-free velocity fields then smoothing the results since convolution and derivatives commute (up to issues at the boundary).

Algorithm 1 summarises the main steps of the method. Implementation of this algorithm is described in the following section. A more thorough description of the derivations of the previous equations can be found in [Mansi et al., 2011a].

Algorithm 3 iLogDemons: Incompressible Elastic LogDemons Registration

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

- 1: **loop** {over n until convergence}
 - 2: Compute the update velocity: $\delta \mathbf{v}^n$ (see [Mansi et al., 2011a]).
 - 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: $\mathbf{v}^n \leftarrow Z(\mathbf{v}^{n-1}, \delta \mathbf{v}^n)$ (see [Vercauteren et al., 2008]).
 - 5: Elastic-like regularisation: $\mathbf{v}^n \leftarrow G_{\sigma, \kappa} \star \mathbf{v}^n$ (see [Mansi et al., 2011a]).
 - 6: Solve: $\Delta p = \nabla \cdot \mathbf{v}^n$ with 0-Dirichlet boundary conditions.
 - 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 $\phi^{-1} = \exp(-\mathbf{v})$.
-

4.2.2 Implementation

The algorithm has been implemented using ITK and the open source implementation of the LogDemons algorithm [Dru and Vercauteren, 2009]. The Poisson equation (which is

Input parameters:	Echo	Cine	Tag
Multi-resolution levels (frame-by-frame registration)	3	2	2
Multi-resolution levels (refinement step)	2	1	1
Number of iterations / level	100	100	100
Sigma (update field) in <i>mm</i>	0.5	0.5	0.5
Kappa (update field) in <i>mm</i>	0	0	0
Sigma (stationary velocity field) in <i>mm</i>	0.5	2	2
Kappa (stationary velocity field) in <i>mm</i>	1	1	1
Incompressibility update field (0-Disable,1-Enable)	0	0	0
Incompressibility velocity field (0-Disable,1-Enable)	1	1	1

solved at the incompressible domain) is discretised on the image grid using finite difference schemes [Simard and Mailloux, 1988] as the incompressible domain Γ may be of irregular shape. Image gradients are computed with periodic boundary conditions over the entire image domain [Dru and Vercauteren, 2009] and the Gaussian filters are implemented with ITK recursive filters.

Despite the additional constraints, the complexity of the algorithm remains reasonable with respect to the LogDemons algorithm. Demons update velocity is computed at each voxel. The elastic-like filter is computed using Gaussian convolutions, therefore 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 .

The algorithm 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 divergence and gradient operators are linear in the number of voxels. The Poisson Equation is solved at each iteration using iterative solvers like GMRES [Saad, 2003].

The codes are written in C++ and require as input the fixed image file and moving image file, as well as optional input of the mask image file, and registration parameters. The parameters used in the registration are summarised in the table below. These values were chosen based on tests performed on similar data-sets that concluded that the key parameter of interest is σ , which defines the weight of the Gaussian smoothing of the velocity field (in mm). The original voxel size of the images are $0.67 \times 0.68 \times 0.58$ for echocardiography, $1.25 \times 1.25 \times 8$ for cine-MR and $0.96 \times 0.96 \times 0.96$ for tagged-MR. The choice of σ is generally based on the voxel size to be around 1-2 times the largest original voxel size. Given the large difference in voxel size for cine-MR σ was a trade-off between the largest and smallest voxel size. More levels were used for the echocardiography sequences to speed up convergence of the simulation. This could also be done for the cine-MR sequences but was not considered necessary in this case. For the tagged-MR sequences, increasing the number of multi-resolution levels can remove the tags from under-sampling.

4.2.3 Image Pre-Processing

In order to apply the algorithm to the different data types, some pre-processing is needed to prepare the data. The method is defined in a way such that the user can give as input a region (which we define as a binary image with value 1 in the incompressible region and value 0 outside) where the incompressibility constraint is imposed. This region is defined at one time-frame only (end diastole). If no input is given the entire image is constrained to be incompressible, otherwise the user can turn off the incompressibility constraint (giving the standard LogDemons algorithm). Therefore, in order to use the iLogDemons algorithm, we need to define the region where we impose the incompressibility constraint by delineating the left ventricle myocardium using image segmentation tools (since in this case we are interested in the deformation of the left ventricle). Note that for the cine-MR sequences we segmented also the right ventricle since it is clearly visible in all the images and provides added information to the registration.

Myocardium Segmentation to Define the Incompressible Region For each image sequence we used an interactive 3D segmentation tool that builds a 3D mask image and mesh. Control points are added by the user to define the inside, outside, and border of the region, from which a 3D mesh is constructed using an implicit variational surfaces approach. The tool is included within the CardioViz3D software package available for download¹. For further details on the tool see [Mansi, 2010]. We segmented the LV endocardium and the LV epicardium and then applied arithmetic tools to obtain the LV myocardium image. We then dilate the resulting mask to ensure that the full myocardium is covered and to avoid possible boundary effects. The incompressibility domain is shown in yellow for each of the the imaging modalities (see Fig. 4.1). A screenshot of the segmentation tool is shown in Fig. 4.2.




Figure 4.1: The incompressibility domain shown on a cine-MR image (left), tagged-MR image (center) and echocardiography image (right). This domain defines where the incompressibility constraint is enforced in the registration algorithm.

¹<http://www-sop.inria.fr/asclepios/software/CardioViz3D/>

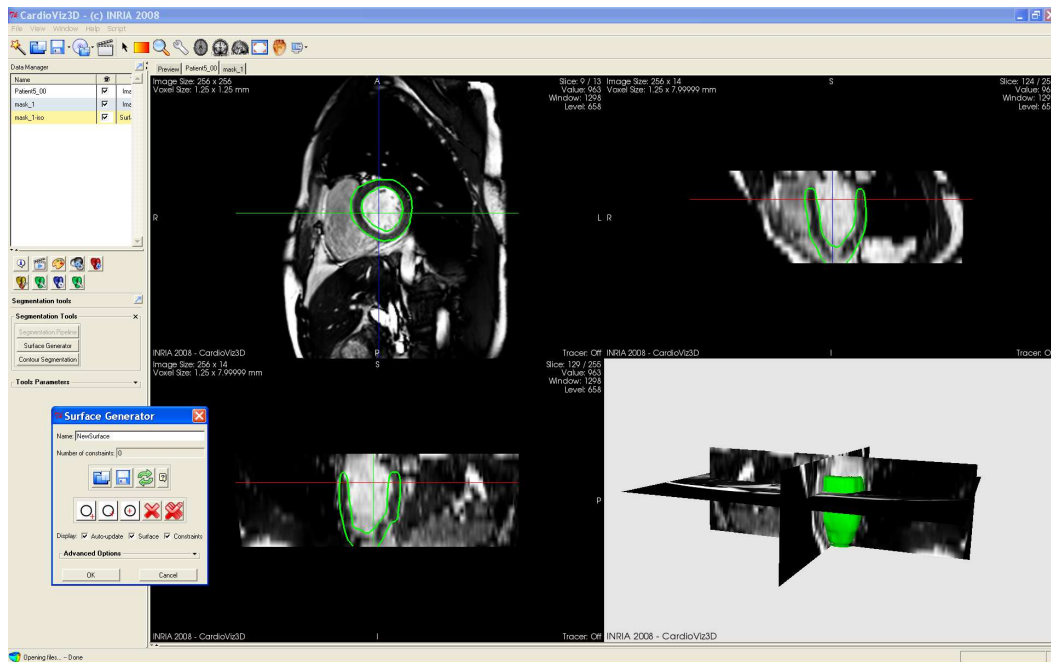


Figure 4.2: A screenshot of the interactive segmentation tool in CardioViz3D which can be downloaded from <http://www-sop.inria.fr/asclepios/software/CardioViz3D/>. The tool requires the user to place control points, from which a surface is built using implicit variational surfaces approach.

Isotropic Resampling The cine-MR images have anisotropic voxel sizes. To correct for this, we re-sampled the voxels to be isotropic in all directions. Isotropic voxel size improves the registration since the transformation is defined on a grid with enough resolution to avoid "aliasing" effects (as is true for any demons algorithm). The echocardiography and tagged-MR image sequences had already isotropic voxels.

Contrast Enhancement To enhance the image contrast we clamped the tails of the grey-level histogram to exclude the 1st and 99th quantiles. The grey level intensities were then normalised for each slice using a fixed scale. This was done for each image in the sequence independently. An example of the before and after image is shown in Fig 4.3. This processing also reduced the effects of tag fading, thus further improving registration results.

4.2.4 Application to Challenge Data

The algorithm was applied to a data-set of 15 volunteers and one phantom, each with cine-MR, tagged-MR and echocardiography images. The resulting deformation fields for each modality are included in the motion tracking challenge. To demonstrate the performance of this method we show the results for each modality for one patient from the data-set of 15 volunteers as well as the results for the phantom data. Similar results were obtained for the remaining volunteers.




Figure 4.3: Original image (left) and processed image (right) after histogram clamping and normalisation to improve image contrast.

4.2.4.1 Results for Echocardiography Sequences

The method was first applied to echocardiography image sequences. In this case, the images show well the endocardium (inside the heart) but the epicardium is difficult to see, particularly in the free wall. However, the motion is more apparent in the echocardiography sequences than in cine-MR due to the speckle that is "stitched" to the muscle and thus follows it as the heart deforms, though this speckle is consistent only between few time frames. Figure 4.4 (first two rows) shows one patient image at full contraction (systole) with the mask propagated using the deformation field computed in the registration overlaid on the image and similarly for the phantom. The masked deformation field is shown on the patient and phantom at full contraction to illustrate the direction and magnitude of motion. In each case the registration captures the expected longitudinal contraction, and circumferential twisting of the ventricle. We can also observe that, although it is difficult to distinguish clearly the epicardium for this modality, the algorithm is able to produce reasonable strain curves, as shown in Fig 4.5.

4.2.4.2 Results for Cine-MR Image Sequences

The algorithm was applied to the short-axis cine-MR images. These images show clearly the myocardium, though there is little information in the apex due to too few slices in the through plane. The algorithm is able to capture a realistic motion of the myocardium, as shown in the middle two rows of Fig 4.4. The strain curves for cine-MR are underestimated mainly due to lack of texture information in the images but show the expected trends (increase in strain towards peak systole, followed by decrease at rest (see Fig 4.5).

4.2.4.3 Results for Tagged-MR Image Sequences

As expected, the tagged-MR registration captures the twisting motion of the myocardium very well, this is particularly evident in the phantom (see Fig 4.4 bottom row second to the right), as well as the longitudinal contraction. The strain curves for the tagged-MR data

shown in Fig 4.5 show a reasonable trend, however the standard deviation over the given regions is high in this case.

4.2.5 Strain Estimation

The strain curves in each of the 17 AHA regions in each of the radial, circumferential and longitudinal directions were computed for each of the modalities. The strain was computed using the 3D Lagrangian finite strain tensor

$$E(x) = \frac{1}{2}[\nabla\mathbf{u}(x) + \nabla\mathbf{u}^T(x) + \nabla\mathbf{u}^T(x)\nabla\mathbf{u}(x)] \quad (4.4)$$

for the estimated displacement $\mathbf{u}(x)$ from the iLogDemons registration at the spatial positions x . The computed strain tensors were then projected onto a local prolate coordinate system as described in [Mansi et al., 2011a].

The strain curves for each modality in each direction are shown in Fig 4.5. The curves show a good consistency between the modalities in respect to curve trends, with the strain rising to a peak in the middle of the cycle at peak systole, and decreasing back towards zero (note that the curves are not temporally synchronised). The curves for the echocardiography sequence show a good agreement to those previously found for cine-MR and tagged-MR presented in [Mansi et al., 2011a]. However, the curves for the cine-MR sequence show less consistency with previously published results in [Moore et al., 2000], as they are under-estimated in all directions. Possible reasons for this could be too much smoothing, a lack of texture information, poor image resolution or errors in the tracking. The curves of the standard deviation among the zones shown are similar to the mean curves shown in green, which displays the consistency among the AHA regions which is expected in healthy subjects with synchronised movement among the regions. Note that here we exclude the apical regions since the apex is not clearly visible in all images.

4.2.6 Discussion

In general, this method provides reasonable results for tracking the myocardium in the three modalities. In particular, the method gives good results for the echocardiography sequences for both the tracking and estimation of strain even given data with poor visibility and little structural information. The method is particularly useful for cardiac motion tracking due to the fact that it can be applied to the imaging modalities that are most commonly used in cardiology.

4.2.6.1 Incompressibility Constraint

We discuss here the advantages and disadvantages of enforcing the incompressibility constraint in the myocardium. The constraint was integrated into the LogDemons method initially to be used on cine-MR sequences, which are known to exhibit only apparent motion in the image. For this reason, it seemed natural to constrain the myocardium to be incompressible to reduce the number of unknowns to force a circumferential and longitudinal deformation when there is a radial contraction/expansion. In the case of echocardiography sequences and tagged-MR sequences, there is more texture information in the




Figure 4.4: Top row: Long axis and short axis views of echocardiography images for one patient (left two columns) and the phantom (right two columns) shown at full contraction overlaid with the mask deformed by the deformation computed using iLogDemons. Second row: Two views of the computed deformation field (normal of intensities and vectors) shown only in the mask region for one patient image (left columns) and the phantom (right columns). Similarly for cine-MR (third and fourth rows) and tagged-MR (fifth and sixth rows). For each modality a realistic motion is obtained (rows one, three and five), as well as the desired direction and magnitude of motion (rows two, four and six), particularly for the phantom. In particular, the longitudinal motion captured by the algorithm can be seen by the vectors pointing downwards towards the apex in the long axis views of rows two, four and six, and the circumferential motion can be seen in the short axis views where the vectors appear to be wrapping around the muscle to an extent.




Figure 4.5: Strain curves in the radial (left), circumferential (middle) and longitudinal (right) directions for echocardiography sequence (top row), cine-MR sequence (middle row) and tagged-MR sequence (bottom row) for one subject. Mean (solid line) and standard deviation (dashed line) are shown for one patient in green, and the mean and standard deviations for systolic strain reported in [Moore et al., 2000] are shown in blue. Note that the curves have not been temporally synchronised. We can see that the magnitudes are under-estimated, though the curve trends are consistent with what is expected with a peak strain at peak contraction in the radial direction, minimum strain at peak contraction for circumferential and longitudinal strain.

image that aids in capturing this motion (speckles in echocardiography images, tag grids in tagged-MR). Nonetheless, for the purposes of the challenge we applied the method to all modalities to analyse the results.

Myocardium Segmentation Since the method requires a mask of the myocardium, we segmented these prior to running the algorithm using the tool described in Section 4.2.3. The incompressibility constraint relies heavily on the accuracy of the segmentation, therefore errors in the tracking can arise due to mis-segmentation of the tissue. This is a problem in particular for the echocardiography and tagged-MR images, which are known to be hard to segment given the poor image quality, poor visibility of the myocardium and noise from the top of the cone in echo images. However, this is the case for any method using localised incompressibility constraint to track the myocardium. In this work, we used a binary mask

for the myocardium. To avoid possible problems related to the boundary conditions, we dilated the mask by 2 voxels.

Constrained Incompressibility vs. Compressibility A common point of discussion for constraining the myocardium to be incompressible is that the myocardium is not in fact fully incompressible. In the literature, the myocardium is observed to have a volume change of around 5% [Glass et al., 1991]. In the case of the LogDemons algorithm, there is no constraint on the compressibility of the myocardium, which results in up to 30% volume change, compared to less than 7% of numerical volume change for the iLogDemons algorithm (see Fig 4.6). Therefore, while the iLogDemons algorithm may under-estimate the volume change in general, with the unconstrained LogDemons algorithm it can be greatly-overestimated. Furthermore, improved strain curves were obtained in [Mansi et al., 2011a] compared to those computed from the LogDemons algorithm. Hence, the incompressibility constraint is a useful prior for cardiac motion tracking, though penalising rather than constraining the compressibility may be more physiologically realistic.




Figure 4.6: Average jacobian determinant in each of the 17 AHA regions for the LogDemons algorithm (top row) and iLogDemons (bottom row) for each of the modalities (echo-left column, cine-MR - centre column, tagged-MR - right column). The iLogDemons algorithm constrains the compressibility to be less than 7% for each modality compared to up to 30% compressibility for LogDemons.

4.2.6.2 Field of View

In some of the sequences in the challenge data-set, the myocardium was on or very close to the border of the image, particularly in the tagged-MR sequences which have a very narrow field of view. How the image and the deformation are treated at the boundary of the image (extrapolated to invisible data) is a key problem in most registration algorithms. Currently the iLogDemons algorithm works in such a way that the intensities on the border of the image are extrapolated outside the image in a given region.

4.2.7 Conclusion

The iLogDemons algorithm was applied to a data-set of 15 subjects and one phantom each with an echocardiography, cine-MR and tagged-MR image sequence. This method was developed for the heart to model elasticity of the tissue and incompressibility in the myocardium. The results show that given few changes in the input parameters, the method is able to retrieve realistic motion of the heart as well as reasonable strain curves for each of the three modalities and is thus a versatile registration algorithm for cardiac motion tracking. However, future work is needed to further analyse the incompressibility prior, possibly including a change in the way the prior is incorporated into the model by means of a penalisation of the compressibility rather than the current method of constraining the velocity fields to be divergence-free.

4.3 Application to Synthetic Echocardiography Sequences

In this section, we evaluate the iLogDemons algorithm for the STACOM 2012 cardiac motion tracking challenge. This algorithm was previously applied to the STACOM 2011 cardiac motion challenge to track the left-ventricle heart tissue in a data-set of volunteers. Even though the previous application showed reasonable results with respect to quality of the registration and computed strain curves; quantitative evaluation of the algorithm in an objective manner is still not trivial. Applying the algorithm to the STACOM 2012 synthetic ultrasound sequence helps to objectively evaluate the algorithm since the ground truth motion is provided. Different configurations of the iLogDemons parameters are used and the estimated left ventricle motion is compared to the ground truth motion. Using this application, quantitative measurements of the motion error are calculated and optimal parameters of the algorithm can be found.

4.3.1 Introduction

Understanding cardiac motion dynamics through the heart beat is fundamental for providing useful insights into cardiac diseases. Analysing medical images is one way to better understand the complex dynamics of the heart and in recent years, cardiac motion tracking algorithms have been developed to attempt to estimate the observed motion. We refer the reader to [Mansi et al., 2011a] for the state of the art on cardiac motion tracking. A cardiac motion tracking challenge was introduced in the STACOM 2011 MICCAI workshop which allowed participants to apply algorithms to a given data-set of healthy volunteers with cine-magnetic resonance, ultrasound, and tagged-magnetic resonance image sequences. In this work we describe the application of the incompressible log-domain demons algorithm (iLogDemons for short) to a set of synthetic ultrasound image sequences for which the ground truth deformation is known and provided for training within the STACOM 2012 MICCAI cardiac motion tracking challenge. From this we are able to compute the error between the ground truth and the estimated deformation for the training data.

[Note: the description of the iLogDemons method from [Prakosa et al., 2012a] has been removed as it is already described in Sec. 4.2]

4.3.1.1 Cardiac Motion Tracking Strategy

We initialise the registration of the template image $T_i(\mathbf{x})$ at frame i to the reference image $R(\mathbf{x})$ with the concatenation of the previous frame ($i - 1$) to reference velocity field $\mathbf{v}_{T_{i-1} \rightarrow R}$ and the current-to-previous frame velocity field $\mathbf{v}_{T_i \rightarrow T_{i-1}}$ by $Z(\mathbf{v}_{T_{i-1} \rightarrow R}, \mathbf{v}_{T_i \rightarrow T_{i-1}})$ with Z is the BCH operation, as a strategy to track the myocardium (cf. Fig. 4.7) [Mansi et al., 2011a]. The final registration is always calculated to the same end diastolic reference image $R(\mathbf{x})$.

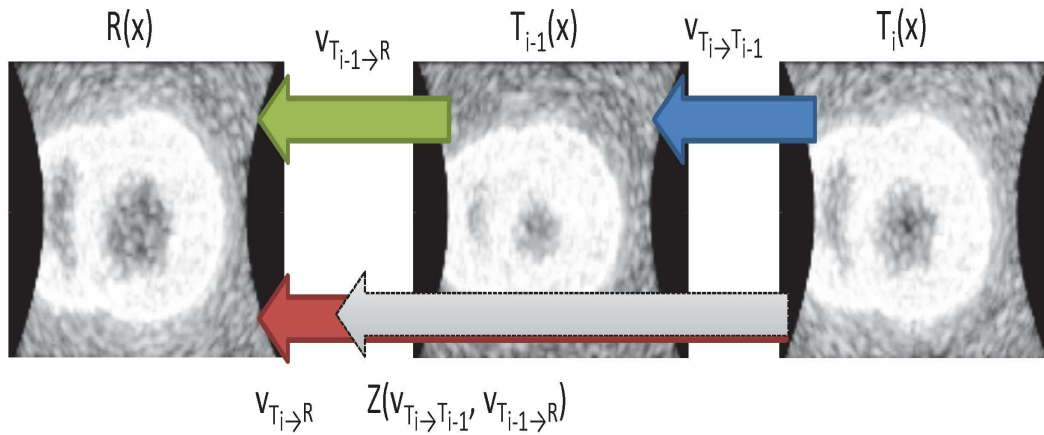


Figure 4.7: The concatenation of the velocity field $\mathbf{v}_{T_i \rightarrow T_{i-1}}$ and $\mathbf{v}_{T_{i-1} \rightarrow R}$ using the BCH formula is used to initiate the registration of the template image $T_i(\mathbf{x})$ to the reference image $R(\mathbf{x})$.

4.3.2 Application to Challenge Data

4.3.2.1 Algorithm Parameter Setting

We used the standard parameters that were used previously in [McLeod et al., 2012a]. However, since the ground truth motion is available for the synthetic ultrasound sequence provided, we also tested different parameters of the iLogDemons as described in Table 5.5.1.1.

iLogDemons non-rigid registration was previously applied to the STACOM 2011 challenge data-set [Tobon-Gomez et al., 2013, McLeod et al., 2012a]. It showed reasonable results in term of the alignment of the registered frames in the cardiac sequence with the reference end diastolic image. Using the estimated transformations, it could also track the myocardium along the cardiac cycle. The calculated strain curve was also comparable to literature for healthy strain values [Moore et al., 2000].

Input parameters:	Value
Multi-resolution levels (frame-by-frame registration)	3
Multi-resolution levels (refinement step)	2
Number of iterations / level	100
σ_f update field in <i>mm</i>	0.5
κ_f update field in <i>mm</i>	0
σ stationary velocity field in <i>mm</i>	1 or 1.5 or 2
κ stationary velocity field in <i>mm</i>	1
Incompressibility update field (0-Disable,1-Enable)	0
Incompressibility velocity field (0-Disable,1-Enable)	1 or 0

Table 4.1: iLogDemons parameters used in the application.

4.3.2.2 Simulated Ultrasound Cardiac Sequence Data

The simulated data-set consisted of 10 synthetic ultrasound sequences with 23 frames per case, with image spatial resolution of $267 \times 355 \times 355$, and isotropic voxel size of 0.33 mm. For each sequence, the left ventricle (LV) is almost fully visible while the right ventricle is only partially visible in the ultrasound acquisition cone. To compensate for the part of the LV which is out-of-window region, we artificially expanded the acquisition pyramid. The boundary voxels were copied to fill this region and additional noise was also added. The data-set contains different motion and deformation patterns (normal, LBBB, RBBB, pacing) with the ground truth deformation provided as the deformation of volumetric meshes in a cardiac cycle (See [De Craene et al., 2013] for further details on the synthetic data-set).

4.3.2.3 Application to the Synthetic Data

In order to find the optimal parameters of the algorithm that are able to handle large deformations, we processed the first case of the ultrasound synthetic data-set since it simulates normal heart motion with large contraction. We launched the parameters that were used previously in [McLeod et al., 2012a] to the full resolution data-set. We also applied our algorithm on down-sampled images to reduce the computational time. We down-sampled the data to a resolution of $88 \times 117 \times 117$ with isotropic voxel size of 1.02 mm. The computation time of the whole sequence processing was reduced from the order of days to hours. The current implementation can be optimised to handle large volumes by improving the memory access scheme since the addition of computation time of current implementation is not caused by the addition of computational complexity. One configuration of parameters was tested for both the full and down-sampled data to verify the accuracy of the down-sampled registration compared to the full-resolution registration and found very small differences in the results (cf. Fig. 4.8). Other configurations of the key parameters were tested on the down-sampled data.




Figure 4.8: The registration error (calculated using the method described in Section. 4.3.2.4) of the full resolution and down-sampled dataset of the first case are compared. They show relatively small difference.

4.3.2.4 Quantitative Evaluation

Displacement Error To evaluate quantitatively the performance of each set of the parameters used for the iLogDemons with incompressibility on the velocity field set to 0 or 1, we calculated the ground truth displacement vector field from the deformation of the provided simulated meshes. We rasterised the displacement vectors to the image $\mathbf{u}_{GT}(\mathbf{x})$ in order to be able to compare them to the iLogDemons estimated displacement field $\mathbf{u}_e(\mathbf{x})$. The norm of the difference of the two vector fields $\|\mathbf{u}_{GT}(\mathbf{x}) - \mathbf{u}_e(\mathbf{x})\|$ is calculated. The global mean of this values over the whole left ventricle are calculated for each time frame in the cardiac cycle (cf. Fig. 4.9). Based on Fig. 4.9, the parameter $\sigma = 1.5$ without the incompressibility constraint gives the lowest maximum error for the first case. We calculated the LV volume of the ground truth deformed meshes in a cardiac cycle and we observed that the current electromechanical model is not incompressible. Fig. 4.10 shows the mean and standard deviation of the LV myocardium volume change in a cardiac cycle for the whole data-set. There is a 10% change of volume during the maximum contraction. In Fig. 4.11, we compare the ground truth displacement vector for each American Heart Association (AHA) region of the left ventricle. We compare it to the iLogDemons estimated displacement vector and calculated the difference for each AHA segment. Fig. 4.11 also shows the error for the basal (regions 1-6), mid (regions 7-12) and apical (regions 13-17) regions. More error is observed in the apical region since the longitudinal motion of the apex toward the base changes the intensity of the apical region.

The result for the whole data-set processing is shown in Fig. 4.12. As also shown in Fig. 4.11 for the first case, the registration of each frame to its previous frame gives small error which is less than one voxel size. For the frame to reference result, we observe that there is an error accumulation during the maximum contraction.

Strain Estimation From the iLogDemons estimated displacement field $\mathbf{u}(x)$, we computed the strain tensor and projected it to the local radial, circumferential and longitudinal directions. The strain tensor was calculated using the 3D Lagrangian finite strain tensor




Figure 4.9: The mean and standard deviation of the displacement error calculated on the whole left ventricle for varying values of σ for the first case.




Figure 4.10: The mean and standard deviation of the LV volume change of the ground truth deformed meshes during a cardiac cycle. Current electromechanical model is not incompressible since there is a 10% of volume change during the maximum contraction.

$E(x) = \frac{1}{2}[\nabla \mathbf{u}(x) + \nabla \mathbf{u}^T(x) + \nabla \mathbf{u}^T(x)\nabla \mathbf{u}(x)]$. The mean and standard deviation of the strain estimation of the whole data-set is shown in Fig. 4.13. The result using incompressibility has more realistic range of value (from -15% to 25%) of the estimated strain compared to the one without incompressibility (from 150% to 300%).

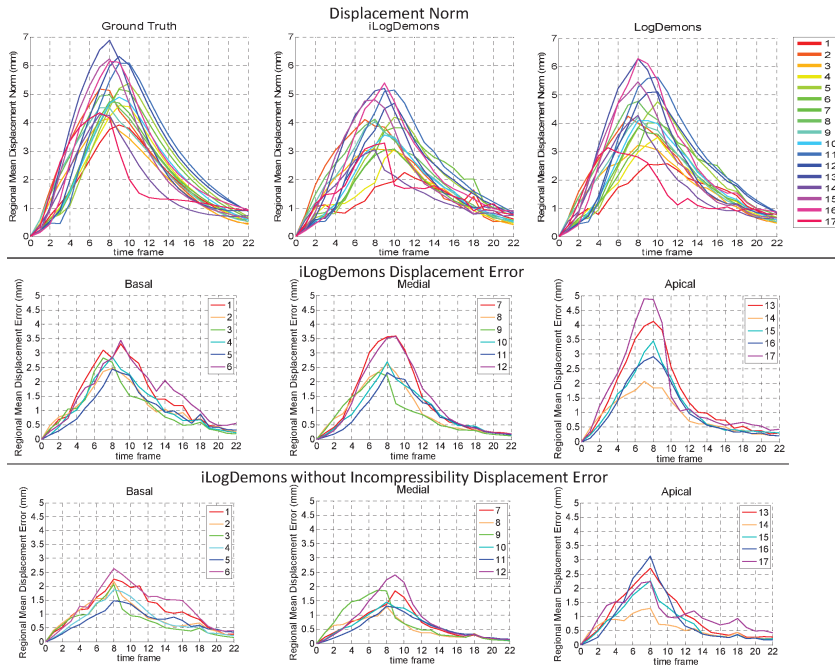


Figure 4.11: The comparison of the ground truth, incompressible and non-incompressible iLogDemons estimated LV displacement norm for the first case on each American Heart Association (AHA) region. In both cases, $\sigma = 1.5$ was used. The mean displacement error is also calculated on each AHA region.

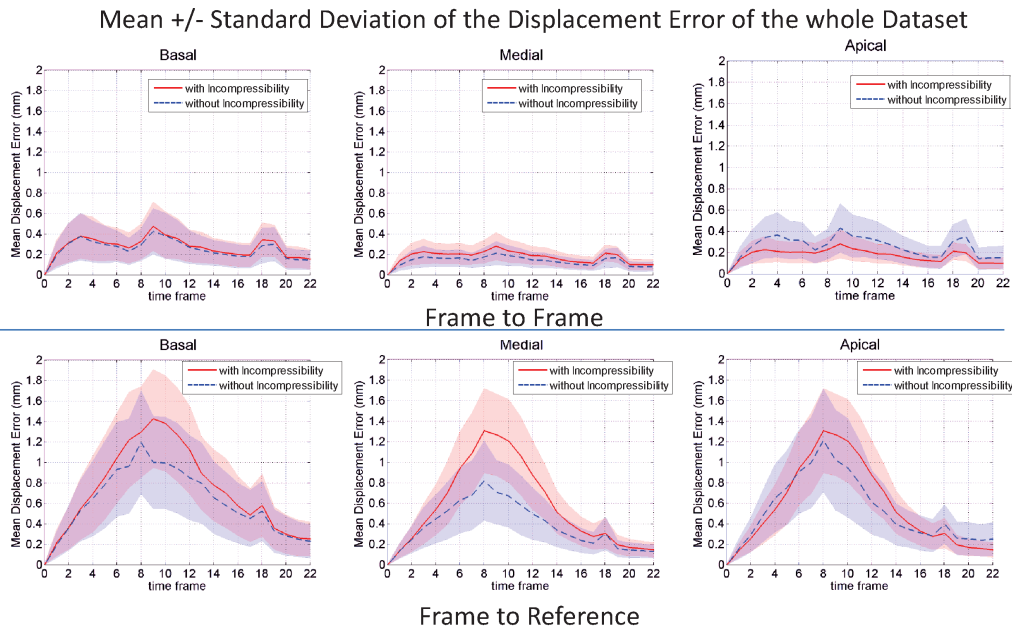


Figure 4.12: The displacement error of the whole training data-set




Figure 4.13: The mean and standard deviation of the estimated strain for the whole training data-set with and without incompressibility constrain. Incompressibility constraint gives more realistic range of value of the estimated strain (from -15% to 25%). This range is shown as black horizontal lines on the result without incompressibility.

4.3.2.5 Myocardium Tracking

Qualitative evaluation of the algorithm is done by comparing the contour of the simulated mesh at the frame with maximum contraction with the deformation of the end diastolic mesh using the iLogDemons estimated displacement field at the same frame for the first case. Reasonable agreement of the contours can be observed in Fig. 4.14, which indicated that the algorithm is able to capture realistic deformations, even in the case of a synthetically simulated sequence.

4.3.3 Discussion

This evaluation shows that the iLogDemons with and without the incompressibility constraint were able to recover the simulated motion in the ultrasound synthetic sequence with reasonable accuracy. It is worth noting that the current electromechanical model is not incompressible, therefore enforcing incompressibility in the registration algorithm naturally does not improve the results, in comparison to the iLogDemons method without the incompressibility constraint. Furthermore, we also found that increasing or decreasing the sigma value does not always improve the result since the best value that we found here is $\sigma_u = 1.5$ while $\sigma_u = 1$ and $\sigma_u = 2$ do not yield significantly better results.




Figure 4.14: Myocardium tracking result for the first case is shown (red for iLogDemons and purple for iLogDemons without incompressibility) and compared to the simulated ground truth (blue) at the time frame 8 which is at the maximum contraction. The tracking result follow the contour of the ground truth, indicating that the algorithm is able to capture reasonably well the dynamics of the motion.

4.3.4 Conclusion

The iLogDemons algorithm was applied to a data-set of synthetic ultrasound sequence with different motion and deformation pattern. The algorithm was able to reasonably estimate the ground truth deformation of the model. Since the provided data-set were created using an electromechanical model which is not incompressible, the incompressibility constraint does not improve the result. However, the incompressibility constraint gives more realistic range of estimated strain value. Future work is needed to deal with the error accumulation during the maximum of contraction.

Anatomically Constrained Regional Motion Model for Cardiac Strain Analysis

Contents

5.1	Chapter Overview	74
5.2	Introduction	74
5.2.1	Cardiac Motion Tracking	75
5.2.2	Polyaffine Image Registration	77
5.2.3	Incompressible Image Registration	77
5.2.4	Aim and Chapter Organisation	78
5.3	Background: Polyaffine LogDemons	79
5.3.1	Log-Euclidean Polyaffine Framework	79
5.3.2	SVF Projection onto the Space of LEPT's	80
5.3.3	Polyaffine Weight Functions	81
5.4	Cardiac-Specific Polyaffine LogDemons	82
5.4.1	Incompressibility Penalisation	83
5.4.2	Neighbouring Region Homogeneity of Parameters	84
5.4.3	Final Cardiac PolyLogDemons Algorithm	85
5.4.4	Cardiac-Specific Priors	85
5.4.5	Image Driven Spatial Priors	87
5.5	Experiments	89
5.5.1	Left Ventricle Motion Tracking	89
5.5.2	Bi-ventricular Cardiac Motion Tracking	94
5.5.3	Myocardial Strain Estimation	94
5.6	Discussion	96
5.6.1	Accuracy of the Method	97
5.6.2	Model Variable Selection	97
5.6.3	Pre-processing	98
5.6.4	Polyaffine Region Definition	98
5.6.5	Computation Time	98
5.7	Conclusion	99
5.8	Appendix: Numerical Implementation	99

In the previous chapter, state-of-the-art methods for cardiac motion tracking were studied. In this chapter we propose a new method for cardiac motion tracking based on the Polyaffine and LogDemons methods. A cardiac-specific motion tracking algorithm is proposed to obtain deformations that describe the motion along the cardiac cycle with few parameters to describe the motion while maintaining reasonable accuracy of the registration. The proposed method is quantitatively analysed on the data-set from Chapter 4, and compared to the iLogDemons method and the TDFFD algorithm applied to the same data-set. This chapter is a starting point towards addressing the second key question of this work on how to model cardiac motion over time at a population level to be able to capture the typical motion dynamics observed for a given population, by making use of a data-set of several subjects drawn from the same population. This chapter is in preparation for submission to a journal and extends on the work published in McLeod et. al [McLeod et al., 2012b].

5.1 Chapter Overview

The PolyLogDemons algorithm is an efficient non-linear registration method for computing the smooth fusion of local affine transformations into a global diffeomorphic transformation using a given stationary velocity field. This method has the advantage of describing a dense transformation with potentially few parameters. However, since the parameters are unconstrained, there is no control over the regional volume change or the degree to which neighbouring regions are allowed to vary. We propose cardiac-specific PolyLogDemons regression with all the advantages of the PolyLogDemons, as well as a regional incompressibility penalisation, regularisation between neighbouring regions and image driven spatial priors on the LogDemons velocity. Furthermore, we propose cardiac-specific weight functions by taking advantage of the prolate spheroidal coordinate system rather than standard Cartesian coordinates. The proposed method was applied to a synthetic cine-MRI sequences with known ground truth deformation to test the effect of each parameter on the displacement error. Cardiac motion tracking was applied to a data-set of cine MRI sequences of 15 healthy volunteers to estimate the 3D myocardium strain of these subjects. The results indicate that the proposed method obtains physiologically reasonable deformations at a level of accuracy consistent with state-of-the-art methods, with a much lower number of parameters.

5.2 Introduction

Cardiac motion tracking is a complex task but can provide useful insight for clinicians to aid in therapy planning. By better understanding the motion dynamics for normal subjects and comparing to the motion observed in pathological cases, more indications could be given for determining the suitability of a given treatment option. In order to be able to make such comparisons in a clinically meaningful way, we require a robust method to describe the motion observed along the cardiac cycle that ideally relies on a small number

of meaningful parameters for clinically relevant regions of the heart.

Cardiac imaging (such as Magnetic Resonance (MR), Echocardiography (ECHO), Computed Tomography (CT)) is becoming more and more widely used in clinical practice as a means of visualising the heart due to the ability to capture 3D motion dynamics. Tracking the cardiac tissue over the cardiac cycle at a regional level can provide useful insight for clinicians to guide with diagnosis, prognosis, and therapy planning [Young, 2006]. However, doing so is not a straightforward task. Manually tracking a given set of points for all frames in an image sequence is not impossible, but can be time consuming and is greatly subject to bias. Furthermore, it is not always easy to determine corresponding points in different frames. Tagged magnetic resonance imaging (tagged MRI) offers a convenient method for cardiac tissue tracking given that the intersecting tagged planes can act as given control points, reducing the problem of tracking from the image level to point-wise tracking [Park et al., 1996, Schaerer et al., 2010]. However, tagged-MRI is not currently routinely used in clinical practice, and accurately tracking control points in an automatic manner is an ongoing topic of research. For this reason, several non-linear registration algorithms have been proposed to compute the deformation along the sequence as a quantitative measure of the motion. Rigid registration, while being simple and easy to implement, is usually insufficient in capturing the complex dynamics, therefore non-rigid registration is more commonly used.

In general, non-linear registration involves estimating the deformation ϕ for a given fixed reference image F and a moving image M ($F, M : \mathbb{R}^3 \rightarrow \mathbb{R}$) that minimises the difference between the fixed image and the deformed moving image using the sum of squared differences for example: $\sum_{x \in \Omega} (F(x) - M(\phi(x)))^2$, for voxel coordinates x in the image domain Ω . Cardiac motion tracking generally involves computing the deformation from all image frames to a reference frame to essentially describe how each image differs from the reference image. In order to preserve the inherent structure of the heart, the deformations should be constrained to be diffeomorphic (to prevent unwanted structural changes such as folding), as well as preserving the volume of the muscle to remain within a physiologically reasonable range. Furthermore, in order to be able to compare normal and pathological motion in a quantitative manner, describing the motion by a small number of parameters can not only improve the robustness of the estimation but also allows meaningful analyses to be drawn between subjects.

5.2.1 Cardiac Motion Tracking

Analysing cardiac deformation at a local and/or global level can help to identify indicators of heart disease. Tracking the tissue with physical markers is invasive and therefore not possible in a clinical setting. Thus, noninvasive methods have been developed in recent years to track the tissue by using non-linear registration.

In response to the growing interest in cardiac motion tracking algorithms, two cardiac motion tracking challenges have been conducted as a part of the STACOM workshop of the MICCAI conference. The first challenge involved phantom and real image sequences from 3D ECHO and tagged and cine MRI [Tobon-Gomez et al., 2013]. The second challenge was focussed on tracking motion in synthetically generated ECHO sequences with

known ground truth [De-Craene et al., 2012]. Each challenge was open to anyone wanting to test and validate their algorithms. Several different methods were applied to the data including those based on B-spline transformations [Wang et al., 2011], [Heyde et al., 2012], free-form deformations (FFD) [De-Craene et al., 2011], [Piella et al., 2012], phase-based registration [Tautz et al., 2011], and optical-flow methods [McLeod et al., 2012a], [Prakosa et al., 2012b], [Somphone et al., 2012], [Alessandrini et al., 2012]. A number of review articles summarising the state-of-the-art in cardiac motion tracking have been presented, the most recent of which is Wang and Amini [Wang and Amini, 2012] (see references therein). Of the different methods currently available, each has their own advantages and disadvantages.

Voxel-based methods, which essentially model the correspondences between voxels in pairs of images, have been widely used for cardiac registration. Of these, the iLogDemons algorithm [Mansi et al., 2011a] was developed specifically for cardiac registration by incorporating physiological constraints in the model such as elasticity and incompressibility of the tissue. This method, as with the LogDemons algorithm [Vercauteren et al., 2008] (which it is based on), models the correspondences between images at a local level and hence the deformation is parameterised at every voxel. This can be on the order of millions for standard cardiac MR images. Similarly, for the method of Tautz et al. [Tautz et al., 2011], which uses phase-based registration with quadrature filters to track cardiac motion from tagged MR and ECHO, the final deformations are represented by dense fields parameterised at the voxel level.

Methods parameterised by B-spline transformations have a lower number of parameters since the deformation is parameterised by the motion of control points. Several spatio-temporal methods have been proposed that consider the transformations as B-splines to register 3D ECHO sequences [De-Craene et al., 2011] and 3D tagged MRI sequences [Chandrashekhara et al., 2004a]. In [Shi et al., 2012] the information from both tagged and non-tagged MR image sequences is used to enhance the image registration accuracy by modelling the transformations by FFD's and spatially weighting the different information before combining to a common framework. While B-spline methods have the advantage of representing the transformation with a reduced number of parameters, these parameters are generally defined at arbitrary locations since the control points of the B-spline model are usually defined on a grid. In [Heyde et al., 2012], the grid was adapted to be more physiologically relevant by using an anatomically oriented coordinate system, rather than a grid in Cartesian coordinates. This improves the accuracy of the motion tracking, though the total number of parameters still relies on the sampling of control points, which is a trade-off between the accuracy and maintaining a low number of parameters. Furthermore, the parameters may not be translatable between subjects to allow direct comparisons between the motion.

We believe that simple models that can accurately represent the deformation with a small number of meaningful parameters could provide a more insightful basis for motion analysis. The Polyaffine model provides such a framework and has been applied in many applications for medical image analysis. A summary of these is described in the following section.

5.2.2 Polyaffine Image Registration

The Polyaffine framework, first proposed by Arsigny et. al [Arsigny et al., 2005], has the nice advantage of characterising the deformation from one image to another by a set of affine transformations. In this way, the deformation is defined at a regional level, though the transformation itself is dense, and therefore defined at every point. The model has been further applied in the field of medical image analysis for a number of different applications including registration of articulated structures [Martín-Fernández et al., 2009], registration of fiber bundles from diffusion MRI [Ziyan et al., 2007], [Wassermann et al., 2011], brain MRI registration [Taquet et al., 2011], cardiac registration [Hansen et al., 2012], as well as for multi-modality registration in cardiovascular imaging [Zhang et al., 2007] and in prostate imaging [Cosse, 2012]. A method for estimating a non-rigid transformation from rigid parts for cervical cord image registration was proposed in [Commowick et al., 2012]. The method of [Hansen et al., 2012] has the nice property of being able to capture large deformations, however the application is limited to 2D image registration, which may miss some of the crucial 3D motion dynamics.

The PolyLogDemons method of [Seiler et al., 2012] takes the advantages of the LogDemons and Polyaffine algorithms to model the correspondences between images at the voxel-level using the LogDemons, and then to represent the resulting deformation by a greatly reduced number of parameters using the Polyaffine model. The Polyaffine regions are defined automatically from the image in a hierarchical manner which allows for a multi-level approach. However, defining the regions in this way may result in inconsistent region definitions between subjects. This method also may result in physiologically unrealistic deformations, with large volume change or great variability between the transformation of neighbouring regions. Since little volume change is observed in biological tissue such as the myocardium, it is desirable to add incompressibility as a property of the registration.

5.2.3 Incompressible Image Registration

The myocardium, as with other biological tissue, is near incompressible, with only a small amount of volume loss caused by the perfusion of blood through the tissue. Though some experimental studies have been conducted to measure the volume change, there is not a consensus in the literature, with a range of reported values between around 4-10% [Vossoughi et al., 1980], [Holzapfel and Ogden, 2009], [Yin et al., 1996]. Nevertheless, since the volume change is low, the tissue is agreed to be near-incompressible. Hence, this important property has been included in a number of registration algorithms, by adding incompressibility constraints on the displacement or divergence-free constraints on the velocities.

The first of these was the method of Song and Leahy [Song and Leahy, 1991], where both a divergence-free and an incompressibility constraint were successfully applied in the context of using optical flow to perform motion analysis from CT images. Since this is a voxel-based method, computing correspondences between voxels in pairs of images, the deformation field is described at the voxel level and hence with a very large number of parameters. In [Mansi et al., 2011a], an incompressibility constraint was added to

the LogDemons algorithm [Vercauteren et al., 2008] by projecting the estimated velocity field to the space of divergence-free velocity fields. In this way, the volume is preserved, however the hard constraint means that the volume is constrained beyond physiological ranges, with less than 1% volume change over the cardiac cycle. This was also the case in [Bistoquet et al., 2007], though this was later transformed into a near-incompressible rather than fully incompressible constraint [Bistoquet et al., 2008]. The disadvantage of [Bistoquet et al., 2008] is that the deformation is defined at a given set of nodes. Given that this number could be low, the number of parameters needed to represent the deformation may be small, however these are highly dependent on the choice of nodes, and thus may not be easily translatable between different patients for the purpose of direct comparison. Incorporating an incompressibility constraint into the registration when B-splines are used to represent the transformation was proposed in [Rohlfing et al., 2003] for breast image registration. This method has a similar disadvantage in that the deformation is defined on a set of control points, typically on an arbitrary grid, and therefore the parameters of the deformation are not meaningful in a way that could be directly compared between different subjects.

5.2.4 Aim and Chapter Organisation

In this chapter we describe the extension of the PolyLogDemons algorithm [Seiler et al., 2012] to cardiac specific registration for motion tracking, with the aim of providing a robust and accurate method for tracking cardiac tissue with a small number of parameters. The proposed method incorporates cardiac specific prior information into the model via the use of clinically relevant region definition for the Polyaffine regions, and anatomically oriented weight functions to follow the shape of the ventricle. This method takes advantage of the efficiency of the LogDemons algorithm to compute the deformation between images, as well as the low-dimensionality of the Polyaffine model to represent this deformation. By incorporating physiological information into the model, we hope to not only improve the overall accuracy of the model, but also to be able to obtain more physiologically plausible deformations. Moreover, since the definition of the regions and weight functions are consistent from one subject to another, the results are reproducible over different subjects. Spatial priors on the input velocity field are added, to give confidence on velocities in regions of high gradient in the image and to reduce the motion in the epicardial regions (since the epicardium moves less than the endocardium). The method is applied to a given set of cardiac-specific regions and fused over these regions using anatomically grounded weight functions.

The main contribution of this method is a framework for motion tracking that provides physiologically reasonable deformations that can be described by a small number of regional parameters (due to the choice of using a small number of Polyaffine regions), where the transformations are defined everywhere regardless of the sparsity of the grid. Moreover, the transformation is driven by a set of intelligible parameters; translation, rotation, strain and shear. An early version of this method was presented at the STACOM MICCAI workshop in 2012 ([McLeod et al., 2012b]) to describe the added incompressibility and regularisation terms to the PolyLogDemons minimisation, and the use of anatomically rel-

evant weight functions was described in [McLeod et al., 2013b]. The contributions of this work include:

- Extended description of the methodology
- Thorough testing of the model variables
- Added spatial priors to give more physiologically reasonable deformations
- The implementation of a multi-resolution scheme
- Extension to bi-ventricular tracking
- Quantitative validation

The remainder of the chapter is organised as follows: Sec. 5.3 reviews the Poly-LogDemons algorithm of Seiler et. al [Seiler et al., 2012] with a brief overview of the Polyaffine framework of Arsigny et. al [Arsigny et al., 2009] and the LogDemons algorithm of Vercauteren et.al [Vercauteren et al., 2008]. In Sec. 5.4, a reformulation of the least squares error solution is described which enables the incompressibility penalisation and neighbouring region regularisation to be incorporated into the linear least squares problem in a simple manner. Cardiac-specific priors are incorporated into the model via curved weight functions to replace the Cartesian-based weights from [Seiler et al., 2012], and consistent and clinically relevant region division of the left ventricle. Image-based spatial priors are included in the model to give weighting to the input velocity vectors depending on the gradient of the image, and to prevent large deformations in the epicardial regions. The method is quantitatively validated on synthetic and real cine MRI sequences in Sec. 8.5 as well as the estimation of myocardial strain from the real sequences obtained from healthy volunteers.

5.3 Background: Polyaffine LogDemons

In this section, we briefly describe the Polyaffine LogDemons algorithm proposed by Seiler et. al [Seiler et al., 2012], which forms the basis of the proposed method. The Polyaffine LogDemons algorithm is designed to estimate a Polyaffine model from an observed stationary velocity field (SVF) by projecting the given field onto the linear space of Log-Euclidean Polyaffine transformations (LEPT). The space of LEPT's is described in Sec. 5.3.1, followed by the regression method in Sec. 5.3.2.

5.3.1 Log-Euclidean Polyaffine Framework

The Polyaffine framework was introduced by Arsigny et. al [Arsigny et al., 2005] in the context of medical image analysis to fuse transformations in a smooth manner by numerically solving ordinary differential equations (ODE). The computation of the solution to the ODE's was improved in Arsigny et. al [Arsigny et al., 2009] and Commowick et. al [Commowick et al., 2008] to obtain a fast Log-Euclidean Polyaffine framework.

The parameters of the Polyaffine model can be defined for points x in homogeneous co-ordinates as

$$\log(T) = \log \begin{pmatrix} A & t \\ 0 & 1 \end{pmatrix} \stackrel{\text{def}}{=} \begin{pmatrix} M \\ 0 \end{pmatrix} \quad (5.1)$$

where \log is a principal matrix logarithm, A is the linear part of the transformation (accounting for the shearing, rotation and scaling), t its translation, and M a 3×4 log-affine transformation matrix. Velocity vector fields can be associated to affine transformations by

$$\vec{v}(x) = M\tilde{x}, \quad (5.2)$$

where $\vec{v}(x)$ is a 3×1 vector, $\tilde{x} = [x \ 1]^T$. For each segment the affine transformations are fused to a global velocity field using the Polyaffine model (see Fig. 5.1):

$$\vec{v}_{poly}(x) = \sum_i \omega_i(x) M_i \tilde{x}, \quad (5.3)$$

where $\omega_i(x)$ is a weight function term defined for each region i at every point x (further described in Sec. 5.3.3).

5.3.2 SVF Projection onto the Space of LEPT's

The affine matrices M_i 's can be estimated by projecting an observed SVF to the space of LEPT's by solving a linear least squares problem with respect to the observed velocity field $\vec{v}(x)$ within a given region of interest Π (represented by vectors in Fig. 5.2). As shown in [Seiler et al., 2012] the log affine parameters M_i can be estimated by the following least-squares regression:

$$C(M) = \int_{\Pi} \left\| \sum_i \omega_i(x) \cdot M_i \tilde{x} - \vec{v}(x) \right\|^2 dx,$$

using the Frobenius inner product $\|W\|^2 = Tr(W^T W)$:

$$\begin{aligned} &= \int_{\Pi} Tr \left(\left(\sum_i \omega_i(x) M_i \tilde{x} - \vec{v}(x) \right)^T \left(\sum_j \omega_j(x) M_j \tilde{x} - \vec{v}(x) \right) \right) dx \\ &= Tr \left(\sum_{i,j} \int_{\Pi} (\omega_i \omega_j \tilde{x}^T M_i^T M_j \tilde{x}) dx - 2 \sum_i \int_{\Pi} \omega_i(x) \tilde{x}^T M_i^T \vec{v}(x) dx \right. \\ &\quad \left. + \int_{\Pi} \vec{v}(x)^T \vec{v}(x) dx \right) \\ &= \|\vec{v}(x)\|^2 - 2 \sum_i Tr(BM_i^T) + Tr(M \Sigma M^T) \end{aligned}$$

where $M = [M_1 M_2 \dots M_3]$, $B_i = \int_{\Pi} \omega_i(x) \cdot \vec{v}(x) \cdot \tilde{x}^T dx$ and $\Sigma_{ij} = \int_{\Pi} \omega_i(x) \cdot \omega_j(x) \cdot x \cdot x^T dx$. For which the optimum is given at $\nabla C_M = 0$ to obtain the following system for M :

$$M = B \Sigma^{-1}. \quad (5.4)$$

The Polyaffine LogDemons algorithm is summarised in 4.




Figure 5.1: The Polyaffine framework to perform a smooth fusion between neighbouring affine transformations to obtain a smooth global transformation. For example, shown here is the Polyaffine fusion of two rotations, obtained from two neighbouring regions of the left ventricle.

5.3.3 Polyaffine Weight Functions

The weights $\omega_i(x)$ play an important role in the Polyaffine model and can greatly affect the results of the registration. They are used to model the spatial effect of a given region. A number of different ways of defining the weights is possible, as long as the fusion is sufficiently smooth to ensure that the final global transformation remains diffeomorphic. They can be defined in a straightforward way by a simple Gaussian function as

$$\omega_i(x) = -\exp\left(\frac{\kappa}{2}(x - \bar{x}_i)^T \phi_i^{(-1)}(x - \bar{x}_i)\right), \quad (5.5)$$

with \bar{x}_i the barycentre (centre point) of zone i and ϕ_i the corresponding covariance matrix, as in [Seiler et al., 2012]. The parameter κ controls how much influence neighbouring




Figure 5.2: The Polyaffine LogDemons framework to project a given stationary velocity field (SVF) to the space of Log-Euclidean Polyaffine transformations (LEPT).

Algorithm 4 Polyaffine LogDemons Algorithm

Require: Initialisation: Let $\vec{v}^0 = v_{poly}(x)$.

loop {over n until convergence}

- Compute the update velocity: $\delta\vec{v}^n$ given \vec{v}^{n-1} .

- Update the correspondence velocity field: $\vec{v}^n \leftarrow Z(\vec{v}^{n-1}, \delta\vec{v}^n)$ using Demons “forces”.

- Estimate affine transformation of each segment from \vec{v}^n by solving (5.4).

- Let $\vec{v}_{poly}(x) \leftarrow \sum_i \omega_i(x) M_i \vec{x}$

return \vec{v} , $\phi = \exp(\vec{v})$ and $\phi^{-1} = \exp(-\vec{v})$.

regions have on each other. This gives independent transformations for $\kappa > 1$ and smoother transitions between regions for $\kappa < 1$.

Alternatively, more sophisticated weight functions can be developed, as in [Martín-Fernández et al., 2009], where the weight functions were modified to be suited for the registration of articulated structures. The above Gaussian weight function was generalised to include modified distance terms suited for articulated registration. In this chapter we use Gaussian weights within the myocardium.

5.4 Cardiac-Specific Polyaffine LogDemons

In this section, we describe the contributions of this work to build on the PolyLogDemons method described in the previous section. Before incorporating new terms into the model, the standard PolyLogDemons solution needs to be re-written in terms of vectors, as described in the following section. Once the solution has been reformulated, an incompressibility penalisation (Sec. 5.4.1) and regularisation term (Sec. 5.4.2) can be easily incorporated into the model. Cardiac specific priors can be incorporated into the model via the choice of meaningful region definitions (Sec. 5.4.4.1), anatomically oriented weight functions (Sec. 5.4.4.2) and image driven confidence on the input velocity (Sec. 5.4.5).

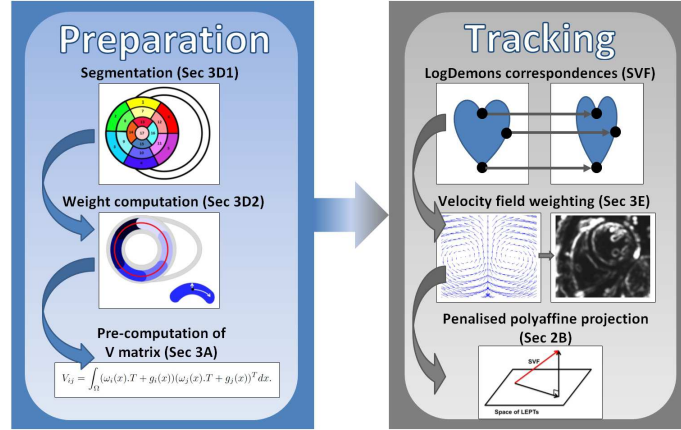


Figure 5.3: The proposed pipeline for cardiac displacement computation divided into two key steps: **preparation** and **tracking**. The **preparation** requires segmentation of the ventricle(s) at the reference frame and division into the Polyaffine regions, followed by computation of the weights and pre-computation of the V matrix (used for the incompressibility penalisation). Cardiac motion **tracking** is performed pair-wise between each frame in the cycle and the reference frame by computing the correspondence field using LogDemons, weighting this field, then projecting that onto the constrained Polyaffine space.

5.4.0.1 Rewriting the Least Squares Minimisation

In order to be able to incorporate the incompressibility and regularisation terms in the PolyLogDemons framework, the solution to the least squares minimisation given in Eq. 5.4 needs to be reformulated in terms of vectors in order to be able to perform the matrix (or rather vector) multiplication needed in the following sections. This can be done simply by vectorising the B and Σ matrices and adding a Kronecker product. Doing so, we obtain the equivalent least-squares minimisation, written in terms of vectors:

$$C(M) = \text{vect}(M)^T (\Sigma \otimes Id_3) \text{vect}(M) - \text{vect}(M)^T \cdot \text{vect}(B) \quad (5.6)$$

with the solution

$$\text{vect}(M) = (\Sigma \otimes Id_3)^{-1} \cdot \text{vect}(B), \quad (5.7)$$

equivalent to the solution in Eq. 5.4, where $\text{vect}(\cdot)$ is the standard matrix vectorisation, \otimes is the Kronecker Product.

5.4.1 Incompressibility Penalisation

In order to control the compressibility of the myocardium to be within physiological ranges, an added penalisation term is needed. Given that a transformation is incompressible if its Jacobian determinant is equal to one (and the log Jacobian is equal to zero), for an infinitesimal transformation $T = Id + \vec{v}_{poly}$ with $\nabla T = Id + \nabla \vec{v}_{poly}$:

$$\det(\nabla T) = \det(Id + \nabla \vec{v}_{poly}) = \text{Tr}(\nabla \vec{v}_{poly}) + \mathcal{O}(\|\nabla \vec{v}_{poly}\|)^2, \quad (5.8)$$

where $\mathcal{O}(\cdot)$ represents higher order terms. Therefore the velocity field \vec{v}_{poly} is locally incompressible if the trace of $\nabla \vec{v}_{poly}$ is zero. A penalisation term can then be derived as:

$$\alpha \int_{\Pi} Tr(\nabla \vec{v}_{poly})^2 dx. \quad (5.9)$$

The parameter α is used to control the strength of the penalisation. Incorporating this term into the least squares minimisation gives the penalised least squares formula:

$$C(M) = \int_{\Pi} \left\| \sum_i \omega_i(x) \cdot M_i \cdot \tilde{x} - \vec{v}_{obs}(x) \right\|^2 + \alpha \int_{\Pi} Tr(\nabla \vec{v}_{poly})^2, \quad (5.10)$$

To incorporate the new term into the least squares computation, (5.9) needs to be reformulated to obtain a quadratic form of $vect(M)$. The partial derivative of the poly-affine velocity field with respect to x is

$$\frac{\partial \vec{v}_{poly}(x)}{\partial x} = \sum_i \left(\omega_i(x) M_i \begin{bmatrix} Id_3 \\ 0 \end{bmatrix} + M_i \cdot \tilde{x} \cdot \frac{\partial \omega_i(x)}{\partial x} \right). \quad (5.11)$$

Using $T = vect([Id_3; 0])$ to extract the diagonal elements from the matrix, we have

$$Tr(\nabla \vec{v}_{poly}(x)) = \sum_i (\omega_i(x) \cdot T^T \cdot vect(M_i) + g_i(x)^T \cdot vect(M_j)), \quad (5.12)$$

with $g_i(x) = vect(\nabla \omega_i(x) \cdot x^T)$. The penalisation term can then be written as:

$$\begin{aligned} \int_{\Pi} Tr(\nabla \vec{v}_{poly})^2 dx &= \sum_{i,j} vect(M_i)^T \cdot V_{ij} \cdot vect(M_j) \\ &= vect(M)^T \cdot V \cdot vect(M) \end{aligned} \quad (5.13)$$

with

$$V_{ij} = \int_{\Pi} (\omega_i(x) \cdot T + g_i(x)) (\omega_j(x) \cdot T + g_j(x))^T dx. \quad (5.14)$$

At first, it may seem possible to simplify V by considering only the first order terms:

$$V_{ij} = \int_{\Pi} (\omega_i(x) \cdot T) (\omega_j(x) \cdot T)^T dx$$

However, this amounts to penalising the trace per region, but does not take into account the directional information. Since neighbouring regions can have high deformations in opposing directions, this may causing problems in the overlap. Therefore, the higher order terms are needed.

5.4.2 Neighbouring Region Homogeneity of Parameters

We can also define a regulariser term to control how neighbouring regions influence one another. The weights $\omega_i(x)$ control how smooth the transition is between two regions, however we would also like to control how similar the affine matrices are, as an addition regularisation. To do this we can add an additional term:

$$\beta \sum_i \Omega_{ij} dist(M_i, M_j), \quad (5.15)$$

where β controls the strength of the regularisation. We define Ω_{ij} as:

$$\Omega_{ij} = \int_{\Pi} \omega_i(x) \omega_j(x) dx, \quad (5.16)$$

which is essentially the Gram matrix of the basis of spatial functions. Defining a matrix Q such that

$$Q = \left[\begin{array}{c|c} Id_3 & 0 \\ \hline 0 & \mu \end{array} \right]$$

allows to account for the different scaling between the rotation/sheering part of the affine matrix and the translation part. A distance term can then be defined to describe the distance between the matrices from two neighbouring regions as:

$$\begin{aligned} dist^2(M_i, M_j) &= Tr[(M_i - M_j)^T \cdot Q \cdot (M_i - M_j)] \\ &= Tr(M_i^T Q M_i) + Tr(M_j^T Q M_j) \\ &\quad - 2Tr(M_i^T \cdot Q \cdot M_j), \end{aligned} \quad (5.17)$$

with

$$\begin{aligned} Tr(M_i^T Q M_j) &= vect(M_i^T Q)^T \cdot vect(M_j) \\ &= vect(M_i)^T \cdot (Q \otimes Id_3) \cdot vect(M_j) \end{aligned} \quad (5.18)$$

Thus we obtain

$$\sum_{i,j} \Omega_{ij} Tr(M_i^T \cdot Q \cdot M_j) = vect(M)^T \cdot \Omega \otimes (Q \otimes Id_3) \cdot vect(M), \quad (5.19)$$

5.4.3 Final Cardiac PolyLogDemons Algorithm

Defining a matrix R as:

$$R = \Omega \otimes (Q \otimes Id_3), \quad (5.20)$$

the final least squares error incorporating both the incompressibility and regularisation term is given by:

$$\begin{aligned} C(M) &= vect(M)^T (\Sigma \otimes Id_3) vect(M) - vect(M)^T vect(B) \\ &\quad + \alpha \cdot vect(M)^T \cdot V \cdot vect(M) + \beta \cdot vect(M)^T \cdot R \cdot vect(M) \end{aligned} \quad (5.21)$$

for which we want to find the optimum by solving $\nabla C_M = 0$.

$$\nabla C_M = (\Sigma \otimes Id_3 + \alpha V + \beta R) vect(M) - vect(B). \quad (5.22)$$

Therefore the solution for M is given by:

$$vect(M) = (\Sigma \otimes Id_3 + \alpha V + \beta R)^{-1} \cdot vect(B) \quad (5.23)$$

5.4.4 Cardiac-Specific Priors

Given the formulation of the PolyLogDemons algorithm, prior information can be incorporated into the model easily through the choice of the Polyaffine regions and weights, as described in the following sections.

5.4.4.1 Anatomically Consistent Region Definition

Computing regional statistics at a population level requires a low-dimensional and consistent division of the ventricle. The American Heart Association (AHA) proposed a standardised method to segment the myocardium which can be used to anatomically define meaningful regions of the ventricle [Cerqueira et al., 2002]. The recommendation given by the AHA is to divide the left ventricle of the heart into 17 regions with six regions for the basal area (1-6), six regions for the mid area (7-12) and five for the apical areas (13-17)(see Fig. 7.1). These definitions are commonly used in clinical practice, therefore applying the method to these regions allows the analysis to be done at a level that is relevant for clinicians. Furthermore, the divisions are performed in a standardised and consistent manner from one patient to another.




Figure 5.4: Left: Bulls-eye plot of the 17 AHA regions of the left ventricle shown in relation to the right ventricle (white). Right: A 3D mesh of the left and right ventricle of one volunteer with the 17 AHA regions of the left ventricle with coloured similarly to the bulls-eye plot, as well as the right ventricle regions (13 regions total: one apical region, three layers with four regions each).

5.4.4.2 Prolate Spheroidal Polyaffine Weights

Polyaffine weights can be easily computed in standard Cartesian coordinates, however for the left ventricle in particular, this gives weights that wrap around the ventricle in a non-anatomical sense. Increasing κ in the weight function above essentially increases the amount of outside area considered in the poly-affine regularisation. Furthermore, the penalisation and regularisation terms described in the previous section rely on an anatomically meaningful description of the weights. Therefore, in order to compute physiologically meaningful weights, the above weight function can rather be computed in prolate spheroidal (PSS) coordinates (see Fig. 5.5). In this way, the weights wrap around the ventricle. These coordinates have already been used in cardiac image analysis to represent the left ventricle for mesh generation [Legrice et al., 2001], [Choi et al., 2008], segmentation [Vallet et al., 2006], for volume and surface computations [Feng et al., 2001] and for deformation reconstruction [O'Dell et al., 1995]. Written in terms of $x, y,$ and $z,$ the prolate

spheroidal coordinates (μ, ν, ϕ) can be expressed using the following equations:

$$x = a \sinh \mu \sin \nu \cos \phi$$

$$y = a \sinh \mu \sin \nu \sin \phi$$

$$z = a \cosh \mu \cos \nu$$




Figure 5.5: Regularisation direction for Cartesian coordinates (left) and prolate spheroidal coordinates (right). The PSS coordinates follow the curvature of the anatomy, compared to Cartesian which go against the curvature.

The prolate spheroidal coordinates for a given image were computed following the method of Toussaint et. al [Toussaint et al., 2013]. These tools can be found online¹. In order to perform the change of coordinates, the mesh needs to be registered to an approximating prolate ellipsoid (computed semi-automatically from the base, apex and LV-RV anterior junction points). The registration is performed using the symmetric log-demons algorithm provided in ITK [Vercauteren et al., 2008]. The PSS coordinates are then computed on the ellipsoid and transported back to the patient-specific geometry via the inverse deformation. Given the PSS coordinate for each voxel, the barycenter, covariance matrix and Gaussian function are computed in prolate spheroidal coordinates, with the output of the Gaussian function at each voxel written out to an image on the standard Cartesian grid. One weight image is computed for each region, as well as an image combining the total weight for all the regions, used for visualisation purposes in order to approximate the most reasonable values of κ . The difference between the Cartesian and PSS weights is portrayed in Fig. 5.6, where increasing the overlap between regions described with Cartesian weights has the effect of averaging out the curvature, whereas in PSS weights, this curvature is followed around the ventricle.

5.4.5 Image Driven Spatial Priors

In our preliminary work [McLeod et al., 2012a], the region of interest Π was defined as a binary mask, where $\Pi = 1$ inside the mask, and $\Pi = 0$ outside. In this way, equal weight is given to the inner part of the tissue, endocardium and epicardium. In cine MR images, however, there is little texture information in the inner part of the tissue. Therefore, the

¹<https://github.com/ntoussaint>




Figure 5.6: The difference between Cartesian and PSS weights shown from above on a spheroid and the directions of the covariance matrix shown on an oval (left - lower right corner) and curved segment (right - lower right corner). The red lines emphasise the different shapes captured by the different weights, where the Cartesian weights sum to a hexagon (left) and the PSS weights sum to a circle (right).

boundary of the tissue is the region that drives the registration, since this is the region of highest gradient. In fact, given the Gaussian nature of the Polyaffine weights, more weighting is given to the centre of the tissue than the boundaries. To account for this, we introduce spatial priors on the LogDemons correspondence field. Firstly, higher confidence is given to the vectors in regions with high gradients in the image. Secondly, we can restrict the displacement in the epicardium, where the registration tends to be dragged by the endocardial motion, though less movement is observed in the epicardium than the endocardium.

The first prior can be easily incorporated into the model by modifying the mask in which we integrate the terms in Eq. 5.4 (II). Rather than using the standard binary mask, we can incorporate a confidence on the input velocity field (the LogDemons correspondence field) by taking the smoothed gradient magnitude image of the fixed and moving images. For the reference image however, the gradient of the mask can be used since the reference mask is already defined. By taking the addition of these two images for the fixed and moving frames, we obtain an image with the added confidence between the fixed and moving images, which can be used to weight the input velocity field to give higher confidence to the velocities in voxels on the boundary of the tissue. An example for one image is shown in Fig. 5.7.

The second prior can be included in the model by down-weighting the input velocity field at the epicardium by a factor τ , where $0 < \tau < 1$ and $\tau = 1$ in regions of high gradient (thus leaving the velocity vectors unchanged). We can define τ as $(1 - (1 - C(x)) \times E(x))$, where $E(x)$ defines the epicardial border, $C(x)$ is the gradient magnitude. In practice, this can be the gradient magnitude image of the epicardial mask. The term $1 - C(x)$ is included to prevent the velocity from being restricted in regions where there is high gradient in the image (and thus in regions where we trust the velocity).




Figure 5.7: The gradient magnitude image (right) of the image on the left used as a spatial image-driven prior on the input velocity field.

5.5 Experiments

The method was tested on two data-sets, the first a synthetically generated cine MR image sequence built from a pathological image sequence. The second was a data-set of 15 healthy volunteer adults. The focus of the experiments was to test the effect of the parameters on motion tracking of the left ventricle and to quantitatively compare the method to state-of-the-art methods (Sec. 5.5.1). Bi-ventricular motion tracking was also tested briefly (Sec. 5.5.2). Myocardial strain was estimated on the healthy volunteers, as described in Sec. 5.5.3.

5.5.1 Left Ventricle Motion Tracking

5.5.1.1 Quantitative Validation on Synthetic Sequences

Quantitatively validating image registration in an unbiased way is a very difficult task given that the underlying ground truth is usually not known. Therefore, for this purpose we made use of the method of Prakosa et. al [Prakosa et al., 2012c] for generating synthetic images sequences. This approach was proposed to generate visually realistic sequences from clinical image sequences. Since the motion is generated synthetically, the ground truth displacement is available and can be directly compared to the modelled displacement to obtain quantitative measures of the displacement error. Using this, the effect of the different parameters was tested.

Image Information A sequence was created from an 84 year old male heart failure patient, with single values for the global contractility mechanical and electrophysiology parameters to simulate the motion dynamics. Steady-state free precision (SSFP) MR images of this patient were acquired using an Achieva MR Philips Medical System scanner with image size of $171 \times 171 \times 98$ voxels and re-sampled image resolution of $1.424 \times 1.424 \times 1.424 \text{mm}^3$ (original slice thickness of 8mm), with 30 frames in the cycle.

Variable	Parameters
Weights	Cartesian, PSS
Incompressibility (α)	0, 0.1, 1, 10, 100, 1000
Regularisation (β)	0, 0.1, 1, 10, 100, 1000
Weight overlap (κ)	0.01, 0.05, 0.1, 0.5, 1
Reg. param. scaling (μ)	0, 0.1, 0.5, 1, 10, 100, 1000, 10000, 100000
Multi-resolution levels	1, 2, 3, 4

Table 5.1: **Summary of Variables.** The different variables of the proposed method and the tested values for each variable

Quantitative Validation and Variable Testing We tested a range of values for the incompressibility - α , the regularisation - β , the overlap of the weights - κ , the scaling of the affine / translation components for the regularisation - μ and the different weight types, and the number of multi-resolution levels. The different values for each variable that were tested are summarised in Table. 5.5.1.1. The error was computed by taking the L^2 distance over all the regions between the ground truth displacement and the computed displacement. As a measure of the mean error over the cardiac cycle, the errors were averaged for all the frames.

The effect of the variables κ and μ was tested by fixing $\alpha = 10$ and $\beta = 10$ and varying κ (fixing $\mu = 1$) and μ (fixing $\kappa = 0.1$). The resulting errors are shown in Fig. 5.8. Four multi-resolution levels were tested ($i = 1 : 4$) for both Cartesian and PSS weights (see Fig. 5.8). For the PSS weights, three multi-resolution levels gives the optimal results. For Cartesian weights, only two levels are needed. In the case of using a multi-resolution, there is a trade-off between the precision and speed of the registration. In general, using more levels means that the registration converges quicker. However, using too many levels can bias the registration when the images are downsampled below the level of detail needed to be captured by the deformation.

In order to try to visualise the potential relationship between α and β , contour plots of the error for each combination of values are given in Fig. 5.9. For both Cartesian and PSS weights, there appears to be a smooth convergence to a global minimum around $\alpha = 1 - 10$ and $\beta = 1 - 10$. In both cases, the method appears to be robust to the choice of α and β . In the case of PSS weights, the added incompressibility (α) and regularisation (β) terms have a large effect on the displacement error. For PSS weights, there is a larger optimum, indicating greater stability with respect to the parameters, and the error is lower than for Cartesian weights.

In order to evaluate the evolution of the errors over times, the mean and standard deviation of the displacement error are plotted over time for $\alpha = 10$ and $\beta = 10$, since these appear to be the optimal values from Fig. 5.9. The errors are plotted in Fig. 5.10. The errors increase over time up until the end systolic frame, which is expected since this is the frame with largest displacement.




Figure 5.8: The effect of κ , μ and the number of multi-resolution levels on the displacement error, tested by fixing $\alpha = 10$, $\beta = 10$, and varying *kappa* with μ fixed at 1 (left) and varying μ with κ fixed at 0.1 (centre) for Cartesian (green) and PSS (blue) weights. The number of multi-resolution levels tested (right) were between 1 – 4. The optimum for *kappa* is at 0.1 and μ around 1, with 3 multi-resolution levels.




Figure 5.9: Contour plots displaying the relationship between α and β with respect to the displacement error computed from the average difference between the ground-truth displacement field and the estimated displacement field for the Cartesian weights (left) and PSS weights (right). The optimum occurs around $\alpha = 1 - 10$, $\beta = 1 - 10$.

Computation Time The computations were performed on a Dell PowerEdge 1950 dual-Xeon 5355 2.66Ghz computer. For the synthetic image sequence with resolution described in Sec. 5.5.1.1, computation of the matrix V took 6.67s. One pairwise registration took 60.24s. Registration of frame n to frame 1 is initialised with the results of the previous frame-to-reference registration (from frame $(n - 1)$) concatenated with the result of the n -to- $(n - 1)$ registration, therefore the tracking is performed sequentially, though the pair-wise registrations can be performed in parallel. Note that the computation time for both the V matrix and the pair-wise registration depends on the size of the image.




Figure 5.10: Displacement error over all the frames for $\alpha = 10$ and $\beta = 10$ showing the mean (solid line) \pm standard deviation (dashed lines) for Cartesian weights (left) and PSS weights (right).

5.5.1.2 Comparison to State-of-the-Art Methods

A first step to better understanding abnormal motion is to understand the motion dynamics observed in healthy subjects. In this way, we can quotient out the normal motion to distinguish abnormal motion patterns. To tackle this first challenge, the proposed method was applied to the STACOM 2011 MICCAI cardiac motion tracking challenge database [Tobon-Gomez et al., 2013]. The primary goal here was to quantitatively compare the method to other state-of-the-art methods.

Image Information and Acquisition Details We illustrate these tools on 15 healthy adults (3 female, mean age \pm SD = 28 ± 5). SSFP MR images were acquired using a 3T scanner (Philips Achieva System, Philips Healthcare) in the short axis view covering entirely both ventricles (12-16 slices; isotropic in-plane resolution: $1.15 \times 1.15 \text{mm}^2$ to $1.25 \times 1.25 \text{mm}^2$; slice thickness: 8mm; 30 frames). For each volunteer, 12 manually annotated landmarks of the left ventricle were tracked along the cardiac cycle by two observers on 3D tagged MR image sequences of the same volunteers. These landmarks were transported to the cardiac MR sequences using the DICOM header information. The 12 landmarks were composed of one landmark per wall (anterior, posterior, lateral, septal) for each of the ventricular levels (basal, mid and apical).

Quantitative Comparison to iLogDemons and TDFFD The STACOM 2011 cardiac motion tracking challenge was open to all willing participants. Of the four competitors that applied for the final challenge, only two submitted their methods to the cine MR sequences. These were the Temporal Diffeomorphic Free Form Deformation (TDFFD) method of De Craene et. al [De-Craene et al., 2011] and the iLogDemons method of Mansi et. al [McLeod et al., 2012a, Mansi et al., 2011a]. The principle behind the TDFFD method is to model the deformations as B-splines. The iLogDemons is rather a gradient-driven optical flow method based on the Demons method. As with the method proposed in this work, both of these methods require a mask of the myocardium as input.




Figure 5.11: Boxplot of errors per subject from the 12 4D annotated landmarks at end diastole and end systole for the Cartesian weights (left) and PSS weights (centre). The average over the population is plotted for the TDFFD algorithm, iLogDemons algorithm, and the proposed method. The proposed method is on the same order as the TDFFD and has lower errors than the iLogDemons method.

Evaluating the accuracy of the provided ground truth landmark tracking is not the objective of this work. Rather, we are interested in quantitatively comparing the registration results of the proposed method to state-of-the-art methods applied to the same sequences. Given that the ground truth computation is subject to errors from the landmark allocation itself (inter- and intra-observer bias), and in alignment between the tagged MR and cine MR sequences (subjects may move between acquisitions), and noting the large slice thickness (8mm), we are interested in evaluating how greatly the errors vary from the other methods, rather than the absolute error itself. Furthermore, we expect higher errors than for ECHO or tagged MR for instance, since there is little or no texture information in the myocardium to drive the registration [Wang and Amini, 2012].

We applied the same workflow for computing the error as described in [Tobon-Gomez et al., 2013] to be able to directly compare the errors. Based on the experiments described in the previous section, the chosen parameters were $\alpha = 1, \beta = 1, \kappa = 0.5, \mu = 1$, with PSS weights, and three multi-resolution levels. The L^2 norm error between the annotated landmarks at end diastole and end systole and the deformed landmarks at the corresponding frames was computed. The median error value of the two time points was 3.07 for the proposed method. This can be directly compared to the median values for the iLogDemons algorithm (4.82) and for the TDFFD algorithm (3.17). Therefore, the errors are of the same order as these methods. The errors are plotted in Fig. 5.11 (left) for each of the 15 volunteers. The error over all the subjects is shown in Fig. 5.11 (right) along with the boxplots for the TDFFD and iLogDemons methods for comparison. The accuracy is equivalent to TDFFD and lower than iLogDemons.

To evaluate the registration qualitatively, the AHA image was deformed by the computed deformation field and the contours of each region is overlaid on the image at end diastole (the reference frame) and end systole, for one volunteer (see Fig. 5.12). Based on these images, the registration appears to capture well the translation of the regions from peak contraction to peak relaxation, though the radial thickening of the muscle is not fully captured in the registration.




Figure 5.12: Regional registration result of volunteer 10 at end diastole (ED) and end systole (ES) for two different views showing the LV AHA regions on one slice annotated on the ED frame, and the warped regions obtained from deforming the AHA regions by the registration results.

5.5.2 Bi-ventricular Cardiac Motion Tracking

Given the connected dynamics of the right and left ventricles, capturing the bi-ventricular motion is important to capture the full dynamics. Given that the PSS weights described in this chapter are specific to the left ventricle, we apply the motion tracking only with the Cartesian weights. The annotated landmarks from the STACOM 2011 cardiac motion tracking challenge were only for the left ventricle, we validate the bi-ventricular motion tracking only qualitatively.

We applied the proposed method to the first volunteer sequence. The right ventricle was divided into 13 regions in a consistent manner to those for the LV (four basal, four mid-section, and five apical). We analysed the results qualitatively by overlaying the deformed mesh on the image sequence. In this way, we can visualise how well the contours of the deformed mesh align with the image sequence. The contours for three frames along the cycle are shown in Fig. 5.13 for three views. The first frame is the reference frame, considered to be the end diastolic frame (end of relaxation), since the sequences were gated from the electrocardiography (ECG) signal. The 10th was chosen manually as the frame corresponding to the end systolic phase of the cycle for the STACOM challenge. The 20th frame is also shown to represent the relaxation period. Since the right ventricle is very thin, there is little information in the image to drive the registration, though the proposed method is able to capture the motion of the right ventricle reasonably well, as shown by all the views of Fig. 5.13.

5.5.3 Myocardial Strain Estimation

Myocardial strain can be used as a measure of function to measure the shortening or lengthening of the tissue in different directions. However, computing the strain directly from cine MR sequences is not straightforward since there is little texture information within the myocardium that can be extracted from the image. Several attempts have been made to estimate strain from non-rigid registration, such as in [Petitjean et al., 2004], [Sinusas et al., 2001], [Ledesma-Carbayo et al., 2006], [Allan et al., 2011], [Mansi et al., 2011a]. These methods typically compute the strain




Figure 5.13: The bi-ventricular mask annotated on the first frame was deformed by the deformation field computed using the proposed method. The image sequence of volunteer 1 at frame 1 (column 1), frame 10 (second column), and frame 20 (third column) of the cycle (of 30 frames total) is shown with the contour of the deformed mask for two views (shown in the different rows). The contours align reasonably well with the image.

from the displacement field via the computation of the finite Lagrangian strain tensor:

$$\mathbf{E}(x) = \frac{1}{2}(\nabla\mathbf{u}(x)^T + \nabla\mathbf{u}(x) + \nabla\mathbf{u}(x)^T\nabla\mathbf{u}(x)), \quad (5.24)$$

where $\mathbf{u}(x)$ is the displacement at x . The strain tensor can then be projected onto a local basis in each of the longitudinal, radial and circumferential directions, following the notation in [Moore et al., 2000].

The average strain for all the volunteers was computed over the cardiac cycle for regions 1-16 (note that region 17 is excluded due to computational issues in defining the directions at the apex). The average strain over each of the AHA regions for $\alpha = \beta = 1$ is summarised in Table 5.2 and plotted over time in each direction in Fig. 5.14. The average radial strain over all subjects and all regions for $\alpha = \beta = 10$ is also plotted in Fig. 5.14 to visualise the effect of increasing the incompressibility and regularisation on strain. It appears that increasing the incompressibility and regularisation decreases the strain, in particular the radial strain. This is expected since increasing the regularisation has the effect of smoothing out local motion (and potentially smoothing out the strain). Note that for $\alpha = \beta = 10$ the displacement error is slightly higher (median error 3.12).

These values can be compared to values reported in the literature. In [Augustine et al., 2013], the authors propose a method for extracting the strain directly from cine MR using a dedicated feature tracking software. The method requires user interaction, and therefore is not optimal for computing strain on a large number of subjects,

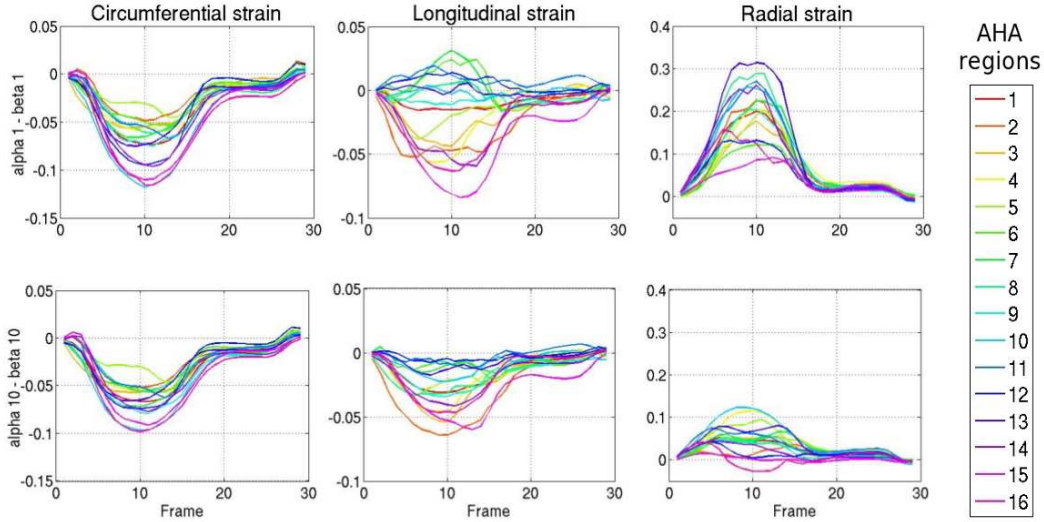


Figure 5.14: Mean regional strain of all volunteers in the circumferential (left), longitudinal (centre) and radial (right) directions over time for $\alpha = \beta = 1$ (top row) and $\alpha = \beta = 10$ (bottom row).

	Radial	Longitudinal	Circumferential
cPolyLogDemons	0.20 ± 0.08	-0.02 ± 0.03	-0.08 ± 0.03
Augustine et. al [Augustine et al., 2013]	0.25 ± 0.06	-0.19 ± 0.03	-0.21 ± 0.03
Moore et. al [Moore et al., 2000]	0.43 ± 0.19	-0.15 ± 0.03	-0.19 ± 0.03

Table 5.2: **Strain comparison.** Strain estimated from the model and compared to the values given in [Augustine et al., 2013, Moore et al., 2000].

and is subject to bias. From the study, the longitudinal, radial and circumferential strain were computed from 145 subjects. Strain was computed globally over the three ventricular layers, though in the longitudinal and radial directions less strain was observed in the apical regions. In [Moore et al., 2000], systolic strain were computed from tagged MR image sequences of healthy subjects. The reported strain values of these studies are summarised in Table 5.2. The computed strain is within the same range as those presented in [Tobon-Gomez et al., 2013] for the TDFFD and iLogDemons methods, but with higher radial strain.

5.6 Discussion

In this chapter we presented a cardiac motion tracking method that allows cardiac motion to be tracked along the cycle with few parameters needed to describe the deformation from each frame to the reference frame. The proposed method extends on the method of Seiler et. al [Seiler et al., 2012] by introducing an incompressibility penalisation by penalising

the trace of the gradient of the Polyaffine velocity field (where the trace of the gradient of the velocity field is directly related to the volume change, thanks to properties of Lie algebra). Added regularisation between neighbouring affine parameters was incorporated by controlling how neighbouring regions differ in terms of their corresponding affine matrices. This regularisation allows for a more homogeneous motion between regions, since we do not expect large variation in motion between connected regions. Cardiac-specific Polyaffine weight functions were used to allow more physiologically meaningful combination of the affine regions, and for the regularisation. Spatial priors on the input velocity vectors (computed using the LogDemons algorithm ([Vercauteren et al., 2008])) were added to give higher confidence to vectors in regions on the boundary of the tissue, where there are high gradients driving the registration. The movement of the epicardium was restricted to prevent the epicardial border from being dragged by the endocardial border that is driven by higher gradients. The accuracy of the method was quantitatively and qualitatively analysed, as discussed further in the following section. A few key points were raised briefly in this chapter including the pre-processing needed to implement the algorithm, and the current computation time. The choice of regions was briefly discussed in the previous sections. These discussions are further developed below.

5.6.1 Accuracy of the Method

Experiments on a synthetic data-set were able to establish a range of optimal values for the model parameters. Using the synthetic data-set, we were able to determine how each parameter effects the precision of the registration. Based on these experiments, we observed that both the incompressibility and regularisation terms improve the accuracy of the registration, and furthermore, that a combination of both of these terms gives better results. Using the PSS weights also gave better results than using Cartesian weights. Using a multi-resolution scheme improves the accuracy and speed of the registration. Using the openly available data-set of healthy volunteers from the STACOM 2011 cardiac motion tracking challenge, we are able to determine that the method is able to track the motion of the heart up to reasonable levels of accuracy consistent with state-of-the-art methods. Preliminary analysis of the method on cardiac motion tracking of both the left and right ventricle show promising results in terms of registration accuracy (based on qualitative analysis). Myocardial strain computed on 15 healthy volunteers indicate that the proposed method is able to capture the circumferential and longitudinal deformation reasonably well, though radial deformation appears to be underestimated with respect to expected radial strain of healthy subjects. Including both the spatial priors give higher strain in all directions than without any spatial priors (results not shown).

5.6.2 Model Variable Selection

The experiments on the synthetic sequence gave indications for a suitable range of values for each variable of the model. Based on these experiments, the model seems to be reasonably insensitive to the choice of values used for the incompressibility and regularisation variables (α and β). The number of multi-resolution levels has a small effect on the accu-

racy after the first level, but little effect between 2-4 levels. The Polyaffine weight overlap variable κ is an important variable in the model as it determines the spatial effect of each region on its neighbouring regions. For future work, a Bayesian approach could allow for more thorough investigation of the optimal values for the variables.

5.6.3 Pre-processing

The current implementation requires a number of pre-processing steps. Using the AHA segmentation to define the Polyaffine regions, we require first segmentation of the myocardium, and then division into the AHA regions is defined from an additional three points (base, apex and LV-RV junction on anterior). The motivation of this choice of region definition is discussed further in the next section. For this work these steps were performed semi-automatically, therefore restricting the large-scale application of this method. However, several methods have been proposed for automatic segmentation of the left ventricle [Rouchdy et al., 2007], and automatic detection of the base, apex and LV-RV junctions. Therefore, these pre-processing steps could potentially be automated to allow for large-scale applications. Computation of the spatial priors is fully automatic.

5.6.4 Polyaffine Region Definition

Using the proposed method, the deformation is described by 204 parameters per frame (12 affine parameters \times 17 Polyaffine regions). A key question raised from this work is whether 17 regions is enough to describe the motion of the left ventricle. In our opinion, given the results of the registration on the healthy volunteer data-set, these regions are sufficient for tracking the left ventricular motion. Other possible choices for the Polyaffine regions are available, such as the possibility to use a multi-scale approach as in Seiler et. al [Seiler et al., 2012] by starting at a coarse scale, and refining up until a desired number of regions (potentially until the voxel level). However, for this application, we want the strain and displacement to be defined in the AHA regions so we can communicate this directly to the clinicians. Given that these regions are already familiar in cardiology, it allows for more clinically meaningful analyses. Furthermore, using more regions means there will be a larger number of parameters required to describe the motion. Naturally there is a trade-off between the desired level of accuracy and the low parameterisation of the motion. We believe that the current region definition provides a suitable trade-off.

5.6.5 Computation Time

The current implementation of the method is mostly based on the ITK framework. Manipulation of the matrices (such as the V and Ω matrices) has not been optimised. Current implementation of the method requires sequential tracking with one frame at a time. Using instead an independent frame-to-reference implementation gives higher errors (median value 3.55 for the STACOM 2011 data-set), though the tracking can be performed much faster since all the registrations can be done in parallel. Given that the interest of this work is not to try to have real-time motion tracking, but rather to perform off-line population analyses of the motion, accuracy, rather than speed is preferred.

5.7 Conclusion

A method for tracking cardiac motion over the cardiac cycle based on a Polyaffine deformation model is proposed. The described method incorporates cardiac-specific prior information and image-driven spatial priors into the model. Cardiac motion tracking can be performed with a good level of accuracy (errors within the same range as current methods), and with few parameters to describe the motion (204 parameters per time frame). Given so few parameters to describe the motion that a consistently defined for all subjects, the method is robust and can be reproduced on large populations for population-based analysis and motion comparison.

5.8 Appendix: Numerical Implementation

The method is implemented fully in ITK following the open-source code of the LogDemons, available from the ITK website². The Polyaffine regression is added as a filter to regularise the correspondence field parameterised by a SVF, to replace the standard LogDemons Gaussian convolution of the SVF. The added terms in the minimisation are included within the filter as extra functions to compute the matrix V and R , and the least squares minimisation is solved with respect to these added terms. The Polyaffine regions are given as an input image with the regions represented by different scalar values to compute the weight images per region. The weight images are computed as described in Sec. 5.4.4.2, also using code written in ITK. These weight images are given as input directly to the PolyLogDemons framework.

5.8.0.1 Pre-Computation of the V Matrix

For the purposes of cardiac motion tracking, the V matrix needs to be computed only once since the regions (and thus the weights) are fixed and defined in such a way so as to encompass all trajectories of the sequence. The matrix is pre-computed and stored as a text file to be read for all the paired image registrations. Since the computation of the V matrix requires the computation of the gradient of each weight image, it is computationally expensive, so pre-computing in this way reduces the total computation time.

5.8.0.2 Multi-Resolution Implementation

Following the multi-resolution implementation of the LogDemons algorithm, a multi-resolution scheme for the proposed method was used. In this way, a coarse-to-fine registration was performed by starting at a coarse level to capture the large-scale deformation, followed by a refinement at a finer level to capture the small-scale deformation. In practice, the multi-resolution implementation converges faster since fewer iterations are needed at the finest level.

The proposed pipeline is outlined in Fig. 5.3 to summarise the key steps involved in the preparation, and tracking stages. The preparation step requires segmentation of the

²www.itk.org

ventricles and subdivision into AHA regions followed by computation of the weights and pre-computation of the V matrix used for the incompressibility penalisation. Tracking is performed by computing the correspondence field using LogDemons to obtain a SVF, weighting this field by the smooth gradient magnitude of the image, then projecting to the Polyaffine space. Tracking is performed by registration between all frames to the reference frame, with each registration initialised with the result from the previous frame. Strain and displacement analysis are performed directly on the affine parameters by considering the first 3×3 block as the strain component, and the last 3×1 column as the displacement component.

Polyaffine Parameter Analysis

Contents

6.1 Chapter Overview	101
6.2 Regional Analysis of Left Ventricle Function using a Cardiac-Specific Polyaffine Motion Model	102
6.2.1 Introduction	102
6.2.2 Regional Analysis of Polyaffine Parameters	103
6.2.3 Discussion and Future Work	104
6.3 Spatio-Temporal Dimension Reduction of Cardiac Motion for Group-Wise Analysis and Statistical Testing	105
6.3.1 Introduction	105
6.3.2 Left-Ventricular Motion Tracking Validation on a Synthetic Sequence	107
6.3.3 Group-wise Statistical Analysis of Motion	108
6.3.4 Discussion and Future Work	110

In the previous chapter, a reduced-order model for cardiac motion tracking was proposed. This model allows the transformation from one image to another in the sequence to be defined by a small number of parameters. The cardiac motion tracking algorithm was developed specifically towards enabling group-wise comparison of motion. This chapter leads towards such an analysis by comparing the transformation parameters of different subjects as a preliminary step towards the full group-wise analysis described in the following chapter. This chapter is based on McLeod et. al [McLeod et al., 2013b] and McLeod et. al [McLeod et al., 2013a] (with the explanation of the cardiac motion tracking model and testing data-set removed here to avoid repetition).

6.1 Chapter Overview

Using the cardiac motion tracking algorithm described in the previous chapter, the estimated transformation parameters are analysed first at a regional level to compare the parameter evolution of healthy and pathological subjects (Sec. 6.3). The analysis is taken a step further to then build group-wise modes of temporal and spatial variation in Sec. 6.3.

6.2 Regional Analysis of Left Ventricle Function using a Cardiac-Specific Polyaffine Motion Model

Given the complex dynamics of cardiac motion, understanding the motion for both normal and pathological cases can aid in understanding how different pathological conditions effect, and are affected by cardiac motion. Naturally, different regions of the left ventricle of the heart move in different ways depending on the location, with significantly different dynamics between the septal and free wall, and basal and apical regions. Therefore, studying the motion at a regional level can provide further information towards identifying abnormal regions for example. The 4D left ventricular motion of a given case was characterised by a low number of parameters at a region level using the cardiac-specific Polyaffine motion model described in Chapter 5. The motion was then studied at a regional level by analysing the computed affine transformation matrix of each region. This was used to examine the regional evolution of normal and pathological subjects over the cardiac cycle. The method was tested on 15 healthy volunteers with 4D ground truth landmarks and 5 pathological patients, all candidates for Cardiac Resynchronisation Therapy. Visually significant differences between normal and pathological subjects in terms of synchrony between the regions were obtained, which enables us to distinguish between healthy and unhealthy subjects. The results indicate that the method may be promising for analysing cardiac function.

6.2.1 Introduction

Cardiovascular disease still remains one of the worlds leading causes of mortality in the developed world, despite recent advances in the field. Research into congestive heart failure is increasingly providing new solutions to extend the life expectancy of heart failure patients. Given that cardiac asynchronous contraction is seen as a key characteristic of heart failure, Cardiac Resynchronisation Therapy (CRT) (where pacing leads are strategically implanted in a given region, or regions) can significantly reduce the risk of death in heart failure patients.

Common procedure for patients with heart disease is to scan the patient to obtain a 4D image sequence (using echocardiography, magnetic resonance (MR) or computed tomography (CT) scanners), for example, and / or electrocardiography (ECG) to measure heart rhythm. Image sequences allow to visualise in 3D the motion dynamics and to potentially identify thickened, ischemic, or scar regions, and asynchronous electrical conduction can be identified with ECG. Cardiac motion tracking from image sequences is a thoroughly studied area of research due to the useful insights it can provide for therapy planning (see [Tobon-Gomez et al., 2013],[Wang and Amini, 2012] and references therein for a sample of such methods).

The aim of the proposed method is two fold: first, differentiating between healthy and unhealthy subjects, and second, identifying abnormal regions of abnormal motion. To do so, a motion tracking scheme was required, with a small number of parameters describing the motion in order to obtain robust and reproducible measures. For this, the method proposed in the previous chapter, which has the advantage of incorporating anatomical and clinically meaningful prior information into the motion tracking through the use of

anatomically oriented weight functions, and clinically relevant region definition, was used. As a clinical application of this study, the method was applied to 5 heart failure patients, to determine if there was a visible difference between the healthy and heart failure subjects. A preliminary analysis of the deformation parameters (the affine matrices) indicated that the method was able to detect meaningful differences between the two groups.

[Note: the description of the cardiac PolyLogDemons method from the previous chapter given in [Mcleod et al., 2013b] has been removed to avoid repetition, as well as the full description of the STACOM 2011 data.]

6.2.2 Regional Analysis of Polyaffine Parameters

6.2.2.1 4D Cardiac MR data sequences

Healthy Volunteers: The experiments performed on the 15 healthy volunteers described in Chapter 5 were used for comparative purposes to give an indication of "normal" parameter evolution.

CRT Heart Failure Patients: To test the methods on pathological cases, we applied the method described in Chapter 5 to five adults with asynchronous contraction of the left ventricle. These patients were candidates for Cardiac Resynchronisation Therapy (CRT), all with a large QRS duration due to a Left Bundle Branch Block (LBBB). Steady-state free precession magnetic resonance images were acquired for each patient in the short axis view covering entirely both ventricles (12-16 slices; isotropic in-plane resolution: $1.42 \times 1.42 \text{mm}^2$ to $1.56 \times 1.56 \text{mm}^2$; slice thickness: 7-9mm; 30-40 frames).

6.2.2.2 Parameter Analysis

Since the transformation T from one image to another is given in terms of the exponential of the affine matrix M (see [Commowick et al., 2008] for further details), i.e $T = \exp(M)$, we can use the property derived from Lie Algebra for the Jacobian determinant: $\det(\exp(M)) = \exp(\text{tr}(M))$. Thus, we can analyse the trace of each matrix directly as a measure of volume change per region. Note that the incompressibility term is a penalisation rather than hard constraint, since we don't want to force full incompressibility. Furthermore, in order to obtain full incompressibility, a larger number of regions would be needed. Therefore, we are interested in reducing, rather than forcing to zero, the compressibility. Plotting the trace per region along time for all the healthy volunteers shows a similar deformation at similar times for all the regions (see Fig. 6.1).

Plotting the same curves for the CRT patients shows abnormal behaviour between the regions compared with the normal cases (see Fig. 6.2). Patients 1, 3 and 4 have asynchrony between the regions, which can be expected given that these patients have ischemic regions. Patients 2 and 5 have Dilated Cardiomyopathy (DCM) resulting in reduced pumping function in the left ventricle.

Using Patient 4 as a specific case study (adult female with heart failure), the trace evolution can be analysed at the scar regions. For this patient, the scar location was segmented

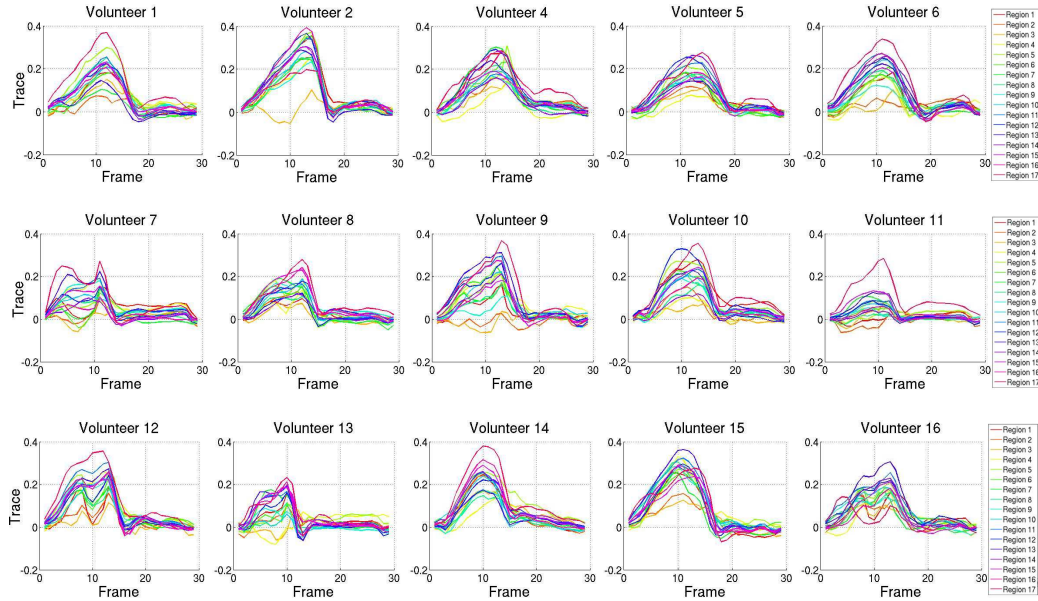


Figure 6.1: Trace per region of each affine matrix for each of the healthy volunteers plotted over time. The figures show abnormal behaviour over time of the trace between all regions for all volunteers in comparison to the normal cases.




Figure 6.2: Trace per region of each affine matrix for the CRT patients. Clearly, there is an asynchrony between the regions for all of the patients.

from the MR images. Significant mechanical asynchrony for the patient was confirmed from echocardiography, with left bundle branch block (LBBB). Two scars were observed, the first in regions 4 and 5, the second in regions 9, 10, 14 and 15 (see Fig. 6.3 left). Plotting the trace separately for each scar (see Fig. 6.3 centre and right), we see that there is a large difference in amplitude between regions 4 and 5 for the first scar, but with a single peak in the middle of the cycle, and smaller differences between the regions for the second scar (see Fig. 6.2). Since the second scar is divided amongst four regions, each region has only a small amount of scar, with mostly healthy tissue.

6.2.3 Discussion and Future Work

A recently proposed method for cardiac motion tracking was applied to a small database of 15 healthy and 5 heart failure subjects. This method presents the advantage of requiring a smaller number of parameters compared to previously reported methods. As the deformation is described with few parameters, they can be used as bio-markers to differentiate




Figure 6.3: Left: Bi-ventricle mesh of a LBBB patient sliced through to show the right ventricle (grey), left ventricle (blue) and scar location (yellow). Trace of the regions where one of the scars are located (centre), and trace of the regions where the other scar is located (right).

between healthy and unhealthy subjects. The results suggest that the proposed method can be used for the comparison of cardiac motion patterns between groups and the identification of abnormal motion areas directly from the images, without requiring any other additional measurements. It was shown that using this methods the motion parameters can be analysed at a regional level to study differences between regions of healthy or unhealthy tissue, though including the percentage of scar per region would provide more depth to this analysis to identify how much the scar region impacts the motion.

6.3 Spatio-Temporal Dimension Reduction of Cardiac Motion for Group-Wise Analysis and Statistical Testing

Given the observed abnormal motion dynamics of patients with heart conditions, quantifying cardiac motion in both normal and pathological cases can provide useful insights for therapy planning. In order to be able to analyse the motion over multiple subjects in a robust manner, it is desirable to represent the motion by a low number of parameters. We propose a reduced order cardiac motion model, reduced in space through a Polyaffine model, and reduced in time by statistical model order reduction. The method is applied to a data-set of synthetic cases with known ground truth to validate the accuracy of the left ventricular motion tracking, and to validate a patient-specific reduced-order motion model. Population-based statistics are computed on a set of 15 healthy volunteers to obtain separate spatial and temporal bases. Results demonstrate that the reduced model can efficiently detect abnormal motion patterns and even allowed to retrospectively reveal abnormal unnoticed motion within the control subjects.

6.3.1 Introduction

With the continuing high incidence of heart disease around the world in all sectors of society, there is a growing interest in cardiac modeling to aid in diagnosis, prognosis, and therapy planning. Models are continually being developed to try to simulate the different

dynamics of cardiac function. Cardiac motion tracking, for instance, is becoming more and more widely studied to try to quantify the different motion dynamics observed between normal and pathological cases.

Over the past few years, a number of methods have been proposed to track the motion of the tissue for different cardiac imaging modalities (ultrasound, cine-magnetic resonance (MR), tagged-MR etc.), with increasingly interesting results (see [Tobon-Gomez et al., 2013] and references therein). The majority of these methods focus on cardiac motion analysis of a single subject, however, population-based motion analysis may provide further information to aid clinicians. Ideally, such analyses should be robust and reproducible for different subjects drawn from different populations (such as healthy vs. pathological subjects), which requires the motion to be described by few parameters. In [Chandrashekhara et al., 2003], a statistical motion model is proposed using principal component analysis on motion fields to define a per phase and per cycle model. However, the models are used for only individual cases. A statistical atlas of motion computed from healthy volunteers was proposed by Duchateau et. al [Duchateau et al., 2011] to quantitatively analyse abnormal motion patterns and aid with therapy planning. This method was effective in characterising abnormal motion, but is limited to 2D ultrasound data. De Craene et. al later extended the method to 4D, but applied the method only to tagged MRI [De Craene et al., 2012]. Recently, a low-dimensional cardiac specific model was proposed by McLeod et. al [McLeod et al., 2012b], that computes the Polyaffine fusion over regions estimated from the log-demons deformation with added incompressibility penalisation and neighbouring region regularisation, but again this method is limited to individual cases. Though this method was able to drastically reduce the number of parameters needed to represent the transformation, we would like to go even further to reduce the number of parameters from several subjects to enable robust population-based statistics on the motion.

To that end, we propose to extend on the method of [McLeod et al., 2012b] by applying model reduction on the Polyaffine model parameters of a single subject (resp. group of subjects) to perform patient (resp. group) -wise statistics on left ventricular motion. The method of [McLeod et al., 2012b] is described briefly in the following section for review. Sec. 6.3.2 describes a validation of the registration on synthetic sequences with known ground truth, as well as patient-based statistics on the computed affine parameters. In Sec. 6.3.3 a method for performing group-wise statistics both spatially and temporally from the Polyaffine model is described and performed on 15 healthy volunteers. Decomposing the motion this way is essentially estimating a tensor decomposition of the spatial and temporal components for multiple subjects, but in a more simplified manner than tensor decomposition (which is not at all straightforward to compute). Similar decomposition was described in [Metz et al., 2012] and [Hoogendoorn et al., 2009], though these required point annotations and temporal correspondences of these points along the cycle. Enforcement of such correspondences can introduce bias since determining such correspondences is not evident. The proposed method requires no such parameterisation and rather using the images directly. The method was able to compute a patient-specific model of the motion with a reduced number of parameters, as well as building separate spatial and temporal bases at a group-wise level to understand the different aspects for a given population and to even detect abnormal motion within that population.

[Note: the description of the cardiac PolyLogDemons method from the previous chapter given in [Mcleod et al., 2013a] has been removed here to avoid repetition, as well as the full description of the STACOM 2011 data.]

6.3.2 Left-Ventricular Motion Tracking Validation on a Synthetic Sequence

In order to validate the accuracy of the registration, the cardiac motion tracking algorithm described in Chapter 5 was applied to a data-set of 5 synthetic cine-MR sequences, with known ground truth were used from the data-set of synthetic image sequences generated by Prakosa et. al [Prakosa et al., 2012c]. All sequences used for the validation were generated from a single patient sequence, by applying a bio-mechanical model to the reference image with different contraction parameters in order to simulate different motion dynamics.

6.3.2.1 Motion Tracking Validation and Comparison to iLogDemons

Given the ground truth deformation from each frame to the reference frame (end diastolic frame), the L^2 -norm error between the computed deformation and the ground truth deformation was computed for each motion simulation over the sequence. The computed errors are shown in Fig. 6.4, comparing against the error of the iLogDemons algorithm [Mansi et al., 2011a] applied to the same sequences. From these curves we can see that the proposed method gives lower errors than the iLogDemons method, which may be because affine motion is more robust.

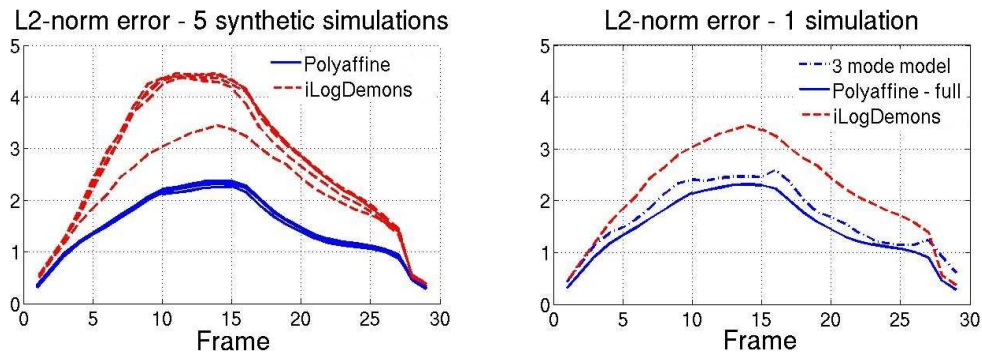


Figure 6.4: Left: Total L^2 -norm error for the proposed method and the iLogDemons method per frame for the 5 simulated sequences. Right: The L^2 norm error for one case comparing the full Polyaffine model (blue solid line), the model computed from only the first 3 modes (blue dashed line), and the iLogDemons algorithm (dashed red line).

6.3.2.2 Singular Value Decomposition on the Polyaffine Parameters

Given the low number of parameters needed to describe the deformation (204 per frame = 17 regions \times 12 affine parameters), a singular value decomposition (SVD) of the parameters can be easily computed. For a single given case, the SVD of the affine matrix is computed as:

$$SVD(M) = U \cdot S \cdot V^T, \quad (6.1)$$

where U and V are unitary matrices, S is a diagonal matrix and V^T denotes the matrix transpose of V . The matrix M can be constructed by stacking the affine parameters for all the regions in a column for all the frames, so that the final matrix will be of size 204×29 .

Using three modes, more than 95% of the variability is captured. Based on this fact, a personalised reduced model was created using the first three modes of each subject. These were then used to re-create the deformation using the Polyaffine model. The resulting error curve is shown in Fig. 6.4 (right) for a single case (similar curves can be shown for the other cases). As expected, there are slightly higher errors than the original Polyaffine model, though the error is still lower than the iLogDemons algorithm. Therefore, we are able to describe the deformation by a reduced number of parameters - now 702 vs. 5916 for the full Polyaffine model, and compared to voxel-based methods on the order of millions (image size in each of the x,y,z directions $\times 3 \times$ number of frames).

6.3.3 Group-wise Statistical Analysis of Motion

The motion tracking was applied to cine-MRI image sequences of 15 healthy volunteers obtained from the STACOM 2011 MICCAI cardiac motion tracking challenge database [Tobon-Gomez et al., 2013]. In order to be able to compare the parameters between subjects, rigid alignment to a common space was first performed using a simple landmark based registration.

6.3.3.1 Group-wise Spatial Mode Decomposition

Given a set of subjects drawn from the same population (in this case healthy volunteers), we are interested in finding a group-wise spatial basis of the motion within this population. This can be done by using again SVD, but this time on the combined matrices of all the subjects (over all the regions, and all the time frames). By stacking the matrices column-wise, we can obtain a group-wise spatial basis, each with patient-specific time loadings (see Fig. 6.5). In this way, we can consider that the expected motion parameters for all the subjects are the same since they are drawn from the same population, but with different offsets (loadings) for the time.




Figure 6.5: Pipeline for computing a group-wise spatial basis by stacking the matrices for all the patients over all the regions and all time column-wise, and performing SVD on the big matrix.

Applying this to the 15 healthy subjects, we find that 3 modes were sufficient for describing 97% of the variability. Patient-specific models can be built using these first three group-specific spatial modes with the subject-specific temporal loadings.

6.3.3.2 Group-wise Temporal Mode Decomposition

We may conversely be interested in finding a temporal basis for a given population. This could, for example, be useful in detecting abnormal motion by those with a temporal trigger significantly different from normal cases. Such a basis can be found by performing SVD on the group of affine matrices, this time stacked row-wise (see Fig. 6.6).




Figure 6.6: Pipeline for computing a temporal basis for a given population. The affine matrices of all the regions, parameters and frames for a given set of patients are stacked row-wise to perform a SVD.

By applying this decomposition to the 15 volunteers, we find that in this case 8 modes are needed to capture 90% of the variability. Given that the sequences were not fully temporally aligned, a larger number of modes is expected. Plotting the time signal for the first 8 modes only (Fig. 6.7, left), we find that the first mode corresponds well to the normal volume curve (Fig. 6.7, centre). Combining the first two modes has the effect of shifting the peak temporally (Fig. 6.7, right).




Figure 6.7: The first three temporal modes computed from the model reduction described in Fig. 6.6 (left) with estimated locations of the cardiac phases (black vertical lines). For comparison, the volume and electrocardiogram curves are shown (centre) (image modified from http://en.wikipedia.org/wiki/File:Wiggers_Diagram.png). The combination of the first two modes (right).

To test the generalisation of the temporal basis to a new subject, a leave-one-out study was performed on the 15 volunteers by learning the model on 14 volunteers and testing on

the 15th. The residual error between the original affine matrices and the model computed from 3 modes was computed for the learning set to determine the range of expected residual values, then on the testing set to evaluate the generalisation of the basis (see Fig. 6.8 left). We observed that there is one outlier in the testing set, and went back to the data to analyse why this subject does not fit the normal model. Interestingly, we found that in fact this patient had trouble during the scan and moved significantly in the scanner due to discomfort. Therefore, the observed difference could in fact be due to motion artifacts. The model was then learned for the 14 normal volunteers (excluding the outlying volunteer) and projected onto two heart failure patients with known abnormal motion. The computed residual errors were much higher than the expected range (see Fig. 6.8 right), therefore the model is able to successfully distinguish abnormal motion in a quantitative manner.




Figure 6.8: Left: Leave-one-out testing of the expected residual from using a 3-mode model rather than full Polyaffine model with learning on 14 volunteers (left box-plot) and testing on the 15th volunteer (right box-plot). Right: Computed residual on the 3-mode model for 15 volunteers used in the learning (box-plot), and residual of the projection of the 3 modes on two pathological cases (blue stars).

6.3.4 Discussion and Future Work

The current work presents patient-specific, as well as group-wise analysis of motion by approximating the motion by a cardiac-specific Polyaffine model, then performing model reduction to the Polyaffine parameters. Since the motion is described by a small number of meaningful parameters, the analysis can be done on a significantly reduced number of parameters and in a meaningful way, by de-coupling the spatial and temporal aspects. By performing the model reduction in this way, rather than directly computing principal component analysis on the affine parameters at all time frames for all the subjects (vectorising the affine and temporal components (represented by the green face of the cube of Fig. 6.6) for all subjects (represented by the orange face of the cube)), we de-couple the temporal and spatial components. We were able to detect irregular motion within the population, as well as distinguish abnormal dynamics from subjects outside the population using the temporal basis. For future work, temporally and spatially aligning the sequences first, as in [Perperidis et al., 2005], may give more meaningful bases with fewer modes needed to describe the variability. This work describes a proof of concept of the proposed method, which needs to be further validated on larger data-sets.

Regional Mean Trajectory Model of the Left Ventricle for Classification and Population-Based Motion Analysis

Contents

7.1	Chapter Overview	112
7.2	Introduction	113
7.3	Transformation Parameter Estimation using a Polyaffine Model	116
7.4	Spatio-Temporal Parameter Alignment	117
7.4.1	Temporal Resampling	117
7.4.2	Regional Temporal Alignment	118
7.4.3	Spatial Alignment	120
7.5	Mean Motion Model Construction	121
7.5.1	Mean Trajectory Model	121
7.5.2	Simulated Motion from the Mean Model	121
7.5.3	PCA Motion Modes	122
7.6	Spatial and Temporal Mode Construction using Tensor Decomposition	123
7.6.1	Tensor Decomposition Using the Tucker Method	124
7.6.2	Tensor Decomposition on Re-scaled Parameters	124
7.7	Healthy vs. Tetralogy of Fallot Motion	126
7.7.1	Mean Normal Motion Model	126
7.7.2	Abnormal Motion Identification in Tetralogy of Fallot Patients	128
7.7.3	Healthy vs. Pathological Spatio-Temporal Bases	128
7.7.4	Tensor Decomposition on P Matrices	131
7.8	Discussion and Perspectives	131
7.8.1	Parameter Realignment	132
7.8.2	Mean Motion Model Construction	133
7.8.3	Statistical Model Reduction	133
7.9	Conclusion	134
7.10	Appendix	134
7.10.1	Tucker Tensor Decomposition using Alternating Least Squares	134

This chapter uses the image registration algorithm described in Chapter 5 to perform statistical analyses of the transformation parameters by extending on the methods described in Chapter 6. We go further towards population-based analyses by proposing a mean motion model, and performing statistical analyses of the parameters to understand the observed motion patterns. Using the methods proposed in this chapter leads to answering the key question of how to model cardiac motion over time at a population level. The aim of such modelling is to be able to capture the typical motion dynamics observed for a given population, by making use of a data-set of several subjects drawn from the same population, which finalises the second part of this thesis on cardiac motion analysis. This chapter is in the process of being submitted to a journal.

7.1 Chapter Overview

Given that heart disease can result in abnormal motion dynamics over the cardiac cycle, understanding and quantifying cardiac motion can provide insight for clinicians to aid in diagnosis, therapy planning, and to determine the prognosis for a given patient. Quantifying motion patterns that are observed in healthy subjects is a first step towards extracting abnormal motion patterns characteristic of a given disease. Since there are common motion patterns observed in patients suffering from the same condition, extracting these patterns can enable a better understanding of a disease. Quantifying cardiac motion at a population level is not a simple task since images can vary largely in terms of image quality, size, resolution and pose. To overcome this, we analyse the parameters obtained from a cardiac-specific Polyaffine motion tracking algorithm. Since the Polyaffine parameters are intelligible (related to the strain / displacement), they can be interpreted directly. The parameters are temporally aligned and resampled to a common space to have the same number of parameters for all subjects, and to align the parameters according to the regional volume peak. Spatial alignment is performed to align all parameters to a common space in a prolate spheroidal sense. Once all parameters are in a common space, a mean motion model of a population (normal or abnormal) can be derived simply by taking the mean of each parameter. The mean can be defined in any space by realigning spatially in the same prolate spheroidal sense. The modes of variation are derived by performing tensor decomposition on the tensor of affine parameters over time for all subjects. Mean motion models of healthy subjects and Tetralogy of Fallot patients are computed, from which the abnormal motion patterns observed in Tetralogy of Fallot patients are extracted. Spatial and temporal bases are derived for each population to determine typical temporal triggers and typical spatial patterns for each group. Significant differences are observed between the motion dynamics of Tetralogy of Fallot patients and healthy subjects in terms of the residual differences between the mean model and the Polyaffine model.

7.2 Introduction

Cardiovascular disease is a worldwide issue, being the leading cause of death and affecting the day-to-day life of billions. Heart diseases can affect the motion dynamics of the heart over the cardiac cycle. Understanding the abnormal dynamics can thus potentially give insight into a pathology and provide information that can aid with diagnosis, therapy planning, and for determining the prognosis for a given patient. Quantitatively defining what are “normal” motion patterns is not straightforward given the complex dynamics and coupling between the ventricles and atria. In the case of patients with ischemia, the motion dynamics will be affected by the poor conductivity in these regions, and there can be differences in motion due to the different tissue properties. In patients with heart conditions that affect the shape of the ventricles, the motion dynamics can be affected by the abnormal shape to give poor pumping function in one or both of the ventricles, causing an increased workload on the heart. For patients with stenosis in the arteries, the motion dynamics can be affected by the heart trying to pump blood to get enough oxygen to the body, and from regurgitated blood causing abnormal motion dynamics. Given the heavy workload on the heart in a lifetime, maintaining healthy heart motion is crucial to maximise longevity.

The motion of the heart can be studied by placing markers on the tissue and tracking the markers over time. However, this kind of measure is invasive and not applicable on a large-scale basis or in all populations. For this reason, medical images are commonly used to non-invasively visualise the motion of the heart as it beats using magnetic resonance imaging (MRI), echocardiography (ECHO), or computed tomography (CT), for example. Using such images, the temporal evolution of the heart can be visualised in 2D or 3D. Such image sequences can be used to qualitatively analyse the motion dynamics. Some medical scanners are able to obtain global measures directly from the images such as ventricular volume, or 1D blood flow measures, however regional measures are not currently available in standard scanners. Also, more in-depth measures of cardiac function cannot be extracted directly from the images.

Therefore, quantitatively tracking the cardiac motion can provide further insights to obtain measures of cardiac function by analysing how efficiently the heart is beating and identifying patterns related to normal (healthy) motion, and patterns characteristic of a given pathology. Myocardial deformation and strain can give early indicators of poor function and disease evolution. A number of methods for tracking cardiac tissue for a single subject have been proposed using image registration techniques to compute the deformation between images in the cardiac cycle. These are further discussed in the following section.

7.2.0.1 Cardiac Motion Tracking using Non-rigid Registration

A number of non-rigid image registration techniques for cardiac motion tracking have been proposed in recent years. Reviews of these are given in [Frangi et al., 2001, Wang and Amini, 2012]. Key features of these methods include ensuring that the deformations are diffeomorphic (smooth transformations that preserve the structure of the material to prevent non-physiological transformations such as folding) and including an in-

compressibility term to prevent large volume change in the myocardium. The challenge in cardiac motion tracking of cMRI images is that the tissue is represented homogeneously in the imaging, giving little texture information to track besides the endo- and epicardial contours. Thus, it is difficult to capture the strain purely from image-driven registration algorithms, in particular the circumferential strain.

Several methods for computing spline-based representations of transformations by a series of free form deformations (FFD) have been proposed. In [Rueckert et al., 1999] a B-spline based method for image registration of breast MR images was proposed and later extended to include an incompressibility constraint [Rohlfing et al., 2003]. A cross correlation method for combining the information from non-tagged and tagged-MRI was proposed in [Shi et al., 2012]. In [De-Craene et al., 2011], a temporal diffeomorphic framework was proposed where the velocities rather than the displacement were represented by FFD's, to allow a temporal consistency in the estimated transformation. Cardiac motion analysis with the objective of abnormality detection was described in [Ledesma-Carbayo et al., 2005, Ledesma-Carbayo et al., 2006], but is limited to 2D analysis. In [Declerck et al., 1998], a 4D polar transformation was used to define the motion of the left ventricle (LV), however the motion is defined only for a given set of points. Spline-based methods in general define the motion of a set of grid nodes, where the size of the grid is variable.

Optical flow methods, originally proposed in the field of computer vision, are driven largely by image intensities, thus assuming consistency in the image intensities along the cycle. For modalities such as tagged-MRI this may not be the case. However, in the case of cMRI registration, this is a reasonable assumption since the intensities remain reasonably constant over the cardiac cycle. Optical flow methods are usually defined at the densely over the image to determine the motion at voxel-wise. A number of methods for estimating transformations based on optical flow principles in the context of cardiac motion tracking have been proposed using phase-based correlation [Tautz et al., 2010] or Demons forces [Mansi et al., 2011a] for example. These methods have the advantage of defining the transformation for every voxel in the image, however this generally means that the deformation is also parameterised at the voxel level. In cMRI sequences, this amounts to millions of parameters per frame.

Polyrigid [Arsigny et al., 2005] and Polyaffine [Commowick et al., 2008, Arsigny et al., 2009] methods have been proposed in recent years to compute a simple transformation at a regional level (rigid or affine) and fuse the regional transformations in a smooth manner to obtain a global transformation. Such methods have the advantage of defining the transformation over the whole image, but with few parameters. Traditionally, a Polyaffine model is derived based on the result of a block-matching strategy for computing the correspondences between images. In [Seiler et al., 2012] this was changed to use a log-domain Demons optical flow velocity field [Vercauteren et al., 2009] instead, which is then projected onto the space of Polyaffine transformations. This was extended in the context of cardiac motion tracking in [McLeod et al., 2013b] to include cardiac-specific prior terms to the model. This method is used in the current work to obtain a representation of the transformation that requires few parameters that can be interpreted and compared amongst different subjects and different populations.

7.2.0.2 Regional Motion Analysis

Assessing cardiac function can be done at a global or regional scale. Global indices can for example be used to measure the pressure in the ventricles, the fraction of blood ejected from the ventricles to the atria, the volume of blood in the ventricles / atria, and the mass of the tissue. Global indices cannot, however, be used to assess regional function of the heart. Regional measures, such as regional stress and strain, can be useful in detecting regions that are functioning poorly due to ischemia for example. Therefore, analysing the function at a regional level can provide vital information for clinicians in determining where to direct treatment for therapy planning purposes. By knowing normal regional patterns, it may be possible to identify abnormal regional patterns to detect regions that are functioning poorly and potentially to classify function based purely on such indices. Using non-rigid image registration and an unscented Kalman smoother, a Bayesian classifier was used to detect regional motion abnormalities in [Punithakumar et al., 2013]. Using this method, abnormal regions are classified for the full cycle, rather than taking into consideration the temporal component. Therefore, some patterns may be averaged out in a sense. Using 2D and 3D ECHO, wall motion was assessed in [Collins et al., 1999] at a regional level by assessing the motion first in 16 segments of the left ventricle using 2D ECHO, then subdividing to 36 regions in 3D ECHO. Regional motion was assessed from rest and stress cMRI in [Suinesiaputra et al., 2011] by delineating the contours of the endo- and epicardium on 2D slices, then dividing into 17 segments consistent with the standardised LV regions defined by the American Heart Association (AHA) [Cerqueira et al., 2002] and rigidly, then non-rigidly registering the contours.

7.2.0.3 Population-based Motion Analysis

In [Qian et al., 2011], myocardial strain extracted from tagged MRI was used to identify and localise regional abnormal cardiac motion patterns, by describing the spatio-temporal motion as a tensor, where the tensor contains regional strain information for 20 regions of the LV. This results in a tensor of size 60×10 (3 strain directions \times 20 regions \times 10 frames from end diastole to end systole) that is used to classify healthy / abnormal motion patterns. The authors go a step further than patient-specific abnormality detection to develop a population-based classifier, trained on the patient-specific tensors. This, however, is limited to strain analysis.

Cardiac motion was compared across subjects in [Rao et al., 2002] by first extracting the motion parameters using a B-spline model then transforming the motion field to a common space by transforming each displacement vector by the inverse Jacobian matrix of the transformation. This was extended to construct a cardiac motion atlas in [Rao et al., 2003] by simply averaging the transformed motion fields from a group of subjects drawn from the same population. A 4D statistical atlas of cardiac anatomy was developed using a B-spline model to represent the transformations in [Perperidis et al., 2004] by first performing spatio-temporal alignment [Perperidis et al., 2003]. This model was used to differentiate between normal motion and motion of patients with hypertrophic cardiomyopathy. A cardiac motion atlas was constructed in [Rougon et al., 2004] by performing principal compo-

ment analysis (PCA) on subject-specific motion fields after first spatio-temporally aligning the fields. Regional measures are then acquired by averaging over the AHA regions. A method for estimating a spatio-temporal atlas from longitudinal datasets was proposed in [Durrleman et al., 2009b] for the detection of developmental delay in children. This framework, is however, limited to the analysis of shapes, rather than images.

7.2.0.4 Aim and Chapter Organisation

Previous methods for performing population-based analyses of motion have been directed towards analysis of dense motion fields, or of regional strain. In the present work, we are rather interested in performing group-wise analysis of a small set of displacement parameters obtained from the cardiac-specific PolyLogDemons algorithm of [McLeod et al., 2013b]. Rather than performing PCA on motion fields as in [Chandrashekhara et al., 2003], or tensor decomposition of regional strain values, as in [Qian et al., 2011], we compute a mean motion model by averaging spatio-temporally aligned Polyaffine parameters, and compute spatial and temporal bases by performing Tucker tensor decomposition on the spatio-temporal parameters of a set of subjects. In contrast to the method proposed in [Xue et al., 2006] that was proposed to represent high-dimensional deformation fields in the context of brain image warping by a 3D statistical representation, we represent the deformation fields directly with the Polyaffine regression.

The Polyaffine parameters are aligned temporally by taking the regional mean peak volume, obtained from taking the trace of the affine matrix per region per subject, as was done in [McLeod et al., 2013b], and aligning to the mean of these region-wise. The parameters are then temporally resampled so that all subjects have the same number of time frames. Spatial alignment is performed on the temporally aligned and resampled parameters by bringing all parameters to a common coordinate space, and realigning spatially in a prolate spheroidal sense. A mean motion model is constructed by averaging the spatio-temporally aligned parameters. Tucker tensor decomposition is performed on the 3D tensor of affine parameters over time for all subjects (with and without removing the mean) to obtain spatial and temporal bases.

The remainder of this chapter is organised as follows. A brief introduction of the cardiac-specific PolyLogDemons algorithm of [McLeod et al., 2013b] that is used to obtain the transformation parameters that are used in the following analyses is given in Sec. 7.3. The spatio-temporal realignment is described in Sec. 7.4. The mean model construction is described in Sec. 7.5, followed by the mode construction via the Tucker tensor decomposition in 7.6. The methods are applied to create a mean motion model of healthy subjects and of Tetralogy of Fallot patients in Sec. 8.5.

7.3 Transformation Parameter Estimation using a Polyaffine Model

In Seiler et.al [Seiler et al., 2012], the authors propose a method to estimate a poly-affine model from a log demons stationary velocity field. Using homogeneous co-ordinates, the

parameters of the poly-affine model can be defined for points $\tilde{x} = [x \ 1]^T$ in Cartesian coordinates as

$$\log(T) \stackrel{\text{def}}{=} \log \begin{pmatrix} A & t \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} M \\ 0 \end{pmatrix} \quad (7.1)$$

where \log is a principal matrix logarithm, A is the linear part of the transformation, t its translation, and M a 3×4 matrix. For each segment the affine deformation fields parameterised by the M_i matrices are fused to a global deformation field using the poly-affine model:

$$\vec{v}_{poly}(x) = \sum_i \omega_i(x) M_i \tilde{x}, \quad (7.2)$$

where ω_i is a parameter controlling the weight of the i^{th} region for each voxel x . Eqn. 7.2 can be estimated by a linear least squares problem with the least squares error with respect to the observed velocity field $\vec{v}(x)$ (in this case computed using the LogDemons algorithm).

In [McLeod et al., 2013b], the authors showed that an incompressibility penalisation and a neighbouring region regularisation can be added to this model by penalising the trace of the gradient of the poly-affine velocity field, and by regularising the similarity between neighbouring regions via a distance term. Both terms were added in the least squares minimisation to obtain a penalised least squares error term, still linear in terms of M . Image spatial priors were included to weight the input velocity field to give more influence on velocity vectors in regions of high gradient in the image (regions that drive the LogDemons registration) and to control the motion in the epicardium to prevent the registration from being dragged by the endocardial motion that is more visible in the image. The Polyaffine regions used are the 17 AHA regions (see Fig. 7.1 left). The Polyaffine weight functions were computed in prolate spheroidal coordinates rather than in the Cartesian frame using the method of Toussaint et. al [Toussaint et al., 2013] (see Fig. 7.1 right). Using this method, a low-dimensional, regional, and consistent definition of the motion can be obtained for a number of subjects.

7.4 Spatio-Temporal Parameter Alignment

Given the variability of the cMRI data, including temporal shifts, different cycle lengths and different coordinate spaces for different acquisitions, spatio-temporal alignment is first performed directly on the transformation parameters to align all the parameters to a common space. A method for temporally resampling the parameters so that all subjects have the same number of frames in Sec. 7.4.1, followed by a method for temporally aligning the parameters is described in Sec. 7.4.2 and finally a method for spatially aligning the parameters to a common space in a prolate spheroidal sense in Sec. 7.4.3. The steps for aligning Polyaffine transformation parameters are summarised in Fig. 7.2.

7.4.1 Temporal Resampling

In cine MR imaging, there is not a fixed number of frames acquired for each subject. In practice, the cycle length (from the R peak to the next R peak) is divided into a given number of frames (typically 10-30), and the cine movie is constructed from the information




Figure 7.1: The 17 AHA regions of the left ventricle (left) used as the Polyaffine regions, and the prolate spheroidal coordinate system used to define the Polyaffine weight functions (right).

gathered over several heart beats (see Fig. 7.3). Such sequences are acquired while subjects hold their breathe, to remove lung motion artifacts. Therefore, since the cycle length differs from one acquisition to another, we need to temporally resample the parameters to a common space to have an equal number of frames for all subjects (in order to be able to directly compare the parameters). This is done simply by interpolating between frames in the case that the number of frames for a given subject is fewer than the reference number of frames. This is represented in Fig. 7.2, step 2.

7.4.2 Regional Temporal Alignment

Cine MR image sequences are generally gated by ECG signals, so that the first frame corresponds to the end diastolic frame, and so that an image sequence covers one cardiac cycle. Standard acquisition is acquired by waiting a given amount of time after the R wave before initiating the sequence (see Fig. 7.4) and images are required between the R-to-R cycle. However, sometimes the first frame does not correspond exactly to the end diastolic frame and therefore the sequence is slightly offset. Furthermore, some sequences do not loop back exactly to the end diastolic frame, where some frames may be missing, or too many frames may be included in the sequence. Also, the ECG signal may be weak in some patients, such as those with large pericardial effusion, making the ECG gating difficult. Subjects may also differ in the duration of the cardiac cycle, therefore the peak frame does not always overlap for a group of patients. To account for this, the transformation parameters were aligned temporally per region, by computing the trace of the affine matrix per region over time, as in [McLeod et al., 2013b]. The trace of the affine matrix corresponds to the volume curve, through properties of Lie algebra, so this essentially amounts to aligning the parameters according to the volume curves. In this way, the mean peak frame can be computed for a group of patients, and all the transformation parameters are aligned to the




Figure 7.2: Spatio-temporal alignment pipeline starting with a temporal resampling step to resample the parameters so that there are an equal number of frames for all subjects. Temporal realignment is then performed regionally to align the parameters by the regional mean peak of the trace of the affine matrix. Spatial alignment is performed to align all parameters to a common coordinate space, by transporting the parameters to be centered at the barycenter of the region, and reoriented in a prolate spheroidal sense.




Figure 7.3: For a cine image sequence, the cardiac cycle is divided into N phases and acquisition is averaged for each phase over several cycles. The number of phases is variable from one acquisition to another.

mean peak frame (see Fig. 7.5). We assume that the transformation starts and finishes at the identity transformation, therefore a temporal shift of $+5$ sets the first 5 frames to identity.




Figure 7.4: The ECG signal is used to detect the R peak. Acquisition is started a fixed amount of time after the R peak, and one cycle is considered as the R-to-R duration. Gating the cine MR sequences in this way can result in different acquisition lengths, especially for patients with low ECG signal.

7.4.3 Spatial Alignment

In order to be able to construct a mean motion model, we first need to transport all the parameters to a common space. Rather than aligning all the images, which would amount to rigidly registering all images to a common space prior to performing the motion tracking, as was done in [Mcleod et al., 2013a], we rather align the transformations after performing the motion tracking. Given the low dimensional representation of the motion given from the registration algorithm summarised in the previous section, this is a matter of simply transporting the affine matrices to a common coordinate space. By choosing one subject (at random) to be the template subject, we can transport all the other subjects to this space. We are interested in analysing the motion at a regional level, hence we are interested in aligning the matrices at a regional level independently from other regions. Hence, we can




Figure 7.5: Trace of a single region plotted over time for a group of subjects shown before alignment (left) and after alignment (right). Aligning the curves brings the peaks to the same frame, to give the amount of shift needed to align the parameters.

align all matrices by first re-centering each matrix around the origin using the barycenter per region. The regions then need to be re-oriented to be aligned in a common manner. Since the Polyaffine weights are defined in PSS coordinates, it follows that the matrices can be oriented in a prolate spheroidal sense (see Fig. 7.1) by taking the Jacobian matrix computed at the barycenter. The same is done to re-orient the result to match with the template orientation, then bringing the matrix to the template space. A simplification of this pipeline is shown in Fig. 7.2, step 3. The prolate spheroidal coordinates were computed using the method of Toussaint et. al [Toussaint et al., 2013].

7.5 Mean Motion Model Construction

7.5.1 Mean Trajectory Model

Once all the matrices are aligned to a common space for all subjects, the mean and principal modes of these observations can be easily computed using basic arithmetic operations. The mean M_t at time t can be computed by:

$$\bar{M}_t = \frac{1}{N_k} \sum_{i=1}^{N_k} M_{k,t}, \quad (7.3)$$

where k is the subject index, N_k the number of subjects and $M_{k,t}$ the transported matrix for subject k at time t . By stacking all the matrices of a given subject to form a column vector (as shown in Fig. 7.7) of [time \times number of regions \times affine parameters], we can similarly compute the mean trajectory \bar{M} by taking the average of these vectors.

7.5.2 Simulated Motion from the Mean Model

Given the estimated mean trajectory, the motion can be simulated on a new subject, given that the subject is in the same coordinate space as the template, by simply applying the mean to the first image of the sequence: $I_t = M_t * I_0$ (see Fig. 7.6). In general, new subjects




Figure 7.6: The mean motion model applied to an image requires simply applying the mean transformation per frame to the reference image.

will not be in the same coordinate space as the template. Therefore, the mean trajectory \bar{M}_t first needs to be transported to the subject space. This can be done in a similar way to the pipeline for aligning the subjects. Given the barycenter and Jacobian matrix for each region of the new subject, \bar{M}_t can be transported to the subject space by:

$$M_r^k = T_r^k \cdot (\bar{T}_r)^{-1} \cdot \bar{M}_r \cdot \bar{T}_r \cdot (T_r^k)^{-1}, \quad (7.4)$$

where M_r^k is the mean motion model in region r in the subject-specific space for subject k , T_r^k is the matrix:

$$T_r^k = \left[\begin{array}{c|c} J_r^k & b_r^k \\ \hline 0 & 1 \end{array} \right]$$

with b_r^k the barycenter of region r for the k th subject (in Cartesian coordinates), J_r^k the Jacobian matrix at the barycenter in the local (prolate spheroidal) coordinates. \bar{T}_r is the mean matrix at the barycenter of region r in the template space.

7.5.3 PCA Motion Modes

The modes of variation around the mean can be computed simply by first forming a matrix M containing all the subject trajectories, and performing PCA on this matrix (as shown in Fig. 7.7).

Creating the matrix M is essentially a matricisation of the tensor \mathcal{T} containing all the parameters (see Fig. 7.8). This differs to the matricisation given in [Mcleod et al., 2013a], where the matrices of the affine parameters \times time were stacked block-wise by row or by column and singular value decomposition (SVD) was then applied to the stacked block matrices. Performing the decomposition in this way gives modes that may be difficult to interpret, in the sense that it is difficult to de-couple the modes related to the temporal or spatial components. To address this issue, we can rather perform the decomposition directly on the tensor, as described in Sec. 7.6.




Figure 7.7: Formulation of the matrix M of all parameters over time and space for all subjects. Performing PCA on this matrix gives the mean and modes of variation around this mean.

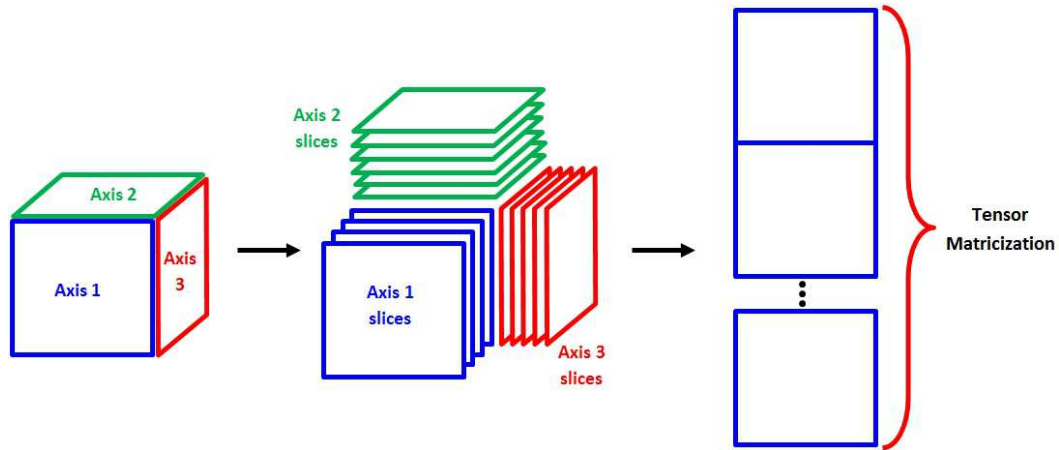


Figure 7.8: One method for matricising a three-way tensor by stacking the matrices block-wise.

7.6 Spatial and Temporal Mode Construction using Tensor Decomposition

Rather than performing PCA on stacked matrices of all parameters for each subject as described in Sec. 7.5.1 or SVD on parameters stacked either temporally or spatially as in [McLeod et al., 2013a], the tensor \mathcal{T} of all parameters can be decomposed directly. A number of methods for tensor decomposition have been proposed, the most common being the CANDECOMP/PARAFAC (CP) decomposition [Carroll and Chang, 1970, Harshman, 1970] and the Tucker decomposition [Tucker, 1963, Tucker, 1966] (see Fig. 7.9). A review of these methods as well as other tensor decomposition methods is given in [Kolda and Bader, 2009]. The CP decomposition is the generalisation of SVD to tensors in which a tensor is factored into the sum of rank-one tensors (tensors that can be written as the outer product of vectors). The Tucker decomposition (also known as higher-order PCA) rather decomposes the tensor into a core tensor and factor matrices corresponding to each axis (see Fig. 7.9). In this work, we focus on the Tucker decomposition, to give




Figure 7.9: Visual representation of the Tucker decomposition (top right) and CP decomposition (bottom right) on a three-way tensor.

decoupled “loadings” of the modes (whereas in the case of the CP decomposition, there is a single loading for each axis, meaning that the n^{th} modes are coupled in some way).

7.6.1 Tensor Decomposition Using the Tucker Method

The Tucker decomposition (a.k.a higher-order PCA / SVD, N-mode SVD / PCA) computes an orthonormal space associated with each axis of the tensor \mathcal{T} . For a three-way tensor, the Tucker decomposition can be written:

$$\begin{aligned} \mathcal{T} \approx \mathcal{G} \times_1 A \times_2 B \times_3 C &= \sum_{k=1}^K \sum_{t=1}^T \sum_{r=1}^R g_{ktr} a_k \circ b_t \circ c_r \\ &= [[\mathcal{G}; A, B, C]]. \end{aligned} \quad (7.5)$$

Computing such a decomposition is not straightforward. One method for doing so is to matricise the tensor (in the three-way case), similar to the decomposition described in the previous section. Another method, known as the Tucker1 method, is to compute optimal components to capture the variation in the n^{th} axis, independently from the other axes. Alternatively, the factors can be estimated using an iterative alternating least squares strategy (ALS), an alternating slice-wise decomposition (ASD), non-linear conjugate gradient method (NCG) and using a direct multi-linear decomposition (DMLD). Further details on the ALS strategy for computing a Tucker decomposition is given in the Appendix.

7.6.2 Tensor Decomposition on Re-scaled Parameters

Rather than performing the tensor decomposition directly on the parameters, we can perform tensor decomposition on re-scaled rotation, stretch, shearing and translation paramete-

ters. In this way, rather than performing the decomposition on the affine matrices M_r^k :

$$M_r^k = \begin{bmatrix} a_{1,1} & a_{2,1} & a_{3,1} & t_1 \\ a_{1,2} & a_{2,2} & a_{3,2} & t_2 \\ a_{1,3} & a_{2,3} & a_{3,3} & t_3 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

the analysis can be performed on vectors $P_{k,t,r}$ made up of the rotation, stretch, shearing and translation components for subject k at time t in region r . The rotation component is given by:

$$Rotation : \begin{bmatrix} R_x \\ R_y \\ R_z \end{bmatrix} = \begin{bmatrix} 1/2(a_{2,3} - a_{3,2}) \\ 1/2(a_{1,3} - a_{3,1}) \\ 1/2(a_{1,2} - a_{2,1}) \end{bmatrix} .$$

The stretch component is given by:

$$Stretch : \begin{bmatrix} S_x \\ S_y \\ S_z \end{bmatrix} = \begin{bmatrix} a_{1,1} \\ a_{2,2} \\ a_{3,3} \end{bmatrix} .$$

The shearing component is given by:

$$Shear : \begin{bmatrix} S_{xy} \\ S_{xz} \\ S_{yz} \end{bmatrix} = \begin{bmatrix} 1/2(a_{1,2} + a_{2,1}) \\ 1/2(a_{1,3} + a_{3,1}) \\ 1/2(a_{2,3} + a_{3,2}) \end{bmatrix} .$$

The translation component is given by:

$$Translation : \begin{bmatrix} T_x \\ T_y \\ T_z \end{bmatrix} = \begin{bmatrix} t_1 \\ t_2 \\ t_3 \end{bmatrix} .$$

The vector $P_{k,t,r}$ can then be expressed as:

$$P_{k,t} = \begin{bmatrix} R_x \\ R_y \\ R_z \\ S_x \\ S_y \\ S_z \\ S_{xy} \\ S_{xz} \\ S_{yz} \\ T_x \\ T_y \\ T_z \end{bmatrix} .$$

The elements of $P_{k,t,r}$ can then be scaled by $\sigma = [\sigma_{R_x}, \sigma_{R_y}, \dots, \sigma_{T_z}]$, where

$$\sigma_{[i]} = \sqrt{\frac{1}{N_k N_t N_l} \sum_{k,t,l} (p_{k,t,l})_{[i]}^2} \tag{7.6}$$

Tensor analysis can then be performed on $q_{k,t,r} = [P_{k,t,r}]_{[i]} / \sigma_{[i]}$

7.7 Healthy vs. Tetralogy of Fallot Motion

In order to apply the methods described in the previous sections, 4D sequences of two populations were tested. The first, a data-set of healthy controls, to establish normal motion patterns. The second, a data-set of patients with Tetralogy of Fallot, to establish motion patterns related to the pathology. This section is organised in the same order as the previous section to test first the mean motion model on both populations (Sections 7.7.1,7.7.2), then computing temporal bases on the stacked matrices to compare to tensor decomposition on the parameters (with and without removing the mean) in Sec. 7.7.3. Tensor decomposition on the P matrices for both populations is described in Sec. 7.7.4.

Healthy Volunteers We illustrate these tools on 15 healthy adults (3 female, mean age \pm SD = 28 ± 5) obtained from the STACOM 2011 MICCAI cardiac motion tracking challenge database [Tobon-Gomez et al., 2013]. Steady-state free precession magnetic resonance images were acquired using a 3T scanner (Philips Achieva System, Philips Healthcare) in the short axis view covering entirely both ventricles (12-16 slices; isotropic in-plane resolution: $1.15 \times 1.15mm^2$ to $1.25 \times 1.25mm^2$; slice thickness: $8mm$; 30 frames).

Tetralogy of Fallot Patients To test the methods on pathological cases, we applied the described method to ten patients with repaired Tetralogy of Fallot (ToF) (5 female, mean age \pm SD = 21 ± 7). These patients all had a full ToF repair early in infancy, resulting in the destruction of the pulmonary valves. Steady-state free precession magnetic resonance images were acquired for each patient in the short axis view covering entirely both ventricles (12-15 slices; isotropic in-plane resolution: $1.21 \times 1.21mm^2$ to $1.36 \times 1.36mm^2$; slice thickness: $8mm$; 15-19 frames).

7.7.1 Mean Normal Motion Model

A mean normal (healthy) motion model was computed from the 15 healthy volunteers using the method described in Sec. 7.5. The first volunteer was chosen to be the template, and all other subjects were aligned to this space. To visualise the mean motion model, the model was transported to the space of the second volunteer. The first image of volunteer 2 was deformed by the motion model. In order to analyse the result, we compare the mean model to the Polyaffine model (since we do not assess here the accuracy of the Polyaffine model used to estimate the transformation parameters). The first image in the sequence was thus transformed by the Polyaffine model for comparison. The images at three frames (frame 5, 10 and 20) are shown in Fig. 7.10, along with the iso-contours of each image (in green). The mean motion model gives a reasonable motion pattern, and actually gives smoother motion than the Polyaffine model (not shown).

A leave-one-out experiment was performed to test the ability of the mean model to approximate the motion of a new subject drawn from the same population by using 14 subjects as the training set, and testing on the 15th subject. The residual error (RMSE) between the mean motion model (in the k^{th} subject space) and the original Polyaffine transformation parameters (before realignment) was computed for subjects in the training set as well as




Figure 7.10: Polyaffine Model (top row) and mean motion model (bottom row) at frame 5 (left), 10 (centre) and 20 (right) for subject 2. The mean model captures the key dynamics of the motion, though the radial contraction is underestimated (as seen on the 10th frame).

for subjects in the testing set. The errors are plotted over time in Fig. 7.11 for the training set (left) and the testing set (right). Note that the error is the difference between the affine transformation parameters. The errors are within the same range for the training and testing set, indicating that the model is able to predict the motion of healthy subjects.

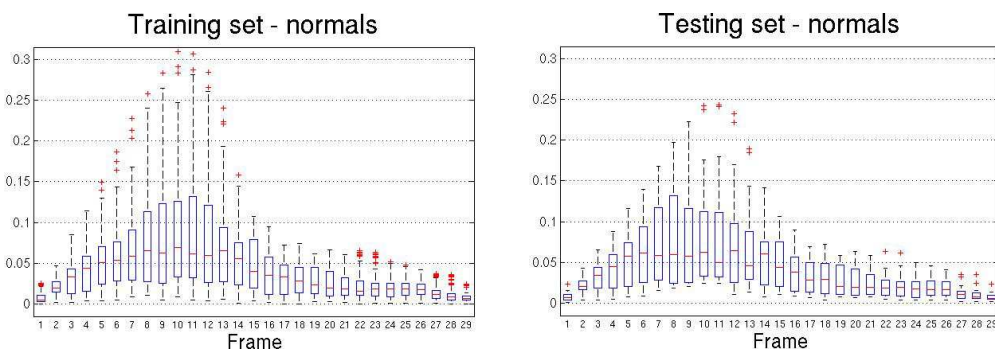


Figure 7.11: Boxplots of the residual norm between the Polyaffine transformation and the normal mean model transformation over time for subjects used in the training set (left), and for subjects not used in the training set (right) for the healthy volunteers.

To test the ability of the model to capture the motion of ToF patients, the mean motion model was applied to each patient, and the residual error between the model and the Polyaffine parameters was computed. The boxplot of the error is given in Fig. 7.12 alongside the boxplots of the training and testing set of the healthy volunteers. As expected the residuals are higher for the ToF patients, indicating that the healthy model is insufficient in capturing the motion dynamics in these patients. This suggests that there are abnormal motion dynamics observed in these patients.




Figure 7.12: Boxplots of the residual norm between the Polyaffine transformation and the normal mean model transformation for subjects used in the training set (left - red), for healthy subjects not used in the training set (centre - blue) and for Tetralogy of Fallot patients (right - green).

7.7.2 Abnormal Motion Identification in Tetralogy of Fallot Patients

The left ventricular motion was tracked in the ToF cases to obtain a baseline motion for these patients. The mean normal model was then applied to each patient to determine whether there are visible regional differences between the obtained matrices and the baseline matrices. In particular, we are interested in determining where the differences lie between the normal and pathological cases. Given the large variability of motion observed in these patients (since the patients received different treatments during their lifetimes, and are at different stages in the progression of the heart remodelling), we expect that a larger number of subjects are needed in the training set to be able to capture the motion dynamics for all subjects. Performing a leave-one-out experiment as was done for the healthy controls, we find that the residual error is higher for the training set than for the testing set, as shown in Fig. 7.13. This is expected since few subjects were used in the training set, and there is a large variability of motion in this population. Therefore, this suggests that the mean ToF motion model is able to capture the motion dynamics of these patients reasonably well.

7.7.3 Healthy vs. Pathological Spatio-Temporal Bases

As a baseline for comparison, the same decomposition as used in [McLeod et al., 2013a] was performed on the realigned parameters for both the healthy controls and for the Tetralogy of Fallot patients (see Fig. 7.14), to compare to the temporal modes computed from the tensor decomposition. There are clear differences between the modes, where the principal mode for the healthy volunteers resembles the typical volume curve of healthy subjects. The principal mode for the ToF patients is less regular, potentially due to the fact that these patients do not have a single peak around the end systolic phase, but have more complex motion dynamics with multiple peaks in contraction.

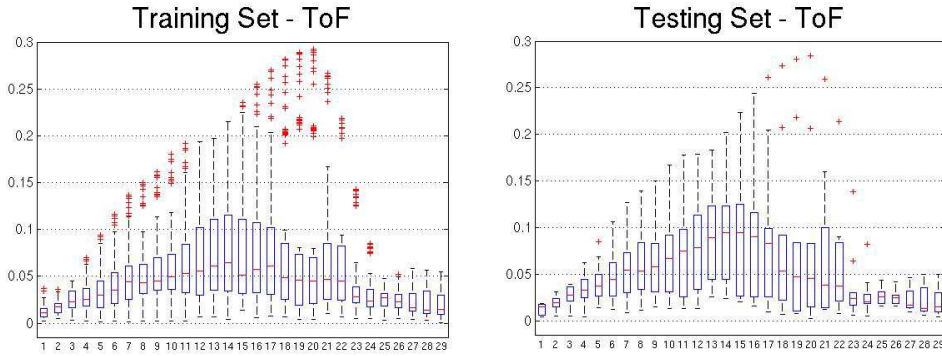


Figure 7.13: Boxplots of the residual norm between the Polyaffine transformation and the normal mean model transformation over time for subjects used in the training set (left), and for subjects not used in the training set (right) for the Tetralogy of Fallot patients.




Figure 7.14: The first 5 temporal modes computed by decomposing the tensor \mathcal{J} block-wise and stacking to one row of blocks, plotted over time (scaled by the eigenvalues) for the healthy controls (left) and the Tetralogy of Fallot patients (right).

The first two spatial modes for the healthy controls are shown in Fig. 7.15 (left) and for the ToF patients (right). The first mode for the healthy controls appears to capture both the radial contraction, and the circumferential motion (shown in block yellow arrows). The ToF modes, on the other hand, appear to capture a translation of the free wall and septal wall towards the right ventricle. This abnormal motion is evident in the image sequences of these patients.

The Tucker tensor decomposition described in Sec. 7.6.1 was performed using the Sandia National Laboratories Tensor Toolbox [Bader et al., 2012], implemented in MATLAB. Using these tools, the three-way tensors of [affine parameters \times time \times subject] of the healthy controls and of the ToF patients were decomposed into factor and core matrices. Using the Tucker decomposition, we obtain temporal bases for each loading related to the spatial bases.

The first three spatial loadings for the healthy (top row) and ToF (bottom row) subjects are shown in Fig. 7.16 for the first five temporal modes (without removing the mean). The first loadings exhibit interesting differences, where the healthy modes resemble the normal volume curve, whereas the first loading for ToF shows slow contraction and show filling.

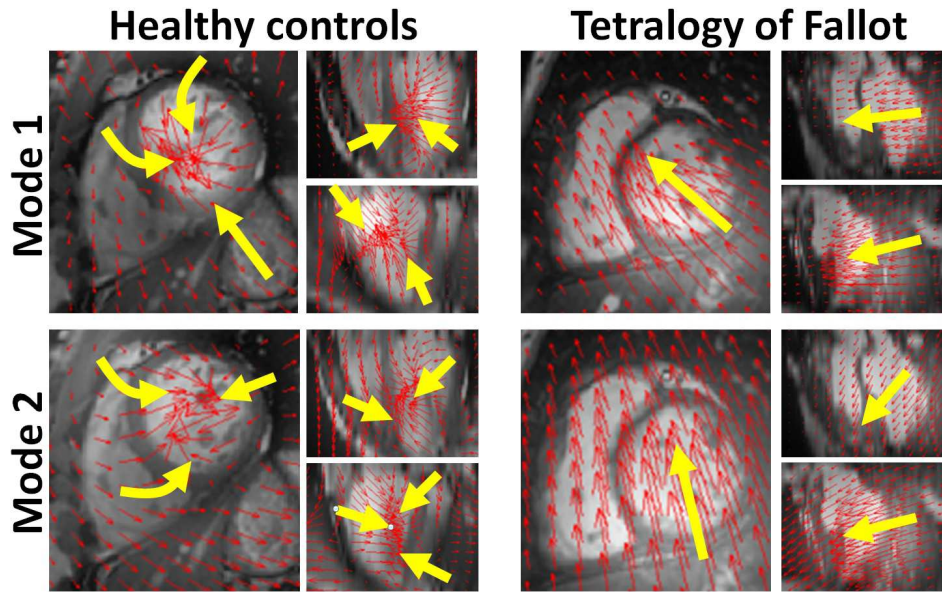


Figure 7.15: Three views of the first(top row) and second (bottom row) spatial modes for the healthy controls (left) and for the Tetralogy of Fallot patients (right). The modes for the healthy controls represent the radial contraction and circumferential motion, whereas the modes for the Tetralogy of Fallot patients represent the translation towards the right ventricle. Yellow arrows indicate direction of motion.

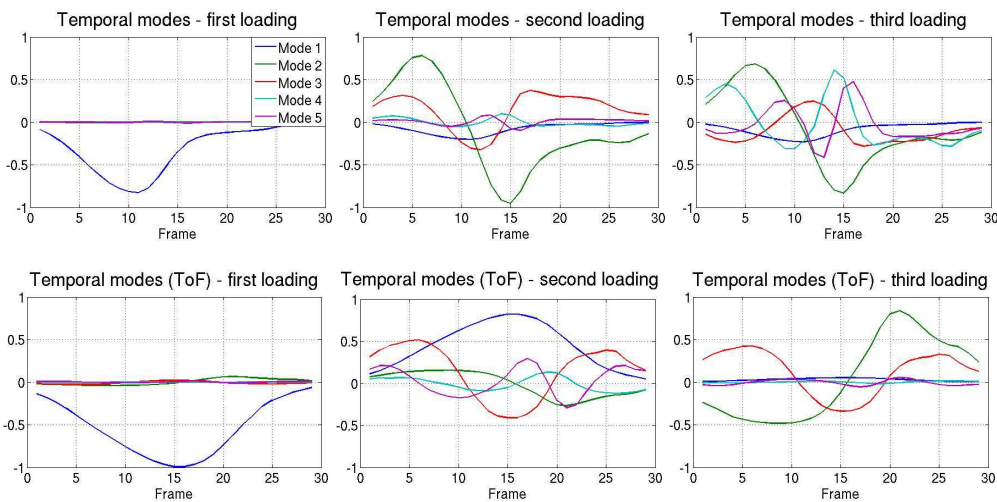


Figure 7.16: Temporal modes for the first (left), second (centre) and third (right) loadings corresponding to the first, second and third spatial modes for the healthy controls (top row), and for the Tetralogy of Fallot patients (bottom row). The first loadings show interesting differences for the ToF patients, with slower contraction and slower filling than for the healthy subjects.

By performing the tensor decomposition on the parameters after removing the mean, the temporal modes show similar behaviour, with multiple peaks, translated to the left (see Fig. 7.17 - left). The curves are flat around the end systolic phase, which is expected since the parameters were aligned to the frame corresponding to the end systolic phase. The peaks occur around the beginning and end of the cycle, and between the 10th and 20th frames, where during the relaxation phase. This may account for the differences in these phases due to the fact that the parameters were temporally aligned only rigidly. In that case, the temporal modes would account for the misalignment during these phases. For the ToF patients (Fig. 7.17 - right), there is a greater number of peaks, which may indicate that there is greater variability temporally in these patients.




Figure 7.17: Temporal modes computed from the tensor decomposition performed on the centered matrices (after subtracting the mean) for the healthy controls (left) and the Tetralogy of Fallot patients (right).

7.7.4 Tensor Decomposition on P Matrices

As described in Sec. 7.6.2, tensor decomposition can be performed on the vectors $P_{k,t,r}$ containing the rotation, stretch, shearing and translation components of each affine matrix. This amounts to performing tensor decomposition on the four-way tensor [matrix components \times time \times region \times subject], which is a tensor of size $[12 \times 29 \times 17 \times N]$ (where $N = 15$ for the healthy subjects and $N = 10$ for the ToF patients). The first affine component mode is summarised in Table 7.1 for each population, highlighting the differences in the motion dynamics. While the healthy controls capture a positive rotation in all directions, the ToF patients have little rotation, and rotate in the opposite direction. Less stretch is observed for the ToF patients than the healthy controls.

7.8 Discussion and Perspectives

In the current work, we proposed to use statistical methods to create a mean motion model, and to extract the modes of variation using standard statistical model reduction techniques, as well as more advanced tensor decomposition methods. The proposed method uses the parameters of the cardiac PolyLogDemons algorithm [McLeod et al., 2013b], which gives

Affine component	Healthy controls	Tetralogy of Fallot
R_x	0.0062	-0.0005
R_y	0.0048	-0.0001
R_z	0.0022	-0.0015
S_x	-0.0054	-0.0022
S_y	-0.0064	-0.0024
S_z	-0.0078	-0.0042
S_{xy}	-0.0093	0.0067
S_{xz}	-0.0083	0.0101
S_{yz}	-0.0100	0.0088
T_x	2.58	-1.085
T_y	3.51	-1.245
T_z	2.73	-1.341

Table 7.1: First principle mode of affine components.

transformations described by a small number of parameters at a regional scale. Spatio-temporal alignment of the parameters was performed prior to analysis. The techniques used are further discussed in Sec. 7.8.1. A mean motion model was constructed by taking the average of the aligned parameters. The mean model construction is further discussed in Sec. 7.8.2. From the obtained Polyaffine transformation parameters, statistical model reduction was applied to the parameters after spatio-temporal alignment. The statistical methods used are discussed in Sec. 7.8.3.

7.8.1 Parameter Realignment

Spatio-temporal alignment was performed by making use of the trace of the affine matrices to align the transformations temporally, and making use of the prolate spheroidal coordinates to align the transformations spatially, in an anatomically meaningful manner. The temporal realignment used in this work was a simple "rigid" transformation of the parameters. The parameters are translated regionally by the mean peak of the trace, since the peak contraction is captured sufficiently well for all regions. However, this may have the effect of removing relevant parameters. Therefore, a non-rigid transformation may allow more intuitive analyses by fixing the first frame and stretch (or shrinking) the parameters to the mean peak of the trace, and stretching (or shrinking) to the relaxation phase, similar to the temporal alignment described in [Perperidis et al., 2005]. Both the temporal and spatial alignment were performed at a regional level. This was due to the fact that we are interested in region-wise analyses, rather than global analyses. However, it is also possible to temporally align the transformations to a global peak frame, and in fact for patients with large dysynchrony between regions, this may be more suitable. Spatially, however, the current formulation relies on the regional construction. The parameters could be spatially aligned by the barycenter of the ventricle, however further data-points would be needed such as the base, apex and LV-RV junction.

7.8.2 Mean Motion Model Construction

From the mean motion model constructed from healthy subjects, we found that the model is able to approximate the motion of other healthy subjects reasonably well, however the motion of Tetralogy of Fallot patients was not fully captured from this model. This was expected given the abnormal motion dynamics observed in ToF patients. Interestingly, the motion model constructed from 9 ToF patients was sufficient to capture the motion dynamics of the 10th subject (based on the results of the leave-one-out experiment). Given that these patients have followed different treatment paths, and are at different stages due to the range of patient age (some patients are children, others adults), we expected to require a larger number of patients to capture the dynamics in this population. However, the related factor between these patients is a missing pulmonary valve, caused by the initial surgery. Therefore, the regurgitated blood from the pulmonary artery may be the key factor affecting the motion dynamics of these patients. Ideally, given more patients with similar treatment history, treatment-specific models could be constructed to distinguish the key differences in the motion dynamics for patients treated with one treatment over another. In future work, the mean motion models could also be used to generate synthetic sequences for new subjects in which only the end diastolic image is available, or as a prior in the cardiac-specific PolyLogDemons of Chapter 5 to regularise the tracking and to reduce errors.

7.8.3 Statistical Model Reduction

A number of statistical model reduction techniques were used in this work to attempt to understand the underlying motion variability of the two populations. We extend on the work in [McLeod et al., 2013a], in which model reduction is performed block-wise on matrices stacked either column- or row-wise to obtain temporal and spatial population bases. In the current work, we perform model reduction on the full tensor using the Tucker decomposition to obtain spatial and temporal bases directly. The advantage of performing model reduction directly on the tensor is that the bases are decoupled and can be recombined independently. However, tensor decomposition is not straightforward and requires numerical methods for estimating the modes, compared to matrix decomposition which could be performed simply with SVD, PCA, POD etc. Due to the coupling between the modes (the fact that a combination of modes contribute to the motion dynamics), we perform model reduction on both the centered data (PCA-like decompositions) as well as on the non-centered data (SVD-like decompositions). Performing PCA-type analyses, we can analyse the bases orthonormal with respect to the variable motion dynamics, rather than to the mean, whereas SVD-type decompositions allow us to analyse the motion with respect to the mean motion patterns. By performing both types of analyses, we can gather more information on the different motion dynamics. We propose to perform tensor decomposition on the affine components (rotation, stretch, shearing and translation) as well as the combined affine matrices, to distinguish population-wide differences in the motion. The tensor decomposition used in this work was the Tucker decomposition, though a number of other methods for decomposing a tensor have been proposed and could be applied to give potentially similar analyses.

7.9 Conclusion

In this work we propose a novel method for constructing a mean cardiac motion model based on image registration and population-based statistics. Using a Polyaffine model to describe the transformation from one image to another, we perform group-wise analysis on the parameters to construct first a mean motion model by averaging spatio-temporally aligned transformation parameters from a group of subjects in the same population (healthy or pathological). The modes of variation around this mean were computed using simple PCA. Further population-based statistical analyses were performed by computing a tensor decomposition on the transformation parameters directly, and on the affine components of these parameters (rotation, stretch, shearing and translation). From the analyses we obtained meaningful differences between the healthy population and the pathological population of Tetralogy of Fallot patients.

7.10 Appendix

7.10.1 Tucker Tensor Decomposition using Alternating Least Squares

The alternating least squares (ALS) method for computing a rank- (R_1, \dots, R_N) Tucker decomposition is described briefly here. Note that the n -Rank of a tensor \mathcal{T} is not the same as the rank of a tensor (the minimum number of rank-one components). R_n is the column rank of the n^{th} axis of \mathcal{T} . In this work, we take $R_n = \text{rank}_n(\mathcal{T})$ to find an exact Tucker decomposition of \mathcal{T} , rather than computing a truncated Tucker decomposition (which gives an inexact decomposition rather than the truncation of an exact decomposition).

The ALS method essentially approximates the factors that minimise the least squares criterion:

$$C(\mathcal{T}) =: \min_{\mathcal{G}, \mathcal{A}^{(n)}} \|\mathcal{T} - [[\mathcal{G}; \mathcal{A}^{(1)}, \mathcal{A}^{(2)}, \dots, \mathcal{A}^{(N)}]]\| \quad (7.7)$$

given $\mathcal{G} \in \mathbb{R}^{R_1 \times R_2 \times \dots \times R_N}$ and $\mathcal{A}^{(n)} \in \mathbb{R}^{I_n \times R_n}$ is column-wise orthogonal for all n . The algorithm for computing a rank- (R_1, \dots, R_N) decomposition is given is Algorithm 5 [Kolda and Bader, 2009].

Algorithm 5 ALS rank- (R_1, R_2, \dots, R_N) Tucker decomposition

Require: Initialisation: $\mathcal{A}^{(n)} \in \mathbb{R}^{I_n \times R}$ for $n = 1, 2, \dots, N$

loop {until convergence}

for $n = 1, \dots, N$ **do**

$y \leftarrow \mathcal{T} \times_1 \mathcal{A}^{(1)T} \dots \times_{n-1} \mathcal{A}^{(n-1)T} \times_{n+1} \mathcal{A}^{(n+1)T} \dots \times_N \mathcal{A}^{(N)T}$

$\mathcal{A}^{(n)} \leftarrow$ the $R^{(n)}$ leading singular vectors of y

$\mathcal{G} \leftarrow \mathcal{T} \times_1 \mathcal{A}^{(1)T} \times_2 \mathcal{A}^{(2)T} \dots \times_N \mathcal{A}^{(N)T}$

return $\mathcal{G}, \mathcal{A}^{(1)}, \dots, \mathcal{A}^{(N)}$

Part III

**STATISTICAL MODELLING OF
CARDIAC BLOOD FLOW**

Group-wise Construction of Reduced Models for Understanding and Characterization of Pulmonary Blood Flows from Medical Images

Contents

8.1	Chapter Overview	138
8.2	Introduction	138
8.3	Computation of an Average Geometry	141
8.3.1	Non-parametric Representation of Surfaces using Currents	141
8.3.2	Surface-to-surface Registration	143
8.3.3	Iterative Estimation of the Atlas	144
8.3.4	Surface Representation of the Atlas	144
8.4	An Atlas-based Reduced Order Model of Blood Flow	145
8.4.1	Reference Blood Flow Simulation and POD Basis	145
8.4.2	Individual Reduced Order Blood Flow Simulation	147
8.5	Numerical Experiments on the Pulmonary Artery of Repaired Tetralogy of Fallot Patients	150
8.5.1	Data Collection	150
8.5.2	Statistical Shape Model of the Pulmonary Arteries	150
8.5.3	Patient-specific Full and Reduced order CFD Simulations	152
8.6	Discussion	162
8.6.1	Results of Reduced Flow Simulations	162
8.6.2	General Remarks and Perspectives	163
8.7	Conclusion	165

Following from the cardiac growth model proposed in Part I, and the cardiac motion analysis described in Part II, we are now interested in understanding how the blood flow dynamics differ in a pathological population compared to a normal population. Given the complexity of subject-specific blood flow simulations, population-based analysis of blood flow dynamics can be difficult. This chapter describes work towards a reduced-order model of blood flow to address the third key question of this work on how to simulate blood flow

through a given organ at a population level to be able to develop a generative blood flow model for a given population. The atlas construction method described in this work in application to constructing an atlas of the pulmonary artery is the same framework as that used to construct the bi-ventricular atlas in Chapter 3. The atlas construction method is described here again to link the notation with the methods that follow in this work. This chapter was accepted for publication to Medical Image Analysis [McLeod et al., 2013] and extends on the work published in McLeod et. al [McLeod et al., 2010a].

8.1 Chapter Overview

3D computational fluid dynamics (CFD) in patient-specific geometries provides complementary insights to clinical imaging, to better understand how heart disease, and the side effects of treating heart disease, affect and are affected by hemodynamics. This information can be useful in treatment planning for designing artificial devices that are subject to stress and pressure from blood flow. Yet, these simulations remain relatively costly within a clinical context. This work aims at reducing the complexity of patient-specific simulations by combining image analysis, computational fluid dynamics and model order reduction techniques. The proposed method makes use of a reference geometry estimated as an average of the population, within a well-posed statistical framework based on the currents representation of shapes. Snapshots of blood flow simulations performed in the reference geometry are used to build a POD (Proper Orthogonal Decomposition) basis, which can then be mapped on new patients to perform reduced order blood flow simulations with specific boundary conditions. This approach is applied to a data-set of 17 tetralogy of Fallot patients to simulate blood flow through the pulmonary artery under normal (healthy or synthetic valves with almost no backflow) and pathological (leaky or absent valve with backflow) conditions to better understand the impact of regurgitated blood on pressure and velocity at the outflow tracts. The model reduction approach is further tested by performing patient simulations under exercise and varying degrees of pathophysiological conditions based on reduction of reference solutions (rest and medium backflow conditions respectively).

8.2 Introduction

Congenital heart disease (CHD), besides decreasing the quality of living and life expectancy of many, is a prominent cause of infant mortality around the world. Abnormal heart rhythm and structure affect blood flow through the heart and into the arteries, which consequently influences the development of the organs and surrounding structures. CHD patients are monitored with several imaging modalities, which are chosen in medical practice based on the pathology and its severity [Puranik et al., 2010]. Anatomy can be easily visualised in the images extracted from typical cardiac imaging modalities such as cine-MR in magnetic resonance (MR) imaging, echocardiography and CT. However current methods for imaging blood flow in a three-dimensional volume are limited. Non-invasive imaging techniques such as phase contrast magnetic resonance imaging (PC-MRI) are used

to quantitatively measure blood flow in clinical practice only if necessary, or are not easy to acquire on multiple locations, especially for CHD patients [Markl et al., 2012]. Compared to PC-MRI, cardiac MR volumetry has been shown to provide non-interchangeable cardiac function assessment in Dobutamine stress test of regurgitant repaired Tetralogy of Fallot (ToF) [Valverde et al., 2011]. Time-resolved 3D contrast-enhanced MR angiography has been proposed to also non-invasively assess lung perfusion in ToF [Tomasian et al., 2009]. 4D-MRI, while potentially providing velocity in a three-dimensional volume over time, is still a subject of research for its acquisition and post-processing [Markl et al., 2012]. Velocity can be measured with Doppler ultrasound machines, which are more often found in clinical practice, but the acquired data is limited to 1D.

This is one reason why computational simulations of blood flow have been developed in recent years in an attempt to better understand blood flow dynamics. The aim of such modeling is to gain insight into how hemodynamics change due to a pathology in terms of flow patterns, pressure, wall shear stress, and so on (see, *e.g.*, [Yeung et al., 2006, Troianowski et al., 2011, LaDisa et al., 2011]). Conversely, pathology developments can depend on abnormal hemodynamics, such as in the absence of a functioning valve, as will be the case in this work. Computer hemodynamic simulations provide a tool to predict hemodynamic changes due to surgical repair [Vignon-Clementel et al., 2010b], explore different scenarios for treatment (see, *e.g.*, [Hsia et al., 2011, Yang et al., 2011]), non-invasively compute indices that are otherwise invasively measured such as fractional flow reserve (FFR) [Koo et al., 2011], and design artificial devices or conduits that are subject to stress and pressure from blood flow (see, *e.g.*, [Prasad et al., 2011, Pant et al., 2011]). Similarly, knowing the anatomy and hemodynamics of a patient, physicians can determine the suitability of a patient for surgical intervention and plan therapy [Morales et al., 2011].

Computational fluid dynamics (CFD) methods are commonly used to simulate blood flow in the larger arteries, by solving numerically the three-dimensional Navier-Stokes equations. However, the large computational cost of numerical simulations is still a relevant issue, especially when dealing with patient-specific geometries. On one hand, the computation time can be reduced by exploiting parallel computing and GPU implementations (see, *e.g.*, [Glowinski et al., 2001, Biswas, 2010]). On the other hand, the issue of complexity has also been addressed by developing so-called model order reduction techniques, which aim at reducing the dimension of the problem by restricting the numerical solution to a pre-defined low order space.

Among others, in [Manzoni et al., 2011], the authors propose to extract a low dimensional parametrisation of the computational domain and solve the parameterised flow equations with the reduced basis method. This work is able to gain speed in computational time but necessitates a shape parametrisation of the object through control points assigned at regions of shape change. Here, we consider a model reduction based on the Proper Orthogonal Decomposition (POD, also known as Karhunen – Loève expansion or Principal Component Analysis – PCA), an approach which aims at finding a low-order basis for the numerical solution starting from a set of pre-computed snapshots. In fluid dynamics, POD was introduced by Lumley in the late 1960s to study turbulent flows. For a more recent presentation of POD, we refer for example to [Volkwein, 1999, Rathinam and Petzold, 2004], among many others.

In this chapter, we extend the algorithm that we briefly presented in [McLeod et al., 2010a]. This method is used to construct reduced models of blood flow through the pulmonary artery by computing a reference POD from a flow simulation on a reference geometry. The reference geometry is obtained by computing a centered atlas of a preliminary population, with a non-parametric method which is thus not dependent on a choice of control points. From a blood flow simulation in the atlas, we generate a POD basis of the flow, which defines a reduced subspace for further simulations. Namely, numerical simulations on (potentially new) patients can be performed directly in the reduced subspace, after transporting the POD basis on the new geometry. As in [McLeod et al., 2010a], this is done using a deformation map computed via 3D shape registration. In the present work, we use the reduced order modeling to compute reduced simulations for a set of individual patients, under different physiological conditions. The first aim is thus to see how for fixed boundary conditions, information computed on a reference geometry can be used to compute hemodynamics on a patient geometry. In particular, we apply the methodology to a data-set of 17 pulmonary artery (PA) geometries with repaired ToF (as shown in Fig. 8.1), where the structures are affected by a known shape abnormality due to the pathology and the initial surgical repair. A schematic diagram showing the pipeline of this method is shown in Fig 8.2. For this




Figure 8.1: A 3D volume rendered image of one patient to visualise the location of the pulmonary artery (PA) with respect to the heart and lungs.

pathology, clinicians are interested in understanding how the artery remodels due to blood regurgitating back into the right ventricle of the heart, and reversely how blood regurgitates due to the deformed artery. Regurgitation results from the absence of a functioning pulmonary valve that maintains one-way blood flow from the right ventricle to the pulmonary artery. Thus the second aim of this chapter is to investigate how the reduced method works for normal as well as pathological flows in the PAs, and if it is robust enough to simulate other physiological states, namely exercise and varying degrees of pathological regurgitation. Previous ToF blood flow simulations have investigated regurgitation with lumped [Kilner et al., 2009] or idealised (geometry and boundary

conditions) three-dimensional [Chern et al., 2008] models, as well as pressure losses for two repair options with either a one-dimensional [Spilker et al., 2007] or a steady three-dimensional [Chai et al., 2010] patient-specific model. But for the aims of this chapter, it is important to combine realistic models of form (three-dimensional geometries) and function (physiological inflow and outlet boundary conditions) as in [Das et al., 2011].

For follow-up treatment planning, clinicians are interested in determining first the optimal method of intervention for each patient, be it surgical or (more preferably) non-invasive, as well as designing artificial devices to act in the same way as the pulmonary valve to enforce one way blood flow. Due to the complex and remodeled pulmonary artery, it is not always possible to insert existing devices in some patients. This work constitutes thus a first step towards the use of atlas-based reduced models to better understand hemodynamics in these patients and efficiently simulate hemodynamics for later device design.

The methodology proposed in this chapter has some similarities with the work presented in [Mcgregor et al., 2008, McGregor et al., 2009]. Nevertheless, in these references the POD basis is used to interpolate a flow field measured on medical images, whereas it is used here to actually solve the fluid mechanics equations.

The chapter is structured as follows. In Sec. 8.3 the construction of the atlas which serves as the reference geometry is described. Sec. 8.4 is dedicated to blood flow simulations, describing the numerical method for the full simulations on the reference geometries and the model order reduction technique for the patient-specific reduced simulations. The procedure is assessed in Sec. 8.5 comparing the reduced method to full CFD simulations in different clinically relevant situations. Finally, in Sec. 8.6 the results and future paths for improvements are discussed, while Sec. 8.7 draws a few concluding remarks.

8.3 Computation of an Average Geometry

We assume that for a set of patients drawn from a given population we have for each patient a surface representation S of the region of interest (such as an organ, or artery) which is defined by delineating the boundaries of the organ on the images. Given the set of such surfaces we would like to compute an average surface representation of the population (an atlas). The geometry of the atlas should be well-defined and centered with respect to the population, and we also require a consistent method for computing the deformations from the individual geometries to the atlas. For this we use the framework proposed in [Durrleman et al., 2009a], as described below.

8.3.1 Non-parametric Representation of Surfaces using Currents

The patient meshes are represented by *currents*, so that we can compute the distance between two meshes in the space of currents, without requiring a point-to-point correspondence between the meshes [Durrleman et al., 2009a]. Currents are used to represent the observations (the patient geometries), the residuals (what is not captured by the model), and the deformations (used to map one geometry to another), in the same common framework. The space of currents forms a vector space, hence we can apply statistical operations such as the mean and variance on surfaces.

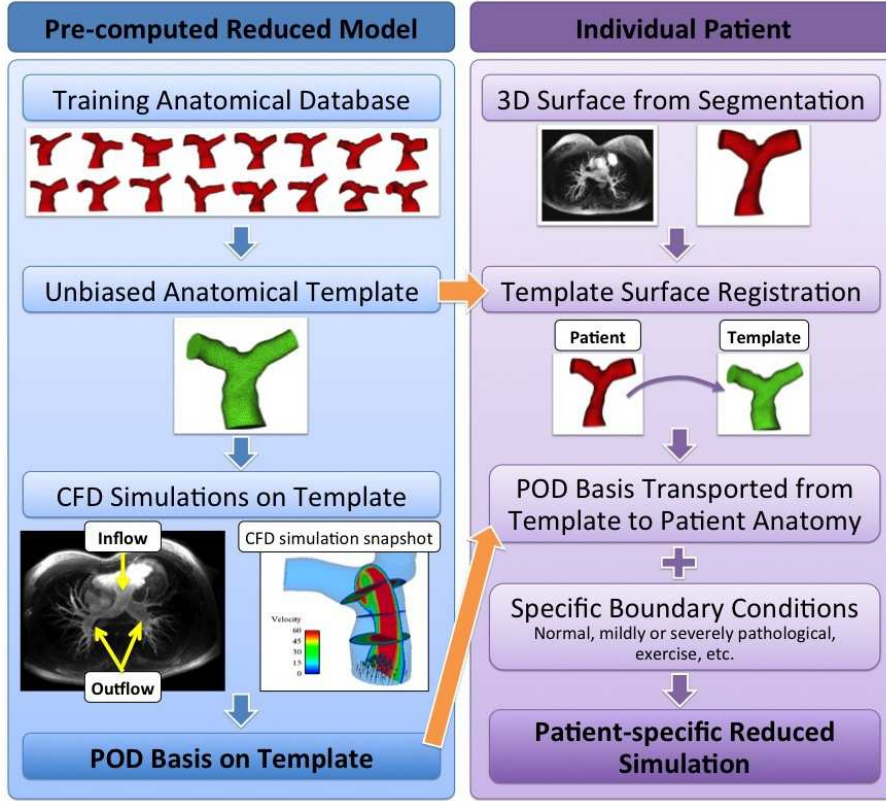


Figure 8.2: Schematic diagram to display the method which consists of an offline pre-computation step by simulating the blood flow on a reference geometry, and then reducing the output of the simulation to obtain a reduced-order basis. An individual flow simulation can be done by transporting the reduced-order reference basis to the individual, and solving the flow equations with the reduced basis in the patient space with patient-specific boundary conditions.

The basic principle of currents is to probe the surface by a set of vector fields in a test space W . In this way, a surface can be characterised by a set of real numbers that represent the solutions of the flux integral

$$S(\omega) = \int_S \omega(x)^t \mathbf{n}(x) d\lambda(x) \quad (8.1)$$

for every vector field ω in W . In (8.1), ω denotes a square integrable 3D vector field, $\mathbf{n}(x)$ is the unit normal of the surface at a point x , and $d\lambda$ is the Lebesgue measure on the surface. The test space W is chosen as the set of convolutions between any square integrable vector field and a smoothing kernel. Hence, W is a Reproducing Kernel Hilbert Space (RKHS). In this work we use a Gaussian kernel, which can be defined, for any points (x, y) , as

$$K_W(x, y) = \exp \frac{-\|x-y\|^2}{\lambda_W^2}, \quad (8.2)$$

where λ_W is the standard deviation. With this choice for the reproducing kernel, we can

then control a metric in the space of currents that allows the distance between two surfaces to be calculated, though this choice is not unique. This space also has the important property that it is the dense span of basis vector fields of the form $\omega(x) = K_W(x, y)\beta$, for any fixed point y and fixed vector β . As a consequence, any vector field ω can be written as an infinite linear combination of the basis elements $K_W(x, y)\beta$. We can also define an inner product in W using the kernel K_W with these basis vectors as

$$\langle K_W(\cdot, x)\alpha, K_W(\cdot, y)\beta \rangle_W = \alpha^t K_W(x, y)\beta. \quad (8.3)$$

This inner product holds for any vector field $\omega(x) = K_W(x, y)\beta$ in W .

The space of currents is defined as the *dual* space of W , denoted W^* , representing the vector space of linear mappings from W to \mathbb{R} . Since W is a Reproducing Kernel Hilbert Space, the evaluation functionals are bounded. Base on the properties of W (the fact that W is densely spanned by the vector fields $\omega(x)$, and has an associated inner product) its dual space W^* is densely spanned by the dual representations of the basis vectors $\omega(x)$, called the Dirac delta currents, defined as:

$$\langle \delta_x^\alpha, \omega \rangle_W = \langle K_W(x, \cdot)\alpha, \omega \rangle_W = \alpha^t \omega(x). \quad (8.4)$$

Given that W^* is a vector space, we can define the sum of two surfaces (represented by currents) C_{S_1} and C_{S_2} as $(C_{S_1} + C_{S_2})(\omega) = C_{S_1}(\omega) + C_{S_2}(\omega)$. In terms of the flux, this means that the flux through the sum of two surfaces is the sum of the flux through each surface. The vector space property of scalar multiplication in this case means that we can scale a current by simply scaling the Dirac delta currents. We can also define an inner product in W^* as:

$$\left\langle \delta_x^\alpha(\omega), \delta_y^\beta(\omega) \right\rangle_{W^*} = \langle K_W(x, \cdot)\alpha, K_W(y, \cdot)\beta \rangle_W = \alpha^t K_W(x, y)\beta. \quad (8.5)$$

A key advantage of the currents representation is that a metric can be defined which does not assume point correspondences between surfaces. Rather than measuring the distances between points on the surface, surfaces are compared at an anatomical level. The distance between two surfaces can then be expressed as the norm of the difference between the surfaces, which is the distance between their currents:

$$\| C_{S_1} - C_{S_2} \|_{W^*}. \quad (8.6)$$

8.3.2 Surface-to-surface Registration

Using currents to represent the meshes, we need a method for computing the deformations. We would like to restrict the transformations to those which preserve the topology of the object and give a one-to-one (invertible) smooth transformation (i.e. a diffeomorphism). Restricting to diffeomorphisms gives non-linear deformations that allow local smooth variations to be captured in the registration. We use a group of diffeomorphisms to allow computations with discrete parametrisation using the Large Deformation Diffeomorphic Metric Mapping (LDDMM) method. LDDMM was used for instance in [Beg et al., 2004] to find the deformations between full 3D images in the context of cardiac anatomy. This

deformation framework can also be used to register surfaces modeled as currents, as shown in shown in [Vaillant and Glaunes, 2005].

The LDDMM framework uses a group of diffeomorphisms constructed through integration of time-varying velocity fields, determined by the initial velocity at time $t = 0$ parameterised by moment vectors, that belong to a RKHS. This yields a diffeomorphism ϕ as well as a differentiable flow of the diffeomorphism ϕ_k for a continuous parameter k within the interval $[0, t]$. At time $k = 0$, we have the identity mapping ϕ_0 . The mapping at time t gives the desired transformation ϕ_t which we require for the atlas estimation. We can then define the path at any point x as $\phi_k(x)$ that leads to the final position $\phi_t(x)$. By following the path that passes through the point x , we can compute easily operations on diffeomorphisms, such as the inverse path from point t to 0. Using this framework, we can minimise the difference between a deformed surface $\phi_t(S_1)$ and another surface S_2 , therefore finding the geodesic path from $\phi_t(S_1)$ to S_2 .

8.3.3 Iterative Estimation of the Atlas

An atlas can be computed from the i patient surfaces S_i using a *forward* approach [Allasonnière et al., 2006] by modeling the observations (the patient meshes) as a noisy deformation of the atlas \hat{S} : $S_i = \phi_i(\hat{S}) + \varepsilon_i$, where ε_i is a residual term that accounts for shape features not represented by the atlas. The deformations ϕ_i can be computed iteratively by minimising the distance from the S_i 's to the atlas, in order to center the atlas with respect to the observations. The shape information is contained within the transformation ϕ_i . Using this approach, statistical analyses can be done on the deformations and the residuals since the framework is well-posed.

The atlas is first initialised by taking the mean of the patient meshes (in the space of currents). This initial atlas is then registered to each of the patients individually. The atlas is then updated through a deformation, in order to minimise the error, for each patient, between the patient mesh and the geometry generated from mapping the atlas onto the patient mesh. Written mathematically, we minimise the error:

$$\varepsilon(\hat{S}) = \|S_i - \phi_i(\hat{S})\|^2 \quad (8.7)$$

for the atlas \hat{S} , for given patient meshes S_i , and deformations (ϕ_i) that map them to the atlas \hat{S} (see [Durrleman et al., 2009a] for details on the minimisation strategy). We then register the updated atlas to the individuals, recompute the atlas and loop until convergence using the forward model (see Algorithm 6).

8.3.4 Surface Representation of the Atlas

Using the atlas-construction method described above, we compute an atlas in the space of currents which can be visualised as the disjoint set of triangles (Dirac currents). This representation is sufficient for computing statistics between objects in the space of currents, however in this case we need a surface representation of the atlas to transport the patient flows into. A general practice for computing this surface is to find the closest patient (in the space of currents) and register this patient to the atlas. Of course, this representation is

Algorithm 6 Atlas Estimation

Require: N segmented patient images (surface meshes).

- 1: Manual rigid alignment of meshes to a reference patient.
- 2: Create initial atlas \widehat{S}^0 as the mean of the patient meshes.
- 3: **loop** {over N until convergence}
- 4: Estimate the transformations ϕ_i that register the atlas \widehat{S}^{N-1} to the individual S_i .
- 5: Update the atlas by minimising the error in 8.7 using the estimated transformations ϕ_i and the atlas \widehat{S}^{N-1}
- 6: **return** Final atlas \widehat{S}^N and the related transformations ϕ_i^N .

biased by which patient is used. To reduce this bias, we followed the minimisation strategy as in [Durrleman et al., 2008a] at a coarse scale, followed by an additional step to average the atlas-to-patient deformations. This average deformation was applied to the closest patient and this mesh was used to initialise the atlas construction pipeline at a finer scale. The initial optimisation is performed at a coarse scale to extract the regional differences in order to obtain a reasonable first estimate of the atlas. This is then refined to capture smaller local shape features (see Algorithm 7).

Algorithm 7 Atlas Refinement

- 1: Compute atlas at a coarse scale for σ_V and σ_W
- 2: Find the closest patient S_j to the atlas and deform the patient to the atlas to get a surface mesh $\widehat{S}_j = \phi_j(S_j)$
- 3: **loop** {over N until convergence}
- 4: Register \widehat{S}_j^N to each patient to get $\phi_{\widehat{S}_j^N, i}^N$
- 5: Average the deformation fields $\phi_{\widehat{S}_j^N, i}^N$ for all i to get $\widehat{\phi}_{\widehat{S}_j^N}^N$
- 6: Apply the average deformation field $\widehat{\phi}_{\widehat{S}_j^N}^N$ to \widehat{S}_j^N
- 7: Use the result from Step 5 as initialisation into the atlas construction at a reduced scale with σ_V and σ_W
- 8: **return** Final refined atlas \widehat{S}^N and the related transformations ϕ_i^N .

8.4 An Atlas-based Reduced Order Model of Blood Flow

8.4.1 Reference Blood Flow Simulation and POD Basis

The reference simulation of blood flow in the pulmonary artery geometries is obtained by solving numerically the incompressible Navier-Stokes equations in the three-dimensional spatial domain of the atlas geometry, denoted by $\widehat{\Omega}$. We assume that the boundary of $\widehat{\Omega}$ is partitioned as

$$\partial\widehat{\Omega} = \Sigma_{\text{in}} \cup \Sigma_{\text{wall}} \cup \Sigma_{\text{out}},$$

which corresponds to the inflow, arterial wall and the outflow boundaries, respectively. The velocity $\widehat{\mathbf{u}} : \widehat{\Omega} \times \mathbb{R}^+ \rightarrow \mathbb{R}^d$ and the pressure $p : \widehat{\Omega} \times \mathbb{R}^+ \rightarrow \mathbb{R}$ are hence described by the

following system

$$\left\{ \begin{array}{l} \rho \partial_t \hat{\mathbf{u}} + \rho \hat{\mathbf{u}} \cdot \nabla \hat{\mathbf{u}} + \nabla \hat{p} - 2\mu \operatorname{div} \varepsilon(\hat{\mathbf{u}}) = 0 \quad \text{in } \widehat{\Omega}, \\ \operatorname{div} \hat{\mathbf{u}} = 0 \quad \text{in } \widehat{\Omega}, \\ \hat{\mathbf{u}} = \mathbf{u}_{\text{in}} \quad \text{on } \Sigma_{\text{in}}, \\ \hat{\mathbf{u}} = 0 \quad \text{on } \Sigma_{\text{wall}}, \\ \boldsymbol{\sigma}(\hat{\mathbf{u}}, \hat{p}) \mathbf{n} = p_{\text{out}} \quad \text{on } \Sigma_{\text{out}}, \end{array} \right. \quad (8.8)$$

where, ρ stands for the density of the fluid and the fluid Cauchy-stress tensor is given by

$$\boldsymbol{\sigma}(\mathbf{u}, p) = -p\mathbf{I} + 2\mu\varepsilon(\mathbf{u}), \quad \varepsilon(\mathbf{u}) \stackrel{\text{def}}{=} \frac{1}{2} (\nabla \mathbf{u} + \nabla \mathbf{u}^T),$$

μ being the dynamic viscosity of the fluid, \mathbf{u}_{in} a given inlet velocity field and p_{out} a given outlet pressure.

At the inlet we prescribe a parabolic velocity profile, whose flow rate varies in time according to a chosen physiological regime. At the outlet, a relationship between pressure and flow is prescribed, in order to represent the pulmonary vessels downstream of the 3D-fluid model. Namely, the outlet pressure p_{out} at each outlet boundary $\Gamma \subset \Sigma_{\text{out}}$ is imposed by a so-called lumped parameter model, in which $p_{\text{out}}(t)$ is related to the outgoing flux at Γ ,

$$Q_{\text{out}}(t) = \int_{\Gamma} \mathbf{u}(t) \cdot \mathbf{n} ds, \quad (8.9)$$

in an analogous way as tension and current are related in electric circuits. In particular, we used a 3-element Windkessel model [Frank, 1899] (see [Vignon-Clementel et al., 2010a] for recent applications in pulmonary artery modeling), in which the outlet pressure and flow are related by the ordinary differential equation:

$$p_{\text{out}} + R_d C \frac{dp_{\text{out}}}{dt} = (R_p + R_d) Q_{\text{out}} + R_p R_d C \frac{dQ_{\text{out}}}{dt}, \quad (8.10)$$

where R_p and R_d model the resistance of the proximal and distal vasculature, respectively, and the capacity C takes into account the deformability of the downstream vessels.

8.4.1.1 Numerical Approximation

Problem (8.8) is discretised in time with a Chorin-Temam projection scheme (see, e.g., [Guermond et al., 2006, Chorin, 1968, Temam, 1968]) in which velocity and pressure are solved separately in two substeps. Namely, let us denote with τ the time-step size, setting $t_n \stackrel{\text{def}}{=} n\tau$ for $1 \leq n \leq N$. For a given initial condition $\mathbf{u}^0 = \mathbf{u}_0$, the time iteration consists of solving the two problems:

1. Velocity step:

$$\left\{ \begin{array}{l} \rho \frac{\hat{\mathbf{u}}^{n+1} - \hat{\mathbf{u}}^n}{\tau} + \rho \hat{\mathbf{u}}^n \cdot \nabla \hat{\mathbf{u}}^{n+1} - 2\mu \nabla \cdot \varepsilon(\hat{\mathbf{u}}^{n+1}) + \nabla \hat{p}^n = 0 \quad \text{in } \widehat{\Omega}, \\ \hat{\mathbf{u}}^n = \mathbf{u}_{\text{in}}(t_{n+1}) \quad \text{on } \Sigma_{\text{in}}, \\ 2\mu \varepsilon(\hat{\mathbf{u}}^{n+1}) \mathbf{n} = 0 \quad \text{on } \Sigma_{\text{out}}, \\ \hat{\mathbf{u}}^n = 0 \quad \text{on } \Sigma_{\text{wall}}. \end{array} \right. \quad (8.11)$$

2. Pressure-Poisson projection step:

$$\left\{ \begin{array}{l} -\frac{\tau}{\rho} \Delta \hat{p}^{n+1} = -\nabla \cdot \hat{\mathbf{u}}^{n+1} \quad \text{in } \hat{\Omega}, \\ \frac{\tau}{\rho} \frac{\partial \hat{p}^{n+1}}{\partial \mathbf{n}} = 0 \quad \text{on } \Sigma_{\text{in}} \cup \Sigma_{\text{wall}}, \\ \hat{p}^{n+1} = p_{\text{out}}^{n+1} \quad \text{on } \Sigma_{\text{out}}. \end{array} \right. \quad (8.12)$$

The 3D-0D coupling on the outlet boundary condition is treated in an implicit fashion (see [Bertoglio et al., 2012]). The discretisation in space of problems (8.11)-(8.12) is performed via continuous piece-wise affine finite elements. We denote by V_h and Q_h the corresponding finite element spaces for the velocity and the pressure, respectively.

8.4.1.2 Proper Orthogonal Decomposition

A proper orthogonal decomposition (POD) of a numerical solution (that is $\{\hat{\mathbf{u}}_h^n\}_{n=1}^N$ and $\{\hat{p}_h^n\}_{n=1}^N$ or, in general, of a given set of data, see [Bergmann et al., 2009, Rathinam and Petzold, 2004] for instance), consists of finding a set of basis functions (orthogonal w.r.t. a given scalar product) which, even containing a small number of elements, can represent sufficiently well the numerical solution. This approach, besides reducing the data size without losing relevant information, allows to perform faster numerical simulations, by restricting the space of the solution to the subspace generated by the POD basis functions.

We computed velocity $\hat{\Phi} = \{\hat{\phi}_i\}_{i=1}^{M_u}$ and pressure $\hat{\Psi} = \{\hat{\psi}_i\}_{i=1}^{M_p}$ POD bases, with $\hat{\phi}_i : \hat{\Omega} \rightarrow \mathbb{R}^3$ and $\hat{\psi}_i : \hat{\Omega} \rightarrow \mathbb{R}$. These POD basis functions are given in terms of their natural decomposition on the finite element basis of V_h and Q_h , respectively. Note that in practice we have $M_u \ll \dim(V_h)$ and $M_p \ll \dim(Q_h)$. Hence, the main idea of model reduction is to perform the spatial approximation of (8.11) and (8.12) on the lower dimensional spaces spanned by $\hat{\Phi}$ and $\hat{\Psi}$, respectively.

For instance, using the velocity POD basis $\hat{\Phi}$, instead of the original finite element basis of V_h , the discrete problem (8.11) can be formulated in terms of the reduced solution

$$\tilde{\mathbf{u}}(\mathbf{x}, t) = \sum_{i=1}^{M_u} \alpha_i(t) \hat{\phi}_i(\mathbf{x}), \quad (8.13)$$

requiring, at each time step, the solution of a linear system of reduced size $M_u \times M_u$.

8.4.2 Individual Reduced Order Blood Flow Simulation

Our goal is to construct a POD basis for the individual patients simulations using the POD basis computed from the flow simulation on the template geometry. The procedure, summarised in Algorithm 8, consists of mapping each 3D-patient mesh onto the template geometry (*3D-shape registration*), then *transporting* the template POD basis onto the individual domain, and finally performing the POD reduction of FE formulation for the patient-specific problem.

8.4.2.1 3D-shape Registration

Let Ω denote the spatial domain of a patient geometry, and let us consider the discretised representations of $\hat{\Omega}$ and Ω , defined by two tetrahedral meshes $\hat{\mathcal{T}}_h$ and \mathcal{T}_h , respectively. Firstly, we are interested in computing a map

$$\mathcal{A} : \hat{\Omega} \rightarrow \Omega \quad (8.14)$$

which maps $\hat{\mathcal{T}}_h$ onto \mathcal{T}_h preserving the mesh topology. In particular, \mathcal{A} defines a one-to-one correspondence between the nodes of the two meshes. In what follows, \mathcal{A} will be called a *3D-shape registration*.

To construct \mathcal{A} , we start from the *surface diffeomorphism* (see section 8.3.2)

$$\phi : \hat{S} \rightarrow S$$

which maps the atlas surface \hat{S} onto an approximation of the patient surface S . Since \hat{S} and S are represented by triangular surface meshes, the diffeomorphism ϕ is equivalent to a piecewise linear deformation defined on the whole surface. Let us denote this deformation as a displacement field $\mathbf{d}_s : \partial\hat{\Omega} \rightarrow \mathbb{R}^3$ such that

$$\partial\Omega = \partial\hat{\Omega} + \mathbf{d}_s(\partial\hat{\Omega}).$$

We aim at defining the map \mathcal{A} by an extension of the \mathbf{d}_s into the three-dimensional domain with a suitable *volume* deformation field

$$\mathbf{d}_v : \hat{\Omega} \rightarrow \mathbb{R}^3,$$

compatible with the surface displacement \mathbf{d}_s on the boundary. In particular, once \mathbf{d}_v is computed, we can define the deformation as

$$\mathcal{A}(\hat{\mathbf{x}}) = \hat{\mathbf{x}} + \mathbf{d}_v(\hat{\mathbf{x}}), \text{ for } \hat{\mathbf{x}} \in \hat{\Omega}. \quad (8.15)$$

In practice, one can consider a harmonic extension of \mathbf{d}_s inside $\hat{\Omega}$, solving the equation

$$\begin{aligned} \Delta \mathbf{d}_v &= 0, & \text{in } \hat{\Omega} \\ \mathbf{d}_v &= \mathbf{d}_s, & \text{on } \partial\hat{\Omega} \end{aligned} \quad (8.16)$$

In general, non-linear extensions of the surface displacement field could also be computed, e.g. considering the volume mesh as an hyperelastic material. This approach is more robust for larger surface deformations. Notice that computing the solution of (8.16) on the reference mesh $\hat{\mathcal{T}}_h$ automatically yields a deformation which preserves the topology of the reference mesh, i.e. such that the inner nodes of the $\hat{\mathcal{T}}_h$ are mapped onto the nodes of the individual volume mesh \mathcal{T}_h .

8.4.2.2 Transporting the Template POD Basis

To construct an individual POD basis for the velocity field on a selected patient, we consider the *inverse Piola transform* of the basis elements of $\hat{\Phi}$, given by

$$\varphi_i(\mathbf{x}) = P(\hat{\varphi}_i)(\mathbf{x}) \stackrel{\text{def}}{=} \frac{1}{J(\hat{\mathbf{x}})} \mathbf{F}(\hat{\mathbf{x}}) \hat{\varphi}_i(\hat{\mathbf{x}}), \quad \hat{\mathbf{x}} = \mathcal{A}^{-1}(\mathbf{x}) \quad (8.17)$$

for $i = 1, \dots, M_u$. Here, \mathbf{F} stands for the deformation gradient, i.e., $\mathbf{F} \stackrel{\text{def}}{=} \nabla \mathcal{A}$, and $J = \det \mathbf{F}$ being the jacobian of \mathbf{F} . From the properties of the Piola transform (see, e.g., [Ciarlet, 1988]), we can infer that

$$J \operatorname{div}_{\mathbf{x}} \varphi_i(\mathbf{x}) = \operatorname{div}_{\hat{\mathbf{x}}} \hat{\varphi}_i(\hat{\mathbf{x}}).$$

Hence, if the POD basis is divergence free in the reference geometry then the transformed basis has the same property in the individual geometry.

Considering the shape registration in the form (8.15), we approximate the gradient

$$F(\hat{\mathbf{x}}) = \mathbf{I} + \nabla \mathbf{d}_v(\hat{\mathbf{x}}),$$

by its projection on the piecewise linear finite element space. In practice, this is done via an L^2 -projection on the finite element space, namely, by inverting the *mass matrix*, whose entries are given by $\hat{M}_{ij} = \int_{\hat{\Omega}} \hat{\mathbf{v}}_i \hat{\mathbf{v}}_j d\hat{\Omega}$, $\hat{\mathbf{v}}_i$ and $\hat{\mathbf{v}}_j$ being basis function of the finite element space in $\hat{\Omega}$. For the pressure, an individual POD basis Ψ can be obtained from the template POD basis $\hat{\Psi}_i$ through a simple coordinate change, namely,

$$\psi_i(\mathbf{x}) \stackrel{\text{def}}{=} \hat{\psi}_i(\hat{\mathbf{x}}), \quad \hat{\mathbf{x}} = \mathcal{A}^{-1}(\mathbf{x}) \quad (8.18)$$

for all $i = 1, \dots, M_p$.

8.4.2.3 Reduced Simulation

With a little abuse of notation, let us denote with $\Phi = (\varphi_1 | \dots | \varphi_{M_u})$ an individual POD basis in matrix form, where each column contains a basis element, obtained with the procedure described in Sections 8.4.2.1-8.4.2.2.

At each time step t_n , let us call A^n and \mathbf{f}^n the matrix and the right-hand side of the linear system associated to the finite element approximation of problem (8.11) in the patient geometry. The reduced model is given by

$$(\Phi^T A^n \Phi) \tilde{\mathbf{U}} = \Phi^T \mathbf{f}^n, \quad (8.19)$$

where $\tilde{\mathbf{U}} = \{\alpha_1 \dots \alpha_{M_u}\}^T$ defines the reduced numerical solution (8.13). Analogously, the reduced system for the pressure-problem (8.12) is built by considering the transported pressure basis $\Psi = (\psi_1 | \dots | \psi_{M_p})$.

Remark 8.4.1 A POD basis transported as in (8.17) and (8.18) is associated to a single patient mesh, and can be employed only to reduce a simulation for this specific geometry. On the other hand, the POD basis is decoupled from other simulation parameters (for example the boundary conditions such as the inlet velocity in (8.11) or model (8.10)). Hence, the same basis can thus be used for different physiological regimes. This will be investigated in Section 8.5.3.1.

Algorithm 8 Individual reduced order simulations

- 1: **Given:** Individual surface S_i and reference POD bases $(\hat{\Phi}, \hat{\Psi})$
 - 2: Compute volume deformation \mathcal{A}_i from the surface diffeomorphism ϕ_i
 - 3: Compute patient-specific deformation gradient \mathbf{F}_i
 - 4: Compute POD bases (from (8.17) and (8.18))
 - 5: Assemble FE formulations and reduce the linear systems (e.g., (8.19) for the velocity)
 - 6: **return** Reduced order solution
-

8.5 Numerical Experiments on the Pulmonary Artery of Repaired Tetralogy of Fallot Patients

8.5.1 Data Collection

Subjects and Image Preparation A data-set of 17 adults with repaired tetralogy of Fallot was used in this study. MRI angiography of the heart was acquired with a 1.5T scanner (Signa excite, GE Medical Systems) with isotropic in-plane resolution $0.703\text{mm} \times 0.703\text{mm}$ and 1 mm thick slices.

Image Segmentation of the Pulmonary Artery In order to extract the surfaces of the pulmonary artery, a user-guided 3D image segmentation tool was applied to delineate the boundaries of the artery. Using this tool, a small number of control points (less than 100) are added by the user to define the inside, outside, and border of the region. With these control points as a guide, a 3D mesh is constructed by an implicit variational surfaces approach. The tool is included within the CardioViz3D software package available for download¹, see [Mansi, 2010] for further details. This tool was applied to each of the patient images to define the artery for at least 3cm of the inflow before the bifurcation and at least 2cm of each outflow branch after the bifurcation (see Fig. 8.3). The variable branch length in the final geometries resulted from the variability of the images used, so in the case when it was possible to include more than 2cm of the outflow, this was included. In particular, images with little image information at the outflow branches resulted in a shorter segmentation.

8.5.2 Statistical Shape Model of the Pulmonary Arteries

The obtained meshes were used to compute a centered atlas as a reference for the population (see Fig. 8.4 and 8.5). The atlas construction pipeline described in Sec. 8.3 requires two parameters to be set to control the ‘stiffness’ of the non-linear deformations (σ_V), where higher values give more global transformations, such as rigid body transformations, and a parameter that characterises the resolution of the currents representation to control how finer deformations are treated as either noise or shape features (σ_W). Since we were mainly interested in the regional ToF alterations related to dilation, valve enlargement, and regional bulging, these parameters were set to $\sigma_W = 30\text{mm}$, $\sigma_V = 5\text{mm}$ for the atlas. With

¹<http://www-sop.inria.fr/asclepios/software/CardioViz3D/>




Figure 8.3: Left: 2D slice of one patient with the segmented outline in yellow. Right: 3D rendering of the same patient with the 3D mesh overlaid (in yellow).




Figure 8.4: The 17 patient meshes (red) and associated atlas mesh (green).

the algorithm described in Sec. 8.3, an atlas was constructed by 4 iterations of the alternate minimisation (the number of iterations needed to reach the convergence criteria). This was sufficient to give a well-centered geometry in which to apply the atlas refinement pipeline described in Sec. 8.3.4.




Figure 8.5: Front view (left), side view (center) and top view (right) of the atlas (blue) and the 17 meshes used to create the atlas (wire-frame red). Though there is a wide shape variability in the population, the atlas is well-centered.

Un-biased Validation of the Atlas To test how biased the atlas is to the population used to create it, we performed a *leave-one-out* validation by creating 17 atlases using 16 patients each. The resulting atlases are shown overlaid on one another in Fig. 8.6. Though there is a large shape variability observed in the population, there is little difference between each of the computed atlases.




Figure 8.6: Front view (left), side view (center) and top view (right) of the atlas created on all the patients (blue) and the 17 atlases created on 16 patients (wire-frame red). There is little difference between the atlases, even given the high variability of shape in the population, which displays the un-biased property of the atlas construction step.

8.5.3 Patient-specific Full and Reduced order CFD Simulations

For the full CFD simulations on the individual geometries, rather than the original patient meshes, we used the geometries obtained by mapping the atlas to each patient via the deformation computed in the 3D-shape registration step (described in Sec. 8.4.2.1). Note that this results in a discrepancy between the original meshes and the atlas-to-patient deformed meshes, since we allow for some noise in the atlas construction step. However, due to the fact that the atlas is well-centered, these differences are small. The simulation approach described in Sec. 8.4 was applied for normal (functioning valve case) and pathological (regurgitant or absent valve) regimes. The aim was to investigate how such a reduced ap-

proach performs for both hemodynamics conditions. For each regime, a full simulation was performed on the atlas, extracting a POD basis containing 30 modes. Then, for each of the individuals, a reduced order simulation was performed with this transported POD and the hemodynamic boundary conditions of that regime. Subsequently, for both regimes, the PODs built under *rest* conditions were tested under *exercise* conditions. Furthermore, the robustness of the method was assessed by investigating if the POD constructed for the reference pathological condition was able to capture other pathological conditions.

To assess the reduced order method, we simulated the flow in the different geometries both with a full FE model and employing the reduced POD basis as described in Sec. 8.4. For each simulation, we monitored three different errors in time indicating the global errors in velocity, pressure and the pressure drop between the right ventricle and the outlets. The latter is particularly interesting from the clinical point of view, and the first two aid in assessing the accuracy of the reduced flow simulation:

- Instantaneous L^2 -norm difference in velocity and pressure, adimensionalised by the maximum L^2 -norm in time:

$$E_{\mathbf{u}} = \frac{\|\mathbf{u}_{\text{reduced}} - \mathbf{u}_{\text{full}}\|_{L^2}}{\max(\|\mathbf{u}_{\text{full}}\|_{L^2})} \quad \text{and} \quad E_p = \frac{\|p_{\text{reduced}} - p_{\text{full}}\|_{L^2}}{\max(\|p_{\text{full}}\|_{L^2})} \quad (8.20)$$

- Instantaneous pressure drop (difference between averaged over the surface inlet and outlet pressures), adimensionalised by the maximum value in time:

$$E_{\Delta p} = \max_{i=1,2} \left(\frac{\|(p_{\text{in}} - p_{\text{out}}^i)_{\text{reduced}} - (p_{\text{in}} - p_{\text{out}}^i)_{\text{full}}\|}{\max(p_{\text{in}} - p_{\text{out}}^i)_{\text{full}}} \right) \quad (8.21)$$

Note that with this number of POD modes, the errors between the full and reduced simulations on the atlas itself are negligible (maximum in time of the velocity, pressure and pressure drop errors below 0.1% in normal conditions). Furthermore, increasing the number of modes does not change the errors between the full and reduced patient simulations.

8.5.3.1 Flow Simulations in Normal and Pathological Regimes

First, we used the pre-computed reduced models to simulate the flow in the individual geometries with different flow conditions. In particular, we distinguished between a *normal cardiac flow condition*, i.e. with very little backflow as is the case in the pulmonary artery of a healthy pulmonary valve, and a *pathological flow condition* with 30% of backflow during diastole, the average amount of expected backflow for patients without pulmonary valves [Schwartz et al., 2011]. Note that flow curves weren't available for all patients, thus the pathological flow curve corresponds to a single patients' measurement, typical of such ToF physiology. The same OD-Windkessel model was applied at both outlets, but the parameters were chosen to obtain in the simulations pressure waveforms typical of a normal pulmonary artery and right ventricle (pathological pressure in the pulmonary artery due to the absence of valve) respectively (see table 8.1). The inlet flow profiles and the resulting

outlet pressures for the atlas simulations for these two cases are depicted in Fig. 8.7 and 8.8. In particular, one can see the characteristic pressure bump at the end of diastole that is observed in the pulmonary artery, as in the right ventricle, for the pathological case. Furthermore, we investigated different degrees of pathological conditions, with 15% and 40% of backflow during diastole (see Fig. 8.8). The errors (8.20)-(8.21) for normal flow




Figure 8.7: Normal boundary conditions: flow rate is prescribed at the inlet (with a mean value of $4.9 \text{ L} \cdot \text{min}^{-1}$) and outlet pressure for the atlas obtained the Windkessel model in Table 8.1. These outlet pressures change slightly from one geometry to another.




Figure 8.8: **Left.** Inlet flow (mean inlet flow rate of $3.5 \text{ L} \cdot \text{min}^{-1}$) and outlet pressure for the atlas under pathological boundary conditions. **Right** pathological inlet flows with different backflows. The prescribed inlet flows with 15% and 40% of backflow respectively present a mean value of 4.3 and $3.1 \text{ L} \cdot \text{min}^{-1}$.

conditions for the different patients are shown in Fig. 8.9. To provide a better picture of the error variability among patients, mean errors and standard deviations are reported in Fig. 8.10. The error in velocity varies over time and among patients between 3% and 42%, with an average curve rising from 7% to 27% and then decreasing slowly down back to 7% (mean over time being 15%). Pressure errors vary over time and among patients from negligible values to 14%, typically rising quickly to its peak value and decreasing equally fast to very low values and remaining low for the rest of the cycle. The mean over time is thus 1%. Pressure loss errors follow similar pattern as the pressure, with a peak over time and patients at 45%, and an average over patients and time of 5%.

Fig. 8.11 shows the errors (8.20)-(8.21) for the different patients under pathological

R_p ($\text{dyn} \cdot \text{s} \cdot \text{cm}^{-5}$)	40	R_p ($\text{dyn} \cdot \text{s} \cdot \text{cm}^{-5}$)	107
C ($\text{cm}^5 \cdot \text{dyn}^{-1}$)	10^{-3}	C ($\text{cm}^5 \cdot \text{dyn}^{-1}$)	0.32
R_d ($\text{dyn} \cdot \text{s} \cdot \text{cm}^{-5}$)	300	R_d ($\text{dyn} \cdot \text{s} \cdot \text{cm}^{-5}$)	308

Table 8.1: RCR Windkessel model values imposed at each outlet for normal (left) and pathological inlet flow conditions (right).




Figure 8.9: Velocity, pressure and pressure drop errors between the full and the reduced order solutions for the same normal boundary conditions.




Figure 8.10: Mean (solid line) and standard deviation (dashed line) for velocity (left), pressure (center) and pressure drop (right) errors between the full and the reduced order solutions with the same normal boundary conditions.

flow conditions. The corresponding mean errors and standard deviations are further reported in Fig. 8.12. We obtained an error in velocity that varies over time and among patients between 10% and 47%, with an average curve going from 18% to 30% with a double-bump waveform (mean in time being 23%). Conversely, pressure errors vary over time and among patients from negligible values to below 5%, also with a double-bump waveform but with a much lower second peak. The mean over time is 1%. Pressure loss errors follow a four-bump pattern, with a negligible minimum value and a peak of 40% over time and amongst patients, and an average over patients and time of 11%.

To better visualise the underlying differences of the 3D fields, velocity magnitude 3D-cuts of the full and the reduced simulations are presented for two representative patients under normal and pathological conditions respectively, at different times in the cardiac cycle. In Fig. 8.13, one can see that the main features of velocity magnitude are well captured in the reduced simulation with respect to the full simulation for patient 7 under normal conditions. There are larger differences during deceleration (2nd row), in coherence with the profile of the L^2 -error (Fig. 8.9). Fig. 8.14 shows that the main features of velocity magnitude are also well captured in the reduced simulation with respect to the full simulation for patient 13 under pathological conditions. In accordance with the L^2 -errors (Fig. 8.11), that do not vary much in time but peak at the beginning of deceleration, we observe the highest disagreement during beginning of deceleration (2nd row), although the four snapshots all show differences. For both normal and pathological flows, flow jets through the main pulmonary artery during peak systole, and presents complex structures in the pulmonary arteries at peak backflow and subsequent diastole.

8.5.3.2 Simulations During Exercise Based on the Resting Conditions POD

The inlet boundary conditions for the reduced simulations (step 5 of Algorithm 8) were modified to simulate normal and pathological conditions under exercise regimes, increasing heart rate and systolic flows (Fig. 8.15). The resulting average flow rates are $7.5 \text{ L} \cdot \text{min}^{-1}$ in the normal case and $6.2 \text{ L} \cdot \text{min}^{-1}$ with 20% backflow in the pathological case. For the reduced simulations, we used the reference POD basis computed in normal conditions,

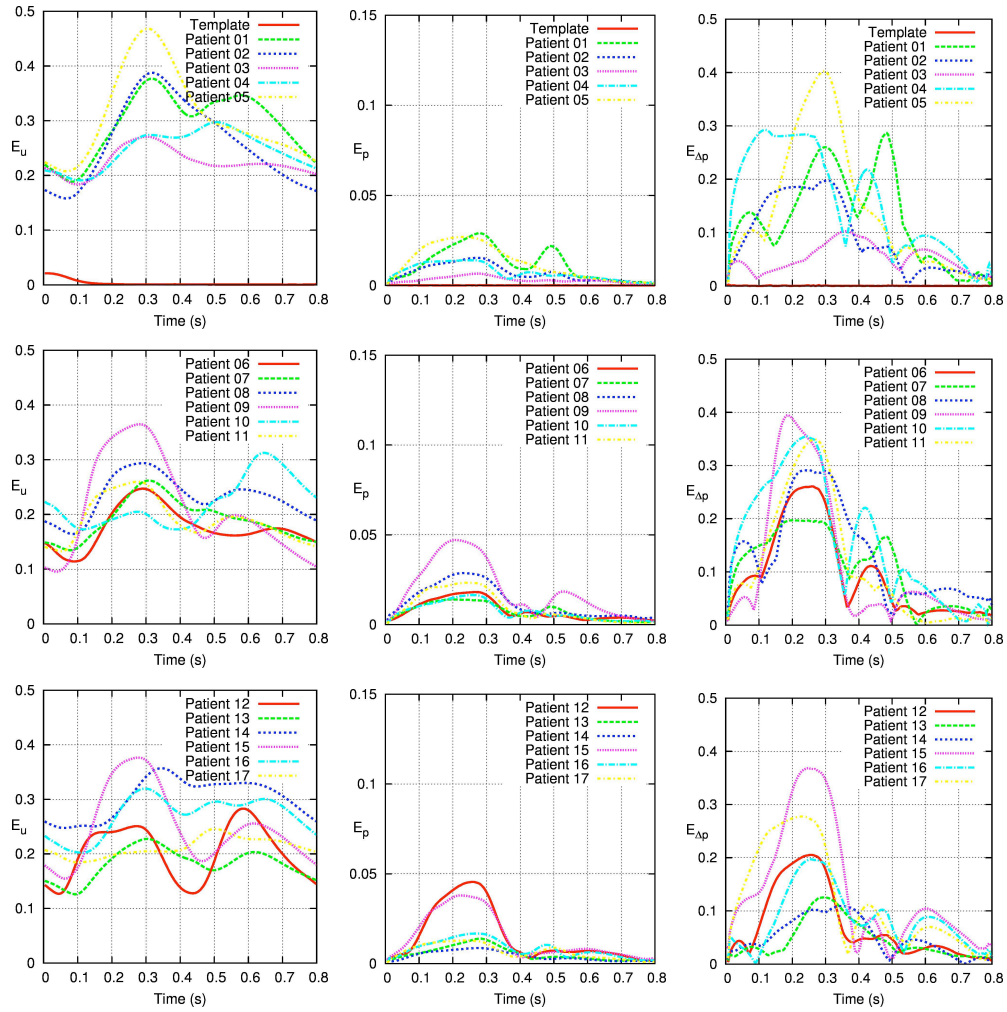


Figure 8.11: Velocity, pressure and pressure drop errors between the full and the reduced order solutions for the same pathological boundary conditions.




Figure 8.12: Mean (solid line) and standard deviation (dashed line) for velocity (left), pressure (center) and pressure drop (right) errors between the full and the reduced order solutions for the same pathological boundary conditions.

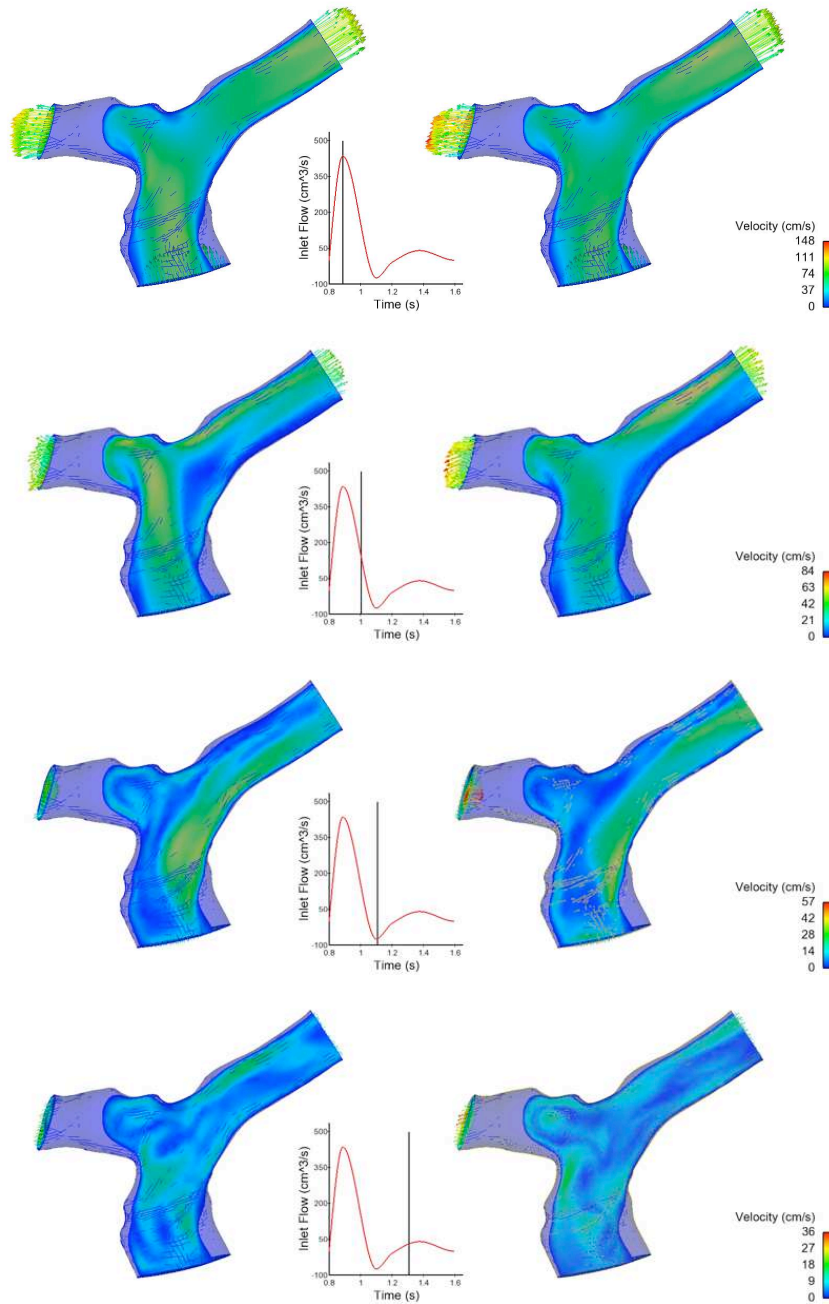


Figure 8.13: Full (left) and reduced (right) velocity fields for patient 07 under normal conditions shown at four instances along the cardiac cycle.

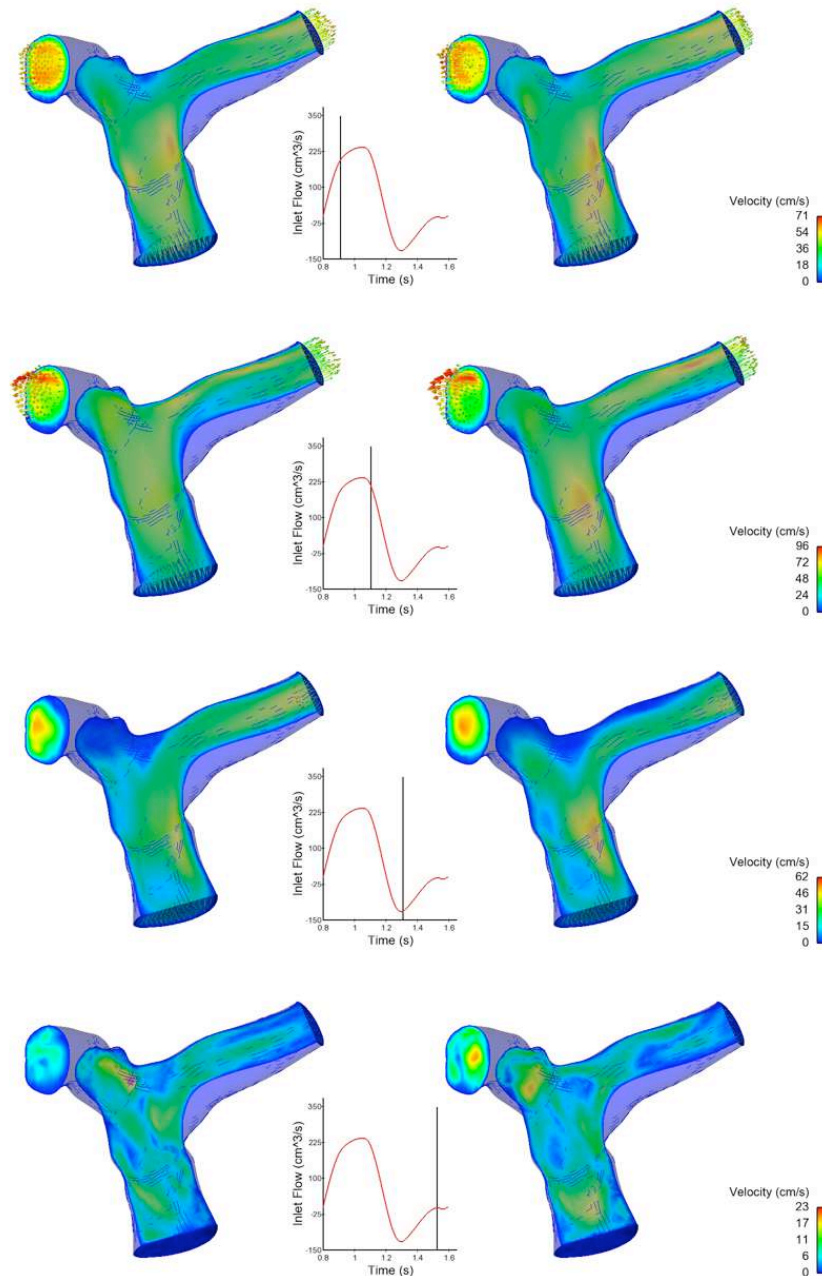


Figure 8.14: Full (left) and reduced (right) velocity fields for patient 13 under pathological conditions shown at four instances along the cardiac cycle.




Figure 8.15: Normal and pathological inlet flows under exercise. Average flow rates are $7.5\text{L} \cdot \text{min}^{-1}$ in the normal case and $6.2\text{L} \cdot \text{min}^{-1}$ with 20% backflow in the pathological case.

thus without performing an additional full computation on the reference geometry with the new boundary conditions. Fig. 8.16 summarises the results for the different patients. The error in velocity on average rises from 7% to 30% and then decreases slowly back to 7%. Pressure errors vary over time on average from negligible values to 7.5%, typically rising rapidly to its peak value and decreasing equally quick to very low values and remaining low for around 60% of the cycle. Pressure loss errors follow similar pattern as the pressure, with a peak at 25%. Overall, the mean and standard deviation time variations and average are very similar to the resting conditions. Figure 8.17 displays the results of the exercise




Figure 8.16: Mean (solid line) and standard deviation (dashed line) for velocity (left), pressure (center) and pressure drop (right) errors between the full and the reduced order solutions under normal exercise boundary conditions (computed with the resting condition POD basis).

reduced simulations for pathological conditions. As in the normal case, we used the POD basis computed on the atlas under rest conditions. In this case, the error in velocity on average rises from 14% to 32%, then decreasing with a small second bump back to 14%. Pressure errors vary over time from negligible values to 3%, remaining very low for around 60% of the cycle. Pressure loss errors follow a four-bump evolution, with a peak at 28%. Overall the mean and standard deviation time variations are similar to resting conditions,

but with less pronounced secondary bumps.




Figure 8.17: Mean (solid line) and standard deviation (dashed line) for velocity (left), pressure (center) and pressure drop (right) errors between the full and the reduced order solutions under pathological exercise boundary conditions (computed with the pathological resting conditions POD basis).

8.5.3.3 Simulations of Different Pathological Conditions Based on the Reference Pathological POD

To further test the robustness of the POD approach, the atlas POD basis computed with 30% backflow was used as reduced basis for the flow in the different patients under other pathological conditions, considering 15% and 40% of backflow. Mean errors and standard deviation results for these pathological conditions are respectively shown in Fig. 8.18 and Fig. 8.19. During systole, the flow conditions are the same for all the pathological con-




Figure 8.18: Mean (solid line) and standard deviation (dashed line) for velocity (left), pressure (center) and pressure drop (right) errors between the full and the reduced order solutions using pathological boundary conditions with 15% of backflow (using original pathological POD bases obtained with 30% of backflow).

ditions, and there is no increase of error during that period. In fact, pressure losses even decrease for the 40% backflow case. During diastole, the 15% flow reversal case has a decrease in all errors compared to the 30% flow reversal case on which the POD was constructed. The velocity error decays at the same rate as for normal flow conditions. The

averages over patients and time are 21% for velocity, 1% for pressure and 9% for pressure loss. For the higher flow reversal case (40%), there is an increase in errors (especially in velocity) after peak flow reversal. The velocity error has a second bump as high as the first one, even though in absolute value peak backflow is not as high as peak forward flow, and similarly for its derivative (re-acceleration versus deceleration). Pressure errors are, on the other hand, only slightly higher. The averages over patients and time are 26% for velocity, 1% for pressure and 10% for pressure loss.




Figure 8.19: Mean (solid line) and standard deviation (dashed line) for velocity (left), pressure (center) and pressure drop (right) errors between the full and the reduced order solutions using pathological boundary conditions with 40% of backflow (using original pathological POD bases obtained with 30% of backflow).

8.6 Discussion

8.6.1 Results of Reduced Flow Simulations

The reduced model framework has been tested with 17 tetralogy of Fallot patients to create the atlas-centered reference geometry. Although the pipeline has been designed to be able to simulate reduced-order blood flow on new patients, in this work we tested the method on the same patients the atlas was created from. This is justified by the fact that the atlas was shown to be unbiased (see the *leave-one-out* validation step of Fig. 8.6), and since the POD basis is computed solely on the atlas geometry, this should have very little effect on the conclusions of the study. From a qualitative point of view, the numerical tests show that the flow dynamics can be well captured by the reduced model in both normal and pathological conditions. However, when looking at velocity magnitude maps at specific times and with their own scales, differences can be more easily seen (Fig. 8.13 and 8.14).

Under normal conditions (Fig. 8.10), the error is lowest for pressure with maximum in general below 10% (Fig. 8.9), and an average over patients and time of 1%. Higher errors were in general obtained for pressure drop. While during diastole both pressure errors are close to zero, the peaks of these pressure errors are located at the beginning of flow deceleration, where flow patterns appear to be more dependent on the geometry. However, the average over time of the mean pressure drop curve (Fig. 8.10) is 5% only. Concerning the velocity field, the error norm in time resulted between 3% and 42%, with relatively

high variability among subjects (Fig. 8.9). But the average over time of the mean curve (Fig. 8.10) is only 15%. In general, note that the error is not proportional to flow, and, as for the pressure and pressure drop, the maximum error occurs during early deceleration. Moreover, the error does not decrease to zero when flow does, but it rather decreases when the flow approaches a steady state, and velocity is lower in the whole domain.

With pathological boundary conditions, and employing a mapped POD basis computed with the same conditions, the peak pressure error is twice lower than in the normal case, with a lower variability between the subjects (lower standard deviation in Fig. 8.12). But the time average remains at 1%. Unlike the case of normal conditions, in this situation the average error follows closely the absolute flow time variation, with the highest bump occurring during high flow (or high pressure, which in this case is almost in phase with flow, see Fig. 8.8). At the same time, the error decrease is less evident than in the normal case, and the error increases again in correspondence to peak flow reversal, which is higher in this pathological condition. In contrast with the pressure field, the velocity and pressure loss errors under pathological flow conditions are larger (increase on average from 15% to 23%, and from 5% to 10% respectively), with more variability among subjects on the magnitude and timing of the error extrema (Fig. 8.9). This suggests that the multiple changes of direction in the flow accentuate geometrical differences and lead to more complex velocity patterns, more sensitive to the individual geometries.

Reduced simulation of exercise conditions were performed with the POD generated under rest conditions. The velocity errors are very comparable on average, with similar time variation and standard deviation as under rest, for both the normal and pathological conditions. This is quite remarkable, since flow peak increases by 50% in the normal case and even by 100% in the pathological case. Concerning the pressure, peak errors in normal conditions generally increase, but remain on the order of 10%, while errors in pressure drop increase slightly. These trends are more pronounced for the pathological regime, with less effect of flow reversal compared to the rest condition, coherently with the fact that there is less backflow under the exercise condition. On average over time however, pressure and pressure drop errors do not change and thus remain low, for both conditions.

8.6.2 General Remarks and Perspectives

The present study shows that the atlas-based reduced order model is able to capture the main hemodynamic features of the flow for a reasonable range of boundary conditions (e.g., rest versus exercise conditions, different degrees of pathologies). This suggests that the main contribution to the error comes from the mismatch between the atlas and individual geometries.

The results also point out the capability of the reduced order method to approximate the pressure fields in different geometries (the time average errors are 1% for pressure and 5-10% for pressure losses). Pressure is mainly set by the Windkessel model which is the same across patients for a given flow regime. This possibly explains why the error in pressure is the lowest. In fact, since pressure represents the main component of the fluid force, this feature could be applied to perform efficient numerical studies of medical device design and stability under varying flow conditions, considering different configurations as

deformations of a reference domain.

Nevertheless, the study of different inlet boundary conditions underlines the sensitivity of the errors to the flow regime (average errors, timings of peak errors), suggesting the importance of using a reference POD basis computed under this regime (normal or pathological conditions respectively). Within a regime however, to simulate exercise with the POD generated at rest worked well. Similarly to simulate different degrees of pathology based on the POD of one pathological case was successful. The importance of the POD basis flow regime has been confirmed by further numerical studies (results not shown), which demonstrated that the POD basis computed with normal boundary conditions does not approximate well the fluid dynamics under pathological conditions. Moreover, we observed that a POD basis combining snapshots of both normal and pathological simulations did not yield any gain (nor loss) in the accuracy (results not shown) of the simulation in the pathological regime.

Another important observation concerns the increase of the error among patients for velocity compared to pressure, especially under complex flow conditions such as in the pathological case. Errors are representative of a geometry, in the sense that, if for a given patient velocity error is among the highest for a given condition, the errors in pressure and for other conditions are among the highest as well.

The influence of the variability of the geometry on the approximation properties of the POD basis remains an open issue. In particular, identifying the geometric parameters at the origin of the highest errors would allow a great advance in the understanding of the flow characteristics and in the construction of better POD bases. However, identifying those that robustly describe flow disturbances is not an easy task [Bijari et al., 2012].

A key point is that the reduced order method has been tested with blood flow solutions, for different regimes, always computed on a single geometry. This surely limits the ability of the transported POD basis to capture the full flow dynamics on different geometries. In order to decrease the error, a possible extension of the method could consider the use of multiple atlas geometries (from different populations) and the computation of different POD bases, which could be combined when building the reduced models for the flow simulations on new patients. A further development could consider generating *intermediate* shapes, interpolating between geometries where flow dynamics is better approximated and the ones where errors were larger. By looking at the point where the error norm increases along these trajectories, we could circumvent the important geometric details that drastically impact the flow. However, more patients are needed to obtain significant clusters.

As a final remark, we observe that the focus of the chapter has been the development of a model order reduction framework, without discussing directly the important issue of computational time. The purpose of reduced-order modeling is obviously to address the reduction of computational time, reducing the *complexity* of the problem (in the presented simulations, the number of degrees of freedom is reduced from order of 10^5 to few tenths). Nevertheless, for a nonlinear problem like the Navier-Stokes equations, it is not trivial to translate a reduction of complexity into a reduction of computational time (see, e.g., [Carlberg et al., 2011]). Indeed, in those cases, the computational cost of the linear system resolution – which is obviously reduced with the number of unknowns – is not necessarily dominant with respect to the computational expense due to the assembly of the

finite element linear system. This issue, which also depends strongly on the structure of the numerical solver, goes beyond the scope of this chapter and will be addressed in future work.

8.7 Conclusion

We presented a model order reduction approach, which combines statistical shape analysis for the computation of a reference geometry and proper orthogonal decomposition (POD). The reduced model framework has been tested with 17 tetralogy of Fallot patients to create the atlas-centered reference geometry. Reduced simulations were computed on these same patients. In order to apply the procedure for the reduced simulation on new patients, the framework requires simply the surface mapping between the atlas and the new geometry, in order to compute the specific reduced basis.

Numerical experiments show that the algorithm yields errors that are acceptable for applications that need to capture pressure and the main velocity features, as, e.g., in the study of medical devices. However, if velocity field or pressure loss are required with higher precision, more accurate approaches may be necessary, which take into account the geometric variability on a smaller scale. This is especially true under pathological conditions. Furthermore, our numerical simulations demonstrate that the reduced model is robust with respect to changes in boundary conditions, in particular to simulate exercise conditions from rest simulations, or different degrees of pathology.

CONCLUSIONS

A Preliminary Approach for a 5D Cardiac Model

Contents

9.1 Chapter Overview	169
9.2 Introduction	169
9.3 Static Growth Modelling using Cross-Sectional Statistics	171
9.4 Dynamic Growth Modelling using Regression	171
9.4.1 Spatio-Temporal Alignment of Transformations	172
9.4.2 Regression of Polyaffine Transformations	172
9.4.3 Cross-Sectional and Longitudinal Model	172
9.5 Coupled Model of Cardiac Growth, Motion and Blood Flow	173
9.6 Perspectives	174
9.6.1 Conclusion	175

This chapter is designed as a discussion for a first step towards tying together the models proposed in each axis of this thesis. We suggest a method that could be used for future development of a coupled model, described here to open the discussion on potential future projects extending on the work in this thesis.

9.1 Chapter Overview

The global objective of this work was to develop a generative growth model of the beating heart, with blood flowing through the ventricles. The purpose of developing such a model is to deform the model to patient-specific geometries to obtain a predictive patient-specific model. Using such a model, we can have an idea of how a heart will evolve over time in terms of anatomy and motion dynamics. A first step towards such a model is to couple the growth and motion models to obtain a 5D model of the beating, growing heart (3D in space +*time* (heart beat) +*time* (years)). In this chapter we briefly discuss an idea for performing such a coupling.

9.2 Introduction

Predictive modelling of heart growth can be used to aid in therapy planning by giving clinicians an idea of how the heart will re-model over time. Given that treatment methods

are continuously being developed to treat heart conditions, simulating the effect of different therapies can give aid in the decision making process. In order to formulate such a model, a method for modelling the growth dynamics, as well as the motion dynamics are first needed. These need to be defined in such a way that the two can then be coupled.

Growth modelling is difficult due to the complex nature of heart development. Different approaches for tackling this issue include modelling techniques such as simulating the heart growth by imposing an underlying mechanical model of the growth, longitudinal modelling from a single subject, and experimental techniques by growing tissue and analysing the result. A theoretical model for stress-modulated growth was proposed in [Lin and Taber, 1995] to study how mechanical load affects the growth and morphogenesis of the developing heart. Longitudinal models require data from several time points over a number of years. Since treatments are proposed on a frequent basis, waiting for years to acquire the data over a life-span of a single patient makes it difficult to analyse the ongoing development of new methods. Furthermore, comparing the data between two patient-specific longitudinal models is not straightforward. In Chapter 3 we discussed a simpler method for predicting the growth using cross-sectional statistical design. This method has the advantage of not requiring longitudinal data of the same patient, and allows a generic model of a given population to be constructed.

The issue of modelling the motion dynamics at a population scale was discussed in Chapter 7. As mentioned, several methods for creating a model of cardiac motion have been proposed using image processing techniques. These include methods using B-spline registration with spatio-temporal alignment motion fields followed by an averaging of the aligned fields [Rao et al., 2003, Perperidis et al., 2004]. A cardiac motion atlas was constructed from tagged-MR images by performing PCA on spatially and temporally aligned motion fields in [Rougon et al., 2004]. These methods, however, require a large number of parameters to describe both the mean motion model, as well as the deformations from the mean motion model to the subject-specific motion models. In contrast, the method proposed in Chapter 7 allows the mean motion model to be defined with a small number of parameters, and furthermore, the subject-specific transformations are described by the same number of parameters. Computing the deviations from the mean requires simple matrix arithmetic operations.

Designing a growth model that also captures the motion dynamics is a challenging task. We propose to use a cross-sectional statistical model to regress the motion of a set of patients against an index of subject growth, in a manner similar to the method described in Chapter 3. To do so, we require a mean motion model to describe the average motion in a population, and a method for computing the deformation from the observations to the mean. Given the formulation of the motion model described in Chapter 7 that requires few parameters, this is used as a basis for developing a coupled model of growth and motion. To the best of our knowledge, there have been no methods proposed that address the coupling of growth with motion to develop a predictive dynamic heart model.

9.3 Static Growth Modelling using Cross-Sectional Statistics

We briefly recap on the static growth model described in Chapter 3, as a basis for developing a dynamic growth model. The model uses statistical techniques to simulate the growth directly from the images rather than implying an underlying biomechanical model.

An atlas of the ventricular surfaces is generated as an “average” of the population. The velocities that describe the deformation from each subject to the atlas are regressed against an index of subject growth (in this case, body surface area (BSA)).

In order to ensure statistically significant results, the dimensionality of the problem is reduced to consider factors related specifically to subject growth. The model reduction technique used for this in Chapter 3 is the partial least squares method (PLS), that has the advantage of being able to compute directly the components most related to BSA.

Modelling the deformations as a function of a growth is a complex problem due to the large number of deformation parameters needed to predict a single growth parameters. Therefore, instead the growth is modelled as a function of the deformations. The relationship can then be inversed using canonical correlation analysis (CCA) to obtain a generative growth model (see Fig. 9.1).




Figure 9.1: The static growth model from Chapter 3, computed from Tetralogy of Fallot patients to capture the pathologically specific growth dynamics.

9.4 Dynamic Growth Modelling using Regression

Using the same concept of regressing observations from different subjects at different stages of the evolution as described in the previous section, we propose to use a similar method to derive a dynamic growth model. Instead of regressing the static observations, we can instead apply regression to transformations computed over the cardiac cycle for each subject, as shown in Fig. 9.2.




Figure 9.2: Left: Linear regression of static observations can be used to develop a growth model as described in Chapter 3. Right: This same concept can be applied to motion transformations, to regress the motion of a set of patients at different stages in the progression of the disease.

9.4.1 Spatio-Temporal Alignment of Transformations

Due to the fact that image sequences are generally aligned differently from one acquisition to another, in order to meaningfully compare the transformations, they need to be first aligned in space and in time. Using the spatio-temporal alignment proposed in Chapter 7, the parameters can be aligned to a common space by resampling the parameters to a common frame, and then realigning the parameters in a rigid manner to fit the mean peak contraction phase (estimated directly from the transformation parameters by taking the trace of the affine matrix per region). Once the parameters are in the same temporal frame, they can be aligned spatially so that all subjects are regionally centered at the same point, and oriented in the same direction. Using the method described in Chapter 7, the reorientation is performed in prolate spheroidal coordinates, to align the regions in an anatomically meaningful manner.

9.4.2 Regression of Polyaffine Transformations

In Chapter 3, the regressed deformations were those obtained by an LDDMM based registration [Durrleman et al., 2009a]. The motion model described in Chapter 5 rather uses Polyaffine transformations. As described in Chapter 7, the mean motion of a population can be estimated by averaging the motion of a set of subjects from the same population (in this case, averaging the Polyaffine transformations after spatio-temporal realignment). Since the transformations of the mean motion model and the subject-specific motion models are described by regional affine matrices, computing the deviation of each subject from the mean requires simple matrix subtraction. As with the static growth model, the variation around the mean could be computed using PLS to obtain factors most related to growth (or in this case, a kind of scaling of the transformations).

9.4.3 Cross-Sectional and Longitudinal Model

The first question that follows from this model is the best way to perform the regression. A simple model using linear regression as in Chapter 3 could be used. Going one step further, we could incorporate data for subjects that have scans at multiple time points by

using multivariate regression or mixed models, for example. Doing so could give better indications of the growth dynamics (see Fig. 9.3). Such methods are discussed in [O'Brien and Fitzmaurice, 2005, Verbeke and Molenberghs, 2009, Diggle et al., 2013].




Figure 9.3: Left: A cross-sectional data-set that implies a negative relationship between age and Y . Right: Assuming the same data was obtained from longitudinal studies with two measurements per subject we may no longer assume a negative relationship. Images obtained from <https://perswww.kuleuven.be/u0018341/documents/longitudinal.pdf>.

9.5 Coupled Model of Cardiac Growth, Motion and Blood Flow

The global objective of this work was towards the development of a coupled model of cardiac growth, motion and blood flow. A first step towards this objective is to couple the growth with the motion to have a growth model of the beating heart was discussed in the previous section (see Fig. 9.4).

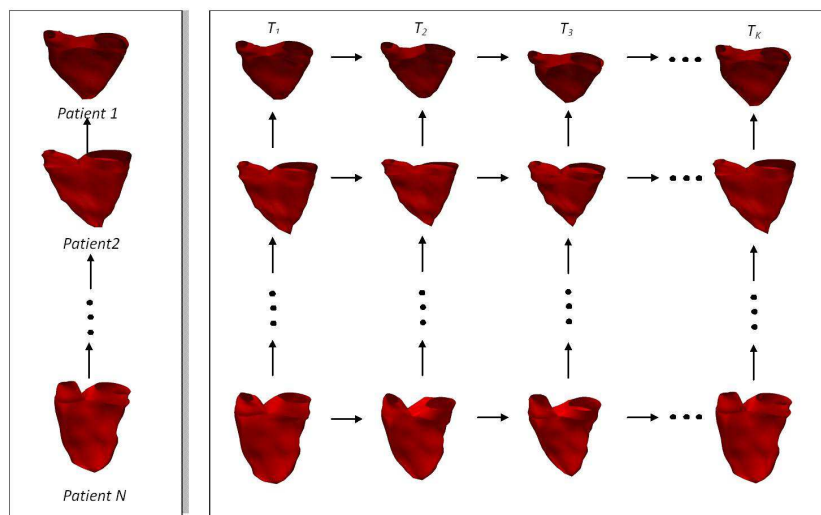


Figure 9.4: Combining the growth and motion dynamics into a common model could give a more informative model of the heart to perform predictive analyses, than using a static growth model. This idea is exemplified here on the right ventricle.

Including the blood flow component is more complex, since the blood flow solutions are still high dimensional, even after model reduction. Moreover, there is an inherent coupling between the anatomy and the blood flow solution. As a first step towards coupling the blood flow model with a growth / motion model, it could be worthwhile first to create a coupled atlas of anatomy / flow by taking snapshots of the flow solution (for example, using streamlines (see Fig. 9.5)).




Figure 9.5: The atlas of the pulmonary artery used in Chapter 8 with extracted streamlines. By coupling the streamlines with the surface, the fluid-structure interaction may be better captured in the atlas construction step.

9.6 Perspectives

Naturally, the primary perspective of this work is to couple the different models to have a combined multiple-factor model of the heart. This chapter touches on the subject of first coupling the models of growth and motion by suggesting a method to perform such a coupling, and then a potential method for also capturing the blood flow dynamics by studying streamlines. However, developing such a model is a very difficult task, hence why few such models have been proposed to address this problem. Having analysed each phenomenon separately, the key questions that need to be addressed in future work in order to approach this problem can now be identified. The most interesting directions to study are summarised below.

Key Questions:

- 1) What is the most appropriate regression method for cross-sectional (and potentially longitudinal) data?
- 2) Should the data be sub-grouped within a population, and if so, how should this be done?

- 3) What are the key features that we want to extract from the model (local shape, regional motion, velocity / pressure)?
- 4) How can the blood flow dynamics be captured in dynamic and growing ventricles?

9.6.1 Conclusion

Throughout this thesis, a primary theme has been population-based modelling of different phenomena. A vital point that was discovered during this PhD was the difficulty in performing population-based analyses on a large number of parameters. This discovery is what motivated the development of the reduced-order models for cardiac growth, motion and blood flow. Naturally, this key point is also valid for coupling models of different phenomena, as the complexity only increases with added model components. The idea of tensor decomposition (used for the mean motion model in Chapter 7) is to have an independent basis for each factor. In this thesis, reduced models (i.e. compact bases) were proposed in order to be able to recombine easily the different phenomena in a kind of tensor form. Tensor decomposition could potentially be performed on the joint tensor of parameters representing the shape / motion / flow patterns. However, exploiting the correlation between factors still remains a key issue for avoiding the combinatorial explosion with the increase in the number of factors.

Conclusions and Perspectives

Contents

10.1 Scope of this Thesis	177
10.2 Main Contributions	178
10.2.1 Bi-ventricular Growth Model	178
10.2.2 Benchmarking of Cardiac Motion Tracking Methods	178
10.2.3 Anatomically Grounded Cardiac Motion Tracking Algorithm	178
10.2.4 Cardiac Motion Pattern Analysis	179
10.2.5 Population-specific Mean Motion Model	179
10.2.6 Population-based Cardiac Blood Flow	179
10.3 Collaborative Work	180
10.3.1 Care4Me and MD-Paedigree European Projects	180
10.3.2 Collaborative Blood Flow Project	180
10.3.3 Cardiac Motion Tracking Challenges	181
10.4 Perspectives and Future Applications	181
10.4.1 Reduced-order Electro-physiological Simulations	181
10.4.2 Reduced-order Modelling with Artificial Devices	181
10.4.3 Inter-modality Cardiac Analysis	182
10.4.4 Longitudinal Reduced-order Models of Other Organs	183
10.5 Conclusion	183

10.1 Scope of this Thesis

Predictive modelling of cardiac growth, motion and blood flow, are challenging areas of research that can provide crucial information for clinicians to aid in diagnosis, classification, therapy planning and prognosis. In this thesis, we presented methods to address each of these areas individually with the aim of having methods that are robust with respect to the data used to build the models, reproducible so that the analyses can be applied easily to a new population and general enough so that new populations can be studied.

The proposed methods were largely based on the analysis of MR images, though the methods are not specific to this modality. Rather, we chose to test the algorithms on this modality due to the large availability of cardiac MR images in the context of congenital heart disease. The methods were applied specifically to the case study of Tetralogy of

Fallot, due to the large shape re-modelling that occurs in these subjects, and due to the clinical challenge of determining the optimal time to intervene, and optimal device type for valve replacement. The methods described in this manuscript can be applied to other heart diseases that undergo significant re-modelling with time due to the pathology.

This thesis was dedicated to better understanding the complex shape, motion, growth and flow dynamics in Tetralogy of Fallot that occur in patients that have damaged or missing pulmonary valves. Due to the high complexity of modelling each component, we applied model reduction techniques to enable robust computation of the growth, motion and flow dynamics. Using reduced-order models we were able to develop population-based models of cardiac growth, motion and blood-flow. Such models can be used to give insight into the healthy growth, motion and blood-flow patterns, to help to identify abnormal patterns related specifically to the pathology.

10.2 Main Contributions

10.2.1 Bi-ventricular Growth Model

In Chapter 3, we extended on the RV growth model of [Mansi et al., 2011b] to couple the bi-ventricular growth dynamics between the left and right ventricles. The coupled growth model was used to predict the growth evolution of repaired Tetralogy of Fallot patients and to establish the modes of shape variation around the mean. By including the left ventricle in the model we were able to capture the coupled growth dynamics between the two ventricles.

10.2.2 Benchmarking of Cardiac Motion Tracking Methods

Working together with the organisers of the 2011 and 2012 STACOM MICCAI workshop cardiac motion tracking challenges, we applied the iLogDemons algorithm to a data-set of real and phantom sequences (in 2011) and to a synthetic data-set (in 2012). This work was aiming towards benchmarking cardiac motion tracking methods to improve the transfer of such methods to the clinical setting by testing state-of-the-art methods on openly available data-sets.

10.2.3 Anatomically Grounded Cardiac Motion Tracking Algorithm

Based on the analysis of state-of-the-art cardiac motion tracking methods, we chose to develop a new method for performing cardiac motion tracking using non-rigid registration, with the transformation from one image to another defined with a low number of intelligible parameters.

The proposed method incorporates cardiac-specific priors into a PolyLogDemons-based model. By incorporating these terms into the model, the accuracy of the registration was improved, compared to using a non-constrained PolyLogDemons model (as tested on synthetic sequences with known ground-truth and real sequences with annotated 4D landmarks). The method was applied to the STACOM 2011 challenge data-set to compare the accuracy of the proposed method to state-of-the-art methods applied to the same data-set.

10.2.4 Cardiac Motion Pattern Analysis

Using the fact that the transformations estimated by the motion tracking method described in Chapter 5 are described with a small number of parameters, we performed a simple analysis of the parameters by first analysing the affine matrices regionally per patient to establish meaningful differences between healthy and pathological cases in terms of synchrony of the trace evolution (where the trace of the affine matrices is related to volume change). We show that the affine parameters can be analysed directly at a regional level to identify regional differences in motion and potentially to detect regions that are functioning poorly.

We proposed a method for analysing the motion over a group of subjects from the same population to determine bases related to the temporal evolution of the LV as well as the spatial patterns through the cardiac cycle. The aim of such analyses is to determine the normal motion patterns observed in healthy patients, to be able to potentially classify, detect, or distinguish abnormal motion in patients with heart disease.

10.2.5 Population-specific Mean Motion Model

By extending further on the regional and population analyses of motion, we proposed a method for constructing a population-specific mean motion model that can be applied to new subjects. The cardiac motion tracking algorithm of Chapter 5 was used to compute the transformations along the cardiac cycle for a group of subjects. We presented a method for spatially and temporally aligning the transformations to a common space so that we can compare and analyse the transformations between subjects. A mean motion model was computed by averaging the aligned transformation parameters. Analysis of the motion patterns was performed by using statistical model reduction techniques on the transformation parameters. Both matrix and tensor-based model reduction was performed to establish dominant spatial and temporal motion patterns within a population.

The mean motion model and statistical model reduction was applied to both healthy and Tetralogy of Fallot subjects to compare the motion patterns in the two groups. From this analysis we were able to identify meaningful differences between the motion dynamics in these two populations.

10.2.6 Population-based Cardiac Blood Flow

A reduced-order model of cardiac blood flow was proposed that makes use of a statistical shape atlas for generating a population-based flow simulation. From the shape atlas, statistical model reduction was performed on the population-based flow simulation, allowing new simulations to be computed on the reduced flow bases. This method was developed as a means of computing a reduced-order “reference” flow solution that can be projected onto new subjects to re-solve the flow solution on the reduced bases. This could potentially be used to reduce the complexity of flow solutions, to reduce the time taken to compute a patient-specific flow simulation, and to obtain population-based analysis of blood flow.

Using the reduced-order model, we showed that it is possible to extract meaningful bases of the pressure, velocity and pressure change. Furthermore, we were able to cap-

ture reasonable flow dynamics from the reduced model. Different degrees of blood flow regurgitation, and different states (rest, exercise) we are examined.

10.3 Collaborative Work

This PhD was partially funded by two European projects (Care4Me and MD-Paedigree). The work conducted during this PhD was in collaboration with several other research teams and with clinicians at hospitals around Europe. Interaction between clinicians and other researchers was an integral part of this PhD, to motivate the development of new methods by everyday clinical challenges. Patient data was provided by clinicians at St. Thomas hospital, London, Necker children's hospital, Paris, OPBG hospital, Rome and from open-source data from the 2011 STACOM MICCAI workshop cardiac motion tracking data-set.

10.3.1 Care4Me and MD-Paedigree European Projects

The Care4Me project (2010-2012) was a project aimed towards improving health-care by advancing medical imaging technology and aiding decision making with the help of medical image analysis techniques. The project brought together 25 partners from 5 countries from academic, industrial and clinical backgrounds. This thesis was focused on the WP4 objective: towards computer-aided detection, quantification and diagnosis, with a specific emphasis on providing decision support with generative models of the heart (leading towards the WP5 objective).

The MD-Paedigree project is a VPH (Virtual Physiological Human)-based initiative to bring together biomedical data, information and knowledge to develop predictive patient-specific models to improve paediatric health-care. This thesis is related to the predictive modelling aspect of the MD-Paedigree project, in regards to modelling of congenital heart disease.

10.3.2 Collaborative Blood Flow Project

The work related to cardiac blood flow simulation was performed as a collaborative project between the REO group at INRIA, Rocquencourt, and the Necker children's hospital in Paris, and the ASCLEPIOS group at INRIA, Sophia Antipolis. This collaboration brought together the different expertise of the group in REO, who have experience in CFD simulation and fluid-structure interactions, and the expertise of the ASCLEPIOS group for image analysis, atlas building, and non-rigid registration.

Patient data from the Necker children's hospital was used to perform patient-specific blood flow simulations in the pulmonary artery, and on these patients an atlas on the pulmonary artery was generated. This work led to collaborative papers; a conference workshop paper presented in Beijing in 2010 [McLeod et al., 2010a] and a journal article submitted to Medical Image Analysis [McLeod et al., 2013].

10.3.3 Cardiac Motion Tracking Challenges

As part of the STACOM MICCAI workshops of 2011 and 2012, there was a large amount of interaction between the workshop challenge organisers and challenge participants to discuss the data, validation, methodology etc. As a result of the 2011 cardiac motion tracking challenge, a collaborative journal article was written between King's College London, Universitat Pompeu Fabra, Philips Research Medisys, Imperial College London, Fraunhofer MEVIS and other clinical / industrial partners [[Tobon-Gomez et al., 2013](#)].

10.4 Perspectives and Future Applications

Future perspectives and conclusions were described in each chapter of this thesis individually. For all of the axes of this work, further validation and testing is anticipated for future work to apply the models on larger data-sets, for a wider range of populations. Each part of this thesis required a large amount of image processing. When possible this processing was performed automatically. Image segmentation was a large burden on this work due to the fact that fully automatic tools were not available during this PhD. Therefore, image segmentation was performed semi-automatically using open-source tools. Naturally, performing the segmentation in this way introduces potential bias into the models. Using an automatic segmentation tool would cut down the processing time significantly. In the current implementation of each method, the codes have not been optimised in any way for time / memory. Since fast computations were not the aim of this work, we did not focus on optimising the implementation, though this could lead towards the methods being more easily transferrable to the clinical setting.

10.4.1 Reduced-order Electro-physiological Simulations

In this work we modelled the growth, motion and blood flow dynamics of the heart. However, the electro-physiological component is also crucial in understanding the heart and analysing cardiac function. Since abnormal electrical conductivity effects the pumping action (and thus the motion) of the heart, studying this component in a similar manner to the other components could provide a more complete picture. Model order reduction could be applied to the electro-physiological bases in a similar manner to the model reduction performed on the blood flow simulations, or to the work described in [[Boulakia et al., 2012](#)] (see Fig. [10.1](#)).

10.4.2 Reduced-order Modelling with Artificial Devices

The models in this work were largely applied to repaired Tetralogy of Fallot patients. The key question for these patients is determining the optimal time to replace the pulmonary valve with an artificial device, the optimal method for implanting the device, and the optimal placement / size of the device (see Fig. [10.4](#)). The next key question for this application is modelling the effects of using different devices, and for predicting the evolution of the heart over time after device implantation. From an anatomical model of the heart, we can

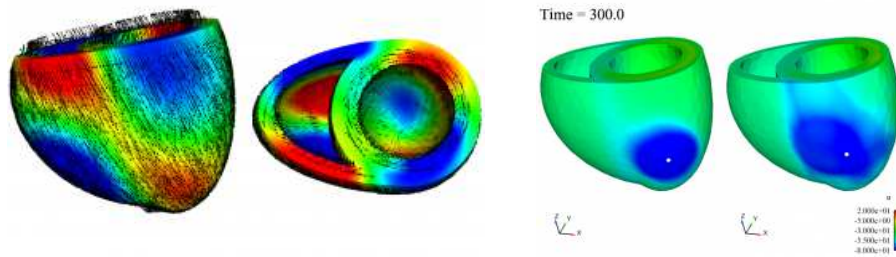


Figure 10.1: Left: An electro-physiological basis computed using a modified Schrodinger operator. Right: Snapshots of a full electro-physiological simulation (left) and a reduced-order simulation performed using a POD method (right). Images obtained from https://team.inria.fr/reo/cardiac_electrophysiology/.

perform virtual surgery on this model. Using the generative bi-ventricular growth model described in Chapter 3, the bi-ventricular heart can be propagated through time to analyse the dynamics over time with respect to the implanted device. The key question would then be whether or not the model is able to predict the shape evolution.




Figure 10.2: 3D image segmentation used to create a 3D model of the pulmonary artery, to create a prototype valve. Image obtained from [Biglino et al., 2012].

10.4.3 Inter-modality Cardiac Analysis

The focus of this work has been on image analysis of cardiac MR images. The tools, however are not specific to this imaging modality. A first step towards multi-modality analysis would be to apply the tools to other modalities to test whether the models yield similar results as those obtained from MR. The bi-ventricular growth model and the reduced-order cardiac blood flow models were applied to surfaces rather than images, therefore these are easily translatable between modalities. However, the cardiac motion model was applied to images directly, therefore there may be some differences between modalities. Given that the cardiac motion tracking algorithm of Chapter 5 is a Demon's based algorithm (and thus driven by image intensities), the results may be different for ECHO, tagged MR and CT. However, the Polyaffine projection is not bound by the use of the LogDemons algorithm

to compute the correspondences between images, therefore this could be adapted to the modality.

For the cardiac motion analyses of Chapters 6 and 7, the analyses are performed on the transformation parameters, rather than the images themselves, therefore the transformation parameters may be easily transferable between different imaging modalities. An interesting extension of the work undertaken in this PhD would be to combine the information from different modalities to build a more informative model of motion (see Fig. 10.3).




Figure 10.3: The Polyaffine transformation parameters could be compared between modalities. Combining the information from different modalities could provide a more robust model of cardiac motion.

10.4.4 Longitudinal Reduced-order Models of Other Organs

This work was focused purely on cardiac modelling, however the tools could be translated to analyse other organs. For example, the cardiac motion tracking algorithm could be used to study longitudinal data on different scales (such as longitudinal scans taken over several years). Longitudinal analysis of brain morphology is a popular area of research, in particular for analysing degenerative brain conditions such as Alzheimer’s disease. To compare the longitudinal evolution between subjects, Lorenzi et. al proposed to transport velocities to a template space using parallel transport [Lorenzi and Pennec, 2012]. In Chapter 7 we transported instead the transformation parameters. This might be more robust than transporting the velocities and requires no pre-alignment of the images.

10.5 Conclusion

The work described in this thesis shows advances in statistical cardiac modelling of the heart towards gaining understanding of the growth, motion and blood flow dynamics specific to a given population. The results indicate that the proposed methods are able to




Figure 10.4: Comparison of the longitudinal evolution of brain scans for patients with Alzheimer's disease. Comparing the evolution of different subjects can be performed by parallel transporting velocities computed from non-rigid registration between image frames. Image courtesy of Marco Lorenzi.

provide useful insight into cardiac disease, that could in future be potentially used to guide therapy planning by providing clinicians with predictive models of cardiac phenomena. The work in this thesis was focused on the application of modelling in Tetralogy of Fallot patients, however the models are not specific to this pathology, and could thus be applied to other pathologies. In order to be able to translate these models into the clinical setting, further research on the pre-processing steps and optimisation of the codes is required to reduce the time needed to prepare the data, and to reduce the running time of the algorithms.

Bibliography

- [Akaike, 1974] Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723. (Cited on page 45.)
- [Alessandrini et al., 2012] Alessandrini, M., Liebgott, H., and Bernard, O. (2012). Monogenic phase based optical flow computation for myocardial motion analysis in 3D echocardiography. In *Proc. STACOM MICCAI Workshop (STACOM'12)*, LNCS. Springer. (Cited on page 76.)
- [Allan et al., 2011] Allan, A., Gao, H., McComb, C., and Berry, C. (2011). Myocardial strain estimated from standard cine MRI closely represents strain estimated from dedicated strain-encoded MRI. In *Proc. Engineering in Medicine and Biological Society (EMBC'11)*, pages 2650–2653. IEEE. (Cited on page 94.)
- [Allasonnière et al., 2006] Allasonnière, S., Amit, Y., and Trouvé, A. (2006). Towards a coherent statistical framework for dense deformable template estimation. *Journal of the Royal Statistical Society: Series B*, 69:2007. (Cited on page 144.)
- [Allasonnière et al., 2007] Allasonnière, S., Amit, Y., and Trouvé, A. (2007). Towards a coherent statistical framework for dense deformable template estimation. *Journal Of The Royal Statistical Society Series B*, 69(1):3–29. (Cited on page 40.)
- [Arsigny et al., 2009] Arsigny, V., Commowick, O., Ayache, N., and Pennec, X. (2009). A fast and log-euclidean polyaffine framework for locally linear registration. *Journal of Mathematical Imaging and Vision*, 33(2):222–238. (Cited on pages 25, 79 and 114.)
- [Arsigny et al., 2006] Arsigny, V., Commowick, O., Pennec, X., and Ayache, N. (2006). A log-euclidean framework for statistics on diffeomorphisms. In *Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI'06), Part I*, volume 4791 of LNCS, pages 924–931. Springer. (Cited on page 55.)
- [Arsigny et al., 2005] Arsigny, V., Pennec, X., and Ayache, N. (2005). Polyrigid and polyaffine transformations: A novel geometrical tool to deal with non-rigid deformations - application to the registration of histological slices. *Medical Image Analysis*, 9(6):507–523. (Cited on pages 77, 79 and 114.)
- [Augustine et al., 2013] Augustine, D., Lewandowski, A. J., Lazdam, M., Rai, A., Francis, J., Myerson, S., Noble, A., Becher, H., Neubauer, S., Petersen, S. E., et al. (2013). Global and regional left ventricular myocardial deformation measures by magnetic resonance feature tracking in healthy volunteers: Comparison with tagging and relevance of gender. *Journal of Cardiovascular Magnetic Resonance*, 15(1):8. (Cited on pages 95 and 96.)
- [Bader et al., 2012] Bader, B. W., Kolda, T. G., et al. (2012). Matlab tensor toolbox version 2.5. Available online. (Cited on page 129.)

- [Beg et al., 2004] Beg, M., Helm, P., McVeigh, E., Miller, M., and Winslow, R. (2004). Computational cardiac anatomy using MRI. *Magnetic Resonance in Medicine*, 52(5):1167–1174. (Cited on page 143.)
- [Beg et al., 2005] Beg, M., Miller, M., Trouvé, A., and Younes, L. (2005). Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International Journal of Computer Vision*, 61(2):139–157. (Cited on page 40.)
- [Bellenger et al., 2000] Bellenger, N., Burgess, M., Ray, S., Lahiri, A., Coats, A., Cleland, J., and Pennell, D. (2000). Comparison of left ventricular ejection fraction and volumes in heart failure by echocardiography, radionuclide ventriculography and cardiovascular magnetic resonance. are they interchangeable? *European Heart Journal*, 21(16):1387–1396. (Cited on page 22.)
- [Bergmann et al., 2009] Bergmann, M., Bruneau, C., and Iollo, A. (2009). Enablers for robust POD models. *Journal of Computational Physics*, 228(2):516–538. (Cited on page 147.)
- [Bertoglio et al., 2012] Bertoglio, C., Caiazzo, A., and Fernández, M. (2012). Fractional-step schemes for the coupling of distributed and lumped models in hemodynamics. *INRIA Research Report*, 7937. (Cited on page 147.)
- [Biglino et al., 2012] Biglino, G., Capelli, C., Binazzi, A., Reggiani, R., Cosentino, D., Migliavacca, F., Bonhoeffer, P., Taylor, A. M., and Schievano, S. (2012). Virtual and real bench testing of a new percutaneous valve device: A case study. *EuroIntervention*, 8(1):120–128. (Cited on page 182.)
- [Bijari et al., 2012] Bijari, P. B., Antiga, L., Gallo, D., Wasserman, B. A., and Steinman, D. A. (2012). Improved prediction of disturbed flow via hemodynamically-inspired geometric variables. *Journal of Biomechanics*, 45(9):1632–1637. (Cited on page 164.)
- [Bistoquet et al., 2007] Bistoquet, A., Oshinski, J., and Skrinjar, O. (2007). Left ventricular deformation recovery from cine MRI using an incompressible model. *IEEE Transactions on Medical Imaging*, 26(9):1136–1153. (Cited on page 78.)
- [Bistoquet et al., 2008] Bistoquet, A., Oshinski, J., and Škrinjar, O. (2008). Myocardial deformation recovery from cine MRI using a nearly incompressible biventricular model. *Medical Image Analysis*, 12(1):69–85. (Cited on page 78.)
- [Biswas, 2010] Biswas, R., editor (2010). *Parallel Computational Fluid Dynamics: Recent Advances and Future Directions*. DEStech Publications. (Cited on page 139.)
- [Bogaert et al., 2012] Bogaert, J., Dymarkowski, S., Taylor, A., and Muthurangu, V. (2012). *Clinical Cardiac MRI*. Medical Radiology / Diagnostic Imaging. Springer. (Cited on page 17.)
- [Bookstein, 1978] Bookstein, F. (1978). *The Measurement of Biological Shape and Shape Change*, volume 24 of *Lecture Notes in Biomathematics*. Springer-Verlag. (Cited on page 35.)

- [Bookstein, 1986] Bookstein, F. (1986). Size and shape spaces for landmark data in two dimensions (with discussion). *Statistical Science*, 1:181–242. (Cited on page 35.)
- [Boulakia et al., 2012] Boulakia, M., Schenone, E., and Gerbeau, J.-F. (2012). Reduced-order modeling for cardiac electrophysiology. Application to parameter identification. *International Journal for Numerical Methods in Biomedical Engineering*, 28:727–744. (Cited on page 181.)
- [Carlberg et al., 2011] Carlberg, K., Bou-Mosleh, C., and Farhat, C. (2011). Efficient non-linear model reduction via a least-squares Petrov-Galerkin projection and compressive tensor approximations. *International Journal for Numerical Methods in Engineering*, 86(2):155–181. (Cited on page 164.)
- [Carroll and Chang, 1970] Carroll, J. D. and Chang, J.-J. (1970). Analysis of individual differences in multidimensional scaling via an n-way generalization of "eckart-young" decomposition. *Psychometrika*, 35(3):283–319. (Cited on page 123.)
- [Cates et al., 2008] Cates, J., Fletcher, P., and Whitaker, R. (2008). A hypothesis testing framework for high-dimensional shape models. In *Proc. of the International Workshop on the Mathematical Foundations of Computational Anatomy (MFCA'08)*, page 170. (Cited on page 42.)
- [Cerqueira et al., 2002] Cerqueira, M., Weissman, N., Dilsizian, V., Jacobs, A., Kaul, S., Laskey, W., Pennell, D., Rumberger, J., Ryan, T., and Verani, M. (2002). Standardized myocardial segmentation and nomenclature for tomographic imaging of the heart. *Circulation*, 105. (Cited on pages 86 and 115.)
- [Chai et al., 2010] Chai, L., Ding, J., and Liu, Y. (2010). *Hemodynamics Simulation of Patient-Specific Surgical Planning For Tetralogy of Fallot*. 2010 3rd International Conference on Biomedical Engineering and Informatics. Yu, W. Zhang M. Wang L. Song Y. (Cited on page 141.)
- [Chandrashekara et al., 2004a] Chandrashekara, R., Mohiaddin, R., and Rueckert, D. (2004a). Cardiac motion tracking in tagged MR images using a 4D B-spline motion model and nonrigid image registration. In *IEEE International Symposium on Biomedical Imaging: Nano to Macro*, pages 468–471. (Cited on page 76.)
- [Chandrashekara et al., 2004b] Chandrashekara, R., Mohiaddin, R. H., and Rueckert, D. (2004b). Analysis of 3-D myocardial motion in tagged MR images using nonrigid image registration. *IEEE Transactions on Medical Imaging*, 23(10):1245–1250. (Cited on page 24.)
- [Chandrashekara et al., 2003] Chandrashekara, R., Rao, A., Sanchez-Ortiz, G., Mohiaddin, R., and Rueckert, D. (2003). Construction of a statistical model for cardiac motion analysis using nonrigid image registration. In *Proc. of Information Processing in Medical Imaging (IPMI'03)*. Springer. (Cited on pages 106 and 116.)

- [Chern et al., 2008] Chern, M.-J., Wu, M.-T., and Wang, H.-L. (2008). Numerical investigation of regurgitation phenomena in pulmonary arteries of Tetralogy of Fallot patients after repair. *Journal of Biomechanics*, 41(14):3002–3009. (Cited on page 141.)
- [Choi et al., 2008] Choi, H. F., Wu, M., Rademakers, F., D’hooge, J., and Claus, P. (2008). Choice of coordinate system for left ventricular FE-mesh generation. In *Proc. European Medical and Biological Engineering Conference (EMBECE’08)*. (Cited on page 86.)
- [Chorin, 1968] Chorin, A. (1968). Numerical solution of the Navier-Stokes equations. *Mathematics of Computation*, 22:745–762. (Cited on page 146.)
- [Ciarlet, 1988] Ciarlet, P. (1988). *Mathematical elasticity. Vol. I*, volume 20 of *Studies in Mathematics and its Applications*. North-Holland Publishing Co. (Cited on page 149.)
- [Clarysse et al., 2002] Clarysse, P., Han, M., Croisille, P., and Magnin, I. (2002). Exploratory analysis of the spatio-temporal deformation of the myocardium during systole from tagged MRI. *IEEE Transactions on Biomedical Engineering*, 149(1):1328–1339. (Cited on page 22.)
- [Clarysse et al., 2011] Clarysse, P., Tafazzoli, J., Delachartre, P., and Croisille, P. (2011). Simulation based evaluation of cardiac motion estimation methods in tagged-MR image sequences. *Journal of Cardiovascular Magnetic Resonance*, 13(Suppl 1):P360. (Cited on page 22.)
- [Collins et al., 1999] Collins, M., Hsieh, A., Ohazama, C. J., Ota, T., Stetten, G., Donovan, C. L., Kisslo, J., and Ryan, T. (1999). Assessment of regional wall motion abnormalities with real-time 3-dimensional echocardiography. *Journal of the American Society of Echocardiography*, 12(1):7–14. (Cited on page 115.)
- [Commowick et al., 2008] Commowick, O., Arsigny, V., Isambert, A., Costa, J., Dhermain, F., Bidault, F., Bondiau, P.-Y., Ayache, N., and Malandain, G. (2008). An efficient locally affine framework for the smooth registration of anatomical structures. *Medical Image Analysis*, 12(4). (Cited on pages 79, 103 and 114.)
- [Commowick et al., 2012] Commowick, O., Wiest-Daesslé, N., and Prima, S. (2012). Automated diffeomorphic registration of anatomical structures with rigid parts: Application to dynamic cervical MRI. In *Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI’12)*, volume 15 of *LNCS*, pages 163–70. (Cited on page 77.)
- [Cosse, 2012] Cosse, A. (2012). Diffeomorphic surface-based registration for MR-US fusion in prostate brachytherapy. In *Proc. Mediterranean Electrotechnical Conference (MELECON’12)*, pages 903–907. IEEE. (Cited on page 77.)
- [Crum et al., 2004] Crum, W., Hartkens, T., and Hill, D. (2004). Non-rigid image registration: Theory and practice. *British Journal of Radiology*, 77(suppl 2):S140–S153. (Cited on page 25.)

- [Crum et al., 2013] Crum, W. R., Modo, M., Vernon, A. C., Barker, G. J., and Williams, S. C. (2013). Registration of challenging pre-clinical brain images. *Journal of Neuroscience Methods*. (Cited on page 23.)
- [D’Agostino et al., 2003] D’Agostino, E., Maes, F., Vandermeulen, D., and Suetens, P. (2003). A viscous fluid model for multimodal non-rigid image registration using mutual information. *Medical Image Analysis*, 7(4):565–575. (Cited on page 25.)
- [Das et al., 2011] Das, A., Gottliebson, W. M., Karve, M., and Banerjee, R. (2011). Comparison of hemodynamic endpoints between normal subject and Tetralogy patient using Womersley velocity profile and MR based flow measurements. *Molecular & Cellular Biomechanics*, 8(1):21–42. (Cited on page 141.)
- [De-Craene et al., 2012] De-Craene, M., Allain, P., Gao, H., Prakosa, A., Marchesseau, S., Hilpert, L., Somphone, O., Delingette, H., Ebeid, S. M., Villain, N., D’hooge, J., Sermesant, M., and Saloux, E. (2012). Synthetic and phantom setups for the second cardiac motion analysis challenge (cMAC2). In *Proc. STACOM MICCAI Workshop (STACOM’12)*, LNCS. Springer. (Cited on page 76.)
- [De Craene et al., 2012] De Craene, M., Duchateau, N., Tobon-Gomez, C., Ghafaryasl, B., Piella, G., Rhode, K., and Frange, A. (2012). SPM to the heart: Mapping of 4D continuous velocities for motion abnormality quantification. In *Proc of IEEE Symposium on Biomedical Imaging: Nano to Macro (ISBI) 2012*. (Cited on page 106.)
- [De Craene et al., 2013] De Craene, M., Marchesseau, S., Heyde, B., Gao, H., Alessandrini, M., Bernard, O., Piella, G., Porras, A., Saloux, E., Tautz, L., Hennemuth, A., Prakosa, A., Liebgott, H., Somphone, O., Allain, P., Ebeid, S., Delingette, H., Sermesant, M., and D’hooge, J. (2013). 3D strain assessment in ultrasound (straus): A synthetic comparison of five tracking methodologies. *IEEE Transactions on Medical Imaging*, (99). (Cited on page 67.)
- [De-Craene et al., 2011] De-Craene, M., Tobon-Gomez, C., Butakoff, C., Duchateau, N., Piella, G., Rhode, K., and Frangi, A. (2011). Temporal diffeomorphic free form deformation (TDFFD) applied to motion and deformation quantification of tagged MRI sequences. In *Proc. STACOM MICCAI Workshop (STACOM’12)*, LNCS. Springer. (Cited on pages 76, 92 and 114.)
- [Declerck et al., 1998] Declerck, J., Feldmar, J., and Ayache, N. (1998). Definition of a four-dimensional continuous planispheric transformation for the tracking and the analysis of left-ventricle motion. *Medical Image Analysis*, 2(2):197–213. (Cited on page 114.)
- [Denton et al., 1999] Denton, E. R., Sonoda, L. I., Rueckert, D., Rankin, S. C., Hayes, C., Leach, M. O., Hill, D. L., and Hawkes, D. J. (1999). Comparison and evaluation of rigid, affine, and nonrigid registration of breast MR images. *Journal of Computer Assisted Tomography*, 23(5):800–805. (Cited on page 24.)

- [Dewey, 2011] Dewey, M. (2011). *Cardiac CT*. Springer Berlin Heidelberg. (Cited on page 17.)
- [Diggle et al., 2013] Diggle, P., Heagerty, P., Liang, K.-Y., and Zeger, S. (2013). *Analysis of longitudinal data*. Number 25. Oxford University Press. (Cited on page 173.)
- [Dru and Vercauteren, 2009] Dru, F. and Vercauteren, T. (2009). An ITK implementation of the symmetric log-domain diffeomorphic demons algorithm. *Insight Journal - 2009 January - June*. (Cited on pages 56 and 57.)
- [Dryden and Mardia, 1991] Dryden, I. and Mardia, K. (1991). Theoretical and distributional aspects of shape analysis. In *Probability Measures on Groups, X (Oberwolfach, 1990)*, pages 95–116, New York. Plenum. (Cited on page 35.)
- [DuBois and DuBois, 1915] DuBois, D. and DuBois, E. (1915). The measurement of the surface area of man. *Archives of Internal Medicine*, 15(5). (Cited on page 46.)
- [Duchateau et al., 2011] Duchateau, N., Craene, M. D., Piella, G., Silva, E., Doltra, A., Sitges, M., Bijmens, B. H., and Frangi, A. F. (2011). A spatiotemporal statistical atlas of motion for the quantification of abnormal myocardial tissue velocities. *Medical Image Analysis*, 15(3). (Cited on page 106.)
- [Durrleman, 2010] Durrleman, S. (2010). *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. (Cited on page 32.)
- [Durrleman et al., 2008a] Durrleman, S., Pennec, X., Trouvé, A., and Ayache, N. (2008a). A forward model to build unbiased atlases from curves and surfaces. In *Proc. of the International Workshop on the Mathematical Foundations of Computational Anatomy (MFCA'08)*. (Cited on pages 40 and 145.)
- [Durrleman et al., 2009a] Durrleman, S., Pennec, X., Trouvé, A., and Ayache, N. (2009a). Statistical models on sets of curves and surfaces based on currents. *Medical Image Analysis*, 13(5):793–808. (Cited on pages 32, 39, 40, 141, 144 and 172.)
- [Durrleman et al., 2009b] Durrleman, S., Pennec, X., Trouvé, A., Gerig, G., and Ayache, N. (2009b). Spatiotemporal atlas estimation for developmental delay detection in longitudinal datasets. In *Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI'09)*, pages 297–304. Springer. (Cited on pages 32 and 116.)
- [Durrleman et al., 2008b] Durrleman, S., Pennec, X., Trouvé, A., Thompson, P., and Ayache, N. (2008b). Inferring brain variability from diffeomorphic deformations of currents: An integrative approach. *Medical Image Analysis*, 12(5):626–637. (Cited on pages 32 and 36.)
- [Ehrhardt et al., 2011] Ehrhardt, J., Werner, R., Schmidt-Richberg, A., and Handels, H. (2011). Statistical modeling of 4D respiratory lung motion using diffeomorphic image registration. *IEEE Transactions on Medical Imaging*, 30(2):251–265. (Cited on page 23.)

- [Feng et al., 2001] Feng, B., Sitek, A., and Gullberg, G. T. (2001). The prolate spheroidal transform for gated SPECT. *Transactions on Nuclear Science*, 48(3):872–875. (Cited on page 86.)
- [Francone et al., 2006] Francone, M., Dymarkowski, S., Kalantzi, M., Rademakers, F. E., and Bogaert, J. (2006). Assessment of ventricular coupling with real-time cine MRI and its value to differentiate constrictive pericarditis from restrictive cardiomyopathy. *European Radiology*, 16(4):944–951. (Cited on page 21.)
- [Frangi et al., 2001] Frangi, A. F., Niessen, W. J., and Viergever, M. A. (2001). Three-dimensional modeling for functional analysis of cardiac images, a review. *IEEE Transactions on Medical Imaging*, 20(1):2–5. (Cited on page 113.)
- [Frank, 1899] Frank, O. (1899). Die grundform des arteriellen pulses. *Zeitung für Biologie*, 37:483–586. (Cited on page 146.)
- [Geva, 2006] Geva, T. (2006). 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. (Cited on page 17.)
- [Glass et al., 1991] Glass, L., Hunter, P., and McCulloch, A. (1991). *Theory of Heart: Biomechanics, Biophysics, and Nonlinear Dynamics of Cardiac Function*. Springer-Verlag. (Cited on page 64.)
- [Glaunès, 2005] Glaunès, J. (2005). *Transport par difféomorphismes de points, de mesures et de courants pour la comparaison de formes et l’anatomie numérique*. Thèse de sciences, Université Paris 13. (Cited on pages 37 and 40.)
- [Glockner et al., 2003] Glockner, J. F., Johnston, D. L., and McGee, K. P. (2003). Evaluation of cardiac valvular disease with MR imaging: Qualitative and quantitative techniques. *Radiographics*, 23(1). (Cited on page 23.)
- [Glowinski et al., 2001] Glowinski, R., Pan, T., Hesla, T., Joseph, D., and Périaux, J. (2001). A fictitious domain approach to the direct numerical simulation of incompressible viscous flow past moving rigid bodies: Application to particulate flow. *Journal of Computational Physics*, 169(2):363–426. (Cited on page 139.)
- [Gorbunova et al., 2012] Gorbunova, V., Sparring, J., Lo, P., Loeve, M., Tiddens, H. A., Nielsen, M., Dirksen, A., and de Bruijne, M. (2012). Mass preserving image registration for lung CT. *Medical Image Analysis*, 16(4):786–795. (Cited on page 23.)
- [Goshtasby, 2012] Goshtasby, A. A. (2012). Image registration methods. In *Image Registration*, pages 415–434. Springer. (Cited on page 23.)
- [Grenander, 1993] Grenander, U. (1993). *General Pattern Theory: A Mathematical Study of Regular Structures*. Oxford University Press Inc., New York, NY. (Cited on page 36.)

- [Guermond et al., 2006] Guermond, J.-L., Mineev, P., and Shen, J. (2006). An overview of projection methods for incompressible flows. *Computer Methods in Applied Mechanics and Engineering*, 195:6011–6045. (Cited on page 146.)
- [Hajnal and Hill, 2010] Hajnal, J. and Hill, D. (2010). *Medical Image Registration*. Biomedical Engineering. Taylor & Francis. (Cited on page 23.)
- [Hansen et al., 2012] Hansen, M. S., Thorup, S. S., and Warfield, S. K. (2012). Polyaffine parametrization of image registration based on geodesic flows. In *Proc. Mathematical Methods in Biomedical Image Analysis Workshop (MMBIA'12)*, pages 289–295. IEEE. (Cited on page 77.)
- [Harshman, 1970] Harshman, R. A. (1970). Foundations of the parafac procedure: Models and conditions for an explanatory multimodal factor analysis. (Cited on page 123.)
- [Helm et al., 2005] Helm, P., Beg, M. F., Miller, M. I., and Winslow, R. L. (2005). Measuring and mapping cardiac fiber and laminar architecture using diffusion tensor MR imaging. *Annals of the New York Academy of Sciences*, 1047(1):296–307. (Cited on page 22.)
- [Heyde et al., 2012] Heyde, B., Barbosa, D., Claus, P., Maes, F., and D'hooge, J. (2012). Three-dimensional cardiac motion estimation based on non-rigid image registration using a novel transformation model adapted to the heart. In *Proc. STACOM MICCAI Workshop (STACOM'12)*, LNCS. (Cited on page 76.)
- [Higgins and Albert de Roos, 2006] Higgins, C. B. and Albert de Roos, M. D. (2006). *MRI and CT of the cardiovascular system*. Wolters Kluwer Health. (Cited on page 17.)
- [Hoffman, 1995] Hoffman, J. (1995). Incidence of congenital heart disease: I. postnatal incidence. *Pediatric Cardiology*, 16(3):103–113. (Cited on page 15.)
- [Hoffman and Kaplan, 2002] Hoffman, J. and Kaplan, S. (2002). The incidence of congenital heart disease. *Journal of the American College of Cardiology*, 39(12):1890–1900. (Cited on page 33.)
- [Holzapfel and Ogden, 2009] Holzapfel, G. A. and Ogden, R. W. (2009). Constitutive modelling of passive myocardium: A structurally based framework for material characterization. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 367(1902):3445–3475. (Cited on page 77.)
- [Hoogendoorn et al., 2009] Hoogendoorn, C., Sukno, F. M., Ordás, S., and Frangi, A. F. (2009). Bilinear models for spatio-temporal point distribution analysis. *International Journal of Computer Vision*, 85(3). (Cited on page 106.)
- [Hsia et al., 2011] Hsia, T. Y., Cosentino, D., Corsini, C., Pennati, G., Dubini, G., Migliavacca, F., and Investigators, M. o. C. H. A. M. (2011). Use of mathematical modeling to compare and predict hemodynamic effects between hybrid and surgical norwood palliations for hypoplastic left heart syndrome. *Circulation*, 124(11 Suppl):S204–10. (Cited on page 139.)

- [Hufnagel, 2010] Hufnagel, H. (2010). *A probabilistic framework for point-based shape modeling in medical image analysis*. Phd thesis, University of Lübeck. (Cited on page 32.)
- [Hundley et al., 1995] Hundley, W. G., Li, H. F., Lange, R. A., Pfeifer, D. P., Meshack, B. M., Willard, J. E., Landau, C., Willett, D., Hillis, L. D., and Peshock, R. M. (1995). Assessment of left-to-right intracardiac shunting by velocity-encoded, phase-difference magnetic resonance imaging a comparison with oximetric and indicator dilution techniques. *Circulation*, 91(12):2955–2960. (Cited on page 23.)
- [Hunold et al., 2005] Hunold, P., Schlosser, T., Vogt, F. M., Eggebrecht, H., Schmermund, A., Bruder, O., Schüller, W. O., and Barkhausen, J. (2005). Myocardial late enhancement in contrast-enhanced cardiac MRI: Distinction between infarction scar and non-infarction-related disease. *American Journal of Roentgenology*, 184(5):1420–1426. (Cited on page 22.)
- [Hutton and Braun, 2003] Hutton, B. F. and Braun, M. (2003). Software for image registration: Algorithms, accuracy, efficacy. *Seminars in Nuclear Medicine*, 33(3):180–192. (Cited on page 24.)
- [Johnson and Christensen, 2002] Johnson, H. J. and Christensen, G. E. (2002). Consistent landmark and intensity-based image registration. *IEEE Transactions on Medical Imaging*, 21(5):450–461. (Cited on page 26.)
- [Joshi et al., 2004] Joshi, S., Davis, B., Jomier, M., and Gerig, G. (2004). Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImage*, 23(Supplement 1):S151 – S160. Mathematics in Brain Imaging. (Cited on page 36.)
- [Joshi et al., 2006] Joshi, S. H., Kaziska, D., Srivastava, A., and Mio, W. (2006). Riemannian structures on shape spaces: A framework for statistical inferences. In *Statistics and Analysis of Shapes, Modeling and Simulation in Science, Engineering and Technology*, pages 313–333. Birkhäuser Boston. (Cited on page 35.)
- [Joshi et al., 2007a] Joshi, S. H., Klassen, E., Srivastava, A., and Jermyn, I. (2007a). A novel representation for riemannian analysis of elastic curves in \mathbb{R}^n . *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 0:1–7. (Cited on page 35.)
- [Joshi et al., 2007b] Joshi, S. H., Klassen, E., Srivastava, A., and Jermyn, I. (2007b). Removing shape-preserving transformations in square-root elastic (SRE) framework for shape analysis of curves. In *EMMCVPR'07*, pages 387–398. (Cited on page 35.)
- [Juh and Suh, 2010] Juh, R. and Suh, T. (2010). Phase correlated image registration to 4D CT for liver stereotactic radiosurgery with image guidance CT. *Medical Physics*, 37:3137. (Cited on page 24.)
- [Kendall, 1989] Kendall, D. (1989). A survey of the statistical theory of shape (with discussion). *Statistical Science*, 4:87–120. (Cited on page 35.)

- [Kilner et al., 2009] Kilner, P. J., Balossino, R., Dubini, G., Babu-Narayan, S. V., Taylor, A. M., Pennati, G., and Migliavacca, F. (2009). Pulmonary regurgitation: The effects of varying pulmonary artery compliance, and of increased resistance proximal or distal to the compliance. *International Journal of Cardiology*, 133(2):157–166. (Cited on page 140.)
- [Klein et al., 2009] Klein, A., Andersson, J., Ardekani, B. A., Ashburner, J., Avants, B., Chiang, M.-C., Christensen, G. E., Collins, D. L., Gee, J., Hellier, P., et al. (2009). Evaluation of 14 nonlinear deformation algorithms applied to human brain MRI registration. *NeuroImage*, 46(3):786–802. (Cited on page 23.)
- [Knauth et al., 2008] Knauth, A. L., Gauvreau, K., Powell, A. J., Landzberg, M. J., Walsh, E. P., Lock, J. E., del Nido, P. J., and Geva, T. (2008). Ventricular size and function assessed by cardiac MRI predict major adverse clinical outcomes late after Tetralogy of Fallot repair. *Heart*, 94(2):211–216. (Cited on page 18.)
- [Kolda and Bader, 2009] Kolda, T. G. and Bader, B. W. (2009). Tensor decompositions and applications. *SIAM Review*, 51(3):455–500. (Cited on pages 123 and 134.)
- [Koo et al., 2011] Koo, B.-K., Erglis, A., Doh, J.-H., Daniels, D. V., Jegere, S., Kim, H.-S., Dunning, A., DeFrance, T., Lansky, A., Leipsic, J., and Min, J. K. (2011). Diagnosis of ischemia-causing coronary stenoses by noninvasive fractional flow reserve computed from coronary computed tomographic angiograms. Results from the prospective multicenter DISCOVER-FLOW (Diagnosis of Ischemia-Causing Stenoses Obtained Via Noninvasive Fractional Flow Reserve) study. *Journal of the American College of Cardiology*, 58(19):1989–1997. (Cited on page 139.)
- [LaDisa et al., 2011] LaDisa, J. F., Dholakia, R. J., Figueroa, C. A., Vignon-Clementel, I. E., Chan, F. P., Samyn, M. M., Cava, J. R., Taylor, C. A., and Feinstein, J. A. (2011). Computational simulations demonstrate altered wall shear stress in aortic coarctation patients treated by resection with end-to-end anastomosis. *Congenital Heart Disease*, 6(5):432–43. (Cited on page 139.)
- [Lang et al., 2005] Lang, R. M., Bierig, M., Devereux, R. B., Flachskampf, F. A., Foster, E., Pellikka, P. A., Picard, M. H., Roman, M. J., Seward, J., Shanewise, J. S., Solomon, S. D., Spencer, K. T., Sutton, M. S. J., and Stewart, W. J. (2005). 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. (Cited on page 18.)
- [Le and Kendall, 1993] Le, H. and Kendall, D. (1993). The Riemannian structure of Euclidean shape space: A novel environment for statistics. *Annals of Statistics*, 21:1225–1271. (Cited on page 35.)

- [Ledesma-Carbayo et al., 2006] Ledesma-Carbayo, M., Mahia-Casado, P., Santos, A., Perez-David, E., Garcia-Fernandez, M., and Desco, M. (2006). Cardiac motion analysis from ultrasound sequences using nonrigid registration: Validation against Doppler tissue velocity. *Ultrasound in Medicine & Biology*, 32(4):483–490. (Cited on pages 94 and 114.)
- [Ledesma-Carbayo et al., 2005] Ledesma-Carbayo, M. J., Kybic, J., Desco, M., Santos, A., Suhling, M., Hunziker, P., and Unser, M. (2005). Spatio-temporal nonrigid registration for ultrasound cardiac motion estimation. *IEEE Transactions on Medical Imaging*, 24(9). (Cited on page 114.)
- [Leeson et al., 2012] Leeson, P., Mitchell, A., Augustine, D., and Becher, H. (2012). *Echocardiography*. Oxford Medical Publications. OUP Oxford. (Cited on page 17.)
- [Legrice et al., 2001] Legrice, I., Hunter, P., Young, A., and Smaill, B. (2001). The architecture of the heart: A data-based model. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 359(1783):1217–1232. (Cited on page 86.)
- [Lin and Taber, 1995] Lin, I.-E. and Taber, L. (1995). A model for stress-induced growth in the developing heart. *Journal of Biomechanical Engineering*, 117(3):343–349. (Cited on page 170.)
- [Lombaert et al., 2012] Lombaert, H., Peyrat, J., Croisille, P., Rapacchi, S., Fanton, L., Cheriet, F., Clarysse, P., Magnin, I., Delingette, H., and Ayache, N. (2012). Human atlas of the cardiac fiber architecture: Study on a healthy population. *IEEE Transactions on Medical Imaging*, 31(7):1436–1447. (Cited on page 24.)
- [Lorenzi and Pennec, 2012] Lorenzi, M. and Pennec, X. (2012). Geodesics, parallel transport & one-parameter subgroups for diffeomorphic image registration. *International Journal of Computer Vision*. (Cited on page 183.)
- [Makela et al., 2002] Makela, T., Clarysse, P., Sipila, O., Pauna, N., Pham, Q. C., Katila, T., and Magnin, I. E. (2002). A review of cardiac image registration methods. *IEEE Transactions on Medical Imaging*, 21(9):1011–1021. (Cited on pages 19 and 24.)
- [Mansi, 2010] Mansi, T. (2010). *Image-Based Physiological and Statistical Models of the Heart, Application to Tetralogy of Fallot*. Thèse de sciences (phd thesis), Ecole Nationale Supérieure des Mines de Paris. (Cited on pages 32, 58 and 150.)
- [Mansi et al., 2009] Mansi, T., Durrleman, S., Bernhardt, B., Sermesant, M., Delingette, H., Voigt, I., Lurz, P., Taylor, A. M., Blanc, J., Boudjemline, Y., Pennec, X., and Ayache, N. (2009). A statistical model of right ventricle in Tetralogy of Fallot for prediction of remodelling and therapy planning. In *Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI'09)*, volume 5761 of *Lecture Notes in Computer Science*, pages 214–221. Springer. (Cited on page 32.)

- [Mansi et al., 2011a] Mansi, T., Pennec, X., Sermesant, M., Delingette, H., and Ayache, N. (2011a). iLogDemons: A demons-based registration algorithm for tracking incompressible elastic biological tissues. *International Journal of Computer Vision*. (Cited on pages [54](#), [55](#), [56](#), [61](#), [64](#), [65](#), [66](#), [76](#), [77](#), [92](#), [94](#), [107](#) and [114](#).)
- [Mansi et al., 2011b] Mansi, T., Voigt, I., Leonardi, B., Pennec, X., Durrleman, S., Sermesant, M., Delingette, H., Taylor, A. M., Boudjemline, Y., Pongiglione, G., and Ayache, N. (2011b). A statistical model for quantification and prediction of cardiac remodelling: Application to Tetralogy of Fallot. *IEEE Transactions on Medical Imaging*. (Cited on pages [7](#), [17](#), [32](#), [45](#) and [178](#).)
- [Manzoni et al., 2011] Manzoni, A., Quarteroni, A., and Rozza, G. (2011). Model reduction techniques for fast blood flow simulation in parametrized geometries. *International Journal for Numerical Methods in Engineering*. (Cited on page [139](#).)
- [Markl et al., 2012] Markl, M., Geiger, J., Jung, B., Hirtler, D., and Arnold, R. (2012). Noninvasive evaluation of 3D hemodynamics in a complex case of single ventricle physiology. *Journal of Magnetic Resonance Imaging*, 35:933–937. (Cited on page [139](#).)
- [Martín-Fernández et al., 2009] Martín-Fernández, M., Martín-Fernández, M., and Alberola-López, C. (2009). A log-euclidean polyaffine registration for articulated structures in medical images. *Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI'09)*, pages 156–164. (Cited on pages [77](#) and [82](#).)
- [McGregor et al., 2009] McGregor, R., Szczerba, D., Muralidhar, K., and Székely, G. (2009). A fast alternative to computational fluid dynamics for high quality imaging of blood flow. *Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI'09)*, pages 124–131. (Cited on page [141](#).)
- [Mcgregor et al., 2008] Mcgregor, R., Szczerba, D., von Siebenthal, M., Muralidhar, K., and Székely, G. (2008). Exploring the use of proper orthogonal decomposition for enhancing blood flow images via computational fluid dynamics. *Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI'08)*, pages 782–789. (Cited on page [141](#).)
- [McLeod et al., 2010a] McLeod, K., Caiazzo, A., Fernández, M. A., Mansi, T., Vignon-Clementel, I. E., Sermesant, M., Pennec, X., Boudjemline, Y., and Gerbeau, J.-F. (2010a). Atlas-based reduced models of blood flows for fast patient-specific simulations. In *Proc. STACOM MICCAI Workshop (STACOM'10)*, volume 6364 of *LNCS*, pages 95–104, Beijing. Springer. (Cited on pages [9](#), [11](#), [138](#), [140](#) and [180](#).)
- [McLeod et al., 2013] McLeod, K., Guibert, R., Caizzo, A., Mansi, T., Fernández, M., Sermesant, M., Pennec, X., Vignon-Clementel, I., Boudjemline, Y., and Gerbeau, J.-F. (2013). Group-wise construction of reduced models for understanding and characterization of pulmonary blood flows from medical images. *Medical Image Analysis*. (Cited on pages [9](#), [10](#), [138](#) and [180](#).)

- [McLeod et al., 2010b] McLeod, K., Mansi, T., Durrleman, S., Sermesant, M., and Pennec, X. (2010b). Statistical analysis of the anatomy: From digital patient to digital population. (Cited on page 11.)
- [McLeod et al., 2013a] McLeod, K., Mansi, T., Sermesant, M., Pongiglione, G., and Pennec, X. (2013a). Statistical shape analysis of surfaces in medical images applied to the Tetralogy of Fallot heart. In *Modeling in Computational Biology and Biomedicine*, LNCS, pages 165–191. Springer. (Cited on pages 7, 10 and 31.)
- [McLeod et al., 2012a] McLeod, K., Prakosa, A., Mansi, T., Sermesant, M., and Pennec, X. (2012a). An incompressible log-domain demons algorithm for tracking heart tissue. In *Proc. MICCAI STACOM Workshop (STACOM'11)*, number 7085 in LNCS. Springer. (Cited on pages 8, 11, 53, 66, 67, 76, 87 and 92.)
- [McLeod et al., 2013b] McLeod, K., Seiler, C., Folgoc, L. L., Toussaint, N., Sermesant, M., and Pennec, X. (2013b). Anatomically constrained regional motion model for cardiac strain analysis. *IEEE Transactions on Medical Imaging*. (Cited on pages 10, 114, 116, 117 and 131.)
- [McLeod et al., 2012b] McLeod, K., Seiler, C., Sermesant, M., and Pennec, X. (2012b). A near-incompressible poly-affine motion model for cardiac function analysis. In *Proc. STACOM MICCAI workshop (STACOM'12)*, LNCS. Springer. (Cited on pages 8, 11, 74, 78 and 106.)
- [McLeod et al., 2013a] McLeod, K., Seiler, C., Sermesant, M., and Pennec, X. (2013a). Spatio-temporal dimension reduction of cardiac motion for group-wise analysis and statistical testing. In *Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI'13)*, LNCS. Springer. (Cited on pages 9, 10, 101, 107, 120, 122, 123, 128 and 133.)
- [McLeod et al., 2013b] McLeod, K., Seiler, C., Toussaint, N., Sermesant, M., and Pennec, X. (2013b). Regional analysis of left ventricle function using a cardiac-specific polyaffine motion model. In *Proc. of Functional Imaging and Modeling of the Heart 2013 (FIMH)*, LNCS. Springer. (Cited on pages 9, 11, 79, 101, 103, 116 and 118.)
- [McVeigh, 1996] McVeigh, E. R. (1996). MRI of myocardial function: Motion tracking techniques. *Magnetic Resonance Imaging*, 14(2):137–150. (Cited on page 23.)
- [Metz et al., 2012] Metz, C. T., Baka, N., Kirisli, H., Schaap, M., Klein, S., Neefjes, L. A., Mollet, N. R., Lelieveldt, B., de Bruijne, M., Niessen, W. J., et al. (2012). Regression-based cardiac motion prediction from single-phase CTA. *IEEE Transactions on Medical Imaging*, 31(6). (Cited on page 106.)
- [Miller et al., 2003] Miller, M., Trouvé, A., and Younes, L. (2003). On the metrics and Euler-Lagrange equations of computational anatomy. *Annual Review of Biomedical Engineering*, pages 375–405. (Cited on page 40.)

- [Miller and Younes, 2001] Miller, M. and Younes, L. (2001). Group actions, homeomorphisms, and matching: A general framework. *International Journal of Computer Vision*, 41(1/2):61–84. (Cited on page [36](#).)
- [Mio et al., 2007] Mio, W., Srivastava, A., and Joshi, S. H. (2007). On shape of plane elastic curves. *International Journal of Computer Vision*, pages 307–324. (Cited on page [35](#).)
- [Moore et al., 2000] Moore, C., Lugo-Olivieri, C., McVeigh, E., and Zerhouni, E. (2000). Three-dimensional systolic strain patterns in the normal human left ventricle: Characterization with tagged MR imaging. *Radiology*, 214:453–466. (Cited on pages [61](#), [63](#), [66](#), [95](#) and [96](#).)
- [Morales et al., 2011] Morales, H. G., Kim, M., Vivas, E. E., Villa-Uriol, M. C., Larrabide, I., Sola, T., Guimaraens, L., and Frangi, A. F. (2011). How do coil configuration and packing density influence intra-aneurysmal hemodynamics? *American Journal of Neuroradiology*, 32(10):1935–1941. (Cited on page [139](#).)
- [Motwani et al., 2012] Motwani, M., Fairbairn, T. A., Larghat, A., Mather, A. N., Biglands, J. D., Radjenovic, A., Greenwood, J. P., and Plein, S. (2012). Systolic versus diastolic acquisition in myocardial perfusion MR imaging. *Radiology*, 262(3):816–823. (Cited on page [22](#).)
- [Nusser et al., 2006] Nusser, T., Höher, M., Merkle, N., Grebe, O. C., Spiess, J., Kestler, H. A., Rasche, V., Kochs, M., Hombach, V., and Wöhrle, J. (2006). Cardiac magnetic resonance imaging and transesophageal echocardiography in patients with transcatheter closure of patent foramen ovale. *Journal of the American College of Cardiology*, 48(2):322–329. (Cited on page [21](#).)
- [O’Brien and Fitzmaurice, 2005] O’Brien, L. M. and Fitzmaurice, G. M. (2005). Regression models for the analysis of longitudinal Gaussian data from multiple sources. *Statistics in Medicine*, 24(11):1725–1744. (Cited on page [173](#).)
- [O’Dell et al., 1995] O’Dell, W. C., Moore, C. C., Hunter, W. C., Zerhouni, E. A., and McVeigh, E. R. (1995). Three-dimensional myocardial deformations: Calculation with displacement field fitting to tagged MR images¹. *Radiology*, 195(3):829. (Cited on page [86](#).)
- [O’Donnell et al., 2009] O’Donnell, D. H., Abbara, S., Chaithiraphan, V., Yared, K., Killeen, R. P., Cury, R. C., and Dodd, J. D. (2009). Cardiac tumors: Optimal cardiac MR sequences and spectrum of imaging appearances. *American Journal of Roentgenology*, 193(2):377–387. (Cited on page [21](#).)
- [O’Donnell et al., 2012] O’Donnell, L. J., Wells III, W. M., Golby, A. J., and Westin, C.-F. (2012). Unbiased groupwise registration of white matter tractography. In *Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI’12)*, pages 123–130. Springer. (Cited on page [24](#).)

- [Osorio et al., 2012] Osorio, E. M. V., Hoogeman, M. S., Romero, A. M., Wielopolski, P., Zolnay, A., and Heijmen, B. J. (2012). Accurate CT/MR vessel-guided nonrigid registration of largely deformed livers. *Medical Physics*, 39:2463. (Cited on page 24.)
- [Paetsch et al., 2004] Paetsch, I., Jahnke, C., Wahl, A., Gebker, R., Neuss, M., Fleck, E., and Nagel, E. (2004). Comparison of dobutamine stress magnetic resonance, adenosine stress magnetic resonance, and adenosine stress magnetic resonance perfusion. *Circulation*, 110(7):835–842. (Cited on page 22.)
- [Pant et al., 2011] Pant, S., Limbert, G., Curzen, N. P., and Bressloff, N. W. (2011). Multiobjective design optimisation of coronary stents. *Biomaterials*, 32(31):7755–7773. (Cited on page 139.)
- [Park et al., 1996] Park, J., Metaxas, D., Young, A. A., and Axel, L. (1996). Deformable models with parameter functions for cardiac motion analysis from tagged MRI data. *IEEE Transactions on Medical Imaging*, 15(3):278–289. (Cited on pages 22 and 75.)
- [Pennec, 2006] Pennec, X. (2006). Intrinsic statistics on riemannian manifolds: Basic tools for geometric measurements. *Journal of Mathematical Imaging and Vision*, 25(1):127–154. (Cited on page 36.)
- [Pennec, 2008] Pennec, X. (2008). Statistical computing on manifolds: From riemannian geometry to computational anatomy. In *Emerging Trends in Visual Computing*, volume 5416 of *LNCS*, pages 347–386. Springer. (Cited on page 36.)
- [Perperidis et al., 2004] Perperidis, D., Lorenzo-Valdes, M., Chandrashekara, R., Rao, A., Mohiaddin, R., Sanchez-Ortiz, G., and Rueckert, D. (2004). Building a 4D atlas of the cardiac anatomy and motion using MR imaging. In *Biomedical Imaging: Nano to Macro, 2004. IEEE International Symposium on*, pages 412–415. IEEE. (Cited on pages 115 and 170.)
- [Perperidis et al., 2005] Perperidis, D., Mohiaddin, R. H., and Rueckert, D. (2005). Spatio-temporal free-form registration of cardiac MR image sequences. *Medical Image Analysis*, 9(5):441–456. (Cited on pages 24, 110 and 132.)
- [Perperidis et al., 2003] Perperidis, D., Rao, A., Lorenzo-Valdés, M., Mohiaddin, R., and Rueckert, D. (2003). Spatio-temporal alignment of 4D cardiac MR images. In *Functional Imaging and Modeling of the Heart (FIMH'03)*, pages 205–214. Springer. (Cited on page 115.)
- [Petitjean et al., 2004] Petitjean, C., Rougon, N., Cluzel, P., Prêteux, F., and Grenier, P. (2004). Quantification of myocardial function using tagged MR and cine MR images. *International Journal of Computer Vision*, 20(6):497–507. (Cited on page 94.)
- [Piella et al., 2012] Piella, G., Porras, A. R., De-Craene, M., Duchateau, N., and Frangi, A. F. (2012). Temporal diffeomorphic free form deformation to quantify changes induced by left and right bundle branch block and pacing. In *Proc. STACOM MICCAI Workshop (STACOM'12)*, LNCS. Springer. (Cited on page 76.)

- [Pitiot et al., 2003] Pitiot, A., Malandain, G., Bardinet, E., and Thompson, P. M. (2003). Piecewise affine registration of biological images. In *Biomedical Image Registration*, pages 91–101. Springer. (Cited on page 25.)
- [Prakosa et al., 2012a] Prakosa, A., McLeod, K., Sermesant, M., and Pennec, X. (2012a). Evaluation of ilogdemons algorithm for cardiac motion tracking in synthetic ultrasound sequence. In *Proc. MICCAI STACOM Workshop (STACOM'12)*, LNCS. Springer. (Cited on pages 8, 11, 53 and 66.)
- [Prakosa et al., 2012b] Prakosa, A., McLeod, K., Sermesant, M., and Pennec, X. (2012b). Evaluation of iLogDemons algorithm for cardiac motion tracking in synthetic ultrasound sequence. In *Proc. STACOM MICCAI Workshop (STACOM'12)*, LNCS. Springer. (Cited on page 76.)
- [Prakosa et al., 2012c] Prakosa, A., Sermesant, M., Delingette, H., Marchesseau, S., Saloux, E., Allain, P., Villain, N., and Ayache, N. (2012c). Generation of synthetic but visually realistic time series of cardiac images combining a biophysical model and clinical images. *IEEE Transactions on Medical Imaging*. (Cited on pages 89 and 107.)
- [Prasad et al., 2011] Prasad, A., To, L. K., Gorrepati, M. L., Zarins, C. K., and Figueroa, C. A. (2011). Computational analysis of stresses acting on intermodular junctions in thoracic aortic endografts. *Journal of Endovascular Therapy*, 18(4):559–68. (Cited on page 139.)
- [Punithakumar et al., 2013] Punithakumar, K., Ayed, I. B., Islam, A., Goela, A., Ross, I. G., Chong, J., and Li, S. (2013). Regional heart motion abnormality detection: An information theoretic approach. *Medical Image Analysis*. (Cited on page 115.)
- [Puranik et al., 2010] Puranik, R., Muthurangu, V., Celermajer, D. S., and Taylor, A. M. (2010). Congenital heart disease and multi-modality imaging. *Heart Lung and Circulation*, 19(3):133–144. (Cited on page 138.)
- [Qian et al., 2011] Qian, Z., Liu, Q., Metaxas, D. N., and Axel, L. (2011). Identifying regional cardiac abnormalities from myocardial strains using nontracking-based strain estimation and spatio-temporal tensor analysis. *IEEE Transactions on Medical Imaging*, 30(12):2017–2029. (Cited on pages 115 and 116.)
- [R Development Core Team, 2009] R Development Core Team (2009). R: A language and environment for statistical computing. *R Foundation for Statistical Computing, Vienna, Austria*. (Cited on page 45.)
- [Rao et al., 2002] Rao, A., Sanchez-Ortiz, G. I., Chandrashekhara, R., Lorenzo-Valdés, M., Mohiaddin, R., and Rueckert, D. (2002). Comparison of cardiac motion across subjects using non-rigid registration. In *Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI'02)*, pages 722–729. Springer. (Cited on page 115.)
- [Rao et al., 2003] Rao, A., Sanchez-Ortiz, G. I., Chandrashekhara, R., Lorenzo-Valdés, M., Mohiaddin, R., and Rueckert, D. (2003). Construction of a cardiac motion atlas from

- MR using non-rigid registration. In *Functional Imaging and Modeling of the Heart (FIMH'03)*, pages 141–150. Springer. (Cited on pages 115 and 170.)
- [Rathinam and Petzold, 2004] Rathinam, M. and Petzold, L. (2004). A new look at proper orthogonal decomposition. *SIAM Journal on Numerical Analysis*, 41(5):1893–1925. (Cited on pages 139 and 147.)
- [Risser et al., 2012] Risser, L., Vialard, F.-X., Baluwala, H. Y., and Schnabel, J. A. (2012). Piecewise-diffeomorphic image registration: Application to the motion estimation between 3D CT lung images with sliding conditions. *Medical Image Analysis*. (Cited on page 23.)
- [Roeleveld et al., 2005] Roeleveld, R. J., Marcus, J. T., Boonstra, A., Postmus, P. E., Marques, K. M., Bronzwaer, J. G., and Vonk-Noordegraaf, A. (2005). A comparison of non-invasive MRI-based methods of estimating pulmonary artery pressure in pulmonary hypertension. *Journal of Magnetic Resonance Imaging*, 22(1):67–72. (Cited on page 22.)
- [Rohlfing et al., 2003] Rohlfing, T., Maurer Jr, C. R., Bluemke, D. A., and Jacobs, M. A. (2003). 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. (Cited on pages 78 and 114.)
- [Rouchdy et al., 2007] Rouchdy, Y., Pousin, J., Schaerer, J., and Clarysse, P. (2007). A nonlinear elastic deformable template for soft structure segmentation: Application to the heart segmentation in MRI. *Inverse Problems*, 23:1017–1035. (Cited on page 98.)
- [Rougon et al., 2004] Rougon, N. F., Petitjean, C., and Preteux, F. J. (2004). Building and using a statistical 3D motion atlas for analyzing myocardial contraction in MRI. In *Medical Imaging*, pages 253–264. International Society for Optics and Photonics. (Cited on pages 115 and 170.)
- [Rueckert et al., 2006] Rueckert, D., Aljabar, P., Heckemann, R. A., Hajnal, J. V., and Hammers, A. (2006). Diffeomorphic registration using b-splines. In *Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI'06)*, pages 702–709. Springer Berlin Heidelberg. (Cited on page 25.)
- [Rueckert et al., 2003] Rueckert, D., Frangi, A. F., and Schnabel, J. A. (2003). Automatic construction of 3-D statistical deformation models of the brain using nonrigid registration. *IEEE Transactions on Medical Imaging*, 22(8):1014–1025. (Cited on page 23.)
- [Rueckert et al., 1999] Rueckert, D., Sonoda, L. I., Hayes, C., Hill, D. L., Leach, M. O., and Hawkes, D. J. (1999). Nonrigid registration using free-form deformations: Application to breast MR images. *IEEE Transactions on Medical Imaging*, 18(8):712–721. (Cited on page 114.)
- [Saad, 2003] Saad, Y. (2003). *Iterative methods for sparse linear systems*, volume 73 of *Society for Industrial Mathematics*. PWS. (Cited on page 57.)

- [Schaerer et al., 2010] Schaerer, J., Casta, C., Pousin, J., and Clarysse, P. (2010). A dynamic elastic model for segmentation and tracking of the heart in MR image sequences. *Medical Image Analysis*, 14(6):738–749. (Cited on page 75.)
- [Scheffler and Lehnhardt, 2003] Scheffler, K. and Lehnhardt, S. (2003). Principles and applications of balanced SSFP techniques. *European Radiology*, 13(11):2409–2418. (Cited on page 21.)
- [Schievano et al., 2007] Schievano, S., Migliavacca, F., Coats, L., Khambadkone, S., Carminati, M., Wilson, N., Deanfield, J. E., Bonhoeffer, P., and Taylor, A. M. (2007). Percutaneous pulmonary valve implantation based on rapid prototyping of right ventricular outflow tract and pulmonary trunk from MR data1. *Radiology*, 242(2):490–497. (Cited on page 17.)
- [Schwartz et al., 2011] Schwartz, M., Rome, J., Gillespie, M., Whitehead, K., Harris, M., Fogel, M., and Glatz, A. (2011). Relation of left ventricular end Diastolic pressure to right ventricular end diastolic volume after operative treatment of Tetralogy of Fallot. *American Journal of Cardiology*. (Cited on page 153.)
- [Schwitter et al., 2001] Schwitter, J., Nanz, D., Kneifel, S., Bertschinger, K., Büchi, M., Knüsel, P., Marincek, B., Lüscher, T., and Von Schulthess, G. (2001). Assessment of myocardial perfusion in coronary artery disease by magnetic resonance a comparison with positron emission tomography and coronary angiography. *Circulation*, 103(18):2230–2235. (Cited on page 22.)
- [Seiler et al., 2012] Seiler, C., Pennec, X., and Reyes, M. (2012). Capturing the multi-scale anatomical shape variability with polyaffine transformation trees. *Medical Image Analysis*. (Cited on pages 77, 78, 79, 80, 81, 96, 98, 114 and 116.)
- [Shi et al., 2012] Shi, W., Zhuang, X., Wang, H., Duckett, S., Luong, D. V., Tobon-Gomez, C., Tung, K., Edwards, P. J., Rhode, K. S., Razavi, R. S., et al. (2012). A comprehensive cardiac motion estimation framework using both untagged and 3-D tagged MR images based on nonrigid registration. *IEEE Transactions on Medical Imaging*, 31(6):1263–1275. (Cited on pages 76 and 114.)
- [Simard and Mailloux, 1988] Simard, P. Y. and Mailloux, G. E. (1988). A projection operator for the restoration of divergence-free vector fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10:248–256. (Cited on page 57.)
- [Sinusas et al., 2001] Sinusas, A. J., Papademetris, X., Constable, R. T., Dione, D. P., Slade, M. D., Shi, P., and Duncan, J. S. (2001). Quantification of 3-D regional myocardial deformation: Shape-based analysis of magnetic resonance images. *American Journal of Physiology-Heart and Circulatory Physiology*, 281(2):H698–H714. (Cited on page 94.)
- [Small, 1996] Small, C. (1996). *The Statistical Theory of Shapes*. Springer series in statistics. Springer. (Cited on page 35.)

- [Somphone et al., 2012] Somphone, O., Dufour, C., and Sherif Makram-Ebeid, L. H., Villain, N., De-Craene, M., Allain, P., and Saloux, E. (2012). Motion estimation in 3D echocardiography using smooth field registration. In *Proc. STACOM MICCAI Workshop (STACOM'12)*, LNCS. Springer. (Cited on page 76.)
- [Song and Leahy, 1991] Song, S. M. and Leahy, R. M. (1991). Computation of 3D velocity fields from 3D cine CT images of a human heart. *IEEE Transactions on Medical Imaging*, 10(3):295–306. (Cited on page 77.)
- [Sørensen et al., 2004] Sørensen, T. S., Körperich, H., Greil, G. F., Eichhorn, J., Barth, P., Meyer, H., Pedersen, E. M., and Beerbaum, P. (2004). Operator-independent isotropic three-dimensional magnetic resonance imaging for morphology in congenital heart disease a validation study. *Circulation*, 110(2):163–169. (Cited on page 21.)
- [Spilker et al., 2007] Spilker, R., Feinstein, J., Parker, D., Reddy, V., and Taylor, C. (2007). Morphometry-based impedance boundary conditions for patient-specific modeling of blood flow in pulmonary arteries. *Annals of Biomedical Engineering*, 35(4):546. (Cited on page 141.)
- [Suinesiaputra et al., 2011] Suinesiaputra, A., Frangi, A. F., Kaandorp, T. A., Lamb, H. J., Bax, J. J., Reiber, J. H., and Lelieveldt, B. P. (2011). Automated regional wall motion abnormality detection by combining rest and stress cardiac MRI: Correlation with contrast-enhanced MRI. *Journal of Magnetic Resonance Imaging*, 34(2):270–278. (Cited on page 115.)
- [Taquet et al., 2011] Taquet, M., Macq, B., and Warfield, S. (2011). Spatially adaptive log-euclidean polyaffine registration based on sparse matches. *Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI'11)*, pages 590–597. (Cited on page 77.)
- [Tautz et al., 2010] Tautz, L., Hennemuth, A., Andersson, M., Seeger, A., Knutsson, H., and Friman, O. (2010). Phase-based non-rigid registration of myocardial perfusion MRI image sequences. In *Biomedical Imaging: From Nano to Macro, 2010 IEEE International Symposium on*, pages 516–519. IEEE. (Cited on page 114.)
- [Tautz et al., 2011] Tautz, L., Hennemuth, A., and Peitgen, H.-O. (2011). Motion analysis with quadrature filter based registration of tagged MRI sequences. In *Proc. STACOM MICCAI Workshop (STACOM'12)*, LNCS. Springer. (Cited on page 76.)
- [Temam, 1968] Temam, R. (1968). Une méthode d'approximation de la solution des équations de Navier-Stokes. *Bulletin de la Société Mathématique de France*, 96:115–152. (Cited on page 146.)
- [Thompson et al., 2004] Thompson, P., Miller, M., Ratnanather, J., Poldrack, R., and Nichols, T. (2004). Editorial of the special issue on mathematics in brain imaging. *NeuroImage*, 23(1). (Cited on page 34.)

- [Thompson, 1917] Thompson, W. D. (1917). *On Growth and Form*. Cambridge University Press, England. (Cited on page 36.)
- [Tobon-Gomez et al., 2013] Tobon-Gomez, C., De-Craene, M., McLeod, K., Lautz, T., Shi, W., Hennemuth, A., Prakosa, A., Wang, H., Carr-White, G., Kapetanakis, S., Lutz, A., Rasche, V., Schaeffter, T., Butakoff, C., Friman, O., Mansi, T., Sermesant, M., Zhuang, X., Ourselin, S., Peitgen, H.-O., Pennec, X., Razavi, R., Reuckert, D., Frangi, A., and Rhode, K. (2013). Benchmarking framework for myocardial tracking and deformation algorithms: An open access database. *Medical Image Analysis*. (Cited on pages 10, 53, 66, 75, 92, 93, 96, 102, 106, 108, 126 and 181.)
- [Tomasian et al., 2009] Tomasian, A., Krishnam, M. S., Lohan, D. G., Moghaddam, A. N., and Finn, J. P. (2009). Adult Tetralogy of Fallot quantitative assessment of pulmonary perfusion with time-resolved three dimensional magnetic resonance angiography. *Investigative Radiology*, 44(1):31–37. (Cited on page 139.)
- [Toussaint et al., 2013] Toussaint, N., Stoeck, C. T., Sermesant, M., Schaeffter, T., Kozerke, S., and Batchelor, P. G. (2013). In vivo human cardiac fibre architecture estimation using shape-based diffusion tensor processing. *Medical Image Analysis*. (Cited on pages 24, 87, 117 and 121.)
- [Troianowski et al., 2011] Troianowski, G., Taylor, C. A., Feinstein, J. A., and Vignon-Clementel, I. E. (2011). Three-dimensional simulations in Glenn patients: Clinically based boundary conditions, hemodynamic results and sensitivity to input data. *Journal of Biomechanical Engineering: ASME DC*, 133(11). (Cited on page 139.)
- [Trouvé, 1998] Trouvé, A. (1998). Diffeomorphisms groups and pattern matching in image analysis. *International Journal of Computer Vision*, 28(3):213–221. (Cited on page 36.)
- [Tucker, 1963] Tucker, L. R. (1963). Implications of factor analysis of three-way matrices for measurement of change. *Problems in Measuring Change*, pages 122–137. (Cited on page 123.)
- [Tucker, 1966] Tucker, L. R. (1966). Some mathematical notes on three-mode factor analysis. *Psychometrika*, 31(3):279–311. (Cited on page 123.)
- [Vaillant and Glaunes, 2005] Vaillant, M. and Glaunes, J. (2005). Surface matching via currents. In *Proc. of Information Processing in Medical Imaging (IPMI'05)*, pages 381–392. (Cited on pages 40 and 144.)
- [Vaillant et al., 2004] Vaillant, M., Miller, M., Younes, L., and Trouvé, A. (2004). Statistics on diffeomorphisms via tangent space representations. *NeuroImage*, 23(Supp. 1):S161–S169. (Cited on page 36.)
- [Vallet et al., 2006] Vallet, B., Angelini, E. D., and Laine, A. F. (2006). Variational segmentation framework in prolate spheroidal coordinates for 3D real-time echocardiography. *Proc. of SPIE Medical Imaging (SPIE'06)*. (Cited on page 86.)

- [Valverde et al., 2011] Valverde, I., Parish, V., Tzifa, A., Head, C., Sarikouch, S., Greil, G., Schaeffter, T., Razavi, R., and Beerbaum, P. (2011). Cardiovascular MR Dobutamine stress in adult Tetralogy of Fallot: Disparity between CMR volumetry and flow for cardiovascular function. *Journal of Magnetic Resonance Imaging*, 33(6):1341–1350. (Cited on page [139](#).)
- [van der Wall et al., 2000] van der Wall, E. E., Kayser, H. W., Bootsma, M. M., de Roos, A., and Schalij, M. J. (2000). Arrhythmogenic right ventricular dysplasia: MRI findings. *Herz*, 25(4):356–364. (Cited on page [22](#).)
- [Velec et al., 2012] Velec, M., Moseley, J., Marshall, A., Dawson, L., and Brock, K. (2012). Development of liver SBRT planning and delivery at the mean respiratory position using deformable image registration. *Medical Physics*, 39(6):3853. (Cited on page [24](#).)
- [Verbeke and Molenberghs, 2009] Verbeke, G. and Molenberghs, G. (2009). *Linear mixed models for longitudinal data*. Springer. (Cited on page [173](#).)
- [Vercauteren et al., 2008] Vercauteren, T., Pennec, X., Perchant, A., and Ayache, N. (2008). Symmetric log-domain diffeomorphic registration: A Demons-based approach. In *Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI'08), Part I*, volume 5241 of *LNCS*, pages 754–761. Springer. (Cited on pages [54](#), [55](#), [56](#), [76](#), [78](#), [79](#), [87](#) and [97](#).)
- [Vercauteren et al., 2009] Vercauteren, T., Pennec, X., Perchant, A., and Ayache, N. (2009). Diffeomorphic demons: Efficient non-parametric image registration. *NeuroImage*, 45(1):S61–S72. (Cited on page [114](#).)
- [Vignon-Clementel et al., 2010a] Vignon-Clementel, I., Figueroa, C., Jansen, K., and Taylor, C. (2010a). Outflow boundary conditions for 3D simulations of non-periodic blood flow and pressure fields in deformable arteries. *Computer Methods in Biomechanics and Biomedical Engineering*, 111(3):502–513. (Cited on page [146](#).)
- [Vignon-Clementel et al., 2010b] Vignon-Clementel, I. E., Marsden, A. L., and Feinstein, J. A. (2010b). A primer on computational simulation in congenital heart disease for the clinician. *Progress in Pediatric Cardiology*, 30(1-2):3–13. (Cited on page [139](#).)
- [Volkwein, 1999] Volkwein, S. (1999). Proper orthogonal decomposition and singular value decomposition. Technical report, Karl-Franzens-Univ. Graz & Techn. Univ. Graz. (Cited on page [139](#).)
- [Vossoughi et al., 1980] Vossoughi, J., Vaishnav, R. N., and Patel, D. J. (1980). Compressibility of the myocardial tissue. *Advances in Bioengineering Mow VC*, pages 45–48. (Cited on page [77](#).)
- [Wang and Amini, 2012] Wang, H. and Amini, A. (2012). Cardiac motion and deformation recovery from MRI: A review. *IEEE Transactions on Medical Imaging*, 31(2). (Cited on pages [76](#), [93](#), [102](#) and [113](#).)

- [Wang et al., 2011] Wang, H., Shi, W., Zhuang, X., Duckett, S., Tung, K., Edwards, P., Razavi, R., Ourselin, S., and Rueckert, D. (2011). Automatic cardiac motion tracking using both untagged and 3D tagged MR images. In *Proc. STACOM MICCAI Workshop (STACOM'12)*, LNCS. Springer. (Cited on page 76.)
- [Wassermann et al., 2011] Wassermann, D., Rathi, Y., Bouix, S., Kubicki, M., Kikinis, R., Shenton, M., and Westin, C.-F. (2011). White matter bundle registration and population analysis based on gaussian processes. In *Proc. of Information Processing in Medical Imaging (IPMI'11)*, pages 320–332. Springer. (Cited on page 77.)
- [Xue et al., 2006] Xue, Z., Shen, D., and Davatzikos, C. (2006). Statistical representation of high-dimensional deformation fields with application to statistically constrained 3D warping. *Medical Image Analysis*, 10(5):740–751. (Cited on page 116.)
- [Yang et al., 2011] Yang, W., Vignon-Clementel, I. E., Troianowski, G., Reddy, V. M., Feinstein, J. A., and Marsden, A. L. (2011). Hepatic blood flow distribution and performance in conventional and novel Y-graft Fontan geometries: A case series computational fluid dynamics study. *The Journal of Thoracic and Cardiovascular Surgery*. (Cited on page 139.)
- [Yeung et al., 2006] Yeung, J. J., Kim, H. J., Abbruzzese, T. A., Vignon-Clementel, I. E., Draney-Blomme, M. T., Yeung, K. K., Perkash, I., Herfkens, R. J., Taylor, C. A., and Dalman, R. L. (2006). Aortoiliac hemodynamic and morphologic adaptation to chronic spinal cord injury. *The Journal of Thoracic and Cardiovascular Surgery*, 44(6):1254–1265. (Cited on page 139.)
- [Yin et al., 1996] Yin, F., Chan, C., and Judd, R. M. (1996). Compressibility of perfused passive myocardium. *American Journal of Physiology-Heart and Circulatory Physiology*, 271(5):H1864–H1870. (Cited on page 77.)
- [Young, 2006] Young, A. A. (2006). Assessment of cardiac performance with magnetic resonance imaging. *Current Cardiology Reviews*, 2(4):271–282. (Cited on page 75.)
- [Zhang et al., 2007] Zhang, W., Noble, J., and Brady, J. (2007). Spatio-temporal registration of real time 3D ultrasound to cardiovascular MR sequences. *Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI'07)*, pages 343–350. (Cited on page 77.)
- [Zheng et al., 2008] Zheng, Y., Barbu, A., Georgescu, B., Scheuering, M., and Comaniciu, D. (2008). 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. (Cited on page 42.)
- [Ziyan et al., 2007] Ziyan, U., Sabuncu, M., O'Donnell, L., and Westin, C.-F. (2007). Nonlinear registration of diffusion MR images based on fiber bundles. *Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI'07)*, pages 351–358. (Cited on page 77.)