

A Statistical Dynamic Cardiac Atlas for the Virtual Physiological Human: Construction and Application

Corné Hoogendoorn

Tesi Doctoral UPF / 2013



Supervised by

Prof. Alejandro F. Frangi

Departament de Tecnologies de la Informació i de la Comunicació

Universitat Pompeu Fabra

and

Department of Mechanical Engineering

The University of Sheffield

About the cover

Cover art by the author, using Blender v2.66a and GNU Image Manipulator Program (GIMP) v2.6. The wireframe rendering of the mean shape of the cardiac atlas (for a color surface rendering, see page xlii in the back) symbolizes the complexity of both the heart itself and the construction of the atlas. The large quantity of CT data that the atlas was generated from is symbolized by the background of the cover, showing central slices of the cardiac CT images, illustrating as well the variability of the heart.

Copyright © 2013 Corné Hoogendoorn.

This work was carried out at the Center for Computational Imaging and Simulation Technologies in Biomedicine (CISTIB), Universitat Pompeu Fabra, and Networking Research Center on Bioengineering, Biomaterials and Nanomedicine (CIBER-BBN).



Reading committee members

Prof. Gábor Székely

Swiss Federal Institute of Technology Zürich
Switzerland

Prof. Carlos Alberola López

Universidad de Valladolid
Spain

Dr. Miguel Ángel González Ballester

Universitat Pompeu Fabra
Spain

Reserve committee members

Prof. Joan Martí Bonmati

Universitat de Girona
Spain

Prof. Petia Radeva

Universitat Autònoma de Barcelona
Spain

To Martha

ABSTRACT

The Virtual Physiological Human (VPH) is a *methodological and technological framework that, once established, will enable collaborative investigation of the human body as a single complex system*. As part of realizing this framework, building detailed spatio-temporal digital representations of the human body in general, and of specific organ systems in particular, is considered of prime importance. Such representations should be able to describe the typical anatomical phenotype of populations of individuals and disease states. Ideally, they should also be multi-scale representations amenable to expressing the desired phenotypes at the relevant spatial and temporal scales.

The use of a common reference frame is a necessity to make objective and accurate comparisons between individuals' groups or of a given individual over time. Tied in with the development of such a reference frame are the tools to transform measurements from various studies to the frame and vice versa. Typically the reference frame is taken to be a model of normality of the anatomical phenotype at hand, expressed as an average and a range of statistically plausible variations, obtained from training data drawn from the population. New measurements are then assimilated, yielding corresponding model parameterizations.

Construction of a common reference frame, or atlas, for the (human) heart is fraught with challenges. The mammalian heart has a complex morphology with great variability across populations, it is a highly dynamic morphology due to cardiac function, and it contains a large number of substructures that ultimately enable this electro-fluid-mechanical bio-machine to work efficiently. Thus, a good cardiac atlas is ideally multi-scale, spatio-temporal, multi-structural, and its construction must be very robust to large morphological variation and capable of handling large amounts of data in a highly automated manner.

This thesis is centered on the construction of a cardiac atlas to serve as such a reference frame. The construction covers the entire construction pipeline, starting from a set of 3D+t multislice computed tomography images, then performing a spatial normalization of these images, segmentation of the synthesized mean image, multi-structure meshing, and finally mapping of the mesh back to the population of images. In addition, two applications are presented in this thesis.

First, the atlas is used to frame a spatio-temporal model of cardiac morphology which models the variability along both 'axes' simultaneously. Such a unified approach should be preferable over existing methods, which decouple the two sources of variation and then model them separately, in isolation.

Second, the proposed atlas is applied to develop an acceleration technique for performing personalized simulation of cardiac electrophysiology (EP). The prior knowledge encapsulated in our atlas is used, in conjunction with a numerical solver of cardiac EP, to build a statistical model linking cardiac morphology with the steady states of myocardial cell models that pre condition detailed cardiac EP simulations. This application puts the proposed dynamic cardiac atlas in the context of VPH-related simulations, of which the computational costs are currently greatly in excess of what is acceptable for their adoption in current clinical practice.

RESUMEN

El Humano Fisiológico Virtual (VPH en inglés) es *un marco metodológico y tecnológico que, una vez establecido, habilitará la investigación colaborativa sobre el cuerpo humano como un único sistema complejo*. Como parte de la realización de este marco, está considerado de suma importancia construir representaciones digitales espacio-temporales, de alto nivel de detalle, del cuerpo humano en general y de sus sistemas orgánicos en específico. Tales representaciones deberán facilitar la descripción del fenotipo anatómico típico de poblaciones de individuos y estados de enfermedad. Idealmente, también deben ser representaciones multi-escala que puedan expresar los fenotipos deseados en las escalas pertinentes en espacio y tiempo.

La representación de datos en un marco de referencia común es una necesidad para hacer comparaciones imparciales y exactas entre individuos o de un individuo a lo largo del tiempo. Atado con el desarrollo de tal marco está el desarrollo de las herramientas para transformar las medidas de varios estudios al marco y viceversa. Típicamente, el marco representa un modelo de normalidad del fenotipo anatómico en utilización, expresada por un promedio y un rango de variaciones estadísticamente plausibles, obtenidos de datos de entrenamiento adquiridos de la población. Entonces, las nuevas medidas serán asimiladas, resultando en las parametrizaciones correspondientes del modelo.

La construcción de un marco de referencia, o atlas, del corazón (humano) presenta una multitud de obstáculos. El corazón mamífero tiene una morfología compleja con gran variabilidad dentro de poblaciones, es una morfología altamente dinámica debido a la función cardíaca, y contiene un gran número de subestructuras que últimamente permiten el funcionamiento eficaz de esta máquina electro-fluido-mecánica. Entonces, idealmente, un buen atlas cardíaco es multi-escala, espacio-temporal, multi-estructura, y su construcción deberá ser robusta a grandes variaciones morfológicas y capaz de procesar grandes cantidades de datos de manera altamente automatizada.

Esta tesis está centrada en la construcción de un atlas cardíaco, para servir como tal marco de referencia. La construcción consiste en la trayectoria completa, empezando con un conjunto de imágenes 3D+t de tomografía computacional multi-corte, y entonces hacer una normalización espacial de las imágenes, segmentación de la imagen promedio sintetizada, un mallado multi-estructura, y finalmente la transformación de la malla a la población de imágenes. Adicionalmente, la tesis presenta dos aplicaciones del atlas.

Primero, el atlas se usa para enmarcar un modelo espacio-temporal de la morfología cardíaca que modela la variación a lo largo de ambos 'ejes' simultáneamente. Tal propuesta debe ser preferible sobre otros métodos existentes, los cuales desacoplan las dos fuentes de variación para modelarlas separadamente, en isolación.

Segundo, el atlas está aplicado al desarrollo de una técnica de aceleración para simulaciones personalizadas de electrofisiología (EF) cardíaca. El conocimiento previo encapsulado en nuestro atlas se usa, en conjunto con un *solver* de EF cardíaca, para construir un modelo estadístico conectando morfología cardíaca con los *steady states* de modelos celulares del miocardio que preconditionan a simulaciones detalladas de EF cardíaca. Esta aplicación posiciona el propuesto atlas dinámico cardíaco en el contexto de simulaciones relacionadas al VPH, cuyo costo computacional actual está en gran exceso de lo

aceptable para su adopción en la práctica clínica de hoy en día.

Agradezco la ayuda de Martha Aguilar en traducir el resumen.

PREFACE

That muscle, the heart, was a double pump that had to beat from well before birth to the final moment before death and do so with unbroken rhythm, unwearying strength, under all conditions. It was the greatest heart in the animal kingdom. The heart of no other mammal beat more than a billion times or so before even the most delayed approach of death, but after a billion heartbeats the human being was merely in early middle-age, in the prime of his strength and power. Men and women had lived long enough to experience well over three billion heartbeats.

— Isaac Asimov, "Fantastic Voyage"

CONTENTS

List of Figures	xi
List of Tables	xiii
Acronyms and Abbreviations	xv
1 Introduction	1
1.1 The Heart and its Function(ing)	2
1.1.1 Electrophysiological Function	3
1.1.2 Pumping and Mechanical Function	3
1.1.3 Health Challenges	4
1.2 Virtual Physiological Human	5
1.2.1 Motivation and Objective of the VPH	5
1.2.2 Technical Challenges	5
1.2.3 The Heart in the VPH	6
1.3 Contributions	6
2 A High-Resolution Atlas and Statistical Model of the Human Heart from Multislice CT	9
2.1 Introduction	11
2.2 Related Work	13
2.2.1 Cardiac Statistical Atlases	13
2.2.2 Spatial Normalization	14
2.3 Atlas Construction	15
2.3.1 Imaging Data	16
2.3.2 Reference Selection	16
2.3.3 Global Registration	19
2.3.4 Diffeomorphic Nonrigid Registration	20
2.3.5 Mean Deformation	21
2.3.6 Parameters and Settings	23
2.3.7 Atlas Segmentation	24
2.3.8 Mesh Construction and Processing	25
2.4 Atlas Properties and Statistical Model	27
2.4.1 Structures	28
2.4.2 Statistical Modeling	28
2.5 Validation	31
2.5.1 Registration—Selection of Volumes and Slices	32
2.5.2 Registration—Measurements	33

2.5.3	Registration—Results	36
2.5.4	Bias Removal	37
2.6	Discussion	40
2.6.1	Framework Performance	40
2.6.2	Limitations	41
2.6.3	Future Perspectives	42
3	Bilinear Models for Spatiotemporal Point Distribution Analysis	43
3.1	Introduction	45
3.1.1	Statistical Shape Models and Time	45
3.1.2	Bilinear Statistical Models	46
3.1.3	Clinical Context	46
3.2	Bilinear Statistical Models	47
3.2.1	Asymmetric Models	49
3.2.2	Symmetric Models	49
3.2.3	Factoring Inter-Subject Variability and Cardiac Dynamics .	49
3.3	Experiments	52
3.3.1	Data	52
3.3.2	Experiments	57
3.3.3	Extrapolation	57
3.3.4	Reconstruction	60
3.3.5	Comparison to Direct Approximation	60
3.4	Results	60
3.4.1	Extrapolation	60
3.4.2	Comparison to Direct Approximation	64
3.5	Discussion and Future Work	66
3.6	Conclusion	67
4	An Atlas-Based Dictionary Approach to Initializing Cardiac Electro-physiology Simulations	69
4.1	Introduction	71
4.1.1	Related Work	72
4.2	Methods	74
4.2.1	Generation of Anatomical Surface Meshes	74
4.2.2	Generation of Anatomical Volumetric Meshes	76
4.2.3	Personalized Peripheral Conduction System	76
4.2.4	Myocardial Fiber Architecture	77
4.2.5	Template-to-Anatomy Mesh Adaptation	77
4.2.6	Prediction of Steady State	78
4.3	Experiments and Results	79
4.3.1	Data – Computational Meshes	79
4.3.2	Data – EP Simulations	80
4.3.3	Hard- and Software	81
4.3.4	Baseline Steady State of 3D Models	81
4.3.5	Prediction Times of the Steady States	83
4.3.6	Accuracy of Simulations Using Predicted Initial States . .	83
4.4	Discussion and Conclusions	90

CONTENTS	ix
5 Conclusions and Outlook	93
5.1 Overview	94
5.2 Outlook	95
5.2.1 Atlases, Cardiac and Other	95
5.2.2 Spatio-temporal Models of Shape	96
5.2.3 Recycling Electrophysiology Simulations	96
5.3 The Virtual Physiological Human's Prognosis	97
Bibliography	xvii
Publications	xxxi
Acknowledgments	xxxv
Funding	xxxvii
Biosketch	xxxix
Color Images	xli

LIST OF FIGURES

1.1	Diagram of cardiac anatomy and circulatory function.	2
1.2	Diagram of the cardiac conduction system.	3
2.1	Flow diagram of the mean image synthesis.	18
2.2	Rankings and mean ranking for subjects with a ranking standard deviation below 10.	19
2.3	The development of \tilde{I}_i , for $0 \leq i \leq 7$	22
2.4	Flow diagram of the meshing steps.	27
2.5	The full atlas mesh, corresponding to the synthesized mean image.	29
2.6	Three phases of two subjects, with their bilinear parameter sets derived from the full 134-subject data set.	30
2.7	The first two modes of variation in the linear statistical model of the full heart at end diastole.	31
2.8	The first two modes of variation in the linear statistical model of the left ventricle at end diastole.	31
2.9	The absolute values of the bias, and the kurtosis observed in the mean vector field.	38
2.10	Histograms of vector components per iteration.	38
2.11	The distributions of vector magnitudes in the mean vector fields show convergence of the registration.	39
3.1	The structure of a symmetric bilinear model.	48
3.2	RMS P2S differences between subsequent phases of the left ventricles.	53
3.3	Surface rendering of one of the training hearts.	54
3.4	The effect of the alignment on the CC of the CT images.	55
3.5	Scatter plots of the relationship between heart rate and the phases in which maximum contraction and maximum deceleration are observed.	56
3.6	Overview of the leave- n -out experiments.	58
3.7	The best, worst and mean RMS P2S reconstruction errors for the left ventricular model.	62
3.8	The mean RMS P2S reconstruction errors for the biventricular models and the full heart models.	63
3.9	The cumulative histogram of errors for the left ventricular, biventricular and full heart models.	63
3.10	The mean RMS P2S reconstruction error for each subject.	63

3.11	Reconstructions of the shapes with the median RMS P2S errors, constructed using the highest dimensional models.	64
3.12	The approximation errors when deriving subject parameters from one phase only.	65
4.1	Action potential duration for four nodes of different type, and Purkinje model state variables at rest.	72
4.2	Flow diagram of the pipeline for generating volumetric meshes V and Purkinje networks P with node correspondence.	75
4.3	Surface renderings of the eleven geometries.	76
4.4	Flow diagram of the prediction and its inclusion into the pipeline.	79
4.5	Edge length distributions in the myocardial mesh and the Purkinje network.	80
4.6	Action potential curves for the various node types.	81
4.7	Dependence on k of NRMS errors at time $t = 5$ and time $t = 10$ of the predictions of myocardium and Purkinje model states.	84
4.8	Mean unsigned relative prediction errors at time $t = 5$ color coded on the surface of the mean myocardium.	85
4.9	Mean unsigned relative prediction errors at time $t = 5$ color coded on the Purkinje network corresponding to the mean geometry.	86
4.10	Mean CPU times for generating the predictions	87
4.11	Residual errors in model states of myocardium and Purkinje network with respect to $t = 10$, for the simulations from scratch and from prediction at $t = 5$	87
4.12	Mean unsigned relative differences between the two approaches to reach time $t = 10$ color coded on the surface of the mean myocardium.	88
4.13	Mean unsigned relative differences between the two different approaches to reach time $t = 10$ color coded on the Purkinje network corresponding to the mean geometry.	89

LIST OF TABLES

2.1	Details of data	17
2.2	Registration parameters and values	23
2.3	Parameters and values for non-registration elements	24
2.4	Atlas subparts (3D) and element statistics	26
2.5	Registration and vector field inversion accuracy of selected anatomical landmarks and contoured structures	34
2.6	Registration accuracy quantified by F-score of overlap of selected structures	35
3.1	The parts of the heart shape that have a closed surface representation	53
3.2	The main statistics for the reconstruction errors using the largest training set	61
4.1	CPU Times to Stable State	82

ACRONYMS AND ABBREVIATIONS

2D	two-dimensional
3D	three-dimensional
μCT	micro-CT
AAM	Active Appearance Model
AAMM	Active Appearance Motion Model
ACVT	approximate centroidal Voronoi tessellation
AP	action potential
APD	action potential duration
APD₉₀	approximation of APD
AV NODE	atrio-ventricular node
AVA	aortic valve
CAD	Coronary Artery Disease
CAP	Cardiac Atlas Project
CC	cross-correlation coefficient
CHF	Congestive Heart Failure
CT	Computed Tomography
CVD	Cardiovascular Disease
DTMRI	diffusion tensor MRI
EKG	electrocardiogram
ED	end diastole
EF	ejection fraction
EP	electrophysiology
ES	end systole
FE	Finite Element
GIMIAS	Graphical Interface for Medical Image Analysis and Simulation
GPU	Graphics Processing Unit
GWMI	Groupwise Mutual Information
HU	Hounsfield Unit
IUPS	International Union of Physiological Sciences
IVC	inferior vena cava
LA	left atrium
LV	left ventricle
MC	marching cubes
MI	Mutual Information (when used in a technical context)
MI	myocardial infarction (when used in a clinical context)

MRI	Magnetic Resonance (Imaging)
MSCT	Multi-Slice Computed Tomography
MVA	mitral valve
NRMS	normalized RMS
ODE	Ordinary Differential Equation
P2S	point-to-surface
PA	pulmonary artery
PCA	Principal Component Analysis
PDM	Point Distribution Model
Pk	Purkinje (network or system)
PMJ	Purkinje-myocardial junction
PV	pulmonary vein
PVA	pulmonary valve
QRS COMPLEX	The complex of Q, R, and S waves in the ECG
RA	right atrium
RCA	right coronary artery
RMS	root mean square (error or difference)
RV	right ventricle
SVC	superior vena cava
SVD	Singular Value Decomposition
TPS	Thin Plate Splines
TVA	tricuspid valve
VPH	Virtual Physiological Human

CHAPTER

ONE

INTRODUCTION

1.1 The Heart and its Function(ing)

It is common knowledge that the heart is essential for all known vertebrate animal life forms. It is also commonly known from biology class that we humans—and all other mammals and birds—have a four-chamber heart. Beyond that, the commonality of the knowledge decreases rapidly.

As illustrated in Fig. 1.1, the human heart has two sides with two chambers each: left and right ventricle, and left and right atrium. The right heart contains deoxygenated blood, while the left heart contains oxygenated blood. This is due to the double cycle that is the human cardiovascular system. One loop of this '8' may start at the right ventricle (RV), from which deoxygenated blood is ejected towards the lungs through the pulmonary artery (PA). There, the blood is oxygenated and re-enters the heart at the left atrium (LA) via the pulmonary veins (PVs). After transfer to the left ventricle (LV), it is ejected into the aorta to provide the entire body with oxygenated blood, after which the blood returns to the heart at the right atrium (RA) via the superior and inferior vena cavae (SVC and IVC). It then moves into the RV, and the double loop repeats.

This is a very simplified way of describing the result of an extraordinarily complex chain of events that is repeated roughly every second of our life, or some 2.5 billion times during a lifespan of about 80 years.

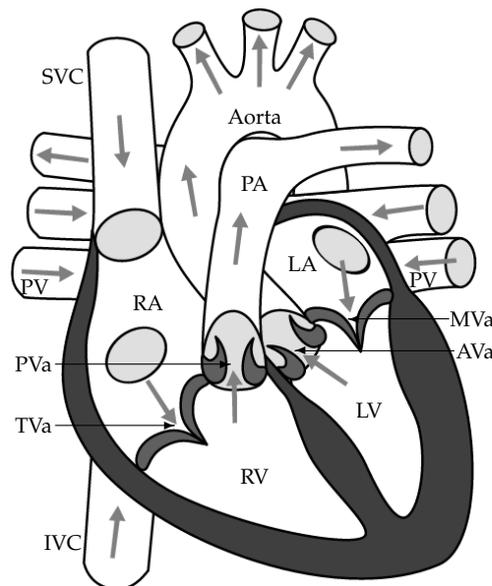


Figure 1.1: Diagram of cardiac anatomy and circulatory function, with left heart structures LA: left atrium; LV: left ventricle; PV's: pulmonary veins; Aorta, and their right heart counterparts RA: right atrium; RV: right ventricle; SVC/IVC: superior/inferior vena cava; PA: pulmonary artery. Also shown are the tricuspid (TVa), pulmonary (PVa), aortic (AVa) and mitral (MVa) valves. Drawing source: Wikimedia Commons.

1.1.1 Electrophysiological Function

The heart is a very intricate machine that works autonomously, although its functioning can be influenced by the sympathetic and parasympathetic nervous system when necessary. It controls the heart rate through the sinus node, a set of specialized cells located in the coronary sinus in the right atrium [7], shown top left in Fig. 1.2. This node is sometimes referred to as the pacemaker of the heart, and in this capacity it synchronizes the cardiac muscle cells, myocytes, to contract in unison. Otherwise, the cells would contract on their own, but in random fashion. The electrical impulse from the sinus node travels through the specialized conduction system of the right atrium, comprised of the terminal crest and Koch's triangle, to the atrioventricular (AV) node, and through Bachmann's bundle to the left atrium. Along the way, both atria are activated (depolarized).

After a short delay, the AV node allows the electrical signal to enter the interventricular septum through the bundle of His, which splits there into the left and right bundle branches. Finally, these bundles branch into an intricate network of fibers known as the Purkinje fibers. This high-velocity conduction system is covered in an insulating sheath up until the terminals, or Purkinje-myocardial junctions, of the Purkinje fibers, which allows the ventricles to depolarize rapidly throughout. As the normal myocytes also propagate the electrical activation at a lower rate, it takes about 100 ms to depolarize the entire healthy human ventricular myocardium. And then it contracts.

1.1.2 Pumping and Mechanical Function

The pumping function of the heart can be felt when one holds a finger on the pulse, on the carotid artery, or simply on one's chest. Usually one feels a single heartbeat as a double thud. This is illustrative of the four-chamber heart's two-phase contraction, separated by the aforementioned delay at the AV node.

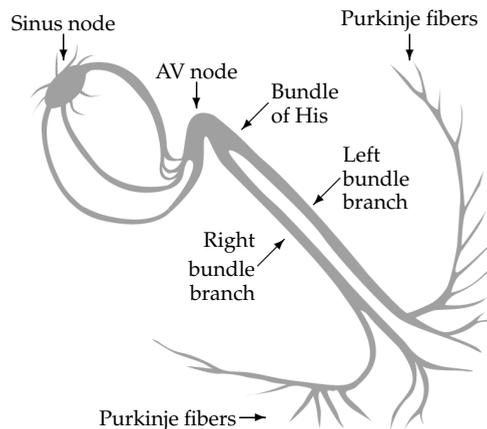


Figure 1.2: *Diagram of the cardiac electrical conduction system. Drawing source: Wikimedia Commons.*

The valves between atria and ventricles (TVa and MVa in Fig. 1.1) are opened when the ventricles are in a contracted state. When the ventricles relax, the pressure differential between atrium and ventricle causes the atrium to empty its contents into the ventricle. The contraction of the atrium squeezes the last bit of blood out. Then, the atrioventricular valves close and the ventricles contract while the aortic and pulmonary valve (AVa and PVa in Fig. 1.1) open, pumping the blood to the body and the lungs.

The timing of these steps is of great importance; together with the shape of the cavities it is optimized to keep the blood going in a fluid motion down into the ventricle, making a 180 degree turn and 'up' into the artery, so as to minimize the actual work required from the myocardium and thus maximize the percentage of blood that is ejected from the ventricle into the artery. As a clinical parameter, this is known as the ejection fraction (EF), computed as $EF = 100 \frac{EDV - ESV}{EDV}$: the end diastolic (before contraction) volume minus the end systolic (after contraction) volume, divided by the end diastolic volume.

EF in a healthy person is usually greater than 50%. However, a contracted cardiac myocyte is only between 10 and 15% shorter than it is in its resting state [48, 103]. The difference in these percentages is resolved through the ingenious arrangement of the myocytes. A helical arrangement that varies throughout the myocardial wall creates a twisting motion of the ventricle when the myocytes contract in unison. At the outside of the wall (epicardium), the fibers are tilted up to 60 degrees from the 'horizontal'. That is, from the true circumferential direction. At the endocardium, the wall of the cavity itself, the fibers are tilted up to -60 degrees. In between, there is a smooth transition [6, 184].

1.1.3 Health Challenges

Worldwide, but particularly in the developed world, the impact of Cardiovascular Diseases (CVDs) is enormous. It is the number one cause of death globally, accounting for roughly one in every three deaths [219] in 2008. In Europe, this percentage is even slightly higher despite recent decrease in the wealthier nations [56]. Its economic cost is estimated to be at nearly €200 billion in the European Union, in health care costs, loss of productivity from mortality and morbidity, and cost of informal care. As populations in the developed world continue to age, and populations in fastly developing countries likely to adopt unhealthier lifestyles [131, 142] as the western world did some decades ago, it is expected that the global mortality rate due to CVDs will continue to rise until at least the year 2030 [219].

This burden emphasizes the importance of all aspects of cardiovascular research. This not only holds true for the classical medical aspects of reducing occurrence through prevention, and reducing morbidity and mortality through early detection and optimized treatment. It seems a natural step to exploit the possibilities offered by present and future computational technologies to improve our understanding of pathogenesis, pathophysiology, and the impact of pharmacological, surgical and other treatments on cardiac physiology.

1.2 Virtual Physiological Human

The Physiome and VPH initiatives [55,84] strive to build a digital representation of an *in silico* human from molecular to organism level. As we could explain our body divided into organ systems, into organs, into organ parts, into cell types, into types of subcellular structures, and finally into molecules and atoms, so too is this human to be described in a multi-scale fashion, and the initiatives represent the concerted efforts to enable such goal.

1.2.1 Motivation and Objective of the VPH

The ultimate objective of the VPH is to enable the computerized simulation of the functioning, in the physiological sense, of the human body or any subset of it. To this end, the various VPH initiatives strive to identify requirements for models and tools, and to reach consensus regarding standards such that these models and tools may interoperate across scale levels and organ boundaries. While the rationale behind crossing the scale boundaries is probably immediately clear, the greater goal of the VPH is motivated more by the somewhat less intuitive need for crossing the frontiers of isolated organs and systems. One of the main keywords here is comorbidities.

As life expectancy at birth in our western society lies around 80 years, the age-related problem of comorbidities, i.e., potentially interrelated health problems in multiple organs or organ systems, is growing. The potential interrelatedness presents a problem in the medical world where specialization has been key to effective delivery of healthcare for over more than a century [216]. Continued increase in the occurrence of comorbidities means that we can no longer see organs as isolated parts and fully understand what their problems mean for the remainder of the body. To this end, the VPH is optimally positioned: by promoting research into enabling multi-scale, inter-system simulation and analysis, the interrelations between otherwise separate systems may be uncovered under both healthy and pathological conditions. In turn, this can be used to improve the various aspects of healthcare: more targeted and cost-effective drug discovery and development through the use of large *in silico* trials; optimization of treatment development and delivery, and improvement of diagnostics and prognostics, through both such trials and through personalization of diagnostic tests and therapeutic planning.

1.2.2 Technical Challenges

The grand technical challenge of the VPH itself is to understand how the various models are interlinked, to each other as well as to actual (clinical) measurements. From this understanding, requirements for new models and tools can be formulated. However, this also translates to requirements for knowledge bases that assist in choosing appropriate models and in the various aspects of the clinical application of the VPH.

Simply put, such a knowledge base would contain a lot of information in a limited number of standardized representations, or reference frames. These reference frames are determined by the type of knowledge they represent, which

in turn translates to the various models to which the specific type of information is relevant.

Such a reference frame is typically a definition of normality for the information it presents, expressed as an average observation and statistically plausible ('normal') deviations from this average, over a population. The derivation of these reference frames, and filling of these knowledge bases, are closely interwoven. A large quantity of data must be processed to establish the average, as dictated by the law of large numbers. While doing so, this means that the knowledge base is being filled, and it implies that a means of mapping observations to the reference frame exists.

In the VPH, the observations can take on a variety of forms as they span the entire spectrum of medical measurements: indexes, measures of length, area, volume, weight, concentration, categorical data, and images are but a few, and these can be either single-shot or have their variation measured over time. That is, over cardiac, pulmonary or dyadic cycles, or over periods measured in weeks, months, or years.

In summary, the complexity of the path towards such an almost sci-fi-like level of computer-assisted healthcare calls for consensus on modeling standards, data acquisition and knowledge generation that the VPH strives to frame.

1.2.3 The Heart in the VPH

A forerunner of the VPH initiatives, the International Union of Physiological Sciences (IUPS), described the human body as a set of twelve organ systems [84], of which the cardiovascular system is one. The heart, although considered an organ within the cardiovascular system, could just as easily be considered a system on its own, based on its complex heterogeneous anatomy and physiology. It is dynamic, and its function at the organ scale is described by three branches of physics (mechanics, fluid dynamics, electrophysiology) in an interconnected manner, to an extent which has so far exceeded the capabilities of current modeling approaches.

A multi-scale model of the heart, thus, links various models of cellular electrophysiology (EP) to cellular or tissue-level models of (electro)-mechanics. These would incorporate also the anisotropic nature of the tissue dictated by the arrangement of cardiomyocytes, commonly known as the myofiber architecture. The model would incorporate cardiac substructures such as the valves and the specialized electrical conduction system described in Sec. 1.1.1, and allow the use of models describing pathological or pharmacological changes at the cellular level. All of this would ultimately be embedded in a macroscopic model of the cardiac anatomy, capable of being linked primarily to the circulatory and pulmonary systems.

1.3 Contributions

This thesis presents three specific contributions, the latter two derived from the first one, with one chapter dedicated to each contribution:

- Chapter Two describes a framework developed for the largely automated construction of a cardiac atlas and statistical model, which can serve as a reference frame in a VPH knowledge base. The processing of a large population, with limitations on per-case optimization and on human interaction/intervention makes this a challenging task.
- Chapter Three describes how the atlas can be used to frame a spatio-temporal model of cardiac shape. The spatial and temporal components are modeled as distinct axes of a bilinear model, which allows us to maintain relationships between model parameterizations and specific time points, which has been one of multiple limitations of previous cardiac modeling approaches.
- The final contribution uses the atlas as reference frame within a knowledge base. As cardiac EP cell models are typically stabilized in isolation (0D), the complex interactions between models when they are embedded in a 3D anatomy mean they must be stabilized further. This leads to a significant computational burden, limiting applicability of such simulations. In Chapter Four this burden is reduced by prediction of the 3D steady states of these cell models using the knowledge base.

CHAPTER

TWO

A HIGH-RESOLUTION ATLAS AND STATISTICAL
MODEL OF THE HUMAN HEART FROM
MULTISLICE CT

Adapted from [75]

C. Hoogendoorn, N. Duchateau, D. Sánchez-Quintana, T. Whitmarsh, M. De Craene, F.M. Sukno, K. Lekadir, and A.F. Frangi
IEEE Transactions on Medical Imaging, vol. 32, no. 1, pp. 28–44, Jan. 2013.

Abstract

Atlases and statistical models play important roles in the personalization and simulation of cardiac physiology. For the study of the heart, however, the construction of comprehensive atlases and spatio-temporal models is faced with a number of challenges, in particular the need to handle large and highly variable image datasets, the multi-region nature of the heart, and the presence of complex as well as small cardiovascular structures.

In this chapter, we present a detailed atlas and spatio-temporal statistical model of the human heart based on a large population of 3D+time multi-slice computed tomography sequences, and the framework for its construction. It uses spatial normalization based on non-rigid image registration to synthesize a population mean image and establish the spatial relationships between the mean and the subjects in the population. Temporal image registration is then applied to resolve each subject-specific cardiac motion and the resulting transformations are used to warp a surface mesh representation of the atlas to fit the images of the remaining cardiac phases in each subject.

Subsequently, we demonstrate the construction of a spatio-temporal statistical model of shape such that the inter-subject and dynamic sources of variation are suitably separated. The framework is applied to a 3D+time data set of 138 subjects. The data is drawn from a variety of pathologies, which benefits its generalization to new subjects and physiological studies. The obtained level of detail and the extendability of the atlas present an advantage over most cardiac models published previously.

2.1 Introduction

Interest in cardiac electrophysiological and mechanical simulation has risen significantly over the past decade, enabled by increased power and availability of computing resources and developments in distributed computing. This interest has been partly structured around initiatives to develop computational physiology models of the (human) body and its systems, collectively known as Virtual Physiological Human (VPH) initiatives [55]. The cardiovascular system is one of twelve systems identified within the International Union of Physiological Sciences (IUPS) Physiome Project [84], and is subdivided further into cardiac and vascular modeling.

Atlases play an important role in computational physiology of any organ, including the heart [27,28,222]. They provide insight regarding the division into and localization of substructures within the body, within an organ, or within a structure. For computational physiology studies to return accurate results, the use of such a ‘map’ is of great importance, as different structures within an organ have different electrical and mechanical properties.

Statistical atlases provide not only an average layout of structures within an encapsulating structure, but also encode deviations from this average. This provides a means to deform the encapsulating structure, within statistically justified bounds, and have the substructures deform and move accordingly based on their statistical correlation. This enables two very important applications of computational physiology: first, simulation studies can be personalized geometrically. This is in addition to incorporating subject-specific measurements to parameterize models of electrical conduction and of mechanics. For geometrical personalization, the atlas can be matched to medical imaging data. Typically, only a subset of atlas structures can be matched explicitly to the image, and the remaining structures are moved by virtue of their statistical correlation. Secondly, the statistics learned from the population can be used to generate populations for virtual population studies. Such a population can be controlled by the user, either to statistically match an existing population, to increase sample size, or to generate more extreme cases.

For both applications, a statistical atlas provides a final advantage in the postprocessing and analysis of any simulation results that are generated. The encoding of population variation implies that spatial relationships between instances are known. This ensures a straightforward and reliable way of warping study outputs into a common reference frame in which they can be further presented, compared and analyzed.

The construction of a statistical atlas from a population of images requires that each of the structures in the atlas be segmented (labeled) in each of the images in the database. To do this manually is generally considered an impossibility for three-dimensional (3D) and 3D+time atlases. To bypass this problem, atlas-based segmentation methods provide a solution. First, one applies *spatial normalization* to the population. This is the synthesis of an average image from the population, usually based on image registration techniques. Additionally, it provides the spatial relationships between the population and this average. As one then *labels this atlas*, one may consider all the population images segmented through these spatial relationships. By representing the atlas as a surface mesh,

one uses point distribution analysis as the approach to *statistical analysis*.

The spatial normalization of a population of 3D+time cardiac images is faced with a number of challenges. Some of these have a counterpart in the construction of brain atlases; others are specific to the cardiac case:

- Cardiac and vascular structures vary significantly in size and geometry, with large global variations observed throughout the populace and both global and local variations due to pathology. In order to extract and represent these accurately, high-resolution image data is required, and the algorithms used must be able to handle this.
- Correct topological relationships between structures are of the utmost importance for the atlas to be used in cardiac simulations. Seamless transitions between structures while maintaining mesh quality are required. Mesh extraction, simplification and smoothing algorithms must be able to provide this.
- High anatomical variability in the heart requires a large population for its statistical modeling. This makes it very desirable to minimize the greatest difference to be resolved, which can be achieved by selecting a suitable initial estimate of the spatially normalized image.
- A large population also means that case-by-case parameter tuning becomes increasingly unfeasible. Consequently, this must be addressed in an automated manner or the algorithms must be robust to potentially sub-optimal parameterization.
- Cardiac motion presents a challenge in statistical analysis, as it means there are at least two sources of shape variation: inter-subject and temporal. Additionally, it introduces an increase in data set size with respect to 3D imaging. Together with the large population requirement and the high resolution requirement, this means that the algorithms must be able to handle very large data sets.

In summary, the algorithms must be robust to suboptimal parameterization and large variations, and able to handle very large data sets. It means that the atlas representation must have the flexibility to permit the statistical modeling as well as the simulation studies. Finally, the statistical modeling must be able to handle spatio-temporal data in a correct manner.

This chapter presents a statistical atlas of the human heart with the detail and flexibility necessary for personalized cardiac simulation, and the framework through which we constructed it. Elements within the framework are freely exchangeable for more advanced techniques, yet in the current form we have opted for some well-established approaches. In addition to the framework for the atlas construction, we present a framework for its validation. Also, we have made the atlas itself publicly available. An additional minor contribution is the automated selection of a reference image as an initial estimate of the mean.

The remainder of this chapter is organized as follows. The next section will provide an overview of related work, on cardiac statistical atlases and on spatial normalization of a population of images. Section 2.3 will delineate the steps

taken in the construction of the atlas, and the considerations implied in the automation of this process at a large scale. Section 2.4 outlines the properties of the atlas and their relationship to the atlas' applications, whereas 2.5 covers experiments regarding the validation of the atlas, and in Sec. 2.6 we discuss possible improvements and extensions.

2.2 Related Work

The traditional atlas in the medical field consists of a collection of commented illustrations—drawings or photos—of the structure of choice, either intact or dissected, as a whole or in close-up. Invariably these are two-dimensional representations of the structures *ex vivo*, without the possibility to see variations, active or passive movement, or points of view other than those presented. Computerized medical imaging techniques provide solutions to each of these issues. In this context, we see related work in the development of digital three- and four-dimensional atlases, as well as efforts in computational anatomy to characterize population variability.

Digital cardiac atlases have been developed for a variety of purposes. Recently, atlas-based segmentation using voxel-based atlases was introduced in the cardiac domain [85, 98, 122, 205]. The construction of a voxel-based atlas of cardiac structures was only outlined in [122], using only three cardiac labels (two ventricular blood pools and left ventricular myocardium).

More recently, voxel-based atlases of cardiac fiber orientation have been constructed, first from canine [158] and later from human [118] *ex vivo* diffusion tensor MRI (dtMRI) data. These atlases provide an important component for simulation of both cardiac electrophysiology and mechanics, despite the limitation of being acquired *ex vivo*.

A special case of cardiac atlases is the Cardiac Atlas Project (CAP) [59]. It adheres more to the traditional idea of an atlas in that it is a collection of maps—a database—rather than a single mean map, or a statistical description of the shapes and their relationships. The construction of a statistical shape model of the left ventricle from part of this data (200 training shapes) was demonstrated in [137].

2.2.1 Cardiac Statistical Atlases

In [61], first a voxel-based and then a surface-based mean atlas was constructed of the same structures used later in [122]. The surface-based atlas, together with deformations obtained during construction of its voxel-based counterpart, was used to generate point correspondence across the population such that a point distribution model could be built. Virtually the same technique was used to generate a four-chamber model in [123].

Similarly, Perperidis used voxel-based atlases of images in the first cardiac phase to generate surface-based training data for a 3D+time statistical model covering both inter-subject and intra-subject (functional) deformations [155], by applying Principal Component Analysis (PCA) first to the mean shapes of each subject, and then to all phases of all subjects, minus their mean shapes. Another

3D+time statistical model of cardiac shape was based on direct bilinear decomposition of the surfaces in the training set [79], using a higher order Singular Value Decomposition. A similar model was used for segmentation of cardiac Magnetic Resonance (MRI) images by Zhu *et al.* [226].

A slightly different approach was followed in [96], where a single surface-represented segmentation was fitted to other segmentations to generate the point correspondence. With the surface sampling depending on the curvature observed in the initial shape, a bias in local mesh density may be introduced. Later, [120] used a similar technique but required full manual segmentation only of the initially chosen image.

In [225], sampling of the curves obtained by intersection of manual delineations with cut planes is employed to generate point correspondence.

A recurring requirement in nearly all of these approaches is the need for (manually) segmented images, as is also apparent from the review in [72]. The same holds true for the various methods suggested to achieve point correspondence on shapes represented by meshes. Overall, this has had the effect of limiting training set sizes in statistical shape modeling. Some effort was made to enlarge such training sets by adding synthetic variations [102], however, one could debate the plausibility of these variations. For further reading on surface-based atlases, including methods that generate surface point correspondence without prior volumetric point correspondence, please see [72].

In this chapter, we extend the works of [61] and [149], in that we construct an atlas to generate correspondences throughout our training population, and we use a surface-based representation so that this correspondence is carried by the surface mesh vertices. Only one volume needs to be segmented explicitly, creating the possibility to use arbitrarily large training sets.

2.2.2 Spatial Normalization

When a synthesized anatomy is used for an atlas, it is usually generated to represent a central tendency of a population. Estimating this central tendency is generally approached as a registration problem. This field has been a very active one since the landmark paper for average image construction by Guimond *et al.* [66]. Over the period since this work, a shift can be observed from image averaging towards transformation averaging, initiated by Christensen *et al.* [34]. Physical constraints have been introduced in the registration approaches, including diffeomorphic and other, more tissue-specific, constraints. Imaging modalities for the construction of atlases have become more varied as development on the mathematics of their deformation and similarity progressed.

Within the field of spatial normalization we can distinguish two main approaches: true groupwise registration on the one hand, and a processing of pairwise registrations on the other hand. The contributions for the latter are mostly strategies to process the outcomes of the pairwise registrations.

Using pairwise registrations starts out from a single reference image. Typically this one is chosen from the population at hand, either by an expert [152], or by an automated method; this can be based on subspace exploration [151] or on an information-theoretical basis [80].

Once the reference has been established, registration of each member of the

population to this reference is carried out, producing a deformation for each population member. The combination of these deformations typically leads to a mean deformation to be applied to the reference image, ideally transforming it into the mean image directly. As demonstrated in [66], this is not necessarily the case. Recently, various interpretations of the use of intermediate means have been proposed. Jongen *et al.* [93] suggested the construction of a mean from a small subset of the population, before registering the entire population to this mean estimate.

This is different from intermediate means used in manifold-based approaches. Jia *et al.* [89] proposed the derivation of a tree to guide registration step by step. The registration of a population member to the population mean is subdivided into steps of registration to neighbors nearer to the mean, thus constructing a path through the population. Child-parent relationships are determined through the stabilization of a clustering method. This contrasts with the approach of Wu *et al.* [220], who computes a Minimum Spanning Tree based exclusively on image intensity differences between images.

Other approaches explore the number of clusters that a population could be divided into. Atlas stratification [18] was an initial approach to only do this clustering, based on k -means. Sabuncu *et al.* [173] introduced the automated atlasing into a similar framework, but used a generalization of Estimation Maximization to discover the clusters.

True groupwise approaches tend to start from an average over the non-deformed images in the population. From that point onwards, Joshi *et al.* [94] update the deformations by explicit minimization of deformation magnitudes with respect to the Fréchet mean of the population, thus obtaining that mean. Lorenzen *et al.* [121] extended this work to handle multi-modality image data, while Fletcher *et al.* [58] replaced the Fréchet mean with the Fréchet median to ensure population membership of the atlas.

The above approaches use an explicit minimization of the average deformations required to deform a population member to the population-representative atlas. Other true groupwise approaches construct the atlas by maximization of a groupwise similarity criterion. This is either similarity of the population to the atlas via information theoretic criteria [15, 134], or a single measure quantifying similarity between all pairs [185]. In either case, an additional constraint is required to enforce zero average deformation.

Our framework can be changed to incorporate any of the above methods. However, in its current form we present it using the standard approach of Guimond *et al.* [66], with multi-scale diffeomorphic B-spline registration [170, 172].

2.3 Atlas Construction

This section covers the steps required to construct a population average image using image registration and the subsequent generation of an atlas represented using a surface mesh. Although various strategies exist for registering a set of images to an initially unknown average [15, 18, 45, 69, 89, 94, 121, 133, 185, 213], we chose the classical approach of registration to a chosen reference, as did Guimond *et al.* [66]. The simplicity of this approach holds some advantages:

1. The process is transparent. Unlike stratification [18] and manifold-based methods [69], it does not introduce more parameters than those required for the pairwise registration and the steps to update the mean image. Additionally, there is no interaction between these sets of parameters.
2. The complexity is kept to only one registration per subject per iteration, rather than a full cross-registration of all or a subset of subjects. Also, the optimization in a groupwise approach is a much higher-dimensional optimization problem, with ramifications for the robustness.
3. The strategies of stratification [18] and of using intermediate local averaging [69, 89, 90, 220] employ distance measures to determine which pairwise registrations must be carried out. It is unclear whether cardiac image volumes would be at an advantage or disadvantage for stratification compared to brain image volumes, given the increased number of disjoint regions of clearly different appearance.

The process is outlined in Fig. 2.1, with parameters listed in the last subsection of this section (2.3.6). It was implemented on a grid computing facility using Fura software (Grid Systems S.A., Palma de Mallorca, Spain) to handle the distribution. The facility was equipped with 12 nodes with 2 64-bit quad core processors each, with 16 Gb of shared memory per node.

Note that we synthesize a mean first cardiac phase image only. This is because of the greater stability of this phase in retrospective gating. Thus, our means are computed over the 138 subjects rather than the 2070 image volumes.

For further reading on cardiac image registration, we refer to Mäkelä *et al.* [128] and Tavakoli *et al.* [191]; for aspects of medical image registration in general, we refer to [74, 101, 127, 160, 182].

2.3.1 Imaging Data

The database of imaging data used in this study was retrospectively collected from a clinical cohort of 138 consecutive patients that underwent a Computed Tomography (CT) examination as part of their routine diagnostic protocol for suspected coronary artery disease, and follow-up. The resulting population distribution, as well as imaging parameters, are outlined in Table 2.1. All information was anonymized before its transfer from the clinic to our group.

While our database is not as large as that of the CAP [59], our image resolution is 2.5 to 5 times higher in-plane [59, 189], and 3 to 4 times higher axially [59, 95, 189]. This enables us to capture more anatomical detail than would be possible with clinical resolution MRI data. Similarly, ultrasound has the advantages of noninvasiveness and higher temporal resolution, but these are undone by the speckle patterns which make automated further processing—specifically inter-subject registration—extremely difficult.

2.3.2 Reference Selection

The selection of a suitable reference is one of multiple approaches to minimize bias in the synthesized mean volume. The first box in Fig. 2.1 is therefore an attempt to select the reference already close to the unknown mean. For this we use

TABLE 2.1: DETAILS OF DATA. CAD: CORONARY ARTERY DISEASE; MI: MYOCARDIAL INFARCTION

<i>Imaging</i>	Machine	Toshiba Aquilion 64 ¹
	Tube voltage	120 kV
	Tube current	400-430 mA
	Contrast agent	Xenetix 350
	- quantity	80-100 ml
	- rate	5 ml/s
<i>Reconstruction</i>	Workstation	Vitrea ²
	Resolution	$0.4 \times 0.4 \times 2.0$ mm
	In-plane grid	512×512 pixels
	Axial grid	65.3 ± 11.3 slices
	Volumes per cycle	15
<i>Population</i>	Size	138
	Men/women	94/44 (64% men)
	Age	59.4 ± 13.0
	Pathology CAD only	20%
	Pathology CAD+MI	20%

¹ Toshiba Medical Systems, Tochigi, Japan² Vital Images Inc., Minnetonka, MN, USA

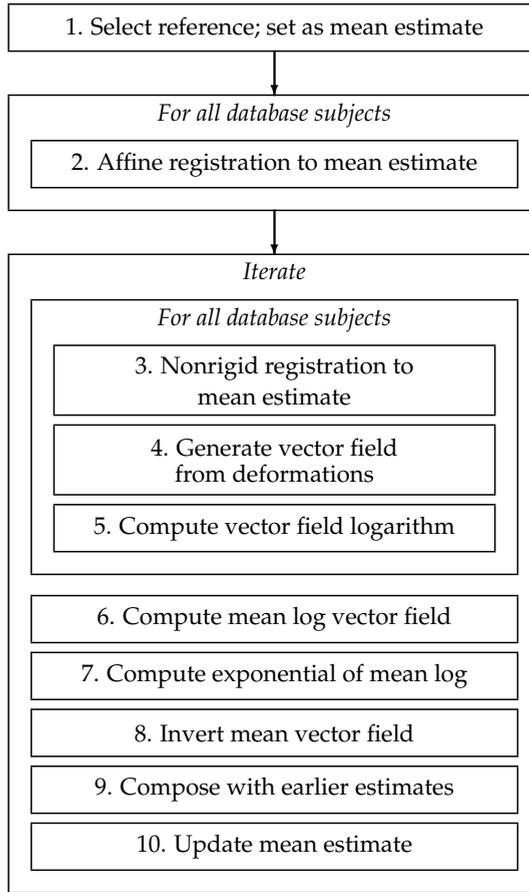


Figure 2.1: Flow diagram of the mean image synthesis. The affine transformation obtained initially is re-used as an initial transform in the non-rigid registration.

a heuristic based on a Groupwise Mutual Information (GWMI) score, an extension of Mutual Information (MI) [125] that describes the amount of information between a single image and a set of images [80]. We compute the GWMI score over affinely registered volumes. The affine registration is required to remove the detrimental influence of pose and size on the score.

In short, the GWMI between image \mathbf{I} and image set \mathcal{J} is defined using the Shannon entropy H over voxel intensities as

$$\text{GWMI}(\mathbf{I}, \mathcal{J}) = H(\mathbf{I}) + H(\mathcal{J}) - H(\mathbf{I}, \mathcal{J}). \quad (2.1)$$

With the probability of an intensity j in the set defined as

$$p(j) = \sum_{\mathbf{J} \in \mathcal{J}} p(j|\mathbf{J})p(\mathbf{J}), \quad (2.2)$$

and with an *a priori* uniform probability distribution over the candidate images

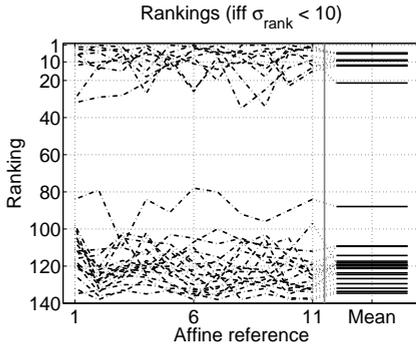


Figure 2.2: *Rankings and mean ranking for subjects with a ranking standard deviation below 10. One could consider the six subjects with both $\sigma_{\text{rank}} < 10$ and $\mu_{\text{rank}} < 10$ to be suitable subjects for the initial reference.*

($p(\mathbf{J}) = \frac{1}{|\mathcal{J}|}$ with $|\cdot|$ denoting cardinality), we are effectively summing the histograms of the images in the set \mathcal{J} defined on the same grid of bins:

$$\text{Histogram}(\mathcal{J}) = \sum_{\mathbf{J} \in \mathcal{J}} \text{Histogram}(\mathbf{J}). \quad (2.3)$$

The joint probabilities of an intensity i in the candidate reference with an intensity j in the set \mathcal{J} are defined analogously as

$$p(i, j) = \sum_{\mathbf{J} \in \mathcal{J}} p(i, j | \mathbf{J}) p(\mathbf{J}), \quad (2.4)$$

also coming down to a summation of the joint histograms:

$$\text{JointHistogram}(\mathbf{I}, \mathcal{J}) = \sum_{\mathbf{J} \in \mathcal{J}} \text{JointHistogram}(\mathbf{I}, \mathbf{J}). \quad (2.5)$$

In the case where \mathcal{J} consists of only one image, this reduces to the standard MI. The method takes one parameter, the number of histogram bins.

The score represents the amount of information each volume carries with respect to the remainder of the volumes, taking into account the variation observed in this set. To make the process even more robust to the reference volume used in the affine registration, we registered all images affinely to 11 randomly selected volumes, and applied GWMI to each of the 11 resulting sets of 138 volumes, obtaining 11 rankings. The reference ultimately chosen for the atlas construction was the volume with the best mean rank. Figure 2.2 illustrates the varying rankings of subjects presenting a rank standard deviation of less than 10 positions with the dash-dot lines. The solid lines (connected to the associated dash-dot lines using dotted lines) illustrate the mean ranks for these subjects. It shows that high and low ranks are relatively stable [80].

2.3.3 Global Registration

With the reference subject chosen from the population as described in the previous section, we register the volume corresponding to each subject's first cardiac phase to the first phase of the reference subject. Corresponding to the second box in Fig. 2.1, the first step of this registration is a global registration, requiring

few parameters to optimize. It has been shown that such an approach increases the robustness of the overall scheme [172].

We aim to model shape variation through a statistical shape model rather than a statistical deformation model (as in [43, 171]). As the inverses of both the global and local transformations are applied to the mesh corresponding to the average image, it is of no importance which global transformation model is used. Hence, we use affine registration to remove global variations, minimizing as much as possible the deformation remaining to be resolved in non-rigid registration.

The small number of degrees of freedom of an affine transformation make the registration suitable for a gradient descent optimization. The similarity metric is MI [125], of which the gradient is computed using Mattes' method [135]. As it was developed for inter-modality image registration, MI is capable of handling varying intensities for corresponding structures, or in other words, to register images with different image intensity profiles. In our dataset, we observed image intensity variations throughout the population in the blood pools, due to varying contrast agent concentrations (note the 25% variation reported in Table 2.1). Additionally, some subjects presented with pacing devices, leading to some very strong local intensity variations.

2.3.4 Diffeomorphic Nonrigid Registration

Following the affine registration, in the third box in Fig. 2.1, nonrigid registration is used to resolve the remaining variations between the reference and population images. With a large number of registrations to be carried out, a fast registration approach is desirable. The multiscale approach using B-splines [172] fulfils this criterion.

A later work by Rueckert *et al.* [170] focused on imposing diffeomorphic constraints on B-splines, meaning that the resulting deformations are smooth and invertible. From an anatomical as well as a technical viewpoint, this is a very important property. Borrowing from the work of Choi and Lee [33], the injectivity of a B-spline-based deformation is guaranteed if the local deformation is limited to 0.4 times the spacing of the control points. As the group of diffeomorphic transformations is closed for the composition operator (\circ) [199], this allows for a combination of the multiscale and the diffeomorphic B-spline approaches. That is, the composition of diffeomorphic B-splines, regardless of equal or varying control point spacings, produces a diffeomorphic transformation. However, it is generally not possible to represent this composition as a single diffeomorphic B-spline.

The multiscale approach is robust and allows for fast registration. The constraints on control point movement are implemented straightforwardly using the Insight Toolkit (Kitware, Inc., Clifton Park, NY, USA) [99] and a bounded quasi-Newton optimizer [25]. Klein *et al.* [101] showed that this optimizer provides a good balance between speed and accuracy when used for registration using B-splines and MI.

In the original formulation of the diffeomorphic B-splines, new splines were added until convergence of the registration. However, we choose to maintain a fixed set of transforms. This results in control over the maximum possible

local deformation, and therefore prevents the construction of outliers that could corrupt the computation of the mean deformation.

Using T for transformations obtained by composition of transforms, A^s for the affine transformation from the reference subject to subject s , B^s for a B-spline based deformation from the reference subject to subject s , $B^{s,\tau}$ for the B-spline based deformation for the nonrigid deformation from frame $\tau - 1$ to frame τ of subject s , the total transformation to frame τ of subject s is given by

$$T_{\text{total}}^{s,\tau} = T_{\text{intra}}^{s,\tau} \circ T_{\text{inter}}^s \circ A^s \quad (2.6)$$

with

$$T_{\text{intra}}^{s,\tau} = \bigcirc_{t=2}^{\tau} B^{s,t} \quad (2.7)$$

and

$$T_{\text{inter}}^s = \bigcirc_{i=1}^{n_b} B_i^s. \quad (2.8)$$

Here, n_b is the number of B-spline transforms we use to compose the deformation from the reference to subject s , and \bigcirc is the big version of the composition operator \circ . After the affine registration has brought the images into a global alignment, by controlling the spline control point spacing and n_b we control the maximum total local deformation. We will show later that we observed a convergence rate similar to that reported by Guimond *et al.* [66].

After the non-rigid registration, we represent the inter-subject transformations T_{inter}^s as displacement vector fields φ^s , thus simplifying further processing in the absence of the composition operators. The use of B-splines to model our deformation reduces the number of degrees of freedom in the optimization problems by multiple orders of magnitude as compared to a parameter-free approach that directly produces a vector field, like the diffeomorphic demons [208].

2.3.5 Mean Deformation

The vector fields $\varphi(s)$ representing $T_{\text{inter}}(s)$, and defined on the same grid as the reference image, belong to the group of diffeomorphic vector fields. This group is closed only for the composition operator (\circ) [199], and therefore the euclidean metrics of addition and division cannot be applied directly. If the deformation is a geodesic in the metric space defined by the deformation model used for the registration, one can use the log-Euclidean framework of Arsigny *et al.* [8] to work in the tangent space of these diffeomorphic vector fields, in which Euclidean metrics can be applied. The mapping to and from the tangent space is provided by the logarithm and exponential operators, respectively. Due to the multi-scale approach to the registration, we should assume that we approximate the geodesic closely enough for the log-Euclidean framework to be applicable. We will later comment on the accuracy of the logarithm computation (Sec. 2.3.6). This leads to the definition of the mean vector field as

$$\bar{\varphi} = \exp \left(\frac{1}{n_s} \sum_{i=1}^{n_s} \log \varphi^s \right), \quad (2.9)$$

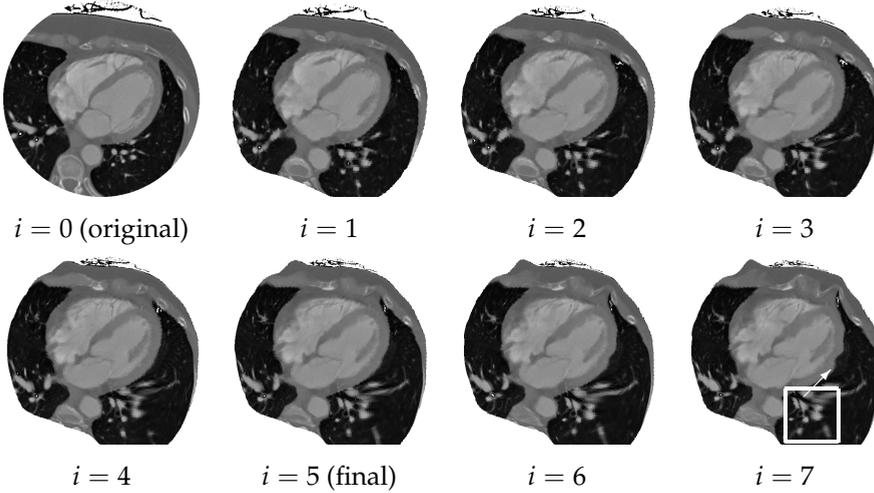


Figure 2.3: The development of $\tilde{\mathbf{I}}_i$, for $0 \leq i \leq 7$. After five iterations, the composed deformations, although diffeomorphic and less biased, become implausible. This is visible in the lateral wall of the left ventricle (the white arrow in $i = 7$, but also visible for $i = 6$).

corresponding to the 5th through 7th boxes of Fig. 2.1. Note again that the mean is computed over the first-phase images only.

Given the mean vector fields obtained in the iterations up to i , and the current approximation of the mean image $\tilde{\mathbf{I}}_{i-1}$, the approximation is updated in boxes 9 and 10 of Fig. 2.1 using

$$\tilde{\mathbf{I}}_i = \bar{\varphi}_{0\dots i}^{-1} \circ \tilde{\mathbf{I}}_0, \quad (2.10)$$

where

$$\bar{\varphi}_{0\dots i}^{-1} = \bar{\varphi}_i^{-1} \circ \bar{\varphi}_{0\dots i-1}^{-1}. \quad (2.11)$$

For completeness, $\bar{\varphi}_{0\dots 0} \equiv \bar{\varphi}_0$. The development of $\tilde{\mathbf{I}}_i$ is illustrated in Fig. 2.3.

Arsigny's algorithm [8] for fast computation of the vector field logarithm is an inverse scaling and squaring method, exploiting the equivalence

$$\log \varphi = 2^n \log \varphi^{2^{-n}}. \quad (2.12)$$

For sufficiently large n , the logarithm can be approximated by

$$\log \varphi^{2^{-n}} \approx \varphi^{2^{-n}} - Id, \quad (2.13)$$

where Id stands for the identity. The square roots are computed using a gradient descent approach, which uses the inverse of the transform. The vector field must extend beyond the image domain and smoothly reduce to zero magnitude to avoid discontinuities in the image domain in the inverse vector field. Effectively, the actual deformations must be contained in a box, rigid under log-

TABLE 2.2: REGISTRATION PARAMETERS AND VALUES

	<i>Parameter</i>	<i>Value</i>
All intersubject	Mask—lower threshold	-1000 HU
	Mask—upper threshold	1500 HU
	Histogram bin count	50
	Voxels per sample	30
Affine	Gaussian smoothing kernel σ	1.0 mm
	Min. optimization step length	0.001
	Max. optimization step length	1.0
32 mm inter (twice)	Gaussian smoothing kernel σ	1.5 mm
16 mm inter	Gaussian smoothing kernel σ	1.0 mm
8 mm inter	Gaussian smoothing kernel σ	0.5 mm
40 mm intra	Histogram bin count	60
	Voxels per sample	30
	Gaussian smoothing kernel σ	0.5 mm
	Mask	warped atlas, 10 mm dilated

arithm, averaging and exponential. The finite support of B-splines provides this without any additional cost in registration complexity.

2.3.6 Parameters and Settings

Each step in the process of synthesizing the mean image volume depends heavily on good parameters, with the image registrations a prime example. However, optimizing each of these parameters for each instance is unfeasible. Therefore it is imperative that only a few, if multiple, distinct configurations are used. As a consequence, for the majority of registrations these configurations are sub-optimal.

In order to obtain the best possible registrations under these conditions, a small set of six first-phase volumes was used to obtain a configuration suitable for all pairwise registrations in the subset, which was the first instance of human intervention. Using a larger subset of 20 first-phase volumes, the useability of this configuration was verified. Table 2.2 lists the values obtained for the registration parameters.

For the inter-subject registrations in the mean image synthesis, image masks are defined based on the image intensities of the fixed and moving images, ranging from air (-1000 Hounsfield Units (HU)) to cortical bone (1500 HU). The mutual information similarity metric is computed from 50-bin histograms generated from voxels randomly sampled within the mask regions, at a rate of one per 30 voxels.

TABLE 2.3: PARAMETERS AND VALUES FOR NON-REGISTRATION ELEMENTS

	<i>Parameter</i>	<i>Value</i>
Reference selection	Histogram bin count	50
Vector field logarithm	Inversion—step size	1.0
	Inversion—stopping criteria	99% error reduction 10 iterations
	Squaring—step size	1.0
	Squaring—stopping criteria	95% error reduction 10 iterations
	Root order (n in Eq. 2.12)	7
Mean vector field inversion	Step size	0.3
	Stopping criterion	99% error reduction 100 iterations

The non-rigid registration uses a composition of B-splines of increasing resolution. Due to the hard limit on the local deformation of each B-spline, we control the maximum local deformation. Using only one each of 32, 16 and 8 mm control point spacing, this would amount to a maximum of 22.4 mm $((32 + 16 + 8) \times 0.4)$. We assumed this to be insufficient, and added an additional allowance of 12.8 mm through an additional B-spline with 32 mm spacing.

The segmentation of the atlas is used to provide the masks for the intra-subject registrations, while the greater histogram resolution is used to improve sensitivity to smaller differences.

An alternative would be to apply a small-range parameter sweep around the values found, for every registration. However, this would increase the computational load prohibitively, even if only one parameter is varied at any given moment.

For the elements of the framework that are not registration steps, we list the parameters in Table 2.3. Error reductions are computed with respect to the error at initialization (identity vector field), and the iteration limits supersede this criterion. For the root order, Arsigny *et al.* [8] empirically reported convergence of the log computation with $n = 7$ in Eq. 2.12. We found this value suitable for our data as well, yielding vector norm errors under 5% for 93.7% of vectors when comparing φ^s to $\exp(\log(\varphi^s))$ in 20 subjects.

Note that the computation of the mean log vector field, the composition of the mean estimate with earlier estimates, and the update of the mean image (boxes 6, 9 and 10 in Fig. 2.1 require no parameters).

2.3.7 Atlas Segmentation

Through registration, we have obtained a mapping from the atlas coordinate system to each of the subjects in our population. Thus, by segmenting the at-

las volume, we obtain a segmentation of each member of the population as well. We represent the segmentation not by a mask image but by a triangulated mesh. This allows us to define other structures in local coordinate systems, structures which are typically obtained through literature description or separate modeling approaches.

By far most human interaction in the procedure is concentrated at this step. An anatomist specialized in cardiac anatomy (D. S.-Q.) segmented the synthesized mean image using GIMIAS v1.2.0b software [107]. In each slice of the image, the structures listed in Table 2.4 were outlined using a free-hand polygon tool, which labeled the voxels on the same grid as the image.

For the coronary arteries, their intersections with the image planes were marked using a pointer tool. Subsequently these sets of points were connected and the arteries themselves were modeled by fitting knotted cones to the segments.

2.3.8 Mesh Construction and Processing

The manual segmentation of the mean image is converted to a mesh representation. There is a sizeable body of literature on the creation of ‘reasonable’ non-manifold meshes from multi-label medical imaging data. These meshes have a smooth surface, high accuracy with respect to the non-smooth (i.e., with staircase artifacts) surface, and with good element quality (usually polygon regularity). In addition to the challenge to achieve this in binary (foreground+background) data, the multi-label property of the input data presents the added challenge of smooth surfaces between materials while maintaining smooth outer surfaces. Examples of approaches to this problem are multi-material extensions to marching cubes (MC) [70, 221], marching tetrahedra [49], dynamic particle systems [139], non-manifold dual contouring [14, 224], and multi-level partitions of unity [51].

We use a more straightforward approach, schematically presented in Fig. 2.4. We first generate a new segmentation by merging labels from the original segmentation. In this case, into three groups: left blood, right blood, and myocardium. The resulting masks are resampled using shape-based interpolation [73, 166], and then triangulated using MC [119]. Distance transforms of the original segmentations are then used to assign labels to the faces of the resulting mesh.

Subsequently, the mesh is coarsened using an approximate centroidal Voronoi tessellation (ACVT) [204]. That is, each Voronoi cell is represented by a set of mesh faces in their entirety, and the seed of each cell coincides with the cell centroid. This approach has the advantage of allowing full control over the number of vertices in the coarse mesh, and as outlined in [204], is guaranteed to converge to a global optimum. Additionally, the resulting mesh is independent from the original MC mesh.

With the coarsening, the mesh is smoothed somewhat due to a lower sampling rate. However, staircase artifacts may remain especially when a large number of vertices is used. Therefore the meshes were submitted to a smoothing step [190]. Overall, the coarsening and smoothing introduced some local distortions, though in general these remained either within acceptable bounds,

TABLE 2.4: ATLAS SUBPARTS (3D) AND ELEMENT STATISTICS

<i>Name</i>	<i>Vertex count</i>	<i>Vertex density</i> (cm^{-2})	<i>Face surface</i> (mm^2)
1 Left Ventricular Myocardium	2032	10.66	4.85 ± 1.88
2 Intraventricular Septum	1492	12.40	4.22 ± 2.19
3 Right Ventricular Myocardium	3454	10.72	4.77 ± 2.34
4 Left Atrial Myocardium	2010	11.68	4.43 ± 1.61
5 Right Atrial Myocardium	1962	10.90	4.80 ± 1.98
6 Aorta	2654	9.40	5.35 ± 1.46
7 Vena Cava (inferior and superior)	742	11.25	4.69 ± 1.81
8 Pulmonary Trunk and Artery	1032	11.46	4.44 ± 1.47
9 Pulmonary Veins (left and right, inferior and superior)	1232	10.27	5.02 ± 1.79
10 Anterior descending artery			
11 Circumflex artery			
12 Diagonal artery			
13 Right coronary artery			
Total/overall	16113	10.44	4.78 ± 1.93

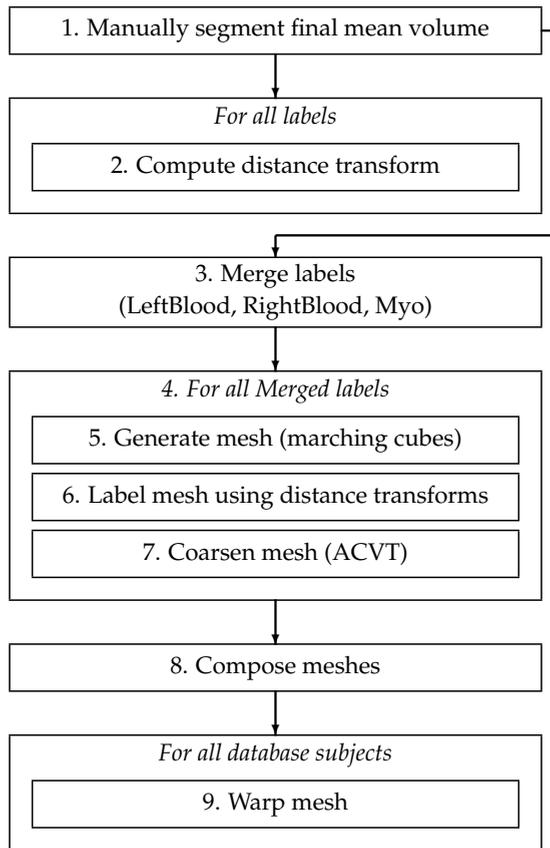


Figure 2.4: Flow diagram of the meshing steps. ACVT: Approximate Centroidal Voronoi Tessellation. Initially the nine labels are merged into three; using the nine distance transforms, the vertices on the marching cubes meshes are labeled before coarsening. The atlas is finally composed by putting the three labeled, coarsened meshes together.

or they smoothed out undesired features which appeared as a result of local inconsistencies in the manual segmentation.

The transformations T_{inter}^s and A^s were used to warp the atlas mesh to the first cardiac phase of each subject s . Subsequently the transformations $T_{\text{intra}}^{s,\tau}$ were applied to these meshes to obtain the meshes pertaining to the remaining 14 phases of each subject.

2.4 Atlas Properties and Statistical Model

We developed the statistical atlas with the aim of using it in a computational physiology setting. In this section we outline the properties of the atlas that

enable this. The atlas is available for downloading at our website¹.

2.4.1 Structures

The atlas has been designed to be instantiated in varying compositions, ranging from the entire heart to a single substructure. This serves the purpose of extracting or generating knowledge from images and simulation studies that only employ the structure of interest, without changing the atlas itself. This in turn enables the seamless incorporation of this new information into the atlas.

The structures outlined are listed in Table 2.4 and colored separately in Fig. 2.5. For the left ventricle, the myocardium and blood pool were segmented, while for both atria, for the right ventricle and for the trunk of the aorta, only the blood pool was segmented. On CT images, the walls of these structures are not discernible and as such are estimated from the boundaries of the blood pool. We did this by extrusion of the surface along its normal direction, to generate the thicknesses reported in the literature. For the right ventricle (RV) wall, the surface was extruded by 4 mm [163], for the aorta this was 3 mm [54] while the atria were extruded by 2 mm [67]. The vertex counts in Table 2.4 already include the extruded parts.

2.4.2 Statistical Modeling

Paramount to the clinical value of computational physiology is the possibility to perform patient-specific simulations, where the computational domain is a patient-specific geometry, or to generate large virtual populations following the model population statistics. Approaches based on statistical shape models to segment patient geometries from image data have shown their worth many times over. For inherently dynamic shapes, naturally one should consider using a dynamic statistical shape model, such as those described in [79, 155, 226]. Here we reiterate some background on the bilinear spatio-temporal statistical shape model, constructed as per [79] from the framework's output.

In Sec. 2.3.8 we described how we obtained the meshes corresponding to the 2070 image volumes we began with. From these, we construct a spatio-temporal statistical model of shapes, decomposing our data along the axes of anatomical variation and motion-induced shape change using a bilinear model as described by Tenenbaum and Freeman [196]. They used the terms *style* and *content* to define the two sources of variation. We could consider here the style to be the temporal element, whereas the content is the anatomical variation that exists between subjects. By using a bilinear model, we maintain temporal information which would be lost if the more traditional PCA [92] had been used.

However, for a rudimentary outlier detection, we used PCA [92] to generate a simple statistical model of the first-phase shapes, explaining 50% of total variation, amounting to five modes. The model parameters of these shapes were used to identify potential misregistration. A shape (and the associated subject) was marked as suspicious if any of its parameters exceeded four standard deviations from the mean. This led to the flagging of four subjects. In each of these

¹<http://www.cistib.upf.edu/cistib/index.php/downloads>, or <https://sites.google.com/site/cornehoogendoorn2013/home/publications/downloads>

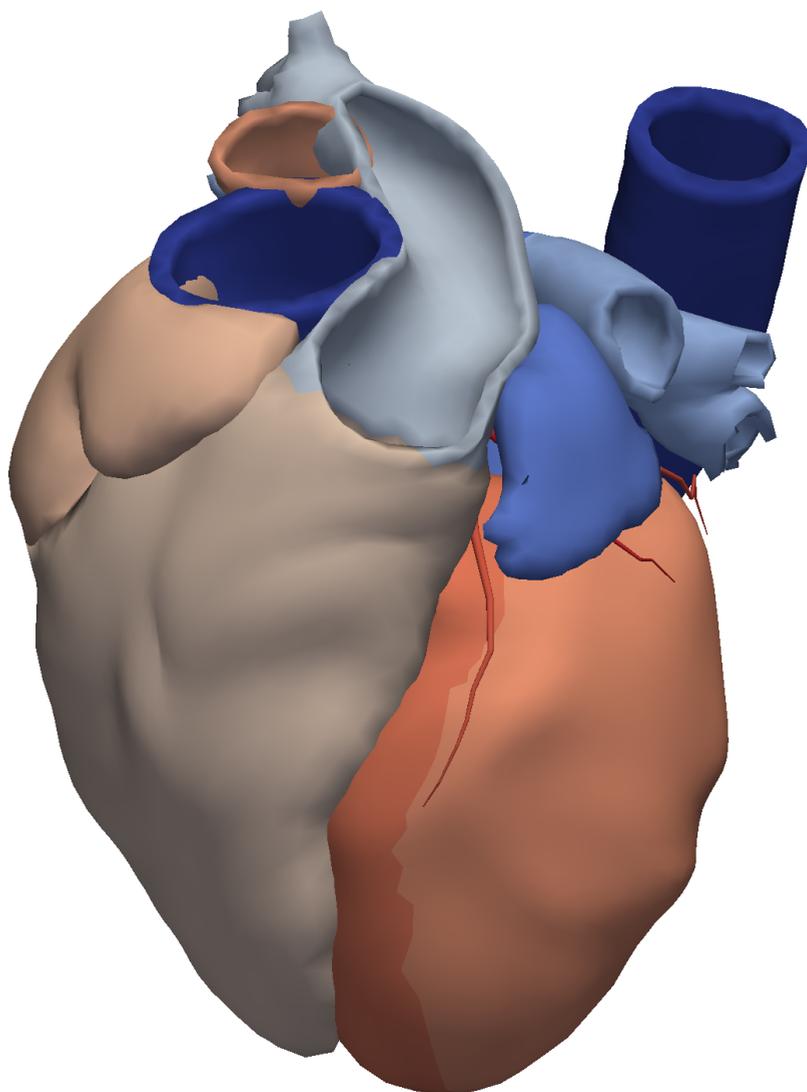


Figure 2.5: *The full atlas mesh, corresponding to the synthesized mean image. Different colors indicate different structures. Reproduced in color on page xlii.*

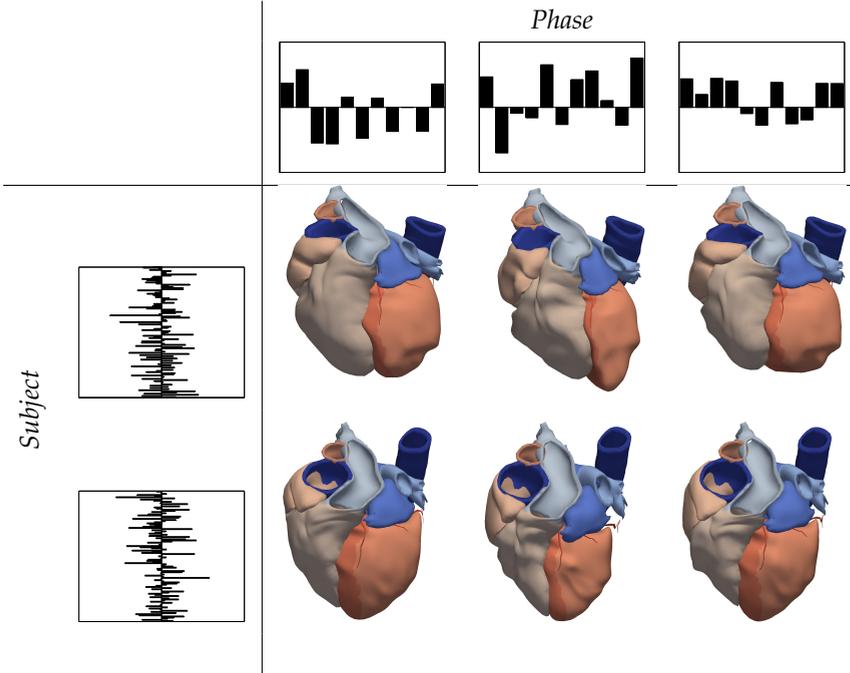


Figure 2.6: Three phases of two subjects, with their bilinear parameter sets derived from the full 134-subject data set. The dimensionality of the phase and subject parameters (11 and 115, respectively) are based on explaining 95% of total variance for each axis.

cases, the affine registration had failed, and thus the subjects were discarded. We generated the spatio-temporal model from the remaining 134 subjects (2010 shapes).

Training of a bilinear shape model is done as described in [79, 196]. With 134 subjects and 15 phases each, the model reduces the data dimensionality to at most 147. For comparison, using Perperidis *et al.*'s approach [155] one would need to reduce the dimensionality from 2025, also from two sets of parameters. For the temporal synchronization, necessary due to the electrocardiogram-based retrospective gating of the CT data, we did use Perperidis' method [154] to identify the end-systolic and end-diastolic phases. More detail on the application of this method to surface meshes is provided in [79].

Figure 2.6 shows three phases of four subjects, together with bar plots for their subject and phase parameters. The applications of cardiac bilinear shape models have already been demonstrated in image segmentation [226] and motion analysis [57, 161]. In earlier works we have also demonstrated pipelines towards electrophysiological [76] and mechanical simulations [214] from surface-based cardiac segmentations.

Figures 2.7 and 2.8 show the first two modes of variation as mean plus and minus one and a half and three standard deviations of both the full heart and the left ventricular endocardium, of the traditional PCA-based shape model.

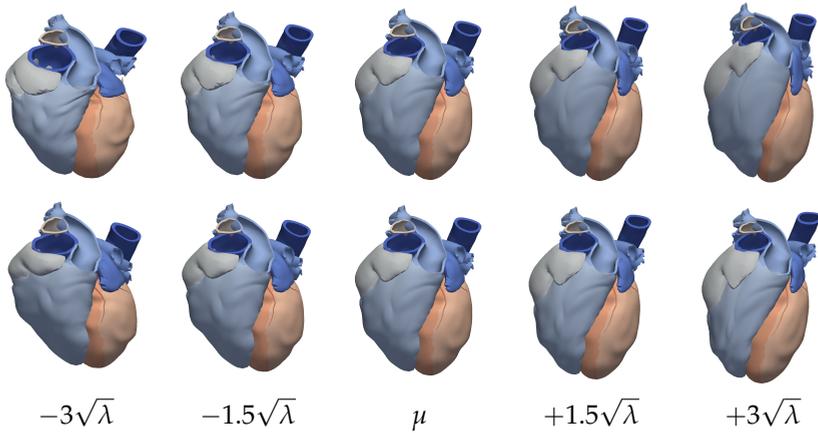


Figure 2.7: The first two modes of variation in the linear statistical model of the full heart at end diastole, showing the mean plus and minus one and a half and three standard deviations.

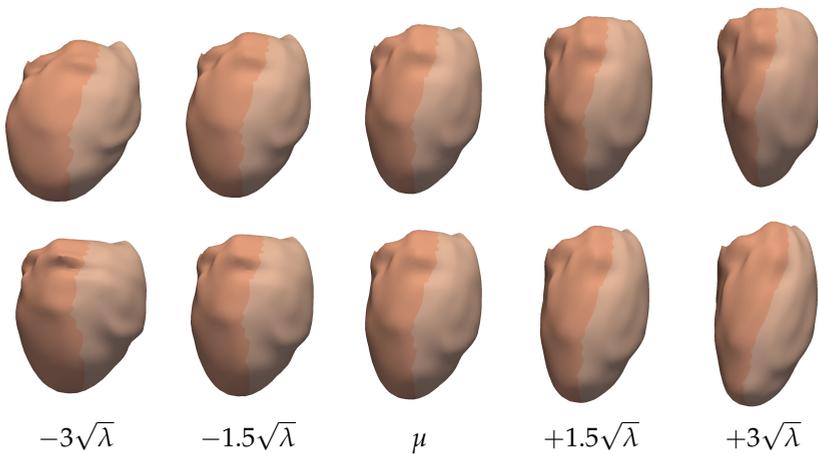


Figure 2.8: The first two modes of variation in the linear statistical model of the left ventricle at end diastole, showing the mean plus and minus one and a half and three standard deviations.

Coincidentally in both models, the first two modes contain a strong component of elongation versus roundness.

2.5 Validation

The quality of the atlas depends on the quality of the registration. In turn, the registration quality depends on each of the basic components of a registration algorithm: similarity metric, transform, and optimizer, and on the initial simi-

larity between volumes [217].

In this section we validate the atlas through evaluation of the registration. This consists of three elements:

- Random selection of volumes and slices. The dataset size does not allow the evaluation of all registrations. It is also undesirable to fully segment volumes for all the selected subjects. Thus, a random selection is made, based on some heuristics which make the subjects and slices therein useful.
- Evaluation of vector field inversion. In order to evaluate everything in one common space to avoid bias to the specific subjects involved, all segmentations and landmarks are warped to the atlas space. This is done using the inverse of the vector fields we obtain during registration. As the inversion is an optimization problem in itself, the inversion error presents a localization uncertainty. We report the inversion error at the locations of the landmarks and contours.
- Distance-based and overlap-based evaluation. After random slice selection and quantification of uncertainty in landmark localization, we report contour-to-surface and point-to-point errors and overlap accuracy to evaluate the registration quality.

Furthermore, we illustrate the convergence of the synthesis of the mean image.

2.5.1 Registration—Selection of Volumes and Slices

Both the optimization and the evaluation of all registrations carried out is unfeasible with large datasets. Therefore it is of importance that a meaningful subset of the data is selected for the evaluation. We used random sampling after identification of volumes and slices satisfying four simple qualification criteria: two for volumes, and two for slices.

A volume qualifies if the image intensity histogram peak pertaining to contrast enhanced blood could be identified, and the modal intensity and width of the peak fall within two interquartile distances from the population median for these two quantities. Additionally, the volume must contain a sequence of at least 20 qualified slices.

A slice qualifies if at least 3.5% of its pixels fall in the intensity range of contrast enhanced blood, as determined via the volume's intensity histogram. Additionally, the slice has to be part of a sequence of at least 20 in the volume.

The resulting sets of qualified slices are nearly guaranteed to intersect the cardiac and supracardiac blood pool. The threshold of 3.5% contrast enhanced blood pixels excludes slices where the enhanced blood is provided exclusively by the vena cava and aorta. The 20 slices cover 38 mm, which is much less than the cardiac blood pool should extend for, but excludes slices that were incorrectly qualified by the first slice criterion.

After determining qualifying volumes and slices, we draw randomly from the volumes. The sequences of qualifying slices in these volumes are divided into three equal parts and we randomly draw a slice from each part.

2.5.2 Registration—Measurements

Contours

From the qualifying subjects, 36 were selected randomly. Six researchers each segmented the slices of six subjects in Graphical Interface for Medical Image Analysis and Simulation (GIMIAS) v1.2.0 [107] using the same tool as used to outline the atlas structures. The set of structures to be segmented correspond either to one or to multiple structures of the final atlas:

1. Aortic root
2. Left ventricular cavity
3. Left ventricular myocardium (constituting the left ventricular free wall and the interventricular septum)
4. Right ventricular cavity
5. Left atrial cavity (including pulmonary vein trunks)
6. Right atrial cavity (including trunk of vena cava)

For each of these structures, the boundary of the segmentation in the slice was extracted and warped to the reference space using the inverse of the vector fields obtained during registration. Subsequently the mean unsigned distance from the warped contour to the corresponding structure's surface in the atlas was recorded. These distances are reported for each of the structures in the top half of Table 2.5.

Overlap

For each of the structures listed above, we also measured the overlap with the atlas segmentations. We warped the voxel-based segmentation into the atlas space using nearest neighbor interpolation. By using three labels (foreground, background in segmented slice, background elsewhere), we could identify and count false negative voxels ('fn') in addition to true and false positive voxels ('tp' and 'fp', respectively). These counts were combined into F-scores for the registrations in Table 2.6. The F-scores are defined as $F = \frac{2pr}{p+r}$, with precision $p = \frac{tp}{tp+fp}$ and recall $r = \frac{tp}{tp+fn}$. Thus, perfect matchings would yield an F-score of 1.

Landmarks

The contour-based evaluations have drawbacks in that they are not measuring the registration accuracy in a truly three-dimensional manner. By measuring the point-to-point accuracy of a set of anatomical landmarks, such evaluation is possible.

From the pool of subjects, 36 were selected for landmark-based evaluation. The same six researchers each processed 12 subjects; each subject was processed

TABLE 2.5: REGISTRATION AND VECTOR FIELD INVERSION ACCURACY OF SELECTED ANATOMICAL LANDMARKS AND CONTOURED STRUCTURES

	<i>Interobserver</i>		<i>Registration error (mm)</i>	<i>Inversion error (mm)</i>
	<i>variability (mm)</i> <i>(image space)</i>	<i>(ref. space)</i>		
<i>Contours</i>				
1	Aorta		2.19 ± 1.81	0.33 ± 0.19
2	LV cavity		1.83 ± 0.64	0.45 ± 0.35
3	LV myocardium		1.59 ± 0.45	0.48 ± 0.38
4	RV cavity		2.45 ± 0.87	0.72 ± 0.68
5	LA cavity		2.31 ± 1.73	0.43 ± 0.45
6	RA cavity		2.76 ± 1.21	0.51 ± 0.39
<i>Landmarks</i>				
1	LV endocardial apex	4.34 ± 2.32	9.28 ± 10.12	0.47 ± 1.04
2	Aortic valve center	4.40 ± 2.88	4.69 ± 3.14	0.38 ± 0.32
3	Mitral valve center	7.31 ± 3.87	7.68 ± 3.84	9.82 ± 12.04
4	Left coronary origin	3.76 ± 2.18	3.40 ± 1.84	5.50 ± 6.86
5	Right coronary origin	2.52 ± 1.30	2.54 ± 1.31	9.22 ± 7.88
				0.89 ± 1.76

TABLE 2.6: REGISTRATION ACCURACY QUANTIFIED BY F-SCORE OF OVERLAP OF SELECTED STRUCTURES

	<i>F-score</i>			
	<i>Initial</i>		<i>Final</i>	
	<i>mean ±</i>	<i>standard deviation</i>	<i>mean ±</i>	<i>standard deviation</i>
	<i>overall</i>	<i>overall</i>	<i>overall</i>	<i>overall</i>
1 Aorta	0.850 ± 0.152	0.833	0.871 ± 0.105	0.867
2 LV cavity	0.925 ± 0.021	0.925	0.916 ± 0.023	0.917
3 LV myocardium	0.773 ± 0.055	0.775	0.773 ± 0.049	0.774
4 RV cavity	0.823 ± 0.070	0.825	0.817 ± 0.043	0.817
5 LA cavity	0.723 ± 0.131	0.723	0.699 ± 0.104	0.700
6 RA cavity	0.643 ± 0.167	0.636	0.628 ± 0.106	0.622
7 Left blood (1+2+5)	0.904 ± 0.056	0.900	0.895 ± 0.042	0.893
8 Right blood (4+6)	0.821 ± 0.079	0.821	0.806 ± 0.058	0.804
Overall		0.838		0.827

by two researchers and each pair of researchers shared either two or three subjects.

For each subject, five landmarks were identified and marked using GIMIAS v1.2.0 [107]: the endocardial apex, the aortic and mitral valve centers, and the origins of the left and right coronary arteries. Like the contours, the landmarks were warped to the reference space. We report the interobserver variations both prior to warping and after warping, and the point-to-point registration error in the reference space. These results are presented in the bottom half of Table 2.5.

Vector field inversion

To warp the contours and landmarks from the subject space to the reference space, we require the inverse of the vector fields obtained during registration. The computation of this inverse is an optimization problem, and as such will introduce some uncertainty in the valuation. We quantify the error in the estimation of the inverse as the vector magnitude of $\varphi^s \circ (\widetilde{\varphi^s})^{-1}$ at the landmark and contour locations in the subject image space. The last column of Table 2.5 lists these errors.

2.5.3 Registration—Results

The contour-based evaluation reported in Table 2.5 provides insight into the agreement on boundaries. The average distance of these contours ranges from 1.59 mm for the left ventricular area to 2.76 mm for the right atrial area. This compares favorably with the slice spacing of 2.0 mm. The difference in performance between the left ventricular area and the atrial and right ventricular areas is easily explained by the low complexity of the left ventricle (LV) compared to the atria, and the much smaller influence of trabeculation on the LV blood-myocardium boundary strength.

The landmark errors compare well with the interobserver variation for the valves and the left coronary origin. The greater error in the apex is due to a sliding along the LV cavity wall; it is on the boundary but not at the apical point. During registration, boundaries are matched yet boundary curvature, by which observers would identify the apex, is not taken into account. The difference between interobserver variation and registration error of the right coronary artery (RCA) origin is due to the size of the RCA. While its position is stable at the base of the aorta, its appearance is often weak, thus offering no boundary for the registration to match.

The left ventricular apex and mitral valve center show a greater interobserver variability than the other landmarks. For the apex, this can be attributed to a dependence on image plane orientation; the observers are free to rotate the image planes to best visualize the landmark to be located, which for the apex tends to be a pure short-axis view. The apparent apex location varies along the endocardial wall with the deviation from this view. The interobserver difference in the mitral valve is due to its structure; whereas the closure of the aortic valve is defined by a curve where the three leaflets meet, the mitral valve's closure is a curved surface on which the valve centroid is not easily identified. In addition,

its oblique location with respect to the image axes results in a high probability that the observers mark the centroid in different slices.

The overlap-based accuracy is reported in Table 2.6 as F-scores in the atlas space. We have listed both the mean and standard deviations per subject per structure as well as an aggregate F-score per structure (column ‘overall’) and over all structures (last row of column ‘overall’). Relatively little change in accuracy between the first (reference) and fifth (final) iterations can be observed. A slight deterioration of the registration quality may be due to the inevitable reduction in image quality when the reference image is warped, which has affected all structures but the aorta, the structure with the most homogeneous blood pool. Seemingly ‘difficult’ structures to match properly have complex geometry (atria) or are thin (myocardium). A difficulty with the aorta was found to be the proximity of the descending part to the spinal column, and of the ascending part to the atria. ‘Bleeding’ into these structures was observed in some of the test subjects.

2.5.4 Bias Removal

The choice of a reference image introduces a bias. The iterative approach to synthesizing a mean image is designed to remove this bias. At full and perfect conversion, the mean vector field with which the mean image is updated should represent an identity transformation. We can therefore quantify the bias by measuring φ_i for each i .

The ideal mean vector field—the identity transformation—would contain exclusively zero vectors. Thus, its vectors’ x , y and z components would follow a zero mean Dirac delta function with the peak at 0. In Fig. 2.9’s left panel we show the deviation from 0 of the mean x , y and z component over seven iterations. It is clear that this deviation decreases strongly in the first iterations and then starts oscillating. In the right panel we plot the sample kurtosis of the distributions over the seven iterations, quantifying the ‘peakedness’ of the distributions. While the Dirac delta distribution’s kurtosis is undefined, one can easily see that kurtosis increases with better approximation of the delta function.

In Fig. 2.10 we show the distributions per component in full over the iterations; after only three iterations the changes are no longer perceivable, which is similar to the findings of Guimond *et al.* [66].

In Fig. 2.11 we show how the distribution of the vector magnitude of φ_i changes over the iterations. The pattern of convergence is clear, yet after five iterations the percentage of longer vectors starts to grow again. This has a detrimental effect on the synthesized mean image and confirms the qualitative observation of increased artifact visibility from Fig. 2.3. Therefore, we consider the synthesized mean image at $i = 5$ to be the final image.

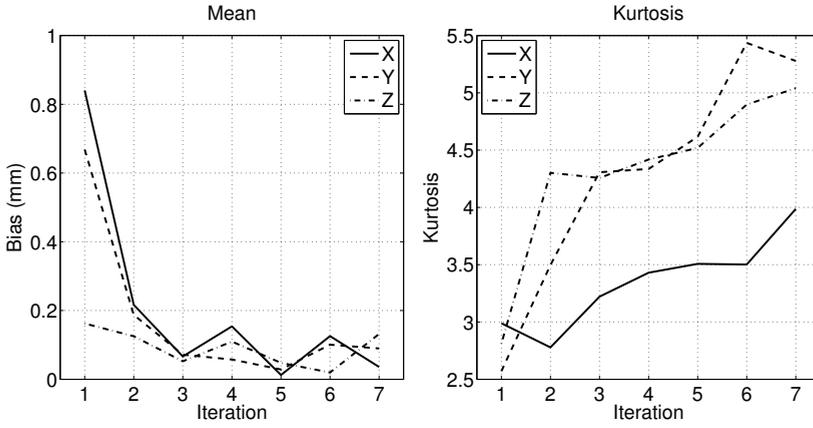


Figure 2.9: The absolute values of the bias (left) observed in the mean vector field decrease with iterations, down to less than half of the shortest side of a voxel in all directions after three iterations. The ‘peakedness’ of the distributions displayed in Fig. 2.10 is quantified through their kurtosis (right). It is clear that this measure continues to increase.

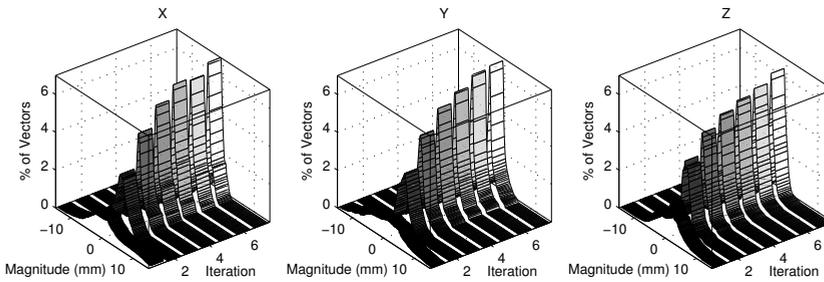


Figure 2.10: Histograms of vector components per iteration. Qualitatively, it shows that the first three iterations resolve the greatest amount of deformation. After this, the majority of change occurs in the y component (sagittal plane).

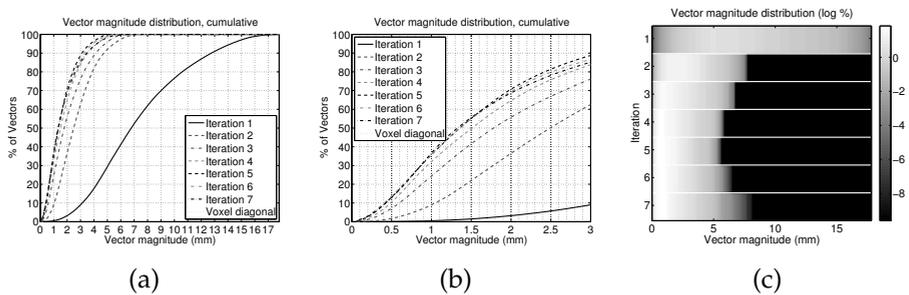


Figure 2.11: The distributions of vector magnitudes in the mean vector fields (displayed as cumulative histograms) show convergence of the registration. In (b) we zoom in on the first three millimeters in (a) to better visualize that after only three iterations the majority of vectors is shorter than the voxel diagonal, and after five iterations around 70 percent of the vectors has a length shorter than the voxel diagonal. In (c) it shows that after five iterations, the tail of the vector magnitude distribution starts to widen. The synthesized mean image at this iteration is considered the end result of the spatial normalization.

2.6 Discussion

We have presented a mesh-represented statistical atlas of the human heart, and a framework for its construction from a large set of 3D+time Multi-Slice Computed Tomography (MSCT) image volumes. We validated the atlas using measures of registration accuracy and evaluated the iterative bias removal.

We have made the atlas publicly available through our group’s website.

2.6.1 Framework Performance

We have demonstrated the construction of a spatio-temporal statistical model of the atlas shape over a population of 134 subjects at 15 cardiac phases. To our knowledge this represents the largest 3D+time population to date to be used for cardiac statistical model generation.

The number of pairwise registrations in the framework scales linearly with the number of subjects. This makes the framework applicable to larger data sets than the set used in this chapter, such as the Cardiac Atlas Project database. The framework derives robustness from automation of the reference selection and the imposition of diffeomorphic constraints on the registrations.

Human interaction is limited to two distinct instances during the application of the framework. The first, tuning of registration parameters, is based on a subset of the population of a user-defined size. In the second instance, the manual labeling of the mean image, the level of detail and the number of structures can be controlled by the user. Existing atlases can be employed to further reduce the human input.

Steps in spatial normalization that were previously typically done by human interaction, such as reference selection and assessment of convergence, were automated. An image-to-group similarity criterion was used to select a suitable reference, while convergence was based on the distribution of vector magnitudes in each iteration’s mean vector field.

We validated the atlas through measures of registration accuracy. Random sampling was employed to balance the need to cover a sufficiently large portion of the population, while maintaining the required human effort at a feasible level. In addition to random selection of subjects, within the subjects the selection of image slices was randomized as well. Domain knowledge was used to encourage the selection of slices presenting cardiac structures, while also enforcing an even distribution along the transversal axis of the body.

To complement the sample-based explicit registration evaluation, a PCA-based outlier detection system was applied to the first-phase shapes. While conceptually simple, it allowed the removal of a very small (3%) subset of the population for which registration had failed, contributing to the robustness of the overall framework.

By applying the spatial normalization only on the first frame of each subject’s volume sequence, we explicitly resolve geometric variations due to inter-subject variation separately from deformation due to cardiac function. This reduced computational requirements with respect to registration using an integrated spatio-temporal transformation model, while improving temporal consistency compared to normalizing the full set of 2070 image volumes.

2.6.2 Limitations

The atlas comprises 13 structures, of which 9 are outlined and 4 are traced; this represents only a small fraction of structures present in the heart but reflects a limitation resulting from the imaging data used. Populations imaged using complementary imaging techniques can be used to provide more structures, further subdivision or level of detail of the current structures, and increase the population size for the statistical modeling. Similarly, the aorta and pulmonary trunk are better modeled based on dedicated imaging studies.

Further limitations due to the imaging data are related to the price of the high spatial resolution of MSCT. Firstly, the temporal resolution is rather low at 15 frames per cycle. This is a well-known tradeoff that exists in all imaging modalities. Secondly, the use of CT data means the subjects were submitted to ionizing radiation, in addition to having contrast agent injected. As a result, all subjects are pathological, which influences both the population mean and the observed variation.

The coronary trees can only be traced over a short section; through the oblique orientation of the heart with respect to the image coordinate system, the coronary arteries quickly disappear 'between slices'. Additionally, the structures are too small and are subject to topological variations, such that accurate registration of these becomes impossible.

Another problem of visibility, though not by obliqueness but by appearance, is that of the thin walls of the atria, right ventricle and trunks of great vessels. The appearance of these structures coincides with that of the pericardium, and thus our atlas only estimates these structures by extrusion of their blood pool boundaries.

A limitation of the framework is that it is currently unable to model the topological variations that are known to occur in the complex of left atrium and pulmonary veins [132], and in the coronary trees [3]. It deserves further study to unify a discrete statistical model of topological variations with the continuous statistical model of shape. The separation of the pulmonary vein trunks from the left atrium is a step towards accommodating such an extension.

The approach of iteratively removing bias from a chosen reference image may require more constraints than only the condition that the registrations produce diffeomorphic deformations. This may be related to the phenomenon we observe in the later iterations in Fig. 2.3. There, even the diffeomorphic transformations start producing seemingly implausible deformations, potentially as a result of shifting from larger to smaller differences to be resolved. As Hamm *et al.* [69] suggested, it may be desirable to explicitly constrain the deformations to generate only images that lie on the manifold generated by the input images.

Furthermore, this manifold may allow discontinuities in the deformations due to different positions of the heart relative to neighboring structures. Consider here the lungs and pulmonary vasculature laterally, the diaphragm inferiorly, and the sternum and rib cage frontally. Such discontinuities cannot be captured using the B-spline deformation model. Providing a region of interest that comprises only the heart may alleviate such a problem, as the B-spline needs only model deformation within this region. This could work under the assumption that the reference image can serve as an initial atlas, or by using an

automated detection and/or rough segmentation method.

By using pairwise registrations to resolve the cardiac motion (the intra-subject registration), the time points at which labeled shapes are obtained are dictated by the imaging protocol. Currently, work on spatio-temporal registration [46, 122, 138] with a spatially and temporally continuous deformation model has not yet been demonstrated to handle high-resolution CT image data. However, the framework allows the pairwise approach to be replaced with a continuous approach when required and feasible.

The high degree of automation of the framework calls for methods, preferably also automated, to quantify accuracy. We used a rudimentary means to identify suspicious shapes by outlier detection. However, as it was rudimentary, it was limited to identifying only very obvious cases, and thus a more sophisticated approach suitable also for early detection is desired.

Finally, the framework presents the user with many decisions to be made; there are many steps involved, one can interchange methods, and each method has its parameters. This combines to a vast number of possible combinations that will not all produce the exact same atlas, and some combinations may even fail. However it is beyond the scope of this chapter to present a quantification of the effect of each decision exhaustively.

2.6.3 Future Perspectives

The various steps in the framework are independent of each other; the methods used in each step can be replaced by other methods according to requirements without affecting the workflow ‘downstream’. One should note stratification [18] as an exception to the latter; while it can be plugged into the framework, it introduces branches in the workflow as it will lead to multiple atlases.

The atlas was developed with the aim of using it in computational physiology studies. We demonstrated the incorporation of fiber orientations in the left ventricular myocardium, which is of paramount importance in the simulation of both electrical and mechanical cardiac function.

The combination of an atlas with a statistical model facilitates personalized computational physiology studies by connecting the image analysis and simulation elements described by Young and Frangi [222].

The flexibility of the statistical model allows fitting the shape to image data to personalize geometry [226]; although the model is built from CT data, other modalities can be used as well with the help of simulated image appearance models [197, 198]. Alternatively, the generative nature of the model can be exploited to generate virtual populations [76, 77, 153] to study the influence of certain parameters—including geometry itself—on simulation results.

Further personalization of the mesh in preparation for simulation can be based on invasive measurements, by taking into account statistical relationships between geometry and fiber orientation [111], or draw upon the literature to introduce normal, abnormal, or expected parameter values.

CHAPTER

THREE

BILINEAR MODELS FOR SPATIOTEMPORAL POINT
DISTRIBUTION ANALYSIS

Adapted from [79]

C. Hoogendoorn, F.M. Sukno, S. Ordás, and A.F. Frangi

International Journal of Computer Vision, vol. 85, no. 3, pp. 237–252, Dec. 2009.

Abstract

In this chapter we describe the usage of bilinear statistical models as a means of factoring the shape variability into two components attributed to inter-subject variation and to the intrinsic dynamics of the human heart. We show that it is feasible to reconstruct the shape of the heart at discrete points in the cardiac cycle. Provided we are given a small number of shape instances representing the same heart at different points in the same cycle, we can use the bilinear model to establish this.

Using a temporal and a spatial alignment step in the preprocessing of the shapes, around half of the reconstruction errors were on the order of the axial image resolution of 2 mm, and over 90% was within 3.5 mm. From this, we conclude that the dynamics were indeed separated from the inter-subject variability in our dataset.

3.1 Introduction

3.1.1 Statistical Shape Models and Time

The use of statistical models of shape has established itself as a popular approach to image analysis problems, in the domain of both natural [150] and medical image analysis [53, 136]. Along the way, much research has been devoted to the development of various types of shape models, as well as to solving problems arising from the construction of such models. Many of those models, however, are essentially spatial models, and extension to the spatiotemporal domain is not as trivial or trivially justifiable as the extension from d to $d + 1$ spatial dimensions.

The analysis of Point Distribution Models (PDMs) in shape space received a significant amount of attention from the mid-1980's to the early 1990's [38, 97, 109, 129]. The most renowned result from this work applied in computer vision is the emergence of the Principal Component Analysis (PCA)-based statistical shape model from Cootes *et al.* [42], who applied PCA to the covariance matrix of their data set in order to extract a set of orthogonal variations of the sampled points. These became well-known as the modes of variation of the shape class.

Efforts to extend linear shape models to the spatiotemporal domain have been made before, for example by Hamarneh and Gustavsson [68], in whose work each sample in the dataset consists of an entire sequence of observations of the same object sampled throughout the temporal exposure window. Mitchell *et al.* [143] and Bosch *et al.* [20] employ the same strategy, extending Active Appearance Models (AAMs) [39] to Active Appearance Motion Models (AAMMs) in order to segment the endocardium in echocardiograms and Magnetic Resonance (MRI) image sequences, respectively. Perperidis [154] constructs two linear models using PCA: one which models the variation across the mean shapes of the subjects, and another one which models the variation within the cardiac cycle. Another approach to add the element of time to 3-dimensional segmentation using statistical shape models was presented by Montagnat and Delingette [144]. After building a PCA-based model treating all subjects and all phases as separate samples, they employ a scheme which segments the cycle as one single object, rather than employing a sequential scheme to segment each frame separately. During segmentation, temporal constraints are introduced in the optimization to limit the differences between segmentations of subsequent frames to reasonable values. Lynch *et al.* [124] model the temporal processes in their level-set approach for cardiac segmentation as a distance function between a set of control points and the level sets.

Statistical models of cardiac left ventricle (LV) deformation only were constructed by Chandrashekhara *et al.* [30], using data from one single subject. A comparable approach to modeling respiratory motion of the liver was taken by Blackall *et al.* [17].

To the best of our knowledge, our previous work [78] was the first to model individual and temporal variations of cardiac shape as two different sources of variability within the same set of data. Before, inter-subject variation and dynamics were not decoupled at all [20, 68, 143]. Later, the application of the model was equipped with constraints to limit the first-order derivative of shape

points over time [144] or by constraining it using a distance prior [124]. By creating a spatiotemporal model of cardiac dynamics that decouples individual and temporal variations, we can extrapolate cardiac phases from the statistical model even when they are not available from the individual measurements. The models most similar to our approach are those from Perperidis [154]. However, our method does not result in a massively higher-dimensional parameterization for intra-subject variability, resulting in a more compact model.

3.1.2 Bilinear Statistical Models

In biometrics, the separation of two (independent) processes that contribute to the overall pattern variability is a well-known problem, which has led to the introduction of bilinear models by Tenenbaum and Freeman in [195]. The reader is referred to [196] for a more detailed description. Dubbing the two sources of variability *style* and *content*, these names can be assigned freely depending on which is most natural given a specific problem. Aside from the examples used in [196], the literature contains examples of the separation of

- face identity and facial expression [1]
- location and content for sparse coding of natural images [65]
- emotion and speech content [35]
- gait (walking characteristics) and viewing conditions [82, 110]
- pairs from the set {identity, action, viewpoint} [44]
- prostate shape between individuals and over time [87, 88]
- face identity and viewpoint [62]
- face identity and illumination [179].

While in Tenenbaum and Freeman [196] the bilinear models and their construction are formulated quite specifically for bilinear decomposition, which we expand on in Section 3.2, they follow the same principles of the multilinear decompositions of higher order tensors presented by De Lathauwer *et al.* [47], which was followed by Vasilescu and Terzopoulos [206] for their trilinear decomposition of expression, identity and illumination.

3.1.3 Clinical Context

The use of statistical models in cardiac image analysis initially focused primarily on segmentation [20, 124, 143, 144, 187]. However, the analysis of cardiac function – dynamics and deformation – has since emerged as a relatively new field of application to which intensive research has been dedicated. In the context of Congestive Heart Failure (CHF) assessment, LV function in particular has been the focus of automated localization and quantification algorithms, employing various types of models. Automated methods for objective modeling and analysis of cardiac morphology and function are therefore desirable and relevant. For an overview of cardiac modeling, we refer to Frangi *et al.* [60].

Most recently, deformation recovery and modeling was done by Bistoquet *et al.* [16], who make use of the near non-compressibility of myocardial tissue to constrain the segmentation of the LV and simultaneously extract the underlying within-tissue deformation. Liu and Shi [117] analyzed cardiac motion under constraints obtained from biomechanical priors.

Lekadir *et al.* [112] model LV deformation as interlandmark motion in a local circular coordinate system in which anomalies present themselves as outliers. Leung and Bosch [114] create a sparse shape variation decomposition using PCA and subsequent orthomax rotations to characterize local abnormal deformations, while Syeda-Mahmood *et al.* [188] use registration and the associated deformation patterns to characterize deformation abnormalities in two-dimensional (2D) ultrasound sequences.

While they may be robust to intersubject variability, none of the above approaches really takes this source of variation into account. By factoring from a dataset the variation attributed to inter-subject differences from that of cardiac dynamics, be it shapes, deformation fields or the images themselves, the contraction pattern of an entire population can be modeled, and then used in various clinical applications. In this chapter we employ bilinear models to establish this factorization on cardiac shapes, and illustrate their power to parameterize the dynamics (or rather a set of discrete phases within the cardiac cycle) and the subject by extrapolating from a small subset of phases to the remaining phases in the cycle.

The remainder of this chapter is organized as follows: we introduce the concept of and techniques behind bilinear models in Section 3.2. Our extrapolation experiments are described in Section 3.3. The results are presented in Section 3.4, followed by a discussion and future research directions in Section 3.5. We conclude this chapter with Section 3.6.

3.2 Bilinear Statistical Models

A bilinear model is a two-factor model which is linear in either factor when the other one is kept constant:

$$y = \vec{a}^T \vec{W} \vec{b} \quad (3.1)$$

where y is a scalar observation, \vec{a} and \vec{b} are parameterization vectors defined by the factors, and \vec{W} is a constant matrix governing the interaction between the factors.

Extending this to the case of multivalued observations, each element y_k^{sc} of a K -sized observation \vec{y}^{sc} in style s and content c can be described by a bilinear model as

$$y_k^{sc} = \sum_{i=1}^I \sum_{j=1}^J w_{ijk} a_i^s b_j^c. \quad (3.2)$$

I and J are the sizes of the parameterization vectors \vec{a} and \vec{b} , respectively. \mathcal{W} is now a 3-dimensional matrix, which forms a mapping from the style and content spaces into observation space and as such is of size $I \times J \times K$. Each w_{ij} is a K -sized base observation, much akin to the eigenface [200] and eigenshape [42]. In

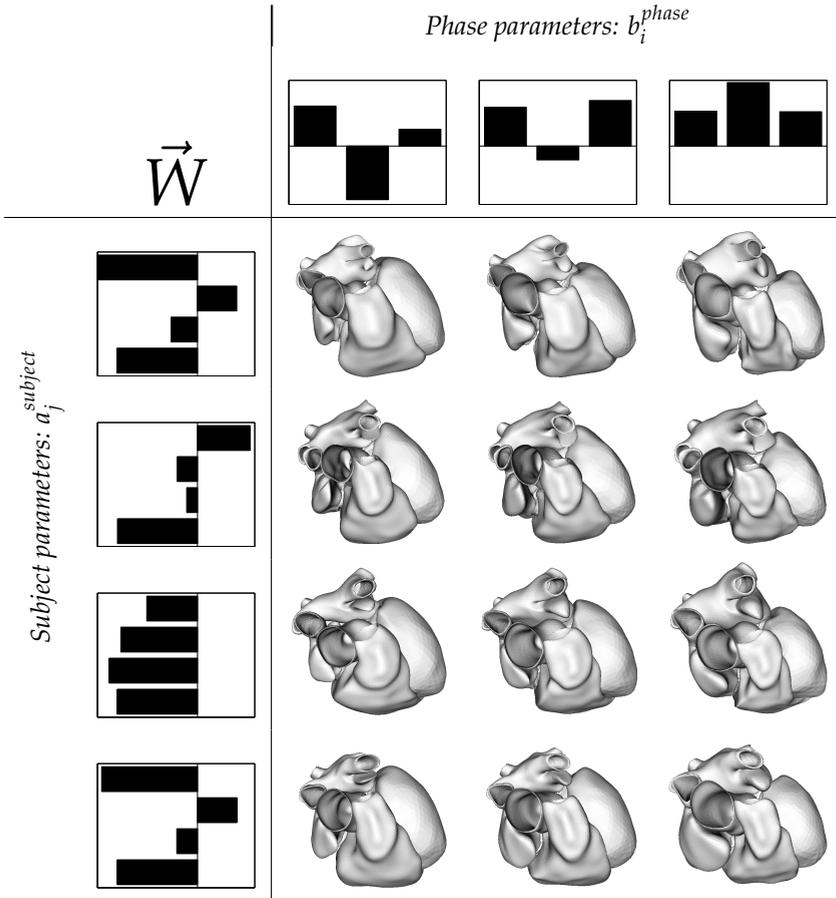


Figure 3.1: The structure of a symmetric bilinear model. Along the horizontal axis, three phases from the cardiac cycle are shown. Along the vertical axis we have four different subjects. The set of basis observations \vec{W} can not be visualized in a meaningful manner.

the case of speech recognition, one could think of them as voice harmonics that need to be combined to form a certain phoneme in a certain accent or intonation. The \vec{a} and \vec{b} vectors provide the information on how to combine those base observations. For the time being, we will adhere to the original nomenclature and call these the content and style vectors, respectively (hence the c and s superscripts). Figure 3.1 illustrates how the model consists of the sets of parameters \vec{a} and \vec{b} , providing a compact representation of the style and the content, and the base observations \vec{W} , derived from the examples in the bottom-right quadrant.

The choice for bilinear models stems from the idea that variations in a set of observations are the consequence of the variation of two independent factors. The examples used to illustrate the usefulness of bilinear models in Tenenbaum and Freeman [196] call upon analyzing the way we manage to recognize known characters, people or phonemes in a font or under viewing circumstances not

observed before, or uttered in an accent not heard before. Somehow, we know the invariants of that character, person, or phoneme, and in an observation we can recognize those irrespective of the variations introduced by the circumstances.

3.2.1 Asymmetric Models

A combining matrix \vec{A} can be the result of contracting $\vec{a}\mathcal{W}$ into a single matrix, leading to an asymmetric model:

$$y_k^{sc} = \sum_{j=1}^J a_{jk}^s b_j^c \quad (3.3)$$

where

$$a_{jk}^s = \sum_{i=1}^I w_{ijk}^s a_i^s. \quad (3.4)$$

This is useful when the mapping from content to observation is dependent on style. The J observations of size K each that make up \vec{A} are then style-specific base observations that can be mixed using the parameterization in \vec{b} . In essence, the model has then become a unilinear model. It is possible to invert the roles of style and content, and thus build a different asymmetric model from the same training data.

3.2.2 Symmetric Models

The symmetric model is the original bilinear model as already presented in (3.2):

$$y_k^{sc} = \sum_{i=1}^I \sum_{j=1}^J w_{ijk} a_i^s b_j^c,$$

where the mapping from the style and content spaces to observation space, \mathcal{W} , is dependent on neither style nor content. The elements of \mathcal{W} are base observations that look like eigenfaces [200] but do not represent an orthogonal basis like eigenelements. The base observations can then be mixed using the \vec{a} and \vec{b} parameterizations to form any element from the training set, and other parameterizations can be used to construct observations of new style and/or new content.

3.2.3 Factoring Inter-Subject Variability and Cardiac Dynamics

Representation of Training Data

The construction of the bilinear models as presented in Tenenbaum and Freeman [196] assumes a vector representation of the observations. This is easily achieved by employing a surface representation as is commonly used for

the construction of PDMs like those in Cootes *et al.* [42]. In d -dimensional Euclidean space, each training shape is landmarked with n_L points – pseudolandmarks [52]¹ – of anatomical correspondence throughout the training set. The set of point coordinates is then concatenated to form an $(n_L d)$ -dimensional shape vector, or a single point in an $(n_L d)$ -dimensional shape space.

After landmarking and vectorization of our observations, we have CS $(n_L d)$ -dimensional shape vectors: C frames per time sequence, S subjects (with one sequence each), n_L landmarks in d dimensions. As is the case in [68], we construct our observation matrix by ‘stacking’ all vectors for one subject onto each other, such that we obtain a $Cn_L d \times S$ shape matrix \vec{Y} :

$$\vec{Y} = \begin{bmatrix} \vec{y}^{11} & \dots & \vec{y}^{1C} \\ \vdots & \ddots & \\ \vec{y}^{S1} & & \vec{y}^{SC} \end{bmatrix}. \quad (3.5)$$

This is the starting point for the construction of both the asymmetric and the symmetric bilinear shape models. Rather than using a 3-dimensional matrix \mathcal{W} , we use \vec{W} which has a format similar to \vec{Y} , with the number of elements limited by the sizes of the parameterization vectors:

$$\vec{W} = \begin{bmatrix} \vec{w}^{11} & \dots & \vec{w}^{1J} \\ \vdots & \ddots & \\ \vec{w}^{I1} & & \vec{w}^{IJ} \end{bmatrix}. \quad (3.6)$$

The modeling consists of minimizing the squared reconstruction error between the original observations and the approximation the model will provide. Denoting the approximation of \vec{y}^{sc} in shorthand notation as

$$\hat{\vec{y}}^{sc} = \begin{cases} \vec{A}^s \vec{b}^c & \text{if the model is asymmetric} \\ \vec{a}^{sT} \vec{W} \vec{b}^c & \text{if the model is symmetric} \end{cases}, \quad (3.7)$$

we minimize

$$E = \sum_{s=1}^S \sum_{c=1}^C \|\vec{y}^{sc} - \hat{\vec{y}}^{sc}\|^2. \quad (3.8)$$

Asymmetric Training

The training of an asymmetric model has a closed-form solution if the number of observations is (nearly) equally distributed over the style and content classes [196]. As the matrix with training data \vec{Y} is the result of the product $\vec{A}\vec{B}$, it suffices to compute the Singular Value Decomposition (SVD) $\vec{Y} = U\vec{S}V^T$. Then, the matrix \vec{B} , containing the phase parameters, can be defined as the first J rows of \vec{V}^T , while \vec{A} will be defined as the first J columns of $U\vec{S}$.

¹In medical image analysis, ‘landmarks’ can refer to either anatomical landmarks or pseudolandmarks.

Symmetric Training

The training of the symmetric model requires the notion of the vector transpose [130]. Unlike the original application of this term, namely the conversion of a column vector into a row vector and vice versa, the vector transpose we use here is a matrix operation.

Given an $IK \times J$ matrix, where each column was constructed by stacking I K -dimensional column vectors onto each other, the vector transpose of this matrix is a $JK \times I$ matrix, with the positions of the K -dimensional column vectors transposed rather than the individual elements. In the case where $K = 1$, the vector transpose is the normal transpose of the matrix. Otherwise, the vector transpose looks like

$$\vec{X}^{\text{VT}} = \left[\begin{array}{c|c} x_{11} & x_{12} \\ x_{21} & x_{22} \\ x_{31} & x_{32} \\ \hline x_{41} & x_{42} \\ x_{51} & x_{52} \\ x_{61} & x_{62} \\ \hline x_{71} & x_{72} \\ x_{81} & x_{82} \\ x_{91} & x_{92} \end{array} \right]^{\text{VT}} = \left[\begin{array}{c|c|c} x_{11} & x_{41} & x_{71} \\ x_{21} & x_{51} & x_{81} \\ x_{31} & x_{61} & x_{91} \\ \hline x_{12} & x_{42} & x_{72} \\ x_{22} & x_{52} & x_{82} \\ x_{32} & x_{62} & x_{92} \end{array} \right] \quad (3.9)$$

for $K = 3$. It is easy to see from (3.9) that the vector transpose operation is invertible: $(\vec{X}^{\text{VT}})^{\text{VT}} \equiv \vec{X}$.

For the symmetric model, an iterative method is required to minimize E in (3.8). To this end, first \vec{A} and \vec{B} are computed. Upon convergence of these computations, \vec{W} is computed using the results.

Analogously to the $K = 1$ case using matrix transpositions, the simplified model equation $\vec{Y} = A\vec{W}B$ can be rewritten as $\vec{Y} = (W^{\text{VT}}\vec{A})^{\text{VT}}B$ as well as $\vec{Y}^{\text{VT}} = (W\vec{B})^{\text{VT}}A$. This leads to two equations that are familiar from the training of the asymmetric model:

$$(\vec{Y}\vec{B}^{\text{T}})^{\text{VT}} = \vec{W}^{\text{VT}}\vec{A} = U\vec{S}V^{\text{T}} \quad (3.10)$$

and

$$(\vec{Y}^{\text{VT}}\vec{A}^{\text{T}})^{\text{VT}} = \vec{W}B = U\vec{S}V^{\text{T}}. \quad (3.11)$$

By iterating over these equations, starting from an initial estimate of \vec{B} using the SVD of \vec{Y} , convergence towards the real \vec{A} and \vec{B} is guaranteed [126]. As was the case in the previous section, dimensionality reduction can be achieved by picking a fixed number of rows of \vec{A} and \vec{B} . Unlike linear models based on PCA, this cannot be left until the end, as a change in dimensionality of one set of parameters will affect the other set of parameters. Therefore this truncation is

done in each iteration. Upon convergence of the computation of \vec{A} and \vec{B} , what is left is the computation of \vec{W} :

$$\vec{W} = \left(\left(\vec{Y} \vec{B}^T \right)^{VT} \vec{A}^T \right)^{VT}, \quad (3.12)$$

which satisfies $\vec{Y} = (W^{VT} \vec{A})^{VT} B$.

In Section 3.1 we already mentioned the relation between the bilinear models as they are formulated in Tenenbaum and Freeman [196] and here, and tensor space decomposition [47]. The two vector transpositions we use correspond to two of the three tensor unfoldings [47] of the third-order tensor we could construct along the axes of subject, phase and pseudolandmark coordinate.

3.3 Experiments

The assumption that shape variation introduced by inter-subject variation is independent from the dynamics is a simplification of reality. However, hearts with equal geometry at rest do not necessarily contract in the same way, while hearts with different geometry may. Factors such as myofiber orientation, local contractility defects like infarction, dilation or hypertrophy, and wall stress influenced by blood pressure and loading all influence the contraction pattern, and none of these are really independent of one another. Therefore we performed experiments to verify the suitability of bilinear modeling to capture the dynamics of the shape of the beating heart. To this end we constructed models factoring the inter-subject variations (as the style) and the dynamics (as the content), as illustrated in Fig. 3.1. Then, the models were used for extrapolating the learnt dynamics to subjects not present in the training set. In addition, we compared predictions based on a single phase to the difference between this shape and phases following it.

3.3.1 Data

Population

Our data consists of 80 full hearts acquired consecutively, with closed surface representations for each of five subparts, listed in Table 3.1 and visualized in Fig. 3.3. The population is distributed as 60% healthy and asymptomatic subjects, 20% subjects with Coronary Artery Disease (CAD) without a history of myocardial infarction (MI), and 20% subjects with CAD and a history of MI. Subject age was 58 ± 8 years, and men made up 56% of the population.

For the experiments we built models of the entire heart as well as models of only the right and left ventricle (biventricular model) and the LV only.

Imaging Parameters

The imaging was performed using a Toshiba Aquilion 64 Multi-Slice Computed Tomography (MSCT) system with a 64-row detector (Toshiba Medical Systems,

TABLE 3.1: THE PARTS OF THE HEART SHAPE THAT HAVE A CLOSED SURFACE REPRESENTATION, WITH THE ASSOCIATED NUMBER OF LANDMARKS n_L . THE PARTS ARE VISUALIZED IN FIG. 3.3.

Symbol	Description	n_L
LV	Left ventricular endo- and epicardium.	2677
RV	Right ventricle without septum, with trunk of pulmonary artery.	7902
LA	Left atrium with trunks of pulmonary veins.	6789
RA	Right atrium with trunks of venae cavae.	7243
Aorta	Trunk of the aorta.	3000

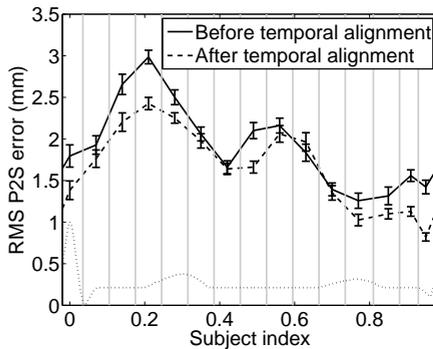


Figure 3.2: RMS P2S differences between subsequent phases of the left ventricles, averaged over all 80 subjects, plus and minus one standard error. The solid line shows the differences before the temporal alignment step, while the dashed line shows the differences after this step (see Section 3.3.1 for details on this process). At time 0, the difference shown is that between phases 15 and 1. The dotted curve at the bottom represents a schematic ECG, the R-peak at the left being the reference for the ECG gating (all acquisitions are at R plus a fraction of the R-R interval). The light grey vertical lines indicate the interval over which each of the phases was sampled.

Tochigi, Japan). Between 80 and 100 ml of contrast material (Xenetic 350) was administered at an injection rate of 5 ml/s. The rotation time of the scanner was, depending on the subject's heart rate, between 400 and 500 ms, and image reconstruction was performed on a Vitrea post-processing workstation (Vital Images, Minnetonka, MN, USA). The resulting dataset consisted of 15 image volumes (temporal phases) obtained using retrospective electrocardiogram (ECG) gating [148] with voxel dimensions of $0.4 \times 0.4 \times 2.0 \text{ mm}^3$ per subject. Figure 3.2 shows the gating sequence over the cardiac cycle, albeit that the ECG used is schematic. Temporal relationships are only preserved with respect to the R-R interval². Overlaid on this sequence are the root mean square (RMS) point-to-surface (P2S) differences between subsequent full heart shapes, using a landmarking scheme as explained hereafter. The value at time 0 reflects the RMS P2S difference between phases 15 and 1.

²The R-R interval is the mean time between the R-peaks – the big peaks – in the ECG, indicating the firing of an electrical impulse to the heart muscle, which then contracts.

Landmarking

The data set was then used for the landmarking as presented in Ordás *et al.* [149], resulting in $n_L = 27611$ points per volume. The point set sampled on the average shape was triangulated using Amira V3.0 (Visage Imaging, Carlsbad, CA, USA) to facilitate visualization of this shape, all shapes in the training set, and all shapes generated using the bilinear model. Table 3.1 lists the number of points sampled uniformly on each of the closed surfaces.

Spatial and Temporal Alignment

The subjects were imaged without sedation, which lead to heart rates of 62.2 ± 11.9 beats per minute. As the time it takes the heart to contract is largely independent of the heart rate, this causes inter-subject phase shifts due to the time points of the phases not being synchronized. Temporal alignment would therefore be desired. To this end, we explored using the method by Perperidis *et al.* [156] for temporal alignment. For each phase, the normalized cross-correlation coefficient (CC) of image intensities between the phase at hand and the first phase in the cycle is computed. As the heart contracts, this scalar reduces as the images become less similar to the first frame. The phase of maximum contraction (end systole (ES)) is identified as the phase with the minimum CC. Subsequently, the phase of maximum deceleration of relaxation (end diastole (ED)) is identified as the phase with the minimum second derivative of the CC *after* the end-systolic phase. Together with the first and last phase, these form the phases which are to be aligned. All phases in between are to be interpolated.

We adapted this approach slightly. We did not use a reference cycle to identify at which phase end-systole and end-diastole should be placed. Instead, we used the mode of the set of identified phases for each, resulting in phase 5 for end-systole and phase 10 for end-diastole. It should be noted, however, that of the 80 subjects, the subjects for which the end-systolic phase was identified as the fifth only slightly outnumbered those for which the end-systolic phase was identified as the sixth. Figure 3.4 shows CC values as the temporal alignment procedure progresses.

It will become apparent, however, that the temporal alignment has its limitations. The spread of the heart rates, together with the low temporal resolution (15 phases, compared to 30 in Perperidis *et al.* [156]), could result in a loss of in-

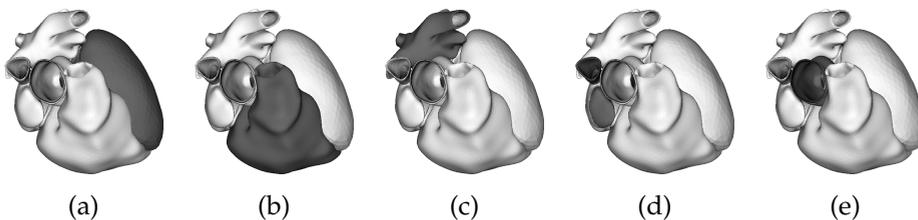


Figure 3.3: Surface rendering of one of the training hearts. In dark (a) the left ventricle, (b) right ventricle, (c) left atrium, (d) right atrium, and (e) aorta.

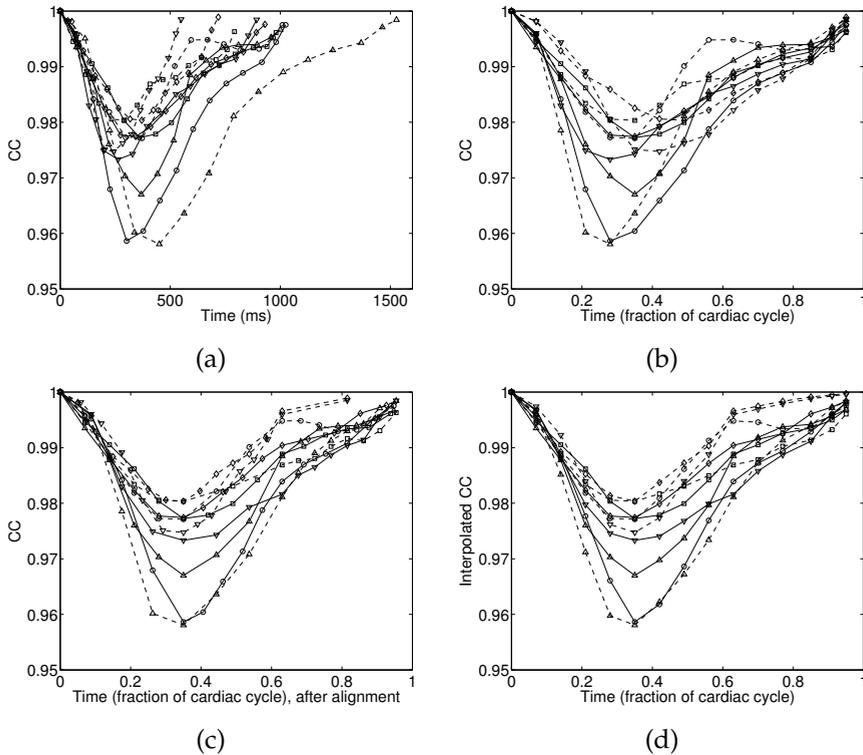


Figure 3.4: The effect of the alignment on the CC of the CT images. For clarity, only values for 10 subjects are shown, including the subjects with shortest and longest R-R interval. The volume of the left ventricle follows a very similar curve [156]. (a) CC against time in milliseconds before alignment. Note the spread along the temporal axis. (b) CC against time in percentage of the cardiac cycle before alignment. This corresponds to aligning the first phase of the cycle and the first phase of the next cycle. (c) CC against time in percentage of the aligned canonic cycle. Positions in time of the original acquisitions are interpolated. (d) Interpolated resampled CC after full alignment. In practice, the shapes were aligned by linear interpolation of corresponding landmarks.

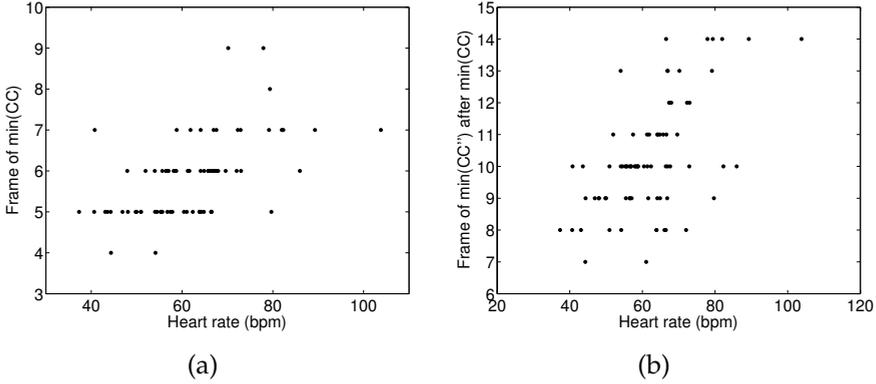


Figure 3.5: Scatter plots of the relationship between heart rate and the phases in which (a) maximum contraction (minimum CC) and (b) maximum deceleration (minimum CC'') are observed. Note that a higher heart rate means a shorter R-R interval and thus the largely rate-independent duration of the beat itself taking up a larger part of this interval.

formation in some of the sequences due to linear interpolation, which may not be compensated for by the resulting correspondence. In some extreme cases, five phases used information from three original phases, while elsewhere in the same cycle, five phases were the result of interpolating between nine. Figure 3.5 shows how the phases identified as end-systolic and end-diastolic depended to some extent on the heart rate, but that even for equal heart rates a significant spread in identified phases may be observed.

Since the identification of the end-systolic and end-diastolic phases was based on the intensities of the original images, we extracted the left ventricle and the biventricular meshes at this point, as the result of temporally aligning those shapes separately would not be influenced.

After the temporal alignment, we also aligned the shapes spatially. With n_L landmarks describing each of our shapes, we denote the point set in \mathbb{R}^3 equivalent to shape vector \vec{x} as the $n_L \times 3$ matrix \vec{X} . The mean phase \vec{X}_{input}^j of each subject j was translated to position its center of gravity at the origin. Then, it was rescaled to have unit norm. Subsequently, Procrustes alignment [63] was performed on these shapes, which, for each subject j , yielded a translation \vec{t}_j , rotation \vec{R}_j and isotropic scaling s_j . Then, for each phase i ,

$$\vec{X}_{\text{output}}^{ji} = s_j \times \vec{X}_{\text{input}}^{ji} \vec{R}_j + \vec{t}_j. \quad (3.13)$$

This way, the spatial relationships within the cycles were preserved. Removal of those relationships would reduce the overall variation observed in the data set, which in turn would lead to smaller reconstruction errors. By retaining the relationships instead, we possibly sacrificed some extrapolation accuracy in exchange for clinical meaningfulness of the results.

3.3.2 Experiments

We performed leave- n -out experiments with increasing training set sizes ($n = \{40, 20, 10, 5\}$). Thus, we divided our set of subjects into $\frac{80}{n}$ disjoint subsets, where each subset plays the role of test set once and forms part of the training set for the remainder of cases, as illustrated on the left side of Fig. 3.6. For each of the subjects from the test subset, five phases were used to derive the subject parameterization using the constructed model, shown at the top right in Fig. 3.6. The remaining ten phases were then approximated (bottom right in Fig. 3.6) and the reconstruction errors recorded. Errors were grouped by phase and model dimensionality, and by phase set. The four sets of five phases covered systole, diastole, rest, and a combination of these (2 each from systole and diastole, 1 from rest), respectively.

3.3.3 Extrapolation

What we wish to do amounts to extrapolation, which is only one of several applications of bilinear models, as is shown in Tenenbaum and Freeman [196]. For a set of data \vec{Y}_{new} defined over an incomplete set of content classes associated with $\vec{B}_{\text{inc}} \subset \vec{B}$, in a new style with unknown parameterization \vec{a}_{new} , the missing elements need to be reconstructed using the characteristics of those contents learned from the training data (other styles) and the style characteristics that are to be derived from the newly presented observations.

The models we built were symmetric models so that we could reduce the dimensionality of both sets of parameterizations to eliminate noise and improve robustness. Knowing which are the phases we have for deriving the new subjects' parameters, we used the phase parameters from the model associated with these known phases:

$$\vec{C}_1^{\text{VT}} = W\vec{V}^T B_{\text{inc}}^T. \quad (3.14)$$

What remains is a linear system of equations

$$\vec{Y}_{\text{new}} = \vec{C}_1 \vec{a}_{\text{new}} \quad (3.15)$$

which, when solved, gives us a parameterization \vec{a}_{new} of the new subject. This system has many more equations than unknowns (the ratio running into the tens of thousands). Therefore, a least squares approximation is found.

Since computing both the subject parameters \vec{a}_i and phase parameters \vec{b}_i in the models involves SVDs, the last column of both \vec{A} and \vec{B} should be a zero vector. In practice we have found that this is not always the case, possibly due to limitations in numerical precision. Additionally, the last few columns typically contain but noise as observed in the data. We used the dimensionality reduction as described in Section 3.2.3 to remove these columns. In the first iteration, we use the associated singular values to determine the number of dimensions required to capture at most 95% of the variance observed, both between subjects and between phases. The resulting dimensionalities are then fixed throughout the remainder of the training procedure.

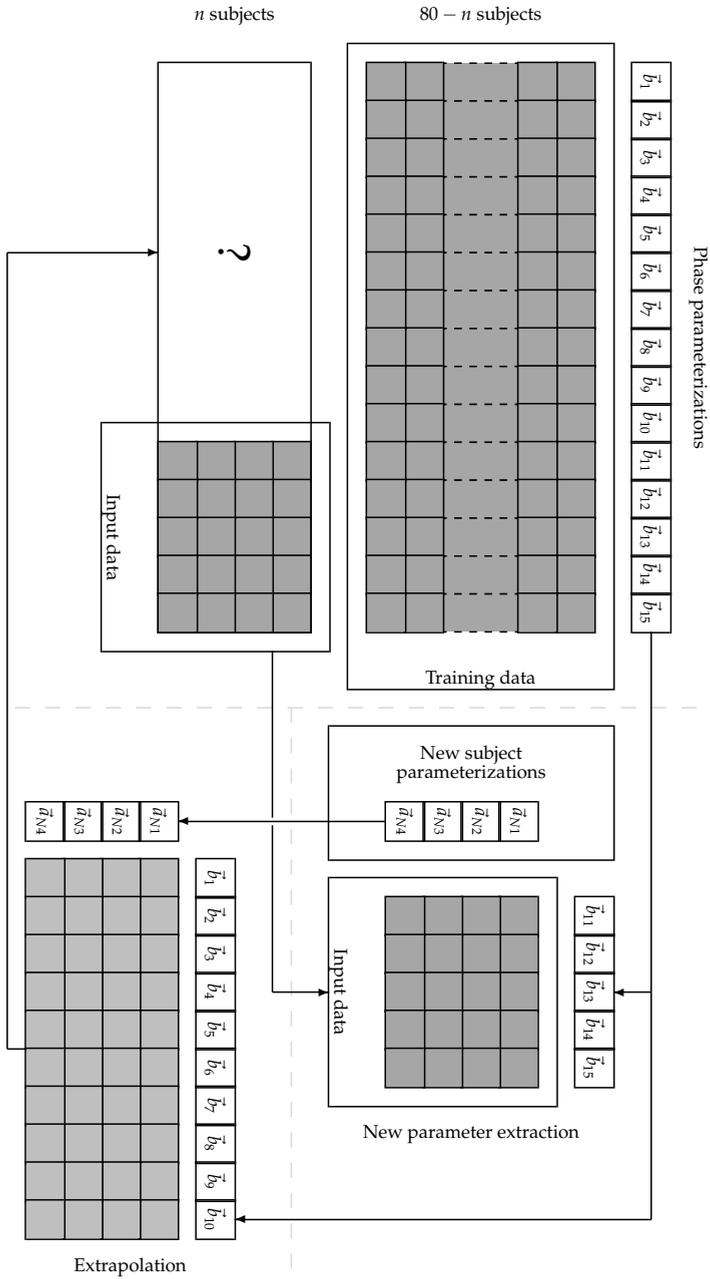


Figure 3.6: Overview of the leave- n -out experiments. (See next page)

Figure 3.6: (See previous page) Overview of the leave- n -out experiments. Dark grey blocks represent known shape data. White blocks with a $\vec{b}_\#$ and $\vec{a}_{N\#}$ denote phase parameterizations from the model and new subject parameterizations, respectively. The subscript N is used to indicate that this is a newly derived parameterization. Arrows with open heads indicate that the elements on either side are the same, yet transported for clarity of the figure. Each of the $\frac{80}{n}$ disjoint sets of n subjects (left side) is used as a test set once, with the other sets combining to form the training data. For each of the subjects in the test set, a set of shapes, corresponding to an equally sized known set of phases (in the figure this is the set $\{11, \dots, 15\}$), is used to derive the subject parameterizations $\vec{a}_{N\#}$ (top right), using \vec{W} (not in the figure) and the phase parameterizations from the model. The resulting subject parameterizations are then combined with \vec{W} and the remaining phase parameterizations to extrapolate the shapes to the other phases for the test subjects (bottom right). The resulting shapes are represented in the figure by light grey blocks, and correspond to the block with the question mark (left side).

3.3.4 Reconstruction

Using the resulting \vec{a}_{new} from (3.15), we approximate the remaining set of phases \vec{Y}_{rem} by

$$\vec{C}_2^{\text{VT}} = \vec{W}^{\text{VT}} \vec{B}_{\text{rem}}^{\text{T}}, \quad (3.16)$$

$$\vec{Y}_{\text{rem}} = \vec{C}_2 \vec{a}_{\text{new}}, \quad (3.17)$$

with

$$\left(\vec{B}_{\text{rem}} \subset \mathbf{B} \right) \wedge \left(\vec{B}_{\text{rem}} \cap \mathbf{B}_{\text{inc}} = \emptyset \right). \quad (3.18)$$

The reconstruction error is then recorded as the RMS P2S error between each predicted shape $\vec{y}_{\text{rem},i}$ in \vec{Y}_{rem} and its corresponding ground truth shape:

$$\text{RMS}(\vec{y}) = \sqrt{\frac{\sum_i d_i^2}{n_L}}, \quad (3.19)$$

where n_L again is the number of landmarks used to describe each of the shapes, and d_i is the distance from the i -th landmark of \vec{y} to the surface defined by the ground truth shape \vec{y} of the corresponding phase and subject.

3.3.5 Comparison to Direct Approximation

In image sequence segmentation, it is not uncommon that the segmentation result of a phase is used as the initialization for the next phase. We believe that the bilinear models can provide an alternative in offering an initialization for a larger part of the sequence, or even the entire sequence, based on a (preliminary) result on the first phase. Subsequently the subject parameters would be optimized to complete the segmentation.

In order to compare the initialization for multiple phases based on one phase, we used an experiment similar to the one described in the previous sections. However, this time we extrapolated five phases from one, for each of the three cardiac periods: systole, diastole, and rest. We compare the extrapolation errors – initialization errors in the segmentation setting – to the errors we would obtain should we use the one phase directly for the initialization of these phases.

3.4 Results

3.4.1 Extrapolation

The results are presented for four sets of phases $\{1 \dots 5\}$, $\{6 \dots 10\}$, $\{11 \dots 15\}$ and $\{4, 5, 9, 10, 15\}$ used for the derivation of a new subject's parameters and thus for the four groups of 10 reconstructed remaining phases. The results were produced using the temporally and spatially aligned data, unless stated otherwise.

The results of the LV models, the biventricular models and the full heart model results were highly correlated. Therefore we refrain from reporting the results on the biventricular and full heart models in the same detail as those from the LV models. Overall results are summarized in Table 3.2.

TABLE 3.2: THE MAIN STATISTICS FOR THE RECONSTRUCTION ERRORS USING THE LARGEST TRAINING SET. STATISTICS WERE COMPUTED USING ALL 3200 RECONSTRUCTIONS FOR EACH MODEL TYPE.

	<i>Structure</i>	<i>Median</i>	<i>Min</i>	<i>Max</i>
With temporal alignment	LV	1.869 mm	0.978 mm	7.224 mm
	LV+RV	2.212 mm	1.084 mm	7.940 mm
	Full heart	2.344 mm	1.380 mm	6.642 mm
Without temporal alignment	LV	1.999 mm	0.820 mm	8.557 mm
	LV+RV	2.359 mm	1.102 mm	8.923 mm
	Full heart	2.436 mm	1.401 mm	7.473 mm

In Figure 3.7, we present the mean RMS error of the worst and best reconstructed phases in the LV models, as well as the mean over all phases, against the size of the training sets (and consequently against the number of folds). It can be noted that larger training set sizes generally produced more accurate reconstructions. Additionally, the tendency of the curves suggests that some performance improvement could be achieved by further increasing the training set size. The errors are comparable to the differences between subsequent phases, reported in Figure 3.2. Figure 3.8 shows the mean reconstruction errors and the standard errors for each set of phases used for extracting the subject parameters, for the biventricular and full heart models. The same tendency can be observed here, although it is more obvious in the results for the more complex structure.

From Figure 3.7, it is possible to conclude that, in general, the worst and best reconstructed phases retain that status irrespective of the training set size. Thus, the quality of the reconstruction depends to some extent on the reconstructed phase, yet quite likely more on the combination of the reconstructed phase and the phases used for parameter extraction. The set of phases used for parameter extraction alone has some influence, but difference between the results for phase set 4 ($\{4, 5, 9, 10, 15\}$) and the other phase sets, presented in Figures 3.7 (d) and 3.8, would best be attributed to the spread of the phases instead of the phases themselves.

Figure 3.9 shows the RMS P2S errors against the percentage of reconstructions, both for the data set that was aligned spatially and temporally, and for the set that was only aligned spatially. It gives an idea of the distribution of the error sizes. For each of the structures, the median error lies around 2 mm, while an error of 3 mm or more is well past the 80th percentile.

Another factor that has a certain degree of influence is which subjects' shapes are being reconstructed. Figure 3.10 shows the mean approximation error for each subject in the 16-fold experiments. The order of the subjects in the plot follows that of the folds. It is obvious that the shapes of certain subjects turn out to be reconstructed poorly at every phase, regardless of the phase set used for deriving the subject parameters. Thus, the model did not generalize well to these subjects.

Finally, the local errors of the reconstructions are shown in Figure 3.11. The errors are color-coded on the surface representations of the reconstructions with

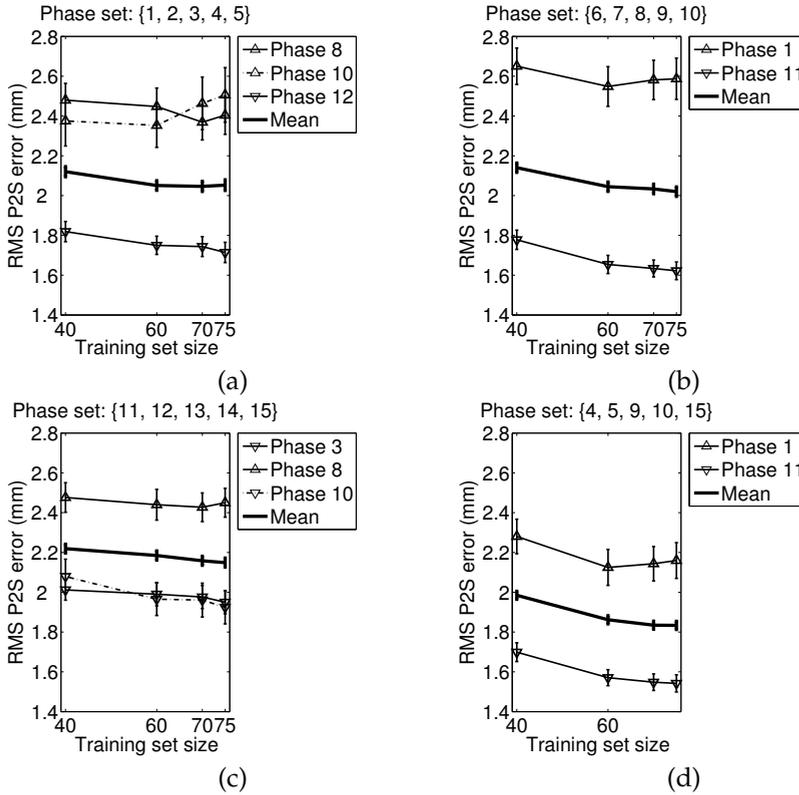


Figure 3.7: The best, worst and mean RMS P2S reconstruction errors for the left ventricular model plus and minus one standard error in mm, after deriving the new subjects' parameters using (a) phases 1 through 5, (b) 6 through 10, (c) 11 through 15, and (d) phases 4, 5, 9, 10 and 15, using models built with increasing training set sizes.

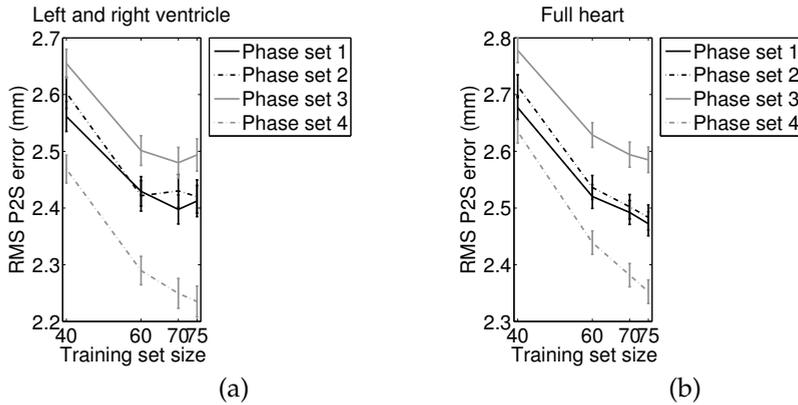


Figure 3.8: The mean RMS P2S reconstruction errors for (a) the biventricular models and (b) the full heart models, plus and minus one standard error in mm, using models built with increasing training set sizes. Each line represents a set of phases used for the derivation of the new subjects' parameters, numbered in the order as they occur in Figure 3.7.

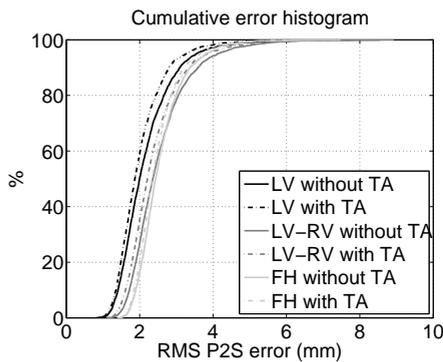


Figure 3.9: The cumulative histogram of errors for the left ventricular (black), biventricular (dark grey) and full heart (light grey) models, using the data pre-processed without (solid line) and with (dash-dot) the temporal alignment (TA) step. Reconstructions included were all those obtained from the 16-fold experiments.

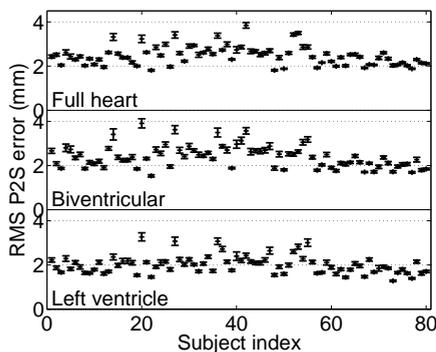


Figure 3.10: The mean RMS P2S reconstruction error for each subject, plus and minus one standard error. It shows that some subjects have a worse reconstruction over the full range of structures and phases, whereas other subjects' shapes were relatively easy to approximate.

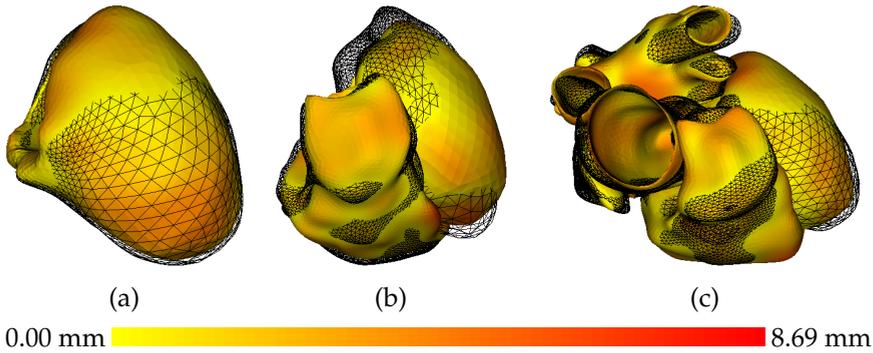


Figure 3.11: Reconstructions of the shapes with the median RMS P2S errors, constructed using the highest dimensional models. (a) Left ventricle (1.869 mm), (b) biventricular (2.212 mm) and (c) full heart (2.344 mm). The wireframe shows the ground truth for these shapes. In (b) and (c), the difference in point density between the left ventricle and the other structures is clearly visible. Reproduced in color on page xliii.

the median RMS P2S error, while the wireframe mesh shows the ground truth shape. Figure 3.11 (c) shows that some error can accumulate in the most complex subparts, namely the atria, and especially in Figure 3.11 (b) it shows that the greater number of points on the right ventricle can result in lower accuracy for the left ventricle.

3.4.2 Comparison to Direct Approximation

We present the comparison between the approximation errors obtained using phases 1, 6 and 11 as approximations for the 5 following phases, and the approximation errors obtained when the subject parameterizations are derived from these same phases, each one separately, is presented in Figure 3.12 (a), (b) and (c) for the LV, biventricular and full heart model, respectively. In the set of phases approximated using the first phase, we can see that the bilinear approach clearly outperforms the direct approximation. This set of phases shows a dramatic change in shape as the heart goes from a state of rest to full contraction. The bilinear approach maintains the mean error within 3 mm. Less dramatic is the shape change in the second set of phases, approximated using phase 6. This is diastole, during which the heart does not return entirely to its full resting state. Recall from section 3.3.1 that the end of diastole was identified as the phase of greatest deceleration; the heart is still relaxing after this point. In this group, the bilinear approach is more favored as the distance from the original phase increases. For the final set of phases, the direct use of phase 11 is favored, although the difference with the bilinear approach is not always significant.

In Figure 3.12 (d) we show how the two approaches compare overall, with respect to the number of phases separating the approximated phase and the phase used for the approximation or subject parameter extraction. It is clear to see that the big differences from the first set of phases dominate.

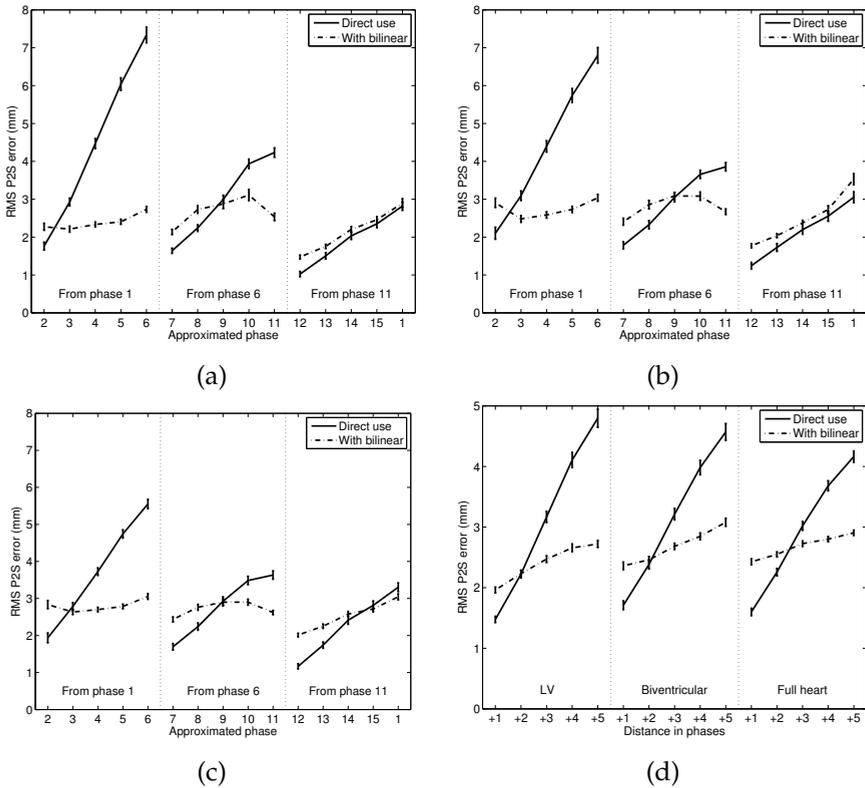


Figure 3.12: The approximation errors plus and minus one standard error when deriving subject parameters from one phase only, for (a) the LV, (b) biventricular model and (c) the full heart model. These would be the initialization errors if used in a segmentation setting. Note that phase 1 is approximated using phase 11, as if it were the 16th phase. Plot (d) shows the errors per distance from the one phase that was used to derive the subject parameters, for each of the three model types: left ventricular, biventricular and full heart.

3.5 Discussion and Future Work

In the previous section we reported extrapolation errors obtained using bilinear models of the heart of varying complexity. We compared the performance of the models built from spatially and temporally aligned shape data with merely spatially aligned data, and we compared the extrapolation errors with the differences between the ground truth of phases, simulating the initialization for image sequence segmentation, which is our intended future use of these models.

Errors of up to 3 mm, as observed in both experiments, should be considered very acceptable, given that the approximated shapes were not seen before. This contrasts with the generalization ability measure commonly used to evaluate PCA-based shape models [186], where the shape to be approximated is known. When providing the initialization for image segmentation, it can offer a better initialization than using a mean shape or the segmentation result of the previous phase, and it can accumulate evidence to improve the initialization as one progresses through the sequence.

We confined ourselves to test data with phases matching those that the model was built with. As is demonstrated in [196], this is not a prerequisite. We are keen to look into the performance of the extrapolation regarding both phase and subject, and to include this in the application of the models to image sequence segmentation.

Although we find the performance of the models very acceptable, there are certainly some limitations regarding both the modeling and the temporal alignment which may have influenced the results. For one, the acquisition scheme as shown in Figure 3.2 is geared towards maximizing visibility of the coronary arteries. Therefore, phases 14 and 15 both have a short time span. The bulk of cardiac motion, however, lies just behind the The complex of Q, R, and S waves in the ECG (QRS complex), resulting in the first five to six phases to be integrated over a time interval with larger spatial changes than the rest. A different temporal sectioning might therefore influence the results. More specifically, higher frequency sampling should result in improved performance both without the temporal alignment as well as *of* the temporal alignment. Additionally, sedation of patients with higher heart rates would have reduced the spread in heart rate and the resulting mismatch of the phases (up to 5 phases for end-systole and up to 7 phases for end-diastole, as shown in Figure 3.5) before the temporal alignment step.

Secondly, the poor approximations of certain subjects' shapes, as seen in Figure 3.10, may be due to the fact that both training and testing data were mixes of healthy and pathological heart shapes. The ratios of these shapes (3:1:1) may have resulted in a bias towards the healthy hearts, resulting in poor derivations of the subjects' parameters in case of pathology. While it was out of the scope of this chapter, whether the bilinear models are powerful enough to separate these groups may be a topic of further research. Also, whether the results would be better if model and test data consisted of only a single class – healthy, or one specific pathology – is an issue that warrants further looking into. When this point is resolved by patient selection, bilinear statistical shape models may provide a means to homogenize multiple acquisition protocols, which could greatly

facilitate retrospective studies.

As the statistics in Table 3.2 and the reconstructions in Figure 3.11 suggest, it is not a given that the more complex shape will always be more difficult to reconstruct. There seems to be a tradeoff between complexity and the amount of information that can be extracted from the example shapes used for the derivation of the new subjects' parameters. However, considering the greater variability that can be observed in the atrial area, and especially around the trunks of the attached vessels and arteries, one would have expected the biventricular models to outperform the full heart model.

How to align sequences of shapes is a question that is not answered entirely. We believe to have taken the most straightforward approach by aligning the mean shapes of each cycle *after* the temporal alignment step, such that the mean would be computed over the same phases for each subject. It would deserve preference, however, to use a more unified approach to align the shapes both temporally and spatially at the same time. Nevertheless, Figure 3.9 does show that employing the temporal alignment indeed improves the extrapolation performance.

We reported the RMS P2S errors of the reconstructions. While these numbers were not unsatisfactory, they do not necessarily have a clinical meaning. Future experiments will focus on the prediction of clinical parameters such as Left Ventricular Volume and Ejection Fraction, Wall Thickening and Wall Motion, and classification of the dynamics based on the parameters acquired using the models.

Finally, the reconstruction errors could decrease further if the training set were larger, which can clearly be concluded from Figures 3.7 and 3.8. Assuming that an appropriate number of training samples for PCA-based models of a complex shape, such as the human heart, easily runs into the triple digits [149], there is room for improvement by increasing the number of training samples for the bilinear model as well.

3.6 Conclusion

We have shown in this chapter how to construct bilinear models of the human heart, and how to derive parameters for new subjects using these models and a limited amount of data from these new subjects. For the construction, simple SVDs and vector transpositions of the data matrix were employed to establish parameterizations for the training subjects and phases, after which a mixing matrix was computed by solving a linear system. The extraction of subject parameters was also reduced to a linear system.

Subsequently, we showed that with such parameters we could predict the shape of the heart over the previously unseen remaining two thirds of cardiac phases with a median RMS P2S error around 2 mm, and that 90% of predictions returned an error below 3.5 mm. Additionally the extraction of parameters from one phase and extrapolation over five following phases returned errors that compare favorably with the differences between subsequent phases. This suggests that a bilinear factorization of the heart shape may be appropriate to separate inter-subject variation from dynamics.

CHAPTER

FOUR

AN ATLAS-BASED DICTIONARY APPROACH TO
INITIALIZING CARDIAC ELECTROPHYSIOLOGY
SIMULATIONS

Adapted from
C. Hoogendoorn, R. Sebastián, J.F. Rodríguez, K. Lekadir, and A.F. Frangi
Under review

Abstract

We present a novel procedure for the initialization of computational simulations of cardiac electrophysiology. By exploiting both volumetric and surface point correspondence, we derive the correlation between domain geometry and results from earlier simulations with similar choices of models. Upon introduction of a new geometry with point correspondence, a prediction of the simulation result is derived from the prior knowledge. We show that the prediction error is typically less than 5% for all model variables, with most variables showing even greater accuracy. When initializing simulations with the predicted model states, it is demonstrated that simulation times can be cut by at least 50% and potentially more. Overall, these results increase the clinical applicability of detailed computational electrophysiology studies.

4.1 Introduction

Computational simulation of electrophysiology (EP) is an increasingly important field in the understanding of cardiac pathophysiology. However, current multiscale approaches require the integration of a large number of variables to account for the dynamic behavior from molecular to organ scale. The trend to incorporate more biological data as it becomes available has a large impact on modeling techniques used to simulate cardiac electrophysiology. The cost associated with running a simulation on a high-resolution personalized anatomy, in terms of computation time and human interaction, limits its translation into a clinical environment and in large scale trial studies.

One of the most important drawbacks of this methodology is the highly nonlinear nature of the models that describe the dynamic behavior of ion channels at the cellular level. This nonlinearity requires a stabilization process before any simulation is carried out. For their use in a three-dimensional anatomy, it is not sufficient to do this stabilization on its basic zero-dimensional building blocks, i.e., isolated cells only. It must be done in the three-dimensional anatomy, allowing the action potentials (APs) of individual cells to be affected by electrotonic currents, specifically the AP amplitude and relative differences in transmembrane currents [32]. The accurate reproduction of AP morphologies is of paramount importance for simulating the complex patterns of tissue activation that may arise with cardiac problems at the cellular level [31, 32, 193].

Specifically this requirement of stabilization prior to introducing a pathological and/or pharmacological modification to the cell behavior can be an obstacle, requiring the simulation of thousands of time steps, which can take days or even weeks to complete in high-resolution meshes of realistic anatomies. In practice, this has led researchers to obtaining a steady state in one reference anatomy or in single isolated cells, and using this for multiple studies. This stands in the way of one of the most rudimentary kinds of personalization of the simulation, namely that of the anatomy.

Additionally, researchers have used the action potential duration (APD) to define correspondence; typically a change of less than 1%, or between 2 and 3 ms, is used. In Fig. 4.1 we can see that for the simulations we run, the APD_{90} (shown in the left panel for one case) ceases to change after only one single heart beat. In contrast, the right panel shows that some variables in the Stewart Purkinje model [183] continue to change more than 10%, averaged over ten cases, for up to nine heart beats. This would suggest not only that the APD_{90} alone is not a sufficiently accurate indicator of convergence, but also that the number of heart beats to be simulated for stabilization may exceed 10 when using models more complex than those employed in this chapter.

In this chapter, we capitalize on a method to establish node correspondence between cardiac ventricular anatomies, to transfer cellular states from one or multiple anatomies (for which electrical simulations have already been carried out) into a new anatomy. The node correspondence allows us to use the notion of similarity between anatomies for identifying a most suitable steady state of cell models in realistic anatomies from the set of available steady states. We show that this data-driven prediction is accurate enough to significantly reduce the time to achieve the steady state in the new anatomy. As more examples be-

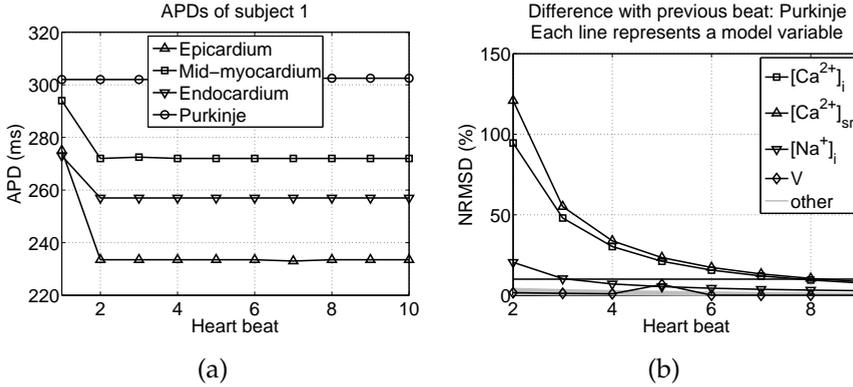


Figure 4.1: (a) Action potential duration for four nodes of different type, and (b) Purkinje model state variables at rest, at each heart beat. It appears that stabilization of the action potential duration does not imply stabilization of the model state.

come available, accuracy should improve and consequently so should the benefit from using the prediction.

4.1.1 Related Work

Simulations have been used for several decades to gain insights into the physiological behavior of living tissues, organisms and populations. More recently, applying computational physiology to interventional or surgical planning has garnered a great deal of interest. This aim imposes additional challenges in terms of patient-specific modeling, to provide individualized treatment, and computational efficiency of the simulations, to be applicable in the clinical setting. Our group focuses on this latter challenge.

While a large body of literature in computational mechanics has focused on optimizing the computational demands of multiscale models, to the best of our knowledge, there has been very little work focused on how to accelerate the convergence to the steady state of the cardiac EP cell model that is required to simulate different therapeutic scenarios and their outcome. We briefly review the two areas where most of the progress in this acceleration has been made:

Mesh Morphing for Simulation Personalization

Finite Element (FE) modeling was introduced in biomedical research over 40 years ago, with an application to orthopaedics [23]. Since then, FE modeling has taken hold in many areas of biomedical research, even though the complexity of anatomical structures has presented a challenge for the construction of quality volumetric meshes for this purpose.

Depending on mesh requirements, quality mesh construction can be a computationally expensive task. In addition, each anatomy, under the same requirements and with the same meshing approach, will return a different mesh in

numbers of nodes and elements, and connectivity. This impedes the comparison of simulation results between anatomies. Recently, the question was raised whether the mesh for an anatomy could be a morphed version of an existing mesh of a different anatomy.

To our knowledge, mesh morphing for FE analysis was first introduced for blood flow studies. This was done using either mesh-to-image registration [9, 218] or using the deformations obtained from image-to-image registration [10]. Later, meshes for orthopaedic studies were morphed based on either an energy minimizing mapping to and from the surface of a low complexity constructive solid geometry, or on the manual identification of matching landmarks and interpolating the deformation using Thin Plate Splines (TPS) [180]. Later, the same group demonstrated two other methods, using parametric descriptions of the template domain [181].

Recently, mesh morphing was also introduced in simulations of cardiac mechanics through the matching of the surfaces of a template volumetric mesh and a patient-specific surface mesh [214, 215], as well as image-to-image registration [105, 106]. In both cases, the application to meshes of cubic Hermite elements requires re-calculation of material properties at the mesh nodes after morphing.

For cardiac electrophysiological simulation, in a previous study, we used TPS to morph a template volumetric mesh of the human left ventricle [76, 77, 153]. The simulations were based on solving the Eikonal equations with fast marching methods, providing only local activation times without regard for physiology at the cellular level. As a result, low-resolution volumetric meshes sufficed to produce mesh independent results.

In this chapter, the TPS approach is applied to high-resolution meshes for detailed multiscale biophysical simulations. The node correspondence resulting from the morphing is exploited to estimate the resting state of ionic model variables on patient-specific meshes.

Faster Cardiac Electrophysiology Simulations

Cell models for cardiac electrophysiology simulation typically employ non-linear Ordinary Differential Equations (ODEs). Solving these on the millions of nodes that typically make up the computational mesh is the most resource intensive of patient-specific simulations, both in hardware and in time. However, many recent developments on the numerics side have contributed to the general notion that cardiac EP simulation will have a future in the clinic.

Firstly, there is the development of solvers specifically for cardiac electrophysiology and electromechanics [22, 71, 159, 211]. Additionally, more and improved use is made of parallelization [5, 104, 207] and Graphics Processing Units (GPUs) [91, 116].

Lately, reduced basis methods [4, 147] have received increased attention for biomedical applications [141, 165]. The objective is to discover the manifold of solutions to a parameterized problem, under the assumption that this manifold will be of lower dimensionality than the problem, facilitating more efficient solution in new cases. However, its application to nonlinear reaction-diffusion equations in general [64], and cardiac electrophysiology simulation in particu-

lar [21], is still very limited.

In this chapter, we use a data-driven method to initialize simulations of cardiac electrophysiology. Rather than computing a regression model using the electrocardiogram as in [223], we use the notion of morphological similarity between cardiac anatomies to choose one or multiple anatomies from a database, and use their steady cell model states for a new anatomy as initial steady-state cell precondition of a new cardiac model outside the database. As the database can be computed off-line and be as extensive as desired, this enables a good trade-off between off-line and on-line computations to make complex simulations more amenable to clinical applications. The database is generated from a statistical cardiac atlas learned from a database of human cardiac images.

4.2 Methods

The methods employed in this chapter comprise a full pipeline for generating volumetric meshes with node correspondence, based on surface meshes with point correspondence. This enables the construction of a database of completed simulations, of which the EP cell model states can be transferred to a new anatomy in a straightforward manner. Once transferred, these states can accelerate the simulation in the new anatomy. The overall scheme to obtain the node correspondence and to populate the database is illustrated in Fig. 4.2. Although the pipeline has been built making specific choices regarding deformation, cell, and fiber models, and choosing a specific numerical solver, these are merely concrete design choices to illustrate our proposed approach in a concrete application but do not reflect any constraints imposed by the pipeline. The various elements of the pipeline are outlined in the following subsections.

4.2.1 Generation of Anatomical Surface Meshes

The pipeline as implemented here assumes point or surface correspondence between surface meshes, without imposing any restrictions on the number of points used to describe the anatomies. Additionally, although for easy distinction from volumetric meshes we use the term ‘surface mesh’, points do not actually need to lie on the surface. Furthermore, the methodology is entirely independent of the mesh element type (triangles, quadrangles, hexagons or a mixture). In this chapter, we use triangulated surface meshes obtained using the aforementioned atlas-based approach [75].

At this stage, a representative anatomy is to be derived from a set of surfaces \mathcal{S} . Given that there is point or surface correspondence between the elements of \mathcal{S} , one could consider generating a mean anatomy S_μ , or one could pick the element minimizing the total dissimilarity to the remainder of the set. For simplicity, we continue to use S_μ to indicate this representative anatomy.

We used eleven human cardiac left ventricle (LV) anatomies obtained using an atlas-based approach as described in [75], and shown in Fig. 4.3. The anatomies were drawn from a larger population of subjects who had undergone a multislice Computed Tomography (CT) scan (Toshiba Aquilion 64, Toshiba Medical Systems, Tochigi, Japan) as part of a standard clinical protocol. The CT

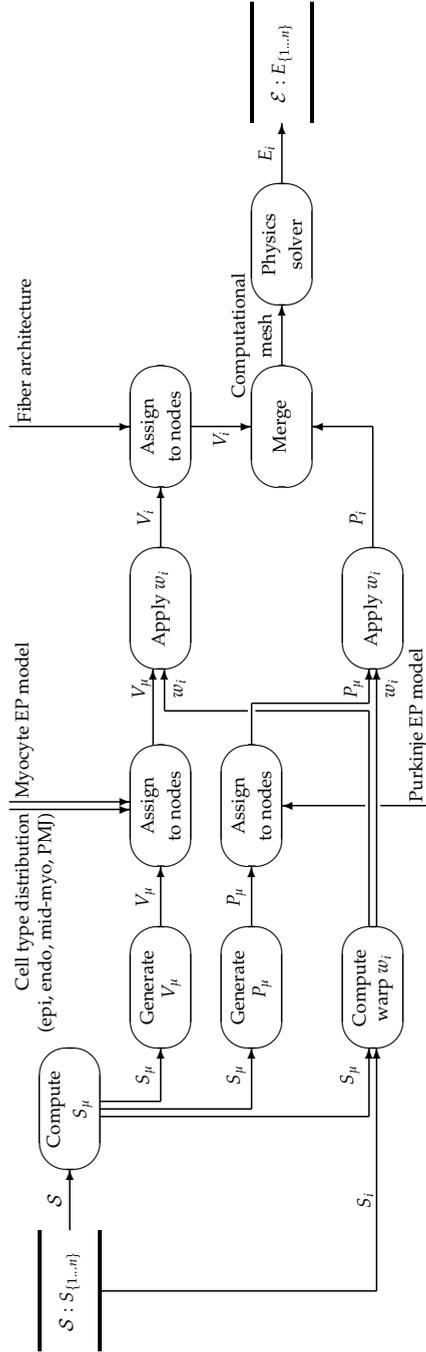


Figure 4.2: Flow diagram of the pipeline for generating volumetric meshes V and Purkinje networks P with node correspondence over a population S of geometries, with surface meshes S_i , in order to obtain simulation result E_i . μ denotes association with the mean or reference geometry.

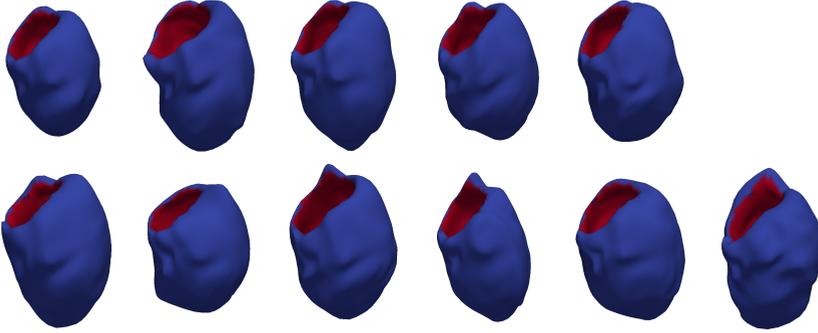


Figure 4.3: Surface renderings of the eleven geometries, with endocardium (red) and epicardium (blue). Reproduced in color on pages xliii.

image resolution was $0.4 \times 0.4 \times 2.0 \text{ mm}^3$, and scans covered the full heart. The LVs were described by a closed triangulated surface mesh with 3,577 vertices and 6,778 faces.

While the first ten anatomies already show clear variations in shape, the eleventh anatomy was chosen specifically to differ from S_μ twice as much as the farthest among the first ten, with the aim to study whether greater distance from the reference anatomy would strongly influence the prediction accuracy.

4.2.2 Generation of Anatomical Volumetric Meshes

From the representative anatomy S_μ , the template volumetric mesh V_μ is generated. The requirements for this volumetric mesh depend on the equations to be solved and the dynamics of the particular ionic model. However, one should keep in mind that these requirements should remain satisfied under deformation of the volumetric mesh. As is the case with the surface meshes, the pipeline itself does not impose any specific resolution, element type or element dimensionality.

In this work, we centered a regular grid of hexahedra, with edge lengths of 0.40 mm, on the template surface mesh, covering it entirely. Each hexahedron with its center inside the surface mesh was assigned to the volumetric mesh. The remaining hexahedra were discarded. The resulting mesh consisted of 2,481,028 nodes and 2,318,956 elements.

4.2.3 Personalized Peripheral Conduction System

A vital part of cardiac anatomy for the synchronized activation of the ventricles is the Purkinje system [37,168,210]. In this work, a single Purkinje network P_μ is generated by applying the tree growing algorithm proposed in [177] to S_μ . The method comprises three steps, each providing a specific part of the network, building on the previous step.

First, the Bundle of His is generated starting from the basal section of the septal wall, running down to the apex and back up again, with a maximum

length constraint. The bundle branches break away from the His Bundle in the upper third of the septal wall, curving round the anterior and posterior walls to the approximate location of the papillary muscle insertions. Some redundant connectivity bundles are generated to complete the first step.

The Purkinje Backbone Branches are then generated based on Lindenmayer systems (L-systems) [83, 115], presenting rules by which construction of further branches proceeds. Distributions of the variables involved are derived from a morphological characterization of stained bovine tissue samples, as in [177].

Finally, the Purkinje Terminal Branches are generated also using L-systems, with a different set of parameter values set to mimic the morphology of these branches upon visual inspection. The Purkinje network produced is completely independent from the underlying ventricular model used; no nodes are shared, and connections between the two exist only as Purkinje-myocardial junctions (PMJs).

4.2.4 Myocardial Fiber Architecture

In cardiac electrophysiology, the myocardial fiber architecture is an important component; the anisotropic conductivity of cardiac muscle cells has been studied extensively [12, 26, 81, 167, 174, 203]. Our fiber orientation model is based on the measurements reported by Streeter *et al.* [184], which were fit to linear and quadratic equations [178] to define a helix angle $\alpha_h(x)$ and a transmural angle $\alpha_t(x)$ at a location x as

$$\alpha_h(x) = 1.90w(x) + 0.86 \quad (4.1)$$

and

$$\alpha_t(x) = 0.215\phi(x)^2 + 0.0089\phi(x) - 0.0093, \quad (4.2)$$

given the normalized distance $w(x)$ of x to the endocardial surface, and the polar angle $\phi(x)$ characterizing the distance between x and the base of the ventricle. The resulting fiber angle between epicardium and endocardium typically rotates from -60° to 60° , corresponding to the observations in [184].

The pipeline allows for elements to be kept out of the morphing, or more precisely, to be added to the myocardial and Purkinje meshes after they have been morphed. In principle, the fiber architecture could be defined prior to morphing. However, as the architecture comprises vector-valued elements, morphing the architecture with nonrigid deformations would require reorientation of the vectors. As we are using a rather simple Streeter-like model, both approaches should yield very similar fiber architectures, but at different computational costs.

4.2.5 Template-to-Anatomy Mesh Adaptation

The pipeline illustrated in Fig. 4.2 was previously used to morph low-resolution volumetric meshes for simulation of electrical activation based on Eikonal equations and fast marching methods [76]. The morphing is done using Thin Plate Splines [19]. TPS is a specific instance of a kernel or landmark-based transform; the displacement of a set of landmarks is defined, and a kernel matrix is computed, which interpolates the deformation between these landmarks. If one

makes no assumptions about the means with which the surface correspondence between template and new surface comes about, this is to be considered one of the most natural deformation models.

In our population, surface correspondence was obtained through an atlas-based approach [75]. However, other approaches to segmenting the structure of interest may also produce surface or point correspondence. Specifically the popular Active Shape Model and Active Appearance Model [39, 40] provide point correspondence automatically. In cases where no correspondence is provided automatically, surface-to-surface registration may be used either to obtain correspondence, or to directly derive a warp for use in the pipeline.

4.2.6 Prediction of Steady State

Cellular models for electrophysiology simulation typically consist of a set of non-linear ODEs. Those equations are usually pre-stabilized as an individual cell, i.e., when a single-cell simulation is run, each cardiac cycle will be the same from the start. The stabilization of a single cell might require a simulation of a few hundred stimuli, which is feasible in a 0D model but not for a large 3D model. Additionally, when a tissue sheet or block is considered or when we consider patient-specific anatomies, cells in the region or organ will interact with one another in a complex manner that will incorporate both knowledge on the type of interacting cells and their distribution in the region or organ. The steady state attained in such configuration is thus the result of complex interactions and cannot be assumed to be the steady state of the individual isolated cellular models. E_{new} in Fig. 4.2 denotes this newly attained steady state; the goal of this work is to estimate this steady state without requiring simulations lasting several cardiac cycles, or alternatively to use such an estimated steady state as the initial cell model state to attain the steady state for a new anatomy in fewer cardiac cycles of simulation.

In Fig. 4.4, we illustrate how the prediction may be included in the pipeline of Fig. 4.2. Again, the diagram is generic regarding the type of predictor used.

The most basic approach to predicting the steady state of the cell models from the LV anatomy is by exploiting measures of morphological similarity in the high-dimensional cartesian space inhabited by the anatomies. Interpolation or label transfer from the k nearest neighbors is a well-known technique in classification. In manifold learning, it forms the basis of various approaches ([11, 50, 169, 176]). When sufficient examples are available, the mapping between the anatomy manifold and the manifold of steady states could be learned. Note that ‘sufficient’ here would likely mean hundreds or thousands of cases.

Let \mathbf{x}_0 be an anatomy represented as a single-column vectorization of its surface mesh vertex coordinates, and $\mathbf{X} = [\mathbf{x}_1 \cdots \mathbf{x}_n]$ a matrix constructed through a side-by-side arrangement of a set of n anatomies in the same representation, with known steady states represented similarly as $\mathbf{Y} = [\mathbf{y}_1 \cdots \mathbf{y}_n]$. Then \mathbf{X}' consists of the k columns of \mathbf{X} such that, treating matrices as sets of column vectors, $(\mathbf{x}_i \in \mathbf{X}' \wedge \mathbf{x}_j \in \mathbf{X} \setminus \mathbf{X}') \Rightarrow d(\mathbf{x}_0, \mathbf{x}_i) \leq d(\mathbf{x}_0, \mathbf{x}_j)$ for some distance or dissimilarity function d . Then, a vector \mathbf{w}_0 of weights is computed as the least-squares

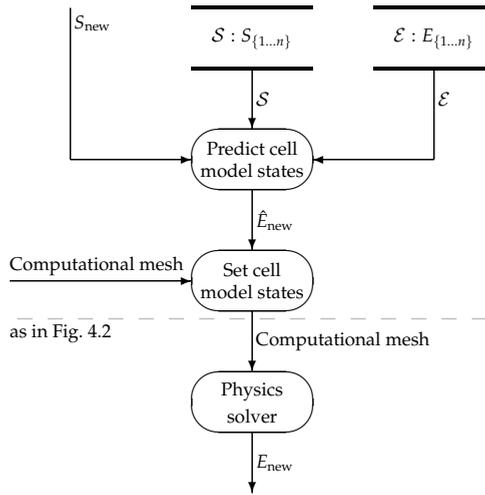


Figure 4.4: Flow diagram of the prediction and its inclusion into the pipeline for the automated construction of volumetric and computational meshes.

solution to $\mathbf{x}_0 = \mathbf{w}_0 \mathbf{X}'$. A predicted steady state $\hat{\mathbf{y}}_0$ is then generated using

$$\hat{\mathbf{y}}_0 = \mathbf{w}_0 \mathbf{Y}', \quad (4.3)$$

where \mathbf{Y}' contains the k columns of \mathbf{Y} corresponding to the k columns of \mathbf{X} that make up \mathbf{X}' . In addition to the unconstrained set of weights \mathbf{w}_0 we also used normalized sets of weights

$$\mathbf{v}_0 = \frac{\mathbf{w}_0}{\sum_{i=1}^k w_{0i}},$$

to ensure that predicted model state values fall within a plausible range. When using \mathbf{v}_0 with $k = 1$, this reduces to a dictionary approach.

4.3 Experiments and Results

The benefits of the method can be assessed in terms of both time gain and its accuracy, compared to obtaining the steady states by simulation using the conventional initialization from isolated-cell steady states. This holds for both the prediction of the initial cellular steady state and its use to initialize simulations of cardiac electrophysiology. We used a leave-one-out approach in this assessment.

4.3.1 Data – Computational Meshes

The volumetric mesh V_μ of the mean anatomy S_μ was generated to obtain regular hexahedra with edge lengths of 0.40 mm, producing 2,481,028 nodes and 2,318,956 elements, as reported before. Note that the hexahedral shape was no longer regular after applying the mesh morphing. Edge lengths of the meshes

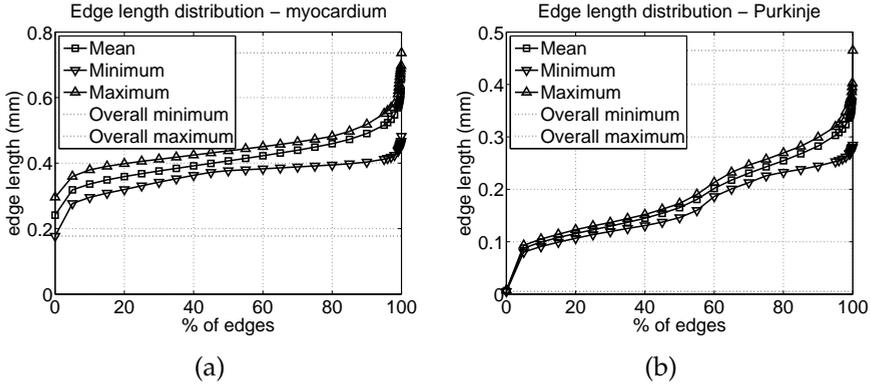


Figure 4.5: Edge length distributions in (a) the myocardial mesh and (b) the Purkinje network, showing minimum, mean and maximum percentile values over the ten subjects.

after morphing are summarized in Fig. 4.5(a). After mesh adaptation, edge lengths did not exceed 0.74 mm in any segment of the first ten anatomies. The maximum edge length in the eleventh anatomy was 0.90 mm.

In V_μ , nodes within 0.80 mm from the epi- and endocardial surfaces (defined in S_μ and colored blue and red, respectively, in Fig. 4.3) were labeled as epi- and endocardial nodes, accounting for 260,541 and 167,013 nodes, respectively. The remainder was labeled as mid-myocardium. This labeling was fixed throughout the datasets, i.e., not re-assigned after morphing.

The Purkinje network P_μ was generated on S_μ , with edge lengths not exceeding 0.27 mm, yielding edge lengths after morphing below 0.5 mm as shown in Fig. 4.5(b). The resulting network consisted of 21,004 1D elements, and of its 20,898 nodes, 4,664 were terminals. The Purkinje network was coupled to the myocardium by connecting these terminals to those nodes in V_μ within a distance of 0.5 mm. These PMJs were modeled as 0D resistances.

The TPS warps were computed from the 3,524 vertices of the LV surface meshes. In order to achieve further reduction in total processing time, the number of landmarks could be reduced [77]. However, the influence of this reduction on detailed biophysical simulations such as those performed here has not been assessed at this point and its analysis is beyond the scope of this work. Computation of the TPS took 2 hours and 7 minutes for each case. Details of the hardware for this step are provided in Sec. 4.3.3.

We should note that we used the same mean anatomy S_μ throughout the experiments, computed over the first ten anatomies. Thus generally the test anatomies were included for the computation of this mean. This was done to maintain a single reference frame for comparison of the results.

4.3.2 Data – EP Simulations

In order to study the method as a proof of concept, we chose to use a cell model with few parameters, that still allows for sufficiently accurate simulations. Such

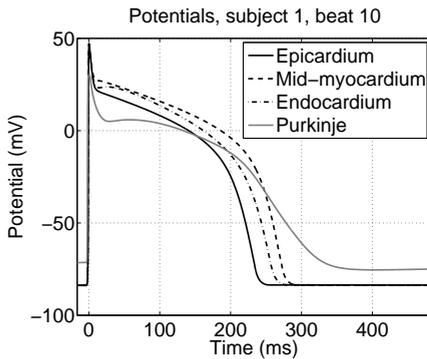


Figure 4.6: Action potential curves for the various node types: epi- and endocardial, mid-myocardial and Purkinje, for one subject, in the final beat. The curves have been shifted on the X axis to have the peak potentials coincide at 0 ms, and the horizontal axis was truncated to 500 ms, half of the basic cycle length.

candidates would be the minimal Bueno-Orovio model [24], which has four parameters, or the six-parameter model presented in [13]. We chose to use the minimal Bueno-Orovio model. Figure 4.6 illustrates the action potential curves in each of the tissue types (epicardial, endocardial, and mid-myocardial).

Specific cell models for the human Purkinje system have been developed recently [175, 183]. Despite their much greater dimensionality compared to the minimal model, we chose to use the Stewart model [183] for appropriateness. The higher resolution of the Purkinje network mesh is due to this increase in model complexity with respect to the minimal model used for the myocardium. The action potential curve of Purkinje nodes is also depicted in Fig. 4.6, in grey.

Cell coupling at tissue level was performed using the monodomain model. The accuracy of such a tissue model was demonstrated to be similar to the bidomain model [162], and with lower computational requirements.

4.3.3 Hard- and Software

Three types of hardware were employed in these experiments; the TPS warps and the predictions were computed on a desktop computer equipped with an Intel Core i7 CPU at 2.67 GHz, with 6 GB of memory, under the Windows 7 operating system. The implementation was in C++ using the Visualization Toolkit [100]. As mentioned before, the warps were computed in about two hours.

Simulations were carried out on two computing clusters using the Elvira solver [71] employing a semi-implicit integration scheme with fixed time steps of 0.02 ms. The first cluster (UPF) is equipped with 24 blades, each containing two Intel Xeon 5355 64-bit quad-core processors running at 2.66 GHz and 16 GB of memory. The second cluster (UZ) is heterogeneous, yet the subcluster used features 16 blades with two quad-core Intel Xeon 5520 64-bit processors at 2.27 GHz and 24 GB of memory each. Inter-blade communication in both clusters is via a high-speed infiniband network. Where relevant, distinction between the clusters will be made in the results reported further on.

4.3.4 Baseline Steady State of 3D Models

On the models as described in the previous section, 10,000 ms of cardiac activity were simulated with step sizes of 0.02 ms and with the heart rate set to 60 beats

per minute. In all cases both the myocardial tissue and the His-Purkinje system showed normal activation sequences with total activation times within physiological ranges. The state after each beat was recorded, and the normalized RMS (NRMS) difference between consecutive beats was computed as

$$\text{NRMS}(\mathbf{y}_{t-1}, \mathbf{y}_t) = \frac{100 \times \text{RMS}(\mathbf{y}_{t-1}, \mathbf{y}_t)}{\max(\mathbf{y}_{t-1}) - \min(\mathbf{y}_{t-1})}.$$

As such, the difference is expressed as a percentage of the range of values in the cell model states. This range was recomputed for each model variable. While APD_{90} , the usual indicator of convergence, showed a steady state after a single beat, variables in the cell models required much longer simulation times (see Fig. 4.1). After an average of 8.64 beats, all variables in the myocardium and Purkinje model showed an NRMS change of less than 10%. The exact number of beats for each case is shown in Table 4.1.

TABLE 4.1: CPU TIMES TO STABLE STATE

<i>case</i>	<i>cluster</i>	<i>cores</i>	<i>beats</i>	<i>CPU time (days)</i>	
				<i>per core</i>	<i>total</i>
1	UPF	8	8	17.85	142.81
2	UPF	8	9	18.36	146.84
3	UPF	8	8	25.47	203.79
4	UPF	8	8	14.26	114.05
5	UPF	8	9	23.46	187.66
6	UPF	8	9	15.72	125.74
7	UPF	8	9	19.10	152.77
8	UPF	8	9	16.78	134.25
9	UPF	8	8	17.27	138.15
10	UZ	16	9	2.52	40.31
11	UZ	16	9	1.41	22.55
Mean UPF					149.56
Median UPF					142.81
Mean UZ					31.43
Mean					128.08
Median					138.15

As our gold standard we use the simulation states after ten heart beats, using the conventional initialization. The accuracy of the predictions is measured as NRMS errors.

Prediction errors on the eleventh anatomy at time $t = 5$ were invariably greater than the errors on the first ten anatomies. They were statistical outliers for one myocardial cell model variable, and for all Purkinje model variables but

one. However, in light of the good overall results, the results on the eleventh geometry have been included in the aggregated results presented from here on.

In Fig. ?? we report the model state NRMS errors for the myocardial and Purkinje models for the various k (number of nearest neighbors), based on the weights \mathbf{w}_0 in the left-hand panels of the top and middle rows and on the weights \mathbf{v}_0 in the right-hand panels and the entire bottom row. Despite the use of different ranges for the vertical axis, it is clear that the normalized weights \mathbf{v}_0 yield better accuracy than \mathbf{w}_0 . It also shows that, although with $k = 1$ the predictions are slightly worse than with $k > 1$, no clear improvement with increasing k is observed, with best performance, in most cases observed around $k = 5$, seldom more than 1% better than the $k = 1$ result.

Figures 4.8 and 4.9 illustrate the spatial distribution of the prediction errors with $k = 1$; unsigned relative errors were averaged over all subjects. The figures show the distributions for each of the myocardial cells' parameters, and for a selection of the Purkinje cells' parameters, corresponding to the ionic concentrations and the transmembrane potential. In addition, histograms of the error distributions are shown to illustrate that the mean unsigned relative error is well below 1% for the vast majority of nodes.

The surface renderings in Fig. 4.8 show that the larger errors are concentrated near Purkinje terminal sites. As the connection between the Purkinje network and myocardium was done after morphing, and was based on a node-to-node distance, inevitably a variability in this connectivity was introduced in the population. Yet, although variable v shows a large maximum error, only very few nodes (1.36%) actually have an average error of 10% or greater.

Prediction errors in the Purkinje cell model states were almost exclusively below 1%. For variable $x_{r,1}$ (not shown), at 16.04% of Purkinje network nodes, the error fell in the 1% to 10% range.

4.3.5 Prediction Times of the Steady States

For the predictions that had their accuracy reported in Fig. ??, computing times were recorded and presented in Fig. 4.10. The main determinant is the amount of data taken up by the cell model states rather than the nearest neighbor search. Consequently, computing time scales linearly with k for $k > 1$. It is clear that this process took mere seconds, peaking at just under 20 seconds for $k = 9$.

The amount of time spent on the nearest neighbor search depends on the complexity of the surface meshes, the size of the database and on whether a data structure for increased search efficiency already exists. With ten neighbors to search and only a few thousand landmarks to describe the surface meshes, its influence on the computing times here was negligible.

4.3.6 Accuracy of Simulations Using Predicted Initial States

We used the predictions with $k = 1$ at the fifth beat to initialize simulations; while not a steady state with respect to the Purkinje model parameters, it allows the assessment of usefulness of the predictions in reaching the steady state faster.

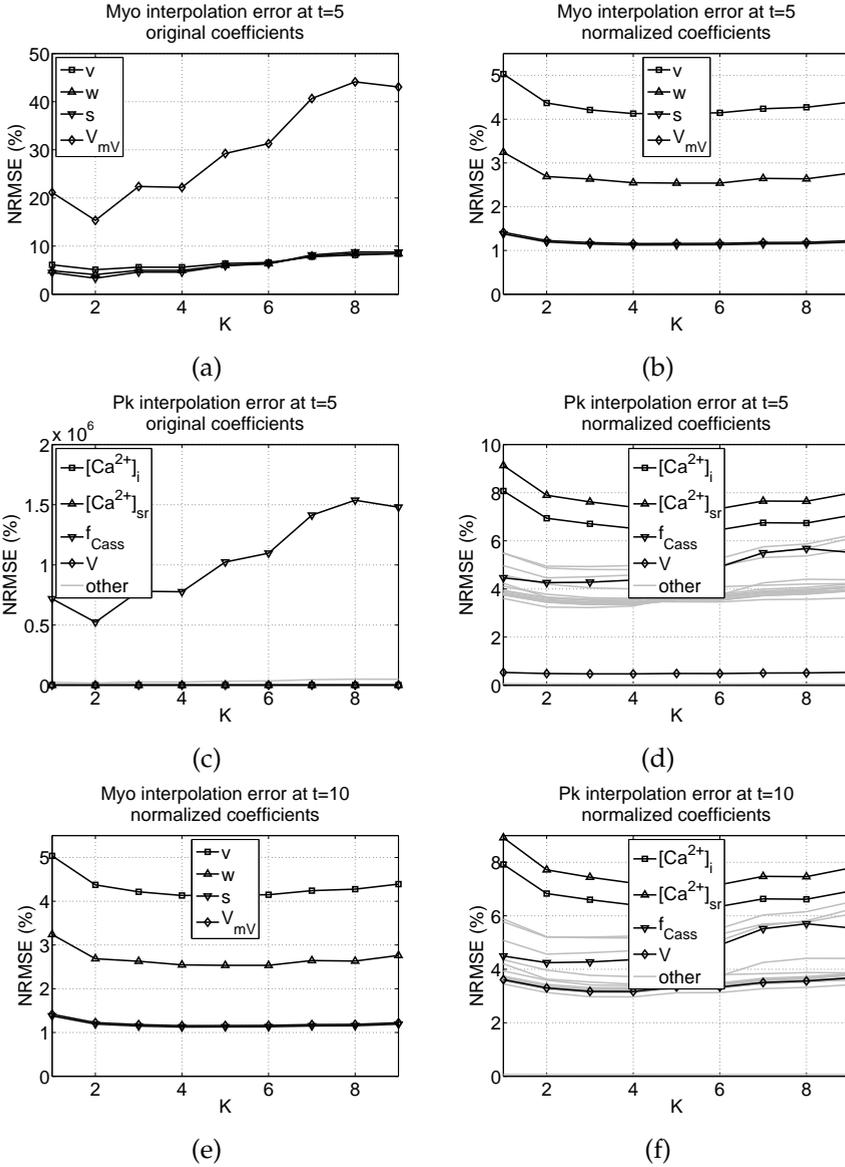


Figure 4.7: Dependence on k (number of nearest neighbors) of NRMS errors at (a-d) time $t = 5$ and (e-f) time $t = 10$ of the predictions of (a,b,e) myocardium and (c,d,f) Purkinje (Pk) model states, with (a,c) original weights (w_0) and (b,d,f) normalized weights (v_0). Highlighted Purkinje model variables are the voltage variable, or outliers in any of the plots shown here.

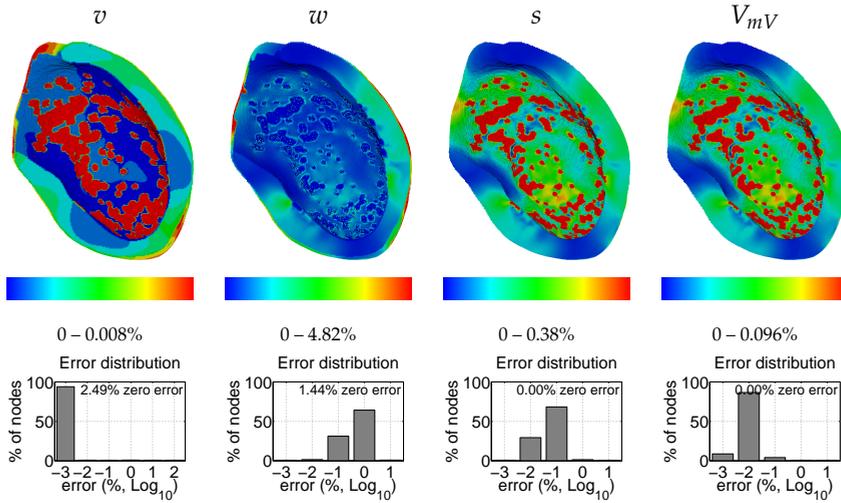


Figure 4.8: Mean unsigned relative prediction errors at time $t = 5$ color coded on the surface of the mean myocardium. Each pair of a surface rendering and histogram corresponds to a model state variable. Color bar upper limits were set to the 95th percentile. Reproduced in color on page xlv.

In Fig. 4.11, we show using the solid lines the mean residual error for each variable as it converges toward the state after 10 beats. The dashed lines running from the fifth beat onwards correspond to the same variables, initialized with the predictions at $t = 5$. The left panel shows that the 5% prediction errors for the myocardial models are largely removed after a single beat. For the Purkinje models, most variables are predicted close to the correct state at five beats, and the convergence from the prediction appears very similar to that of the ground truth. This suggests that the predictions are very effective as an initial cell model state.

After running the simulations from $t = 5$ for an additional five beats, the mean unsigned relative differences are color-coded on the myocardial and Purkinje mesh surfaces in Figs. 4.12 and 4.13. While the maximum error has not been reduced, the portion of myocardial nodes with errors in variable v greater than 10% has dropped to 0.0014% (35 nodes), while the cell model states in 86.28% of nodes are the same between the two approaches to machine precision.

For the Purkinje system, the error ranges have been compressed, and now all nodes' errors for all variables are below 1%, with only three variables (including $[Na^+]_i$ and V , shown in Fig. 4.13) presenting errors in the 0.1% to 1% range, for 0.46% ($[Na^+]_i$), 2.12% and 0.019% (V) of nodes, respectively.

Combining these results with those reported in Sec. 4.3.5, we have achieved nearly identical cell model states in only 50% of the time. The speed with which the myocardial cell model state variables converge to the ground truth, and the initial accuracy of the Purkinje cell model state prediction (see Fig. 4.11), suggest that prediction at $t = 9$ could yield a 90% reduction in computing time.

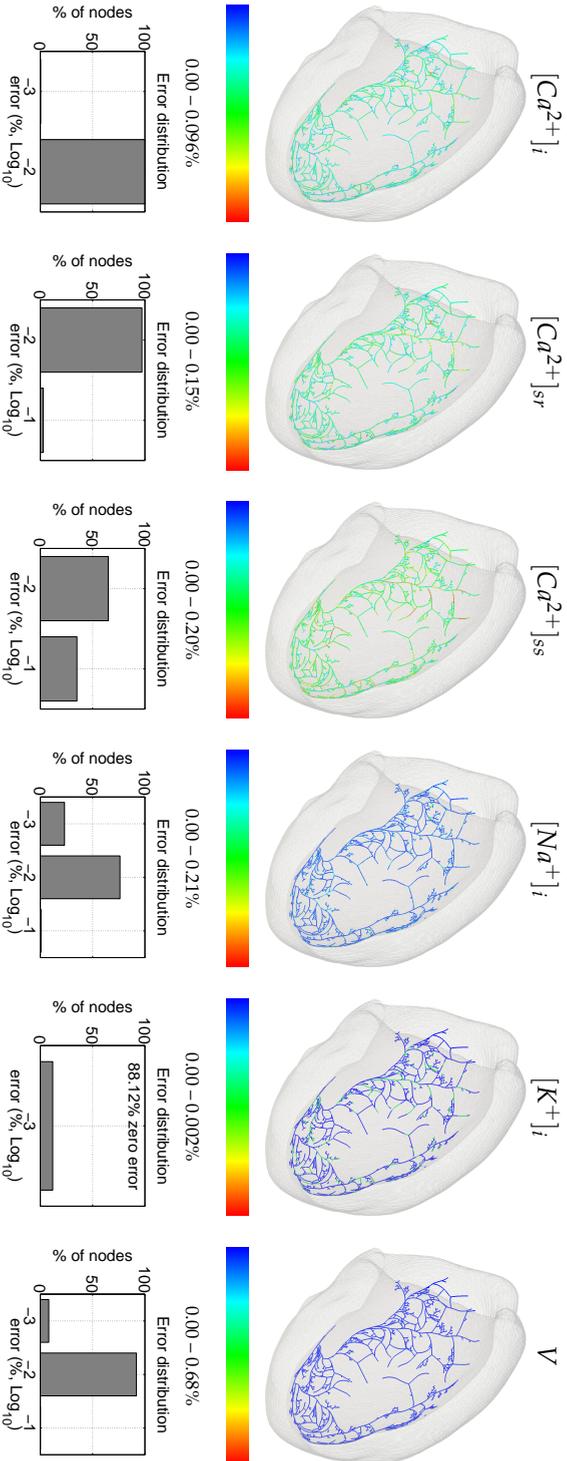


Figure 4.9: Mean unsigned relative prediction errors at time $t = 5$ color coded on the Purkinje network corresponding to the mean geometry. Each pair of a tree rendering and histogram corresponds to a model state variable. Only the ionic concentrations and the voltage V are shown. Reproduced in color on page xlv.

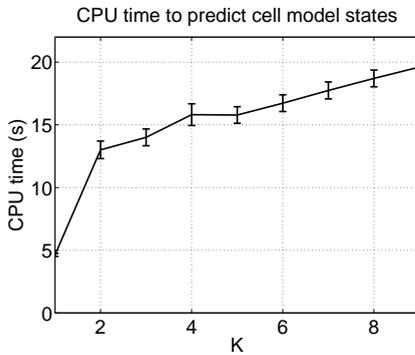


Figure 4.10: Mean CPU times (\pm standard error) for generating the predictions, depending on k , using the normalized coefficients v_0 .

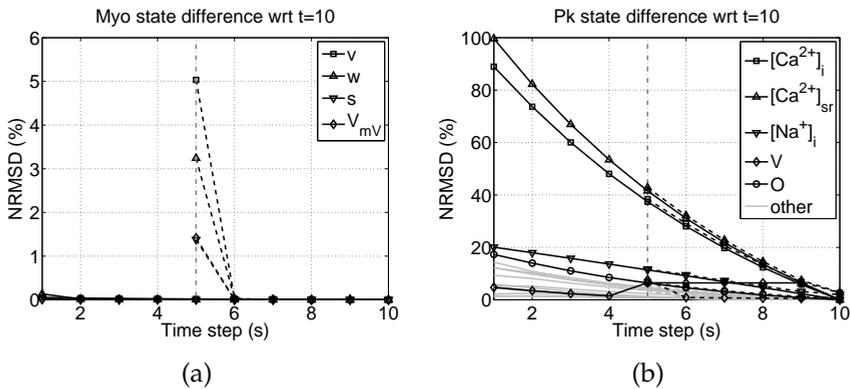


Figure 4.11: Residual errors in model states of (a) myocardium and (b) Purkinje (Pk) network with respect to $t = 10$, for the simulations from scratch (solid lines) and from prediction at $t = 5$ (dashed lines). Each line represents one model state variable, averaged over the subjects. Lines for corresponding variables have equal markers. Highlighted Purkinje model variables are the voltage variable, and the four slowest converging variables.

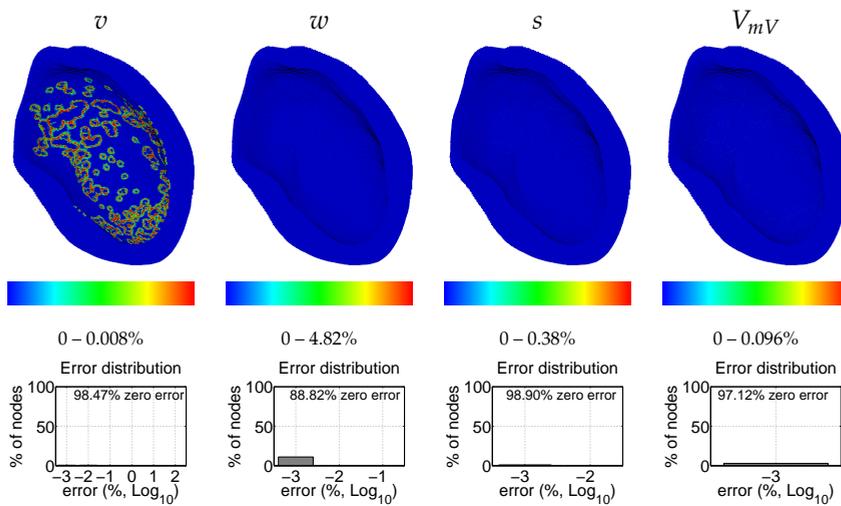


Figure 4.12: Mean unsigned relative differences between the two approaches to reach time $t = 10$ color coded on the surface of the mean myocardium. Each pair of a surface rendering and histogram corresponds to a model state variable. Color bar ranges are the same as in Fig. 4.8. Reproduced in color on page xlvi.

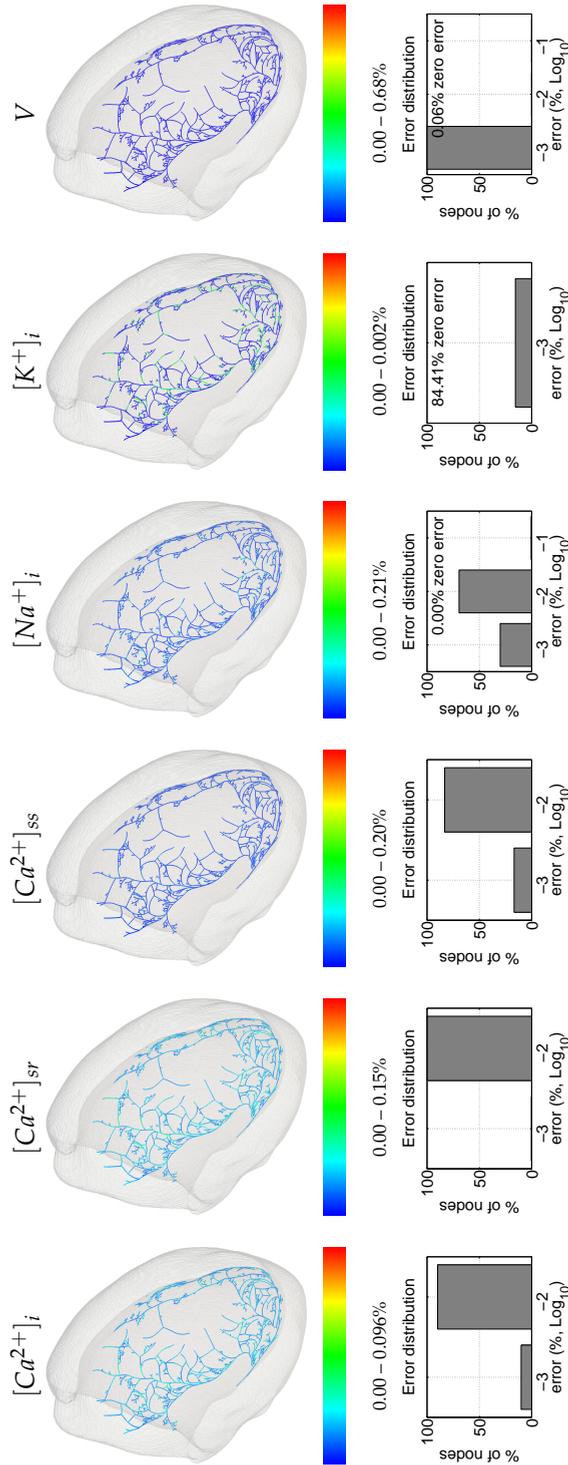


Figure 4.13: Mean unsigned relative differences between the two different approaches to reach time $t = 10$ color coded on the Purkinje network corresponding to the mean geometry. Each pair of a tree rendering and histogram corresponds to a model state variable. Only the ionic concentrations and the voltage V are shown. Color bar ranges are the same as in Fig. 4.9. Reproduced in color on page xlvi.

4.4 Discussion and Conclusions

We have presented a method to automate the construction of computational meshes for cardiac electrophysiology with node correspondence. We have also proposed a method to predict cellular steady states to initialize cardiac EP simulations on high-resolution meshes of realistic cardiac anatomies, driven by data obtained from simulations on similar anatomies.

With a computation time for the prediction best expressed in seconds, we have demonstrated that the time required to obtain a steady state of the cell models can be reduced by at least 50% and possibly even up to 90%, thus further enabling the use of personalized EP simulations in the clinic, as well as studies on larger populations than previously have been used.

Good performance was obtained by employing a k -nearest neighbor approach and linear interpolation of cell model states, with prediction errors varying only slightly between $k = 1$ and $k = 5$. This could allow optimization of cell state stabilization in a larger population of anatomies through the construction of a processing tree. With larger populations, the effect of k on the prediction accuracy may also be better assessed, and more elaborate (non-linear) manifold learning techniques may be employed.

The implementation of the framework in Fig. 4.2 as used in our experiments was based on surface mesh point correspondence. This is not a requirement, as long as a dense deformation to be applied to the volumetric mesh and Purkinje network can be obtained. This is possible either during segmentation in an atlas-based segmentation approach, or in absence of an inherent point or surface correspondence, can be obtained using surface-to-surface registration. Note that the role of the template surface mesh in the pipeline might change.

Although the results are encouraging, some elements of this work provide some clear avenues for future work.

Convergence of the simulations initialized using the predicted steady state is not guaranteed, an issue which becomes more pertinent as the model complexity increases. However, our results show that the predictions for the Purkinje model are a better approximation of the target state than the cell models' steady state in isolation. Consequently, the likelihood of convergence is expected to increase rather than decrease.

One could consider that the use of one single solver and a predominant type of mesh element may lead to biased results. However, the use of hexahedral elements provides a hard limit on the number of neighboring nodes, which benefits simulation speed. Therefore, time gains on meshes with irregular tessellations are expected to be greater.

A second limitation is that we have used a mesh at normal resolution for the representative anatomy. As the deformation includes local stretching of a number of edges, one may choose the mesh resolution for the representative anatomy based on the expected stretching; to maintain edge lengths below 0.5 mm for the population used here this would require 30% shorter edge lengths, accounting for three times as many mesh nodes. The introduction of an adaptive meshing approach into or immediately after the morphing could reduce the initial resolution requirements. An additional step in the prediction would then encompass the interpolation to newly inserted nodes. As a conse-

quence, this would enable the application of our method to a greater range of anatomical variation.

Conversely, the extensibility and modularity of the framework lends itself to adaptation to more complex simulation scenarios.

Firstly, cell and tissue properties were set to reflect a healthy heart. We believe the method could be extended to cases where local conduction properties may deviate from the normal due to electrical remodeling. When such regions have typical locations and morphology (e.g., ventricular ischemic areas are linked to coronary irrigation; atrial ablation is typically done in the coronary sinus), databases for various templates of abnormal regions could be constructed. In addition, prediction results from various databases could be combined to better match a pattern not captured by one single database.

Additionally, the use of a myocardial fiber model based on Streeter's measurements [184] is a simplistic approach which does not incorporate population variability in the fiber architecture. As this work provides a proof of concept, this was of no consequence. However, development of data-driven models for fiber orientation is already underway [111, 113, 118], and could straightforwardly be inserted in the framework.

Finally, we opted for a low-dimensional cell model in the myocardium. Although this model accurately reproduces myocyte action potentials, it does not explicitly incorporate any elements of cellular anatomy, which could be considered a limitation in studying pathological changes. Thus, extendability of the method to more complex myocardial cell models (e.g., [86, 164, 192, 194]) should be an important avenue for further research. However, in this direction, the results regarding the more complex Purkinje cell models are encouraging.

CHAPTER

FIVE

CONCLUSIONS AND OUTLOOK

5.1 Overview

The challenges that this thesis addressed relate to the processing of large quantities of data for atlas construction, to the computational requirements for both population studies and personalization in the clinical setting of electrophysiological simulations, and to the question of statistical modeling of spatio-temporal shape variation.

In each of the three preceding chapters I presented a method which addresses each of these issues, thus bringing the Virtual Physiological Human (VPH) closer to becoming a reality. The conclusions specific to each of these issues are already covered by the respective chapters' own concluding remarks (Secs. 2.6, 3.6 and 4.4).

The first chapter showed that one can generate an anatomical atlas from high-resolution CT image volume sequences largely automatically. Optimization of the registration parameters was done on a mere 5% of the population. Selection of the initial reference required the use of only 10% of the population as 'dummy' candidates during an initial affine registration step and performing Groupwise Mutual Information (GWMI) rank averaging; the ultimately chosen reference was not one of the 10%. Afterward, the removal of bias toward this reference took five interaction-free iterations.

The second chapter illustrated how one can construct a bilinear model to separate anatomical variation from functional variation in a population of cardiac shapes. One would exploit point correspondence between shapes, and one may generate temporal correspondence between sequences, in order to construct a model consisting of a mixing matrix while at the same time obtaining the subject and phase parameterizations required for reconstructing the training set. The generative power of the model was demonstrated by extrapolating two-thirds of a cardiac cycle from only one third of available shapes, as well as extrapolating one-third of the cycle from a single available shape.

The final contribution took the shape of another predictive model, albeit rudimentary due to the small sample size. The differences in normal electrophysiological function, based only on different cardiac geometry, are sufficiently small to use linear interpolation between shapes to predict the state of cell models used in simulation. Sufficiently small meaning specifically that such predictions can cut simulation time significantly compared to current practice in simulation initialization. Even though we know that the EP processes are highly nonlinear, which is reflected in the model equations [145], [146], [24], [194], [36], [183], [175], the resting states themselves appear to be closely linked with geometry.

Tying the second and third contribution together is the atlas in Chapter Two.¹ The volumetric correspondence obtained through the image registration was translated into point correspondence by constructing a surface mesh representation of the atlas. This correspondence was exploited to construct the spatio-temporal statistical shape model.

¹The shapes used in Chapter Three are not those from Chapter Two. They were from an earlier version of the atlas, constructed using simpler techniques [149]. While this seemed to work, the work done for the preliminary version could not be reproduced. A new atlas was constructed, which is the one presented in Chapter Two.

In the third contribution, the surface correspondence was exploited to be able to receive a new geometry –output from an image segmentation algorithm– and construct a computational mesh with node correspondence. This in turn is exploited in the interpolation step.

5.2 Outlook

As could be expected, the methods presented in the preceding chapters are not definite answers to the questions addressed. Each elicited new questions and opened avenues for future research. In some cases, new research has already treaded those avenues.

5.2.1 Atlases, Cardiac and Other

Unbiased atlas construction and atlas-based segmentation are resource-heavy operations that have only relatively recently gained momentum in the medical imaging field. Even if computational resources will become more and more accessible and computing power continues to increase according to Moore’s law, it will remain resource-heavy. Either we will want to use more cases, use more detailed data, or we come up with a better method that demands yet more resources.

While the atlas served to illustrate the framework, the framework itself is of great importance to enable the construction of other atlases. The framework presented in Chapter Two (Fig. 2.1) allows for the replacement of its elements with other methods for equivalent operations, either simply better, or better suited for the data at hand. For some of these elements, future work can be clear-cut. However, the framework itself is based on the idea that for large databases, we cannot optimize for each case separately. Our database contained 138 cases, but as a prime example the database of the Cardiac Atlas Project (CAP) comes to mind, which already contains around 3,000 image volumes. Of those, only a few hundred have been segmented into three parts [59, 95, 189]. Manual processing of the remaining datasets is not feasible even with three parts.

While the brain has received much attention in atlas construction, and the framework was presented on a cardiac atlas construction problem, the framework could be applied to other structures in the body. Any structure that could be subdivided into smaller parts, logically from an anatomical perspective (the subdivision is anatomically well-defined) as well as a clinical perspective (the parts are actually large enough to be seen on clinical scans). What comes to mind are the liver, where segmentation has focused on the gross anatomy only [72] and the lungs, where lobe segmentation has only recently attracted interest [108, 201]. Additionally, some (musculo-)skeletal complexes are likely to be modeled more holistically using atlases [29, 140].

Another avenue, which has already been explored for atlases of brain structures, is the incorporation of multiple modalities in the atlas construction. These could be any of the clinical imaging modalities, or a clinical modality combined with an *ex vivo* modality such as micro-CT (μ CT) or histology, or currently research-only *in vivo* modalities such as diffusion tensor MRI (dtMRI). As inter-

modality registration already presents a number of challenges [125, 135, 182], some would be exacerbated by the specifics of the modalities involved, like color in histological and microscopy images, field of view and resolution in histological, microscopy, and micro-CT images, and the non-scalar nature of tensors in μ CT and dtMRI. All of these challenges will introduce a need for some level of human interaction; with larger databases, this interaction must be smart and minimal.

Finally, possibly the biggest challenge technically, is to let go of an assumption one invariably encounters: that of fixed topology throughout the population. The topology of structures in the human body may vary. Some structures are more stable in this respect, but many structures, either naturally or pathologically, will exhibit topological differences throughout a population. How to deal with this is a very open question, with a great challenge even in distinguishing between the anatomical and mathematical sense of the term. Integration of state-of-the-art registration techniques in the atlas stratification approach of [18] may be a step in this direction. At the level of shape-based atlases, constituent parts could be modeled separately to be combined at time of segmentation, as was done in [157], and statistical models could be mixture models as illustrated in [41]. Stratification of the training images could automate the model building process for this approach.

5.2.2 Spatio-temporal Models of Shape

The bilinear models of Chapter Three were built after a temporal resynchronization of the input data. For the method of alternating Singular Value Decomposition (SVD), this was a necessity. However, when minimizing the reconstruction error (Eq. 3.8) directly, full temporal synchronicity is not a requirement. The requirement becomes that for each time point to be returned, samples are provided for multiple, not all, subjects. Such an approach may allow for the incorporation of training data from multiple modalities of different temporal resolution. How then to deal with the different spatial detail that may be observed, is another question to look into.

Since the publication of the chapter as an article, few other methods have been proposed to unify modeling of spatial and temporal variations [2]. In image registration, some headway was made in four-dimensional transformation models to generate time-consistent registrations throughout a cardiac cycle [46], and to register entire image sequences [138].

Bilinear models have been used now to segment the LV from Magnetic Resonance (MRI) [226], and also to segment the prostate [88]. Other structures where a spatio-temporal state model could be of use would be lungs, diaphragm and the liver, which exhibit respiratory motion. As the bilinear models do not assume cyclicity of the motion, complexes of joints and muscles are also candidates.

5.2.3 Recycling Electrophysiology Simulations

The VPH's reliance on computer simulation studies can be seen as a bottleneck. While computational power and its availability are still increasing, so is its de-

mand with the desire to include more and more measured data in the simulations. As a result, the computational burden of computer simulations is not decreasing very rapidly. Other ways to achieve significant reduction are a necessity. In Chapter Four, we explored the possibility to 'recycle' simulation results or states. We used a mixture of a simple myocardial cell model [24] and a more complex Purkinje cell model [183]. The results on the Purkinje cells suggest that the method might also work for more complex myocardial cell models [36,175,183,194], and perhaps also to relate our method to the reduced basis approach [64,141,165].

The mesh morphing used in the chapter was rather straightforward, and the resolution of the mean volumetric mesh was set to make element size not an issue after morphing, given the population used. An adaptive meshing step could have alleviated this, but could also have complicated the prediction step. However, for the prediction approach to become more universally applicable, it probably is a necessity.

From the eleven geometries used in the study, with the surfaces described by some 3,500 vertices and the simulation states by some 10 million state variables, one could not hope to approximate the manifold that the surface meshes and the simulation states span. However, it is likely that both populations –generally–span a low-dimensional nonlinear manifold embedded in their respective (in this thesis) 10,000-dimensional and 10,000,000-dimensional spaces. With the size of the populations required by a clinically effective VPH, these manifolds may be learned and mapped to each other, resulting in statistics-based predictive models. These should be expected to provide greater prediction accuracy and robustness in more complex –pathological– cases.

5.3 The Virtual Physiological Human's Prognosis

Seeing the Virtual Physiological Human as a concept rather than as the collective VPH initiatives, it is probably safe to say that it is an inevitable future reality. "Mission impossible yet also mission imperative" is another way it has been characterized [212]. The motivations as laid out in the introduction of the Future and Emerging Technologies research roadmap for the VPH [209], are and will remain valid: we cannot look at parts of the body in isolation, be such parts anatomical subdivisions or single physiological scales.

As comorbidities become more frequent with age, this problem will be exacerbated by the ageing society we live in, and even more so by the ageing society that the developing world will see soon: all four of the so-called BRIC countries² are among the ten most populated in the world, currently accounting for over 40% of the world population [202]. The value of a Virtual Physiological Human as the pinnacle of evidence-based medicine towards optimizing healthcare quality and efficiency is almost impossible to ignore.

²Brazil, Russia, India, China

BIBLIOGRAPHY

- 1] B. ABBOUD AND F. DAVOINE, *Bilinear factorization for facial expression analysis and synthesis*, IEE Proceedings – Vision, Image & Signal Processing, vol. 152, no. 3 (2005), pp. 327–333. CITED ON PAGE 46.
- 2] I. AKHTER, T. SIMON, S. KHAN, I. MATTHEWS, AND Y. SHEIKH, *Bilinear spatiotemporal basis models*, ACM Transactions on Graphics, vol. 31, no. 2 (2012), p. 17. CITED ON PAGE 96.
- 3] E.L. ALDERMAN AND M. STADIUS, *The angiographic definitions of the Bypass Angioplasty Revascularization Investigation*, Coronary Artery Disease, vol. 3, no. 12 (1992), pp. 1189–1207. CITED ON PAGE 41.
- 4] B.O. ALMROTH, P. STERN, AND F.A. BROGAN, *Automatic choice of global shape functions in structural analysis*, AIAA Journal, vol. 16, no. 5 (1978), pp. 525–528. CITED ON PAGE 73.
- 5] J.M. ALONSO, J.M. FERRERO, JR., V. HERNÁNDEZ, G. MOLTÓ, J. SAIZ, AND B. TRÉNOR, *A grid computing-based approach for the acceleration of simulations in cardiology*, IEEE Transactions on Information Technology in Biomedicine, vol. 12, no. 2 (2008), pp. 138–144. CITED ON PAGE 73.
- 6] R.H. ANDERSON, M. SMERUP, D. SÁNCHEZ-QUINTANA, M. LOUKAS, AND P.P. LUNKENHEIMER, *The three-dimensional arrangement of the myocytes in the ventricular walls*, Clinical Anatomy, vol. 22, no. 1 (2009), pp. 64–76. CITED ON PAGE 4.
- 7] R.H. ANDERSON, J. YANNI, M.R. BOYETT, N.J. CHANDLER, AND H. DOBRZYNSKI, *The anatomy of the cardiac conduction system*, Clinical Anatomy, vol. 22, no. 1 (2009), pp. 99–113. CITED ON PAGE 3.
- 8] V. ARSIGNY, O. COMMOWICK, X. PENNEC, AND N. AYACHE, *A log-euclidean framework for statistics on diffeomorphisms*, in Proc. 9th Int. Conf. Medical Image Computing and Computer Assisted Intervention (MICCAI), Copenhagen, Denmark, vol. 4190 of Lecture Notes in Computer Science, 2006, pp. 924–931. CITED ON PAGES 21, 22, AND 24.
- 9] L. BAGHDADI, D.A. STEINMAN, AND H.M. LADAK, *Template-based finite-element mesh generation from medical images*, Computer Methods and Programs in Biomedicine, vol. 77, no. 1 (2005), pp. 11–21. CITED ON PAGE 73.
- 10] D.C. BARBER, E. OUBEL, A.F. FRANGI, AND D.R. HOSE, *Efficient computational fluid dynamics mesh generation by image registration*, Medical Image Analysis, vol. 11, no. 6 (2007), pp. 648–662. CITED ON PAGE 73.
- 11] M. BELKIN AND P. NIYOGI, *Laplacian eigenmaps for dimensionality reduction and data representation*, Neural Computation, vol. 15, no. 6 (2003), pp. 1373–1396. CITED ON PAGE 78.
- 12] A.P. BENSON, O. BERNUS, H. DIERCKX, S.H. GILBERT, J.P. GREENWOOD, A.V. HOLDEN, K. MOHEE, S. PLEIN, A. RADJENOVIC, M.E. RIES, G.L. SMITH, S. SOURBRON, AND R.D. WALTON, *Construction and validation of anisotropic and orthotropic ventricular geometries for quantitative predictive cardiac electrophysiology*, Interface Focus, vol. 1, no. 1 (2011), pp. 101–116. CITED ON PAGE 77.
- 13] O. BERNUS, R. WILDERS, C.W. ZEMLIN, H. VERSHELDE, AND A.V. PANFILOV, *A computationally efficient electrophysiological model of human ventricular cells*, American Journal of Physiology - Heart and Circulatory Physiology, vol. 282, no. 6 (2002). CITED ON PAGE 81.
- 14] M. BERTRAM, G. REIS, R. VAN LENGEN, S. KÖHN, AND H. HAGEN, *Non-manifold mesh extraction from time-varying segmented volumes used for modeling a human heart*, in Proc. Joint EuroGraphics/IEEE VGTC Symp. on Visualization (EuroVis), Leeds, UK, 2005, pp. 199–206. CITED ON PAGE 25.

- 15] K.K. BHATIA, J.V. HAJNAL, B.K. PURI, A.D. EDWARDS, AND D. RUECKERT, *Consistent group-wise non-rigid registration for atlas construction*, in Proc. IEEE Int. Symp. on Biomedical Imaging (ISBI), Arlington, VA, USA, 2004, pp. 908–911. CITED ON PAGE 15.
- 16] A. BISTOQUET, J. OSHINSKI, AND O. ŠKRINJAR, *Left ventricular deformation recovery from cine MRI using an incompressible model*, IEEE Transactions on Medical Imaging, vol. 26, no. 9 (2007), pp. 1136–1153. CITED ON PAGE 47.
- 17] J.M. BLACKALL, A.P. KING, G.P. PENNEY, A. ADAM, AND D.J. HAWKES, *A statistical model of respiratory motion and deformation of the liver*, in Proc. 4th Int. Conf. Medical Image Computing and Computer Assisted Intervention (MICCAI), Utrecht, The Netherlands, vol. 2208 of Lecture Notes in Computer Science, 2001, pp. 1338–1340. CITED ON PAGE 45.
- 18] D.J. BLEZEK AND J.V. MILLER, *Atlas stratification*, Medical Image Analysis, vol. 11, no. 5 (2007), pp. 443–457. CITED ON PAGES 15, 16, 42, AND 96.
- 19] F.L. BOOKSTEIN, *Principal warps: Thin-plate splines and the decomposition of deformations*, IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 11, no. 6 (1989), pp. 567–585. CITED ON PAGE 77.
- 20] J.G. BOSCH, S.C. MITCHELL, B.P.F. LELIEVELDT, F. NIJLAND, O. KAMP, M. SONKA, AND J.H.C. REIBER, *Automatic segmentation of echocardiographic sequences by active appearance motion models*, IEEE Transactions on Medical Imaging, vol. 21, no. 11 (2002), pp. 1374–1383. CITED ON PAGES 45 AND 46.
- 21] M. BOULAKIA, E. SCHENONE, AND J.F. GERBEAU, *Reduced-order modeling for cardiac electrophysiology. Application to parameter identification*, International Journal for Numerical Methods in Biomedical Engineering, vol. 28, no. 6/7 (2012), pp. 727–744. CITED ON PAGE 74.
- 22] C. BRADLEY, A. BOWERY, R. BRITTEN, V. BUDELMANN, O. CAMARA, R. CHRISTIE, A. COOKSON, A.F. FRANGI, T. BABARENDA GAMAGE, T. HEIDLAF, S. KRITTIAN, D. LADD, C. LITTLE, K. MITHRARATNE, M. NASH, D. NICKERSON, P. NIELSEN, Ø. NORDBØ, S. OMHOLT, A. PASHAEI, D. PATERSON, V. RAJAGOPAL, A. REEVE, O. RÖHRLE, S. SAFAEI, R. SEBASTIÁN, M. STEGHÖFER, T. WU, T. YU, H. ZHANG, AND P. HUNTER, *OpenCMISS: A multi-physics & multi-scale computational infrastructure for the VPH/Physiome project*, Progress in Biophysics & Molecular Biology, vol. 107, no. 1 (2011), pp. 32–47. CITED ON PAGE 73.
- 23] W.A.M. BREKELMANS, H.W. POORT, AND T.J.J.H. SLOOFF, *A new method to analyse the mechanical behaviour of skeletal parts*, Acta Orthopaedica Scandinavica, vol. 43, no. 5 (1972), pp. 301–317. CITED ON PAGE 72.
- 24] A. BUENO-OROVIO, E.M. CHERRY, AND F.H. FENTON, *Minimal model for human ventricular action potentials in tissue*, Journal of Theoretical Biology, vol. 253, no. 3 (2008), pp. 544–560. CITED ON PAGES 81, 94, AND 97.
- 25] R.H. BYRD, P. LU, J. NOCEDAL, AND C. ZHU, *A limited memory algorithm for bound constrained optimization*, SIAM Journal on Scientific Computing, vol. 16, no. 5 (1995), pp. 1190–1208. CITED ON PAGE 20.
- 26] B.J. CALDWELL, M.L. TREW, G.B. SANDS, D.A. HOOKS, I.J. LEGRICE, AND B.H. SMAILL, *Three distinct directions of intramural activation reveal nonuniform side-to-side electrical coupling of ventricular myocytes*, Circulation Arrhythmia and Electrophysiology, vol. 2, no. 4 (2009), pp. 433–440. CITED ON PAGE 77.
- 27] O. CAMARA, E. KONUKOGLU, M. POP, K. RHODE, M. SERMESANT, AND A. YOUNG, eds., *Proceedings of the Second International Workshop on Statistical Atlases and Computational Models of the Heart (STACOM): Imaging and Modelling Challenges*, vol. 7085 of Lecture Notes in Computer Science, Springer, 2011. CITED ON PAGE 11.
- 28] O. CAMARA, M. POP, K. RHODE, M. SERMESANT, N. SMITH, AND A. YOUNG, eds., *Proceedings of the First International Workshop on Statistical Atlases and Computational Models of the Heart (STACOM)*, vol. 6364 of Lecture Notes in Computer Science, Springer, 2010. CITED ON PAGE 11.
- 29] J. CARBALLIDO-GAMIO AND S. MAJUMDAR, *Atlas-based knee cartilage assessment*, Magnetic Resonance in Medicine, vol. 66, no. 2 (2011), pp. 575–581. CITED ON PAGE 95.
- 30] R. CHANDRASHEKARA, A. RAO, G.I. SANCHEZ-ORTIZ, R.H. MOHIADDIN, AND D. RUECKERT, *Construction of a statistical model for cardiac motion analysis using nonrigid image registration*, in Proc. 18th Int. Conf. Information Processing in Medical Imaging (IPMI), Ambleside, United Kingdom, vol. 2732 of Lecture Notes in Computer Science, 2003, pp. 599–610. CITED ON PAGE 45.

- 31] E.M. CHERRY AND F.H. FENTON, *Suppression of alternans and conduction blocks despite steep APD restitution: Electrotonic, memory, and conduction velocity restitution effects*, American Journal of Physiology - Heart and Circulatory Physiology, vol. 286, no. 6 (2004), pp. H2332–H2341. CITED ON PAGE 71.
- 32] ———, *A tale of two dogs: Analyzing two models of canine ventricular electrophysiology*, American Journal of Physiology - Heart and Circulatory Physiology, vol. 292, no. 1 (2007), pp. H43–H55. CITED ON PAGE 71.
- 33] Y. CHOI AND S. LEE, *Injectivity conditions of 2D and 3D uniform cubic B-spline functions*, Graphical Models, vol. 62, no. 6 (2000), pp. 411–427. CITED ON PAGE 20.
- 34] G.E. CHRISTENSEN, H.J. JOHNSON, AND M.W. VANNIER, *Synthesizing average 3D anatomical shapes*, NeuroImage, vol. 32, no. 1 (2006), pp. 146–158. CITED ON PAGE 14.
- 35] E. CHUANG AND C. BREGLER, *Mood swings: Expressive speech animation*, ACM Transactions on Graphics, vol. 24, no. 2 (2005), pp. 331–347. CITED ON PAGE 46.
- 36] R.H. CLAYTON AND A.V. PANFILOV, *A guide to modelling cardiac electrical activity in anatomically detailed ventricles*, Progress in Biophysics & Molecular Biology, vol. 96, no. 1-3 (2008), pp. 19–43. CITED ON PAGES 94 AND 97.
- 37] P. COLLI FRANZONE, L. GUERRI, M. PENNACCHIO, AND B. TACCARDI, *Spread of excitation in 3-D models of the anisotropic cardiac tissue. II. Effects of fiber architecture and ventricular geometry*, Mathematical Bioscience, vol. 147, no. 2 (1998), pp. 131–171. CITED ON PAGE 76.
- 38] T.F. COOTES, D.H. COOPER, C.J. TAYLOR, AND J. GRAHAM, *Trainable method of parametric shape description*, Image & Vision Computing, vol. 10, no. 5 (1992), pp. 289–294. CITED ON PAGE 45.
- 39] T.F. COOTES, G.J. EDWARDS, AND C.J. TAYLOR, *Active appearance models*, IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 21, no. 6 (2001), pp. 681–685. CITED ON PAGES 45 AND 78.
- 40] T.F. COOTES, A. HILL, C.J. TAYLOR, AND J. HASLAM, *The use of active shape models for locating structures in medical images*, Image & Vision Computing, vol. 12, no. 6 (1994), pp. 355–366. CITED ON PAGE 78.
- 41] T.F. COOTES AND C.J. TAYLOR, *A mixture model for representing shape variation*, Image & Vision Computing, vol. 17, no. 8 (1999), pp. 567–573. CITED ON PAGE 96.
- 42] T.F. COOTES, C.J. TAYLOR, D.H. COOPER, AND J. GRAHAM, *Active shape models—their training and application*, Computer Vision and Image Understanding, vol. 61, no. 1 (1995), pp. 38–59. CITED ON PAGES 45, 47, AND 50.
- 43] T.F. COOTES, C.J. TWINING, K.O. BABALOLA, AND C.J. TAYLOR, *Diffeomorphic statistical shape models*, Image & Vision Computing, vol. 26, no. 3 (2008), pp. 326–332. CITED ON PAGE 20.
- 44] F. CUZZOLIN, *Using bilinear models for view-invariant action and identity recognition*, in Proc. IEEE Int. Conf. on Computer Vision and Pattern Recognition (CVPR), New York, NY, USA, vol. 2, 2006, pp. 1701–1708. CITED ON PAGE 46.
- 45] M. DE CRAENE, A. DU BOIS D’AISCHE, B. MACQ, AND S.K. WARFIELD, *Multi-subject registration for unbiased statistical atlas construction*, in Proc. 7th Int. Conf. Medical Image Computing and Computer Assisted Intervention (MICCAI), Rennes/St. Malo, France, vol. 3216 of Lecture Notes in Computer Science, 2004, pp. 655–662. CITED ON PAGE 15.
- 46] M. DE CRAENE, G. PIELLA, O. CAMARA, N. DUCHATEAU, E. SILVA, A. DOLTRA, J. D’HOOGHE, J. BRUGADA, M. SITGES, AND A.F. FRANGL, *Temporal diffeomorphic free-form deformation: Application to motion and strain estimation from 3D echocardiography*, Medical Image Analysis, vol. 16, no. 2 (2012), pp. 427–450. CITED ON PAGES 42 AND 96.
- 47] L. DE LATHAUWER, B. DE MOOR, AND J. VANDEWALLE, *A multilinear singular value decomposition*, SIAM Journal on Matrix Analysis and Applications, vol. 21, no. 4 (2000), pp. 1253–1278. CITED ON PAGES 46 AND 52.
- 48] L.M.D. DELBRIDGE AND K.P. ROOS, *Optical methods to evaluate the contractile function of unloaded isolated cardiac myocytes*, Journal of Molecular and Cellular Cardiology, vol. 29, no. 1 (1997), pp. 11–25. CITED ON PAGE 4.
- 49] S.E. DILLARD, J.F. BINGERT, D. THOMA, AND B. HAMANN, *Construction of simplified boundary surfaces from serial-sectioned metal micrographs*, IEEE Transactions on Visualization and Computer Graphics, vol. 13, no. 6 (2007), pp. 1528–1535. CITED ON PAGE 25.

- 50] D.L. DONOHO AND C. GRIMES, *Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data*, Proceedings of the National Academy of Sciences of the United States of America, vol. 100, no. 10 (2003), pp. 5591–5596. CITED ON PAGE 78.
- 51] V. D’OTREPPE, R. BOMAN, AND J.P. PONTHOT, *Generating smooth surface meshes from multi-region medical images*, International Journal for Numerical Methods in Biomedical Engineering, vol. 28, no. 6/7 (2012), pp. 642–660. CITED ON PAGE 25.
- 52] I.L. DRYDEN AND K.V. MARDIA, *Statistical Shape Analysis*, Wiley, 1998. CITED ON PAGE 50.
- 53] J.S. DUNCAN AND N. AYACHE, *Medical image analysis: Progress over two decades and the challenges ahead*, IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 22, no. 1 (2000), pp. 85–106. CITED ON PAGE 45.
- 54] R. ERBEL AND H. EGGBRECHT, *Aortic dimensions and the risk of dissection*, Heart, vol. 92, no. 1 (2006), pp. 137–142. CITED ON PAGE 28.
- 55] EU VPH INITIATIVE, *Virtual Physiological Human Network of Excellence*. CITED ON PAGES 5 AND 11.
- 56] EUROPEAN HEART NETWORK AND EUROPEAN SOCIETY OF CARDIOLOGY, *European cardiovascular disease statistics*, Sept. 2012. CITED ON PAGE 4.
- 57] R.M. FIGUERAS I VENTURA, C. HOOGENDOORN, F.M. SUKNO, AND A.F. FRANGI, *Bilinear point distribution models for heart motion analysis*, in Proc. IEEE Int. Symp. on Biomedical Imaging (ISBI), Rotterdam, The Netherlands, 2010, pp. 476–479. CITED ON PAGE 30.
- 58] P.T. FLETCHER, S. VENKATASUBRAMANIAN, AND S. JOSHI, *The geometric median on Riemannian manifolds with application to robust atlas estimation*, NeuroImage, vol. 45, no. 1 (2009), pp. S143–S152. CITED ON PAGE 15.
- 59] C.G. FONSECA, M. BACKHAUS, D.A. BLUEMKE, R.D. BRITTEN, J.D. CHUNG, B.R. COWAN, I.D. DINOVI, J.P. FINN, P.J. HUNTER, A.H. KADISH, D.C. LEE, J.A.C. LIMA, P. MEDRANO-GRACIA, K. SHIVKUMAR, A. SUINESIAPUTRA, W. TAO, AND A.A. YOUNG, *The cardiac atlas project—an imaging database for computational modeling and statistical atlases of the heart*, Bioinformatics, vol. 27, no. 16 (2011), pp. 2288–2295. CITED ON PAGES 13, 16, AND 95.
- 60] A.F. FRANGI, W.J. NIESSEN, M.A. VIERGEVER, AND B.P.F. LELIEVELDT, *A survey of three-dimensional modeling techniques for quantitative functional analysis of cardiac images*, in Advanced Image Processing in Magnetic Resonance Imaging, L. Landini, V. Positano, and M.F. Santarelli, eds., CRC Press, 2005, ch. 9, pp. 267–342. CITED ON PAGE 46.
- 61] A.F. FRANGI, D. RUECKERT, J.A. SCHNABEL, AND W.J. NIESSEN, *Automatic construction of multiple-object three-dimensional statistical shape models: Application to cardiac modeling*, IEEE Transactions on Medical Imaging, vol. 21, no. 9 (2002), pp. 1151–1166. CITED ON PAGES 13 AND 14.
- 62] J. GONZÁLEZ-MORA, F. DE LA TORRE, R. MURTHI, N. GUIL, AND E.L. ZAPATA, *Bilinear active appearance models*, in Proc. Int. Workshop on Non-rigid Registration and Tracking through Learning (NRTL), Rio de Janeiro, Brazil, 2007. CITED ON PAGE 46.
- 63] C. GOODALL, *Procrustes methods in shape analysis*, Journal of the Royal Statistical Society Series B: Statistical Methodology, vol. 53, no. 2 (1991), pp. 285–339. CITED ON PAGE 56.
- 64] M.A. GREPL, *Model order reduction of parametrized nonlinear reaction-diffusion systems*, Computers & Chemical Engineering, vol. 43 (2012), pp. 33–44. CITED ON PAGES 73 AND 97.
- 65] D.B. GRIMES AND R.P.N. RAO, *Bilinear sparse coding for invariant vision*, Neural Computation, vol. 17, no. 1 (2005), pp. 47–73. CITED ON PAGE 46.
- 66] A. GUIMOND, J. MEUNIER, AND J.P. THIRION, *Average brain models: A convergence study*, Computer Vision and Image Understanding, vol. 77, no. 2 (2000), pp. 192–210. CITED ON PAGES 14, 15, 21, AND 37.
- 67] B. HALL, V. JEEVANANTHAM, R. SIMON, J. FILIPPONE, G. VOROBIOF, AND J. DAUBERT, *Variation in left atrial transmural wall thickness at sites commonly targeted for ablation of atrial fibrillation*, Journal of Interventional Cardiac Electrophysiology, vol. 17, no. 2 (2006), pp. 127–132. CITED ON PAGE 28.
- 68] G. HAMARNEH AND T. GUSTAVSSON, *Deformable spatio-temporal shape models: Extending active shape models to 2D + time*, Image & Vision Computing, vol. 22, no. 6 (2004), pp. 461–470. CITED ON PAGES 45 AND 50.

- 69] J. HAMM, D.H. YE, R. VERMA, AND C. DAVATZIKOS, *GRAM: A framework for geodesic registration on anatomical manifolds*, *Medical Image Analysis*, vol. 14, no. 5 (2010), pp. 633–642. CITED ON PAGES 15, 16, AND 41.
- 70] H. HEGE, M. SEEBASS, D. STALLING, AND M. ZÖCKLER, *A generalized marching cubes algorithm based on non-binary classifications*, Tech. Report ZIB-Report SC 97-05, Konrad-Zuse-Zentrum für Informationstechnik Berlin, 1997. CITED ON PAGE 25.
- 71] E.A. HEIDENREICH, J.M. FERRERO, M. DOBLARÉ, AND J.F. RODRÍGUEZ, *Adaptive macro finite elements for the numerical solution of monodomain equations in cardiac electrophysiology*, *Annals of Biomedical Engineering*, vol. 38, no. 7 (2010), pp. 2331–2345. CITED ON PAGES 73 AND 81.
- 72] T. HEIMANN AND H.P. MEINZER, *Statistical shape models for 3D medical image segmentation: A review*, *Medical Image Analysis*, vol. 13, no. 4 (2009), pp. 543–563. CITED ON PAGES 14 AND 95.
- 73] G.T. HERMAN, J. ZHENG, AND C.A. BUCHOLTZ, *Shape-based interpolation*, *IEEE Computer Graphics and Applications*, vol. 12, no. 3 (1992), pp. 69–79. CITED ON PAGE 25.
- 74] M. HOLDEN, *A review of geometric transformations for nonrigid body registration*, *IEEE Transactions on Medical Imaging*, vol. 27, no. 1 (2008), pp. 111–128. CITED ON PAGE 16.
- 75] C. HOOGENDOORN, N. DUCHATEAU, D. SÁNCHEZ-QUINTANA, T. WHITMARSH, M. DE CRAENE, F.M. SUKNO, K. LEKADIR, AND A.F. FRANGI, *A high-resolution atlas and statistical model of the human heart from multislice CT*, *IEEE Transactions on Medical Imaging*, vol. 32, no. 1 (2013), pp. 28–44. CITED ON PAGES 9, 74, AND 78.
- 76] C. HOOGENDOORN, A. PASHAEI, R. SEBASTIÁN, F.M. SUKNO, O. CÁMARA, AND A.F. FRANGI, *Influence of geometric variations on LV activation times: A study on an atlas-based virtual population*, in Proc. Workshop Statistical Atlases and Computational Models of the Heart (STACOM), Beijing, China, vol. 6364 of Lecture Notes in Computer Science, 2010, pp. 242–251. CITED ON PAGES 30, 42, 73, AND 77.
- 77] ———, *Sensitivity analysis of mesh warping and subsampling strategies for generating large scale electrophysiological simulation data*, in Proc. 6th Int. Conf. on Functional Imaging and Modeling of the Heart (FIMH), New York, NY, USA, vol. 6666 of Lecture Notes in Computer Science, 2011, pp. 418–426. CITED ON PAGES 42, 73, AND 80.
- 78] C. HOOGENDOORN, F.M. SUKNO, S. ORDÁS, AND A.F. FRANGI, *Bilinear models for spatio-temporal point distribution analysis: Application to extrapolation of whole heart cardiac dynamics*, in Proc. 8th Int. Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA), Rio de Janeiro, Brazil, 2007. CITED ON PAGE 45.
- 79] ———, *Bilinear models for spatio-temporal point distribution analysis: Application to extrapolation of left ventricular, biventricular and whole heart cardiac dynamics*, *International Journal of Computer Vision*, vol. 85, no. 3 (2009), pp. 237–252. CITED ON PAGES 14, 28, 30, AND 43.
- 80] C. HOOGENDOORN, T. WHITMARSH, N. DUCHATEAU, F.M. SUKNO, M. DE CRAENE, AND A.F. FRANGI, *A groupwise mutual information metric for cost efficient selection of a suitable reference in cardiac computational atlas construction*, in Proc. SPIE Medical Imaging, San Diego, CA, USA, vol. 7623, 2010. CITED ON PAGES 14, 18, AND 19.
- 81] D.A. HOOKS, K.A. TOMLINSON, S.G. MARSDEN, I.J. LEGRICE, B.H. SMAILL, A.J. PULLAN, AND P.J. HUNTER, *Cardiac microstructure: Implications for electrical propagation and defibrillation in the heart*, *Circulation Research*, vol. 91, no. 4 (2002), pp. 331–338. CITED ON PAGE 77.
- 82] E. HSU, K. PULLI, AND J. POPOVIĆ, *Style translation for human motion*, in Proc. 32nd ACM SIGGRAPH Int. Conf. on Computer Graphics and Interactive Techniques, Los Angeles, CA, USA, 2005, pp. 1082–1089. CITED ON PAGE 46.
- 83] T. IJIRI, T. ASHIHARA, T. YAMAGUCHI, K. TAKAYAMA, T. IGARASHI, T. SHIMADA, T. NAMBA, R. HARAGUCHI, AND K. NAKAZAWA, *A procedural method for modeling the Purkinje fibers of the heart*, *The Journal of Physiological Sciences*, vol. 58, no. 7 (2008), pp. 481–486. CITED ON PAGE 77.
- 84] INTERNATIONAL UNION OF PHYSIOLOGICAL SCIENCES, *IUPS Physiome Project*. CITED ON PAGES 5, 6, AND 11.
- 85] I. IŞGUM, M. STARING, A. RUTTEN, M. PROKOP, M.A. VIERGEVER, AND B. VAN GINNEKEN, *Multi-atlas-based segmentation with local decision fusion—application to cardiac and aortic segmentation in CT scans*, *IEEE Transactions on Medical Imaging*, vol. 28, no. 7 (2009), pp. 1000–1010. CITED ON PAGE 13.

- 86] V. IYER, R. MAZHARI, AND R.L. WINSLOW, *A computational model of the human left-ventricular epicardial myocyte*, *Biophysical Journal*, vol. 87, no. 3 (2004), pp. 1507–1525. CITED ON PAGE 91.
- 87] Y. JEONG AND R.J. RADKE, *Modeling inter- and intra-patient anatomical variation using a bilinear model*, in *Proc. 7th Int. Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA)*, New York, NY, USA, 2006, pp. 76–83. CITED ON PAGE 46.
- 88] Y. JEONG, R.J. RADKE, AND D.M. LOVELOCK, *Bilinear models for inter- and intra-patient variation of the prostate*, *Physics in Medicine and Biology*, vol. 55, no. 13 (2010), pp. 3725–3739. CITED ON PAGES 46 AND 96.
- 89] H. JIA, G. WU, Q. WANG, AND D. SHEN, *ABSORB: Atlas building by self-organized registration and bundling*, *NeuroImage*, vol. 51, no. 3 (2010), pp. 1057–1070. CITED ON PAGES 15 AND 16.
- 90] H. JIA, P.T. YAP, G. WU, Q. WANG, AND D. SHEN, *Intermediate templates guided groupwise registration of diffusion tensor images*, *NeuroImage*, vol. 54, no. 2 (2011), pp. 928–939. CITED ON PAGE 16.
- 91] B. JIANG, A. STRUTHERS, Z. SUN, Z. FENG, X. ZHAO, K. ZHAO, W. DAI, X. ZHOU, M.E. BERENS, AND L. ZHANG, *Employing graphics processing unit technology, alternating direction implicit method and domain decomposition to speed up the numerical diffusion solver for the biomedical engineering research*, *International Journal for Numerical Methods in Biomedical Engineering*, vol. 27, no. 11 (2011), pp. 1829–1849. CITED ON PAGE 73.
- 92] I.T. JOLLIFFE, *Principal Component Analysis*, Springer-Verlag, New York, second ed., 2002. CITED ON PAGE 28.
- 93] C. JONGEN, J.P.W. PLUIM, P.J. NEDERKOORN, M.A. VIERGEVER, AND W.J. NIESSEN, *Construction and evaluation of an average CT brain image for inter-subject registration*, *Computers in Biology and Medicine*, vol. 34, no. 8 (2004), pp. 647–662. CITED ON PAGE 15.
- 94] S. JOSHI, B. DAVIS, M. JOMIER, AND G. GERIG, *Unbiased diffeomorphic atlas construction for computational anatomy*, *NeuroImage*, vol. 23, no. S1 (2004), pp. S151–S160. CITED ON PAGE 15.
- 95] A.H. KADISH, D. BELLO, J.P. FINN, R.O. BONOW, A. SCHAECHTER, H. SUBACIUS, C. ALBERT, J.P. DAUBERT, C.G. FONSECA, AND J.J. GOLDBERGER, *Rationale and design for the defibrillators to reduce risk by magnetic resonance imaging evaluation (DETERMINE) trial*, *Journal of Cardiovascular Electrophysiology*, vol. 20, no. 9 (2009), pp. 982–987. CITED ON PAGES 16 AND 95.
- 96] M.R. KAUS, V. PEKAR, C. LORENZ, R. TRUYEN, S. LOBREGT, AND J. WEESE, *Automated 3-D PDM construction from segmented images using deformable models*, *IEEE Transactions on Medical Imaging*, vol. 22, no. 8 (2003), pp. 1005–1013. CITED ON PAGE 14.
- 97] D.G. KENDALL, *Shape manifolds, procrustean metrics and complex projective spaces*, *Bulletin of the London Mathematical Society*, vol. 16, no. 2 (1984), pp. 81–121. CITED ON PAGE 45.
- 98] H. KIRIŞLI, M. SCHAAP, S. KLEIN, T. VAN WALSUM, S.L. PAPADOPOULOU, A.C. WEUSTINK, N.R. MOLLET, M. BONARDI, C.H. CHEN, E.J. VONKEN, R.J. VAN DER GEEST, AND W.J. NIESSEN, *Evaluation of a multi-atlas based method for segmentation of cardiac CTA data: A large-scale, multicenter and multivendor study*, *Medical Physics*, vol. 37, no. 12 (2010), pp. 6279–6291. CITED ON PAGE 13.
- 99] KITWARE INC., *The Insight Segmentation and Registration Toolkit*. CITED ON PAGE 20.
- 100] ———, *The Visualization Toolkit*. CITED ON PAGE 81.
- 101] S. KLEIN, M. STARING, AND J.P.W. PLUIM, *Evaluation of optimization methods for nonrigid medical image registration using mutual information and B-splines*, *IEEE Transactions on Image Processing*, vol. 16, no. 12 (2007), pp. 2879–2890. CITED ON PAGES 16 AND 20.
- 102] J. KOIKALAINEN, T. TÖLLI, K. LAUERMA, K. ANTILA, E. MATTILA, M. LILJA, AND J. LÖTJÖNEN, *Methods of artificial enlargement of the training set for statistical shape models*, *IEEE Transactions on Medical Imaging*, vol. 27, no. 11 (2008), pp. 1643–1654. CITED ON PAGE 14.
- 103] J.W. KRUEGER, D. FORLETTI, AND B.A. WITTENBERG, *Uniform sarcomere shortening behavior in isolated cardiac muscle cells*, *Journal of General Physiology*, vol. 76, no. 5 (1980), pp. 587–607. CITED ON PAGE 4.
- 104] P. LAFORTUNE, R. ARÍS, M. VÁZQUEZ, AND G. HOUZEAUX, *Coupled electromechanical model of the heart: Parallel finite element formulation*, *International Journal for Numerical Methods in Biomedical Engineering*, vol. 28, no. 1 (2012), pp. 72–86. CITED ON PAGE 73.

- 105] P. LAMATA, S. NIEDERER, D. BARBER, D. NORSLETEN, J. LEE, R. HOSE, AND N. SMITH, *Personalization of cubic hermite meshes for efficient biomechanical simulations*, in Proc. 13th Int. Conf. Medical Image Computing and Computer Assisted Intervention (MICCAI), Beijing, China, vol. 6362 of Lecture Notes in Computer Science, 2010, pp. 380–387. CITED ON PAGE 73.
- 106] P. LAMATA, S. NIEDERER, D. NORDSLETEN, D.C. BARBER, I. ROY, D.R. HOSE, AND N. SMITH, *An accurate, fast and robust method to generate patient-specific cubic Hermite meshes*, Medical Image Analysis, vol. 15, no. 6 (2011), pp. 801–813. CITED ON PAGE 73.
- 107] I. LARRABIDE, P. OMEDAS, Y. MARTELLI, X. PLANES, M. NIEBER, J.A. MOYA, C. BUTAKOFF, R. SEBASTIÁN, O. CAMARA, M. DE CRAENE, B.H. BIJNENS, AND A.F. FRANGI, *GIMIAS: An open source framework for efficient development of research tools and clinical prototypes*, in Proc. 5th Int. Conf. on Functional Imaging and Modeling of the Heart (FIMH), Nice, France, vol. 5528 of Lecture Notes in Computer Science, 2009, pp. 417–426. CITED ON PAGES 25, 33, AND 36.
- 108] B. LASSEN, E.M. VAN RIKKOORT, M. SCHMIDT, S. KERKSTRA, B. VAN GINNEKEN, AND J.M. KUHNIGK, *Automatic segmentation of the pulmonary lobes from chest CT scans based on fissures, vessels, and bronchi*, IEEE Transactions on Medical Imaging, vol. 32, no. 2 (2013), pp. 210–222. CITED ON PAGE 95.
- 109] H. LE AND D.G. KENDALL, *The Riemannian structure of euclidean shape spaces: A novel environment for statistics*, Annals of Statistics, vol. 21, no. 3 (1993), pp. 1225–1271. CITED ON PAGE 45.
- 110] C.S. LEE AND A. ELGAMMAL, *Gait style and gait content: Bilinear models for gait recognition using gait resampling*, in Proc. 6th IEEE Int. Conf. on Automatic Face and Gesture Recognition (FGR), Seoul, Korea, 2004, pp. 147–152. CITED ON PAGE 46.
- 111] K. LEKADIR, B. GHAFARYASL, E. MUÑOZ-MORENO, C. BUTAKOFF, C. HOOGENDOORN, AND A.F. FRANGI, *Predictive modeling of cardiac fiber orientation using the Knutsson mapping*, in Proc. 14th Int. Conf. Medical Image Computing and Computer Assisted Intervention (MICCAI), Toronto, ON, Canada, vol. 6892 of Lecture Notes in Computer Science, 2011, pp. 50–57. CITED ON PAGES 42 AND 91.
- 112] K. LEKADIR, N. KEENAN, D. PENNELL, AND G.Z. YANG, *Shape-based myocardial contractility analysis using multivariate outlier detection*, in Proc. 10th Int. Conf. Medical Image Computing and Computer Assisted Intervention (MICCAI), Brisbane, QLD, Australia, vol. 4792 of Lecture Notes in Computer Science, 2007, pp. 834–841. CITED ON PAGE 47.
- 113] K. LEKADIR, A. PASHAEI, C. HOOGENDOORN, M. PEREAÑEZ, F. POVEDA, AND A.F. FRANGI, *Personalized modeling of cardiac electrophysiology using shape-based prediction of fiber orientation*, in Proc. Workshop Statistical Atlases and Computational Models of the Heart (STACOM), Nagoya, Japan, vol. 8330 of Lecture Notes in Computer Science, 2013, pp. 196–203. CITED ON PAGE 91.
- 114] K.Y.E. LEUNG AND J.G. BOSCH, *Localized shape variations for classifying wall motion in echocardiograms*, in Proc. 10th Int. Conf. Medical Image Computing and Computer Assisted Intervention (MICCAI), Brisbane, QLD, Australia, vol. 4791 of Lecture Notes in Computer Science, 2007, pp. 52–59. CITED ON PAGE 47.
- 115] A. LINDENMAYER, *Mathematical models for cellular interactions in development I. Filaments with one-sided inputs*, Journal of Theoretical Biology, vol. 18, no. 3 (1968), pp. 280–299. CITED ON PAGE 77.
- 116] F. LIONETTI, A. MCCULLOCH, AND S. BADEN, *GPU accelerated solvers for ODEs describing cardiac membrane equations*. NVIDIA GPU Technology Conference, 2009. CITED ON PAGE 73.
- 117] H. LIU AND P. SHI, *State-space analysis of cardiac motion with biomechanical constraints*, IEEE Transactions on Image Processing, vol. 16, no. 4 (2007), pp. 901–916. CITED ON PAGE 47.
- 118] H. LOMBAERT, J.M. PEYRAT, P. CROISILLE, S. RAPACCHI, L. FANTON, F. CHERIET, P. CLARYSSE, I. MAGNIN, H. DELINGETTE, AND N. AYACHE, *Human atlas of the cardiac fiber architecture: Study on a healthy population*, IEEE Transactions on Medical Imaging, vol. 31, no. 7 (2012), pp. 1436–1447. CITED ON PAGES 13 AND 91.
- 119] W.E. LORENSEN AND H.E. CLINE, *Marching cubes: A high resolution 3D surface construction algorithm*, in Proc. 14th ACM SIGGRAPH Int. Conf. on Computer Graphics and Interactive Techniques, Anaheim, CA, USA, 1987, pp. 163–169. CITED ON PAGE 25.
- 120] C. LORENZ AND J. VON BERG, *A comprehensive shape model of the heart*, Medical Image Analysis, vol. 10, no. 4 (2006), pp. 657–670. CITED ON PAGE 14.

- 121] P. LORENZEN, M. PRASTAWA, B. DAVIS, G. GERIG, E. BULLITT, AND S. JOSHI, *Multi-modal image set registration and atlas formation*, *Medical Image Analysis*, vol. 10, no. 3 (2006), pp. 440–451. CITED ON PAGE 15.
- 122] M. LORENZO-VALDÉS, G.I. SANCHEZ-ORTIZ, A.G. ELKINGTON, R. MOHIADDIN, AND D. RUECKERT, *Segmentation of 4D cardiac MR images using a probabilistic atlas and the EM algorithm*, *Medical Image Analysis*, vol. 8, no. 3 (2004), pp. 255–265. CITED ON PAGES 13 AND 42.
- 123] J. LÖTJÖNEN, S. KIVISTÖ, J. KOIKKALAINEN, D. SMUTEK, AND K. LAUERMA, *Statistical shape model of atria, ventricles and epicardium from short- and long-axis MR images*, *Medical Image Analysis*, vol. 8, no. 3 (2004), pp. 371–386. CITED ON PAGE 13.
- 124] M. LYNCH, O. GHITA, AND P. WHELAN, *Segmentation of the left ventricle of the heart in 3-D+t MRI data using an optimised non-rigid temporal model*, *IEEE Transactions on Medical Imaging*, vol. 27, no. 2 (2008), pp. 195–203. CITED ON PAGES 45 AND 46.
- 125] F. MAES, A. COLLIGNON, D. VANDERMEULEN, G. MARCHAL, AND P. SUETENS, *Multimodality image registration by maximization of mutual information*, *IEEE Transactions on Medical Imaging*, vol. 16, no. 2 (1997), pp. 187–198. CITED ON PAGES 18, 20, AND 96.
- 126] J.R. MAGNUS AND H. NEUDECKER, *Matrix Differential Calculus With Applications in Statistics and Econometrics*, Wiley, 1988. CITED ON PAGE 51.
- 127] J.B.A. MAINTZ AND M.A. VIERGEVER, *A survey of medical image registration*, *Medical Image Analysis*, vol. 2, no. 1 (1998), pp. 1–36. CITED ON PAGE 16.
- 128] T. MÄKELÄ, P. CLARYSSE, O. SIPILÄ, N. PAUNA, Q.C. PHAM, T. KATILA, AND I.E. MAGNIN, *A review of cardiac image registration methods*, *IEEE Transactions on Medical Imaging*, vol. 21, no. 9 (2002), pp. 1011–1021. CITED ON PAGE 16.
- 129] K.V. MARDIA AND I.L. DRYDEN, *Shape distributions for landmark data*, *Advances in Applied Probability*, vol. 21, no. 4 (1989), pp. 742–755. CITED ON PAGE 45.
- 130] D.H. MARIMONT AND B.A. WANDELL, *Linear models of surface and illuminant spectra*, *Journal of the Optical Society of America A: Optics, Image Science, and Vision*, vol. 9, no. 11 (1992), pp. 1905–1913. CITED ON PAGE 51.
- 131] D.B. MARK, F.J. VAN DE WERF, R.J. SIMES, H.D. WHITE, L.C. WALLENTIN, R.M. CALIFF, AND P.W. ARMSTRONG, *Cardiovascular disease on a global scale: Defining the path forward for research and practice*, *European Heart Journal*, vol. 28, no. 21 (2007), pp. 2678–2684. CITED ON PAGE 4.
- 132] E.M. MAROM, J.E. HERNDON, Y.H. KIM, AND H.P. MCADAMS, *Variations in pulmonary venous drainage to the left atrium: Implications for radiofrequency ablation*, *Radiology*, vol. 230 (2004), pp. 824–829. CITED ON PAGE 41.
- 133] S. MARSLAND, C.J. TWING, AND C.J. TAYLOR, *Groupwise non-rigid registration using polyharmonic clamped-plate splines*, in *Proc. 6th Int. Conf. Medical Image Computing and Computer Assisted Intervention (MICCAI)*, Montréal, QC, Canada, vol. 2879 of *Lecture Notes in Computer Science*, 2003, pp. 771–779. CITED ON PAGE 15.
- 134] ———, *A minimum description length objective function for groupwise non-rigid image registration*, *Image & Vision Computing*, vol. 26, no. 3 (2008), pp. 333–346. CITED ON PAGE 15.
- 135] D. MATTES, D.R. HAYNOR, H. VESSELLE, T.K. LEWELLEN, AND W. EUBANK, *PET-CT image registration in the chest using free-form deformations*, *IEEE Transactions on Medical Imaging*, vol. 22, no. 1 (2003), pp. 120–128. CITED ON PAGES 20 AND 96.
- 136] T. MCINERNEY AND D. TERZOPOULOS, *Deformable models in medical image analysis: A survey*, *Medical Image Analysis*, vol. 1, no. 2 (1996), pp. 91–108. CITED ON PAGE 45.
- 137] P. MEDRANO-GRACIA, B.R. COWAN, J.P. FINN, C.G. FONSECA, A.H. KADISH, D.C. LEE, W. TAO, AND A.A. YOUNG, *The Cardiac Atlas Project: Preliminary description of heart shape in patients with myocardial infarction*, in *Proc. Workshop Statistical Atlases and Computational Models of the Heart (STACOM)*, Beijing, China, 2010, pp. 46–53. CITED ON PAGE 13.
- 138] C.T. METZ, S. KLEIN, M. SCHAAP, T. VAN WALSUM, AND W.J. NIESSEN, *Nonrigid registration of dynamic medical imaging data using $nd + t$ B-splines and a groupwise optimization approach*, *Medical Image Analysis*, vol. 15, no. 2 (2011), pp. 238–249. CITED ON PAGES 42 AND 96.
- 139] M. MEYER, R. WHITAKER, R.M. KIRBY, C. LEDERGERBER, AND H. PFISTER, *Particle-based sampling and meshing of surfaces in multimaterial volumes*, *IEEE Transactions on Visualization and Computer Graphics*, vol. 14, no. 6 (2008), pp. 1539–1546. CITED ON PAGE 25.

- 140] S.K. MICHPOULOU, L. COSTARIDOU, E. PANAGIOTOPOULOS, R. SPELLER, G. PANAYIOTAKIS, AND A. TODD-POKROPEK, *Atlas-based segmentation of degenerated lumbar intervertebral discs from MR images of the spine*, IEEE Transactions on Biomedical Engineering, vol. 56, no. 9 (2009), pp. 2225–2231. CITED ON PAGE 95.
- 141] R. MILANI, A. QUARTERONI, AND G. ROZZA, *Reduced basis method for linear elasticity problems with many parameters*, Computer Methods in Applied Mechanics and Engineering, vol. 197, no. 51/52 (2008), pp. 4812–4829. CITED ON PAGES 73 AND 97.
- 142] J.J. MIRANDA, S. KINRA, J.P. CASAS, G.D. SMITH, AND S. EBRAHIM, *Non-communicable diseases in low- and middle-income countries: Context, determinants and health policy*, Tropical Medicine & International Health, vol. 13, no. 10 (2008), pp. 1225–1234. CITED ON PAGE 4.
- 143] S.C. MITCHELL, J.G. BOSCH, B.P.F. LELIEVELDT, R.J. VAN DER GEEST, J.H.C. REIBER, AND M. SONKA, *3-D active appearance models: Segmentation of cardiac MR and ultrasound images*, IEEE Transactions on Medical Imaging, vol. 21, no. 9 (2002), pp. 1167–1178. CITED ON PAGES 45 AND 46.
- 144] J. MONTAGNAT AND H. DELINGETTE, *4D deformable models with temporal constraints: Application to 4D cardiac image segmentation*, Medical Image Analysis, vol. 9, no. 1 (2005), pp. 87–100. CITED ON PAGES 45 AND 46.
- 145] D. NOBLE, *A modification of the Hodgkin–Huxley equations applicable to Purkinje fibre action and pacemaker potentials*, The Journal of Physiology, vol. 160, no. 2 (1962), pp. 317–352. CITED ON PAGE 94.
- 146] D. NOBLE AND Y. RUDY, *Models of cardiac ventricular action potentials: Iterative interaction between experiment and simulation*, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, vol. 359, no. 1783 (2001), pp. 1127–1142. CITED ON PAGE 94.
- 147] A.K. NOOR AND J.M. PETERS, *Reduced basis technique for nonlinear analysis of structures*, AIAA Journal, vol. 18, no. 4 (1980), pp. 455–462. CITED ON PAGE 73.
- 148] B.M. OHNESORGE, C.R. BECKER, T.G. FLOHR, AND M.F. REISER, *Multi-slice CT in Cardiac Imaging: Technical Principles, Clinical Application and Future Developments*, Springer-Verlag, 2002. CITED ON PAGE 53.
- 149] S. ORDÁS, E. OUBEL, R. LETA, F. CARRERA, AND A.F. FRANGI, *A statistical shape model of the heart and its application to model-based segmentation*, in Proc. SPIE Medical Imaging, San Diego, CA, USA, vol. 6511, 2007. Article no. 65111K. CITED ON PAGES 14, 54, 67, AND 94.
- 150] M. PANTIC AND L.J.M. ROTHKRANTZ, *Automatic analysis of facial expressions: The state of the art*, IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 22, no. 12 (2000), pp. 1424–1445. CITED ON PAGE 45.
- 151] H. PARK, P.H. BLAND, A.O. HERO III, AND C.R. MEYER, *Least biased target selection in probabilistic atlas construction*, in Proc. 8th Int. Conf. Medical Image Computing and Computer Assisted Intervention (MICCAI), Palm Springs, CA, USA, vol. 3750 of Lecture Notes in Computer Science, 2005, pp. 419–426. CITED ON PAGE 14.
- 152] H. PARK, P.H. BLAND, AND C.R. MEYER, *Construction of an abdominal probabilistic atlas and its application in segmentation*, IEEE Transactions on Medical Imaging, vol. 22, no. 4 (2003), pp. 483–492. CITED ON PAGE 14.
- 153] A. PASHAEL, C. HOOGENDOORN, R. SEBASTIÁN, D. ROMERO, O. CÁMARA, AND A.F. FRANGI, *Infarction and fast electrophysiology of human heart: In-silico study on atlas-based virtual population*, in Proc. 6th Int. Conf. on Functional Imaging and Modeling of the Heart (FIMH), New York, NY, USA, vol. 6666 of Lecture Notes in Computer Science, 2011, pp. 427–436. CITED ON PAGES 42 AND 73.
- 154] D. PERPERIDIS, *Spatio-Temporal Registration and Modelling of the Heart Using Cardiovascular MR Imaging*, PhD thesis, Imperial College London, 2005. CITED ON PAGES 30, 45, AND 46.
- 155] D. PERPERIDIS, R. MOHIADDIN, AND D. RUECKERT, *Construction of a 4D statistical atlas of the cardiac anatomy and its use in classification*, in Proc. 8th Int. Conf. Medical Image Computing and Computer Assisted Intervention (MICCAI), Palm Springs, CA, USA, vol. 3750 of Lecture Notes in Computer Science, 2005, pp. 402–410. CITED ON PAGES 13, 28, AND 30.
- 156] D. PERPERIDIS, R.H. MOHIADDIN, AND D. RUECKERT, *Spatio-temporal free-form registration of cardiac MR image sequences*, Medical Image Analysis, vol. 9, no. 5 (2005), pp. 441–456. CITED ON PAGES 54 AND 55.

- 157] J. PETERS, O. ECABERT, C. MEYER, R. KNESER, AND J. WEESE, *Optimizing boundary detection via Simulated Search with applications to multi-modal heart segmentation*, *Medical Image Analysis*, vol. 14, no. 1 (2010), pp. 70–84. CITED ON PAGE 96.
- 158] J.M. PEYRAT, M. SERMESANT, X. PENNEC, H. DELINGETTE, C. XU, E.R. MCVEIGH, AND N. AYACHE, *A computational framework for the statistical analysis of cardiac diffusion tensors: Application to a small database of canine hearts*, *IEEE Transactions on Medical Imaging*, vol. 26, no. 11 (2007), pp. 1500–1514. CITED ON PAGE 13.
- 159] J. PITT-FRANCIS, P. PATHMANATHAN, M.O. BERNABEU, R. BORDAS, J. COOPER, A.G. FLETCHER, G.R. MIRAMS, P. MURRAY, J.M. OSBORNE, A. WALTER, S.J. CHAPMAN, A. GARNY, I.M.M. VAN LEEUWEN, P.K. MAINI, B. RODRÍGUEZ, S.L. WATERS, J.P. WHITELEY, H.M. BYRNE, AND D.J. GAVAGHAN, *Chaste: A test-driven approach to software development for biological modelling*, *Computer Physics Communications*, vol. 180, no. 12 (2009), pp. 2452–2471. CITED ON PAGE 73.
- 160] J.P.W. PLUIM, J.B.A. MAINTZ, AND M.A. VIERGEVER, *Mutual-information based registration of medical images: A survey*, *IEEE Transactions on Medical Imaging*, vol. 22, no. 8 (2003), pp. 986–1004. CITED ON PAGE 16.
- 161] A.R. PORRAS, G. PIELLA, A. BERRUEZO, C. HOOGENDOORN, D. ANDREU, M. SITGES, J. FERNANDEZ-ARMENTA, AND A.F. FRANGI, *Interventional endocardial motion estimation from electro-anatomical mapping data: Application to scar characterization*, *IEEE Transactions on Biomedical Engineering*, vol. 60, no. 5 (2013), pp. 1217–1224. CITED ON PAGE 30.
- 162] M. POTSE, B. DUBÉ, J. RICHER, A. VINET, AND R.M. GULRAJANI, *A comparison of monodomain and bidomain reaction-diffusion models for action potential propagation in the human heart*, *IEEE Transactions on Biomedical Engineering*, vol. 53, no. 12-1 (2006), pp. 2425–2435. CITED ON PAGE 81.
- 163] R. PRAKASH, *Determination of right ventricular wall thickness in systole and diastole*, *British Heart Journal*, vol. 40, no. 11 (1978), pp. 1257–1261. CITED ON PAGE 28.
- 164] L. PRIEBE AND D.J. BEUCKELMANN, *Simulation study of cellular electric properties in heart failure*, *Circulation Research*, vol. 82, no. 11 (1998), pp. 1206–1223. CITED ON PAGE 91.
- 165] A. QUARTERONI AND G. ROZZA, *Numerical solution of parametrized Navier-Stokes equations by reduced basis methods*, *Numerical Methods for Partial Differential Equations*, vol. 23, no. 4 (2007), pp. 923–948. CITED ON PAGES 73 AND 97.
- 166] S.P. RAYA AND J.K. UDUPA, *Shape-based interpolation of multidimensional objects*, *IEEE Transactions on Medical Imaging*, vol. 9, no. 1 (1990), pp. 32–42. CITED ON PAGE 25.
- 167] D.E. ROBERTS, L.T. HERSH, AND A.M. SCHER, *Influence of cardiac fiber orientation on wavefront voltage, conduction velocity, and tissue resistivity in the dog*, *Circulation Research*, vol. 44, no. 5 (1979), pp. 701–712. CITED ON PAGE 77.
- 168] D. ROMERO, R. SEBASTIAN, B.H. BIJNENS, V. ZIMMERMAN, P.M. BOYLE, E.J. VIGMOND, AND A.F. FRANGI, *Effects of the Purkinje system and cardiac geometry on biventricular pacing: A model study*, *Annals of Biomedical Engineering*, vol. 38, no. 4 (2010), pp. 1388–1398. CITED ON PAGE 76.
- 169] S.T. ROWEIS AND L.K. SAUL, *Nonlinear dimensionality reduction by locally linear embedding*, *Science*, vol. 290, no. 5500 (2000), pp. 2323–2326. CITED ON PAGE 78.
- 170] D. RUECKERT, P. ALJABAR, R.A. HECKEMANN, J.V. HAJNAL, AND A. HAMMERS, *Diffeomorphic registration using B-splines*, in *Proc. 9th Int. Conf. Medical Image Computing and Computer Assisted Intervention (MICCAI)*, Copenhagen, Denmark, vol. 4191 of *Lecture Notes in Computer Science*, 2006, pp. 702–709. CITED ON PAGES 15 AND 20.
- 171] D. RUECKERT, A.F. FRANGI, AND J.A. SCHNABEL, *Automatic construction of 3-D statistical deformation models of the brain using nonrigid registration*, *IEEE Transactions on Medical Imaging*, vol. 22, no. 8 (2003), pp. 1014–1025. CITED ON PAGE 20.
- 172] D. RUECKERT, L.I. SONODA, C. HAYES, D.L.G. HILL, M.O. LEACH, AND D.J. HAWKES, *Non-rigid registration using free-form deformations: Application to breast MR images*, *IEEE Transactions on Medical Imaging*, vol. 18, no. 8 (1999), pp. 712–721. CITED ON PAGES 15 AND 20.
- 173] M.R. SABUNCU, S.K. BALCI, M.E. SHENTON, AND P. GOLLAND, *Image-driven population analysis through mixture modeling*, *IEEE Transactions on Medical Imaging*, vol. 28, no. 9 (2009), pp. 1473–1487. CITED ON PAGE 15.

- 174] J.E. SAFFITZ, H.L. KANTER, K.G. GREEN, T.K. TOLLEY, AND E.C. BEYER, *Tissue-specific determinants of anisotropic conduction velocity in canine atrial and ventricular myocardium*, *Circulation Research*, vol. 74, no. 6 (1994), pp. 1065–1070. CITED ON PAGE 77.
- 175] K.J. SAMPSON, V. IYER, A.R. MARKS, AND R.S. KASS, *A computational model of Purkinje fibre single cell electrophysiology: Implications for the long QT syndrome*, *The Journal of Physiology*, vol. 588, no. 14 (2010), pp. 2643–2655. CITED ON PAGES 81, 94, AND 97.
- 176] L.K. SAUL AND S.T. ROWEIS, *Think globally, fit locally: Unsupervised learning of low dimensional manifolds*, *Journal of Machine Learning Research*, vol. 4 (2003), pp. 119–155. CITED ON PAGE 78.
- 177] R. SEBASTIAN, V. ZIMMERMAN, D. ROMERO, D. SANCHEZ-QUINTANA, AND A.F. FRANGI, *Characterization and modeling of the peripheral cardiac conduction system*, *IEEE Transactions on Medical Imaging*, vol. 32, no. 1 (2013), pp. 45–55. CITED ON PAGES 76 AND 77.
- 178] R. SEBASTIAN, V. ZIMMERMAN, F. SUKNO, B.H. BIJNENS, AND A.F. FRANGI, *Cardiac modelling for pathophysiology research and clinical applications: The need for an automated pipeline*, in *Proc. World Congr. Medical Physics and Biomedical Engineering*, Munich, Germany, vol. 25/4 of IFMBE Proceedings, 2009, pp. 2207–2210. CITED ON PAGE 77.
- 179] D. SHIN, H.S. LEE, AND D. KIM, *Illumination-robust face recognition using ridge regressive bilinear models*, *Pattern Recognition Letters*, vol. 29, no. 1 (2008), pp. 49–58. CITED ON PAGE 46.
- 180] I.A. SIGAL, M.R. HARDISTY, AND C.M. WHYNE, *Mesh-morphing algorithms for specimen-specific finite element modeling*, *Journal of Biomechanics*, vol. 41, no. 7 (2008), pp. 1381–1389. CITED ON PAGE 73.
- 181] I.A. SIGAL, H. YANG, M.D. ROBERTS, AND J.C. DOWNS, *Morphing methods to parameterize specimen-specific finite element geometries*, *Journal of Biomechanics*, vol. 43, no. 2 (2010), pp. 254–262. CITED ON PAGE 73.
- 182] A. SOTIRAS, C. DAVATZIKOS, AND N. PARAGIOS, *Deformable medical image registration: A survey*, *IEEE Transactions on Medical Imaging*, vol. 32, no. 7 (2013), pp. 1153–1190. CITED ON PAGES 16 AND 96.
- 183] P. STEWART, O.V. ASLANIDI, D. NOBLE, P.J. NOBLE, M.R. BOYETT, AND H. ZHANG, *Mathematical models of the electrical action potential of Purkinje fibre cells*, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 367, no. 1896 (2009), pp. 2225–2255. CITED ON PAGES 71, 81, 94, AND 97.
- 184] D.D. STREETER, JR., H.M. SPOTNITZ, D.P. PATEL, J. ROSS, JR., AND E.H. SONNENBLICK, *Fiber orientation in the canine left ventricle during diastole and systole*, *Circulation Research*, vol. 24, no. 3 (1969), pp. 339–347. CITED ON PAGES 4, 77, AND 91.
- 185] C. STUDHOLME AND V. CARDENAS, *A template free approach to volumetric spatial normalization of brain anatomy*, *Pattern Recognition Letters*, vol. 25, no. 10 (2004), pp. 1191–1202. CITED ON PAGE 15.
- 186] M.A. STYNER, K.T. RAJAMANI, L.P. NOLTE, G. ZSEMLYE, G. SZÉKELY, C.J. TAYLOR, AND R.H. DAVIES, *Evaluation of 3D correspondence methods for model building*, in *Proc. 18th Int. Conf. Information Processing in Medical Imaging (IPMI)*, Ambleside, United Kingdom, vol. 2732 of *Lecture Notes in Computer Science*, 2003, pp. 63–75. CITED ON PAGE 66.
- 187] J.S. SURI, *Computer vision, pattern recognition and image processing in left ventricle segmentation: The last 50 years*, *Pattern Analysis & Applications*, vol. 3, no. 3 (2000), pp. 209–242. CITED ON PAGE 46.
- 188] T. SYEDA-MAHMOOD, F. WANG, D. BEYMER, M. LONDON, AND R. REDDY, *Characterizing spatio-temporal patterns for disease discrimination in cardiac echo videos*, in *Proc. 10th Int. Conf. Medical Image Computing and Computer Assisted Intervention (MICCAI)*, Brisbane, QLD, Australia, vol. 4791 of *Lecture Notes in Computer Science*, 2007, pp. 261–269. CITED ON PAGE 47.
- 189] H. TANDRI, S.K. DAYA, K. NASIR, C. BOMMA, J.A.C. LIMA, H. CALKINS, AND D.A. BLUEMKE, *Normal reference values for the adult right ventricle by magnetic resonance imaging*, *American Journal of Cardiology*, vol. 98, no. 12 (2006), pp. 1660–1664. CITED ON PAGES 16 AND 95.
- 190] G. TAUBIN, T. ZHANG, AND G. GOLUB, *Optimal surface smoothing as filter design*, in *Proc. 4th European Conf. on Computer Vision (ECCV)*, Cambridge, United Kingdom, vol. 1064 of *Lecture Notes in Computer Science*, 1996, pp. 283–292. CITED ON PAGE 25.

- 191] V. TAVAKOLI AND A.A. AMINI, *A survey of shaped-based registration and segmentation techniques for cardiac images*, Computer Vision and Image Understanding, vol. 117, no. 9 (2013), pp. 966–989. CITED ON PAGE 16.
- 192] K.H.W.J. TEN TUSSCHER, D. NOBLE, P.J. NOBLE, AND A.V. PANFILOV, *A model for human ventricular tissue*, American Journal of Physiology - Heart and Circulatory Physiology, vol. 286, no. 4 (2004), pp. H1573–H1589. CITED ON PAGE 91.
- 193] K.H.W.J. TEN TUSSCHER AND A.V. PANFILOV, *Alternans and spiral breakup in a human ventricular tissue model*, American Journal of Physiology - Heart and Circulatory Physiology, vol. 290, no. 3 (2006), pp. H1088–H1100. CITED ON PAGE 71.
- 194] ———, *Modelling of the ventricular conduction system*, Progress in Biophysics & Molecular Biology, vol. 96, no. 1-3 (2008), pp. 152–170. CITED ON PAGES 91, 94, AND 97.
- 195] J.B. TENENBAUM AND W.T. FREEMAN, *Separating style and content*, in Proc. Neural Information Processing Systems Conf. (NIPS), Denver, CO, USA, vol. 9, 1996, pp. 662–668. CITED ON PAGE 46.
- 196] ———, *Separating style and content with bilinear models*, Neural Computation, vol. 12, no. 6 (2000), pp. 1247–1283. CITED ON PAGES 28, 30, 46, 48, 49, 50, 52, 57, AND 66.
- 197] C. TOBON-GOMEZ, C. BUTAKOFF, S. AGUADE, F.M. SUKNO, G. MORAGAS, AND A.F. FRANGI, *Automatic construction of 3D-ASM intensity models by simulating image acquisition: Application to myocardial gated SPECT studies*, IEEE Transactions on Medical Imaging, vol. 27, no. 11 (2008), pp. 1655–1667. CITED ON PAGE 42.
- 198] C. TOBON-GOMEZ, F.M. SUKNO, B.H. BIJNENS, M. HUGUET, AND A.F. FRANGI, *Realistic simulation of cardiac magnetic resonance studies modeling anatomical variability, trabeculae, and papillary muscles*, Magnetic Resonance in Medicine, vol. 65, no. 1 (2011), pp. 280–288. CITED ON PAGE 42.
- 199] A. TROUVÉ, *Diffeomorphisms groups and pattern matching in image analysis*, International Journal of Computer Vision, vol. 28, no. 3 (1998), pp. 213–221. CITED ON PAGES 20 AND 21.
- 200] M. TURK AND A. PENTLAND, *Eigenfaces for recognition*, Journal of Cognitive Neuroscience, vol. 3, no. 1 (1991), pp. 71–86. CITED ON PAGES 47 AND 49.
- 201] S. UKIL AND J.M. REINHARDT, *Anatomy-guided lung lobe segmentation in X-ray CT images*, IEEE Transactions on Medical Imaging, vol. 28, no. 2 (2009), pp. 202–214. CITED ON PAGE 95.
- 202] UNITED NATIONS, DEPARTMENT OF ECONOMIC AND SOCIAL AFFAIRS, POPULATION DIVISION, *World Population Prospects: The 2010 Revision, Volume I: Comprehensive Tables*, 2011. CITED ON PAGE 97.
- 203] M. VALDERRÁBANO, *Influence of anisotropic conduction properties in the propagation of the cardiac action potential*, Progress in Biophysics & Molecular Biology, vol. 94, no. 1/2 (2007), pp. 144–168. CITED ON PAGE 77.
- 204] S. VALETTE, J.M. CHASSERY, AND R. PROST, *Generic remeshing of 3D triangular meshes with metric-dependent discrete Voronoi diagrams*, IEEE Transactions on Visualization and Computer Graphics, vol. 14, no. 2 (2008), pp. 369–381. CITED ON PAGE 25.
- 205] E.M. VAN RIKXOORT, I. ISGUM, Y. ARZHAeva, M. STARING, S. KLEIN, M.A. VIERGEVER, J.P.W. PLUIM, AND B. VAN GINNEKEN, *Adaptive local multi-atlas segmentation: Application to the heart and the caudate nucleus*, Medical Image Analysis, vol. 14, no. 1 (2010), pp. 39–49. CITED ON PAGE 13.
- 206] M.A.O. VASILESCU AND D. TERZOPOULOS, *Multilinear subspace analysis of image ensembles*, in Proc. IEEE Int. Conf. on Computer Vision and Pattern Recognition (CVPR), Madison, WI, USA, vol. 2, 2003, pp. 93–99. CITED ON PAGE 46.
- 207] M. VÁZQUEZ, R. ARÍS, G. HOUZEAUX, R. AUBRY, P. VILLAR, J. GARCIA-BARNÉS, D. GIL, AND F. CARRERAS, *A massively parallel computational electrophysiology model of the heart*, International Journal for Numerical Methods in Biomedical Engineering, vol. 27, no. 12 (2011), pp. 1911–1929. CITED ON PAGE 73.
- 208] T. VERCAUTEREN, X. PENNEC, A. PERCHANT, AND N. AYACHE, *Diffeomorphic demons: Efficient non-parametric image registration*, NeuroImage, vol. 45, no. 1 (2009), pp. S61–S72. CITED ON PAGE 21.
- 209] M. VICECONTI AND G. CLAPWORTHY, eds., *VPH-FET Research Roadmap: Advanced Technologies for the Future of the Virtual Physiological Human*, 2011. CITED ON PAGE 97.

- 210] E.J. VIGMOND, C. CLEMENTS, D.M. MCQUEEN, AND C.S. PESKIN, *Effect of bundle branch block on cardiac output: A whole heart simulation study*, Progress in Biophysics & Molecular Biology, vol. 97, no. 2/3 (2008), pp. 520–542. CITED ON PAGE 76.
- 211] E.J. VIGMOND, R. WEBER DOS SANTOS, A.J. PRASSL, M. DEO, AND G. PLANK, *Solvers for the cardiac bidomain equations*, Progress in Biophysics & Molecular Biology, vol. 96, no. 1-3 (2008), pp. 3–18. CITED ON PAGE 73.
- 212] VIRTUAL PHYSIOLOGICAL HUMAN NETWORK OF EXCELLENCE [UCLVPH], *Virtual physiological human*, 2011. CITED ON PAGE 97.
- 213] F. WANG, B.C. VEMURI, A. RANGARAJAN, AND S.J. EISENSCHENK, *Simultaneous nonrigid registration of multiple point sets and atlas construction*, IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 30, no. 11 (2008), pp. 2011–2022. CITED ON PAGE 15.
- 214] V.Y. WANG, C. HOOGENDOORN, G. ENGELBRECHT, A.F. FRANGI, A.A. YOUNG, P.J. HUNTER, AND M.P. NASH, *Unsupervised segmentation and personalised FE modelling of in vivo human myocardial mechanics based on an MRI atlas*, in Proc. IEEE Int. Symp. on Biomedical Imaging (ISBI), Barcelona, Spain, 2012, pp. 1360–1363. CITED ON PAGES 30 AND 73.
- 215] V.Y. WANG, C. HOOGENDOORN, A.F. FRANGI, B.R. COWAN, P.J. HUNTER, AND M.P. NASH, *Automated personalised human left ventricular FE models to investigate heart failure mechanics*, in Proc. Workshop Statistical Atlases and Computational Models of the Heart (STACOM), Nice, France, vol. 7746 of Lecture Notes in Computer Science, 2012, pp. 307–316. CITED ON PAGE 73.
- 216] G. WEISZ, *The emergence of medical specialization in the nineteenth century*, Bulletin of the History of Medicine, vol. 77, no. 3 (2003), pp. 536–574. CITED ON PAGE 5.
- 217] M. WIERZBICKI, M. DRANGOVA, G. GUIRAUDON, AND T. PETERS, *Validation of dynamic heart models obtained using non-linear registration for virtual reality training, planning and guidance of minimally invasive cardiac surgeries*, Medical Image Analysis, vol. 8, no. 3 (2004), pp. 387–401. CITED ON PAGE 32.
- 218] B.J.B.M. WOLTERS, M.C.M. RUTTEN, G.W.H. SCHURINK, U. KOSE, J. DE HART, AND F.N. VAN DE VOSSE, *A patient-specific computational model of fluid-structure interaction in abdominal aortic aneurysms*, Medical Engineering & Physics, vol. 27, no. 10 (2005), pp. 871–883. CITED ON PAGE 73.
- 219] WORLD HEALTH ORGANIZATION, *Factsheet 317: Cardiovascular diseases (CVDs)*, Mar. 2013. CITED ON PAGE 4.
- 220] G. WU, H. JIA, Q. WANG, AND D. SHEN, *SharpMean: Groupwise registration guided by sharp mean image and tree-based registration*, NeuroImage, vol. 56, no. 4 (2011), pp. 1968–1981. CITED ON PAGES 15 AND 16.
- 221] Z. WU AND J.M. SULLIVAN, JR., *Multiple material marching cubes algorithm*, International Journal for Numerical Methods in Engineering, vol. 58, no. 2 (2003), pp. 189–207. CITED ON PAGE 25.
- 222] A.A. YOUNG AND A.F. FRANGI, *Computational cardiac atlases: From patient to population and back*, Experimental Physiology, vol. 94, no. 5 (2009), pp. 578–596. CITED ON PAGES 11 AND 42.
- 223] O. ZETTINIG, T. MANSI, B. GEORGESCU, E. KAYVANPOUR, F. SEDAGHAT-HAMEDANI, A. AMR, J. HAAS, H. STEEN, B. MEDER, H. KATUS, N. NAVAB, A. KAMEN, AND D. COMANICIU, *Fast data-driven calibration of a cardiac electrophysiology model from images and ECG*, in Proc. 16th Int. Conf. Medical Image Computing and Computer Assisted Intervention (MICCAI), Nagoya, Japan, vol. 8149 of Lecture Notes in Computer Science, 2013, pp. 1–8. CITED ON PAGE 74.
- 224] Y. ZHANG, T.J.R. HUGHES, AND C.L. BAJAJ, *An automatic 3D mesh generation method for domains with multiple materials*, Computer Methods in Applied Mechanics and Engineering, vol. 199, no. 5-8 (2010), pp. 405–415. CITED ON PAGE 25.
- 225] Y. ZHENG, A. BARBU, B. GEORGESCU, M. SCHEUERING, AND D. COMANICIU, *Four-chamber heart modeling and automatic segmentation for 3-D cardiac CT volumes using marginal space learning and steerable features*, IEEE Transactions on Medical Imaging, vol. 27, no. 11 (2008), pp. 1668–1681. CITED ON PAGE 14.
- 226] Y. ZHU, X. PAPADEMETRIS, A.J. SINUSAS, AND J.S. DUNCAN, *Segmentation of the left ventricle from cardiac MR images using a subject-specific dynamical model*, IEEE Transactions on Medical Imaging, vol. 29, no. 3 (2010), pp. 669–687. CITED ON PAGES 14, 28, 30, 42, AND 96.

PUBLICATIONS

Articles in international peer-reviewed journals

- J1 **C. HOOGENDOORN**, N. DUCHATEAU, D. SÁNCHEZ-QUINTANA, T. WHITMARSH, M. DE CRAENE, F.M. SUKNO, K. LEKADIR, AND A.F. FRANGI, *A high-resolution atlas and statistical model of the human heart from multislice CT*. IEEE Transactions on Medical Imaging, 32 (2013), pp. 38–44.
- J2 **C. HOOGENDOORN**, F.M. SUKNO, S. ORDÁS, AND A.F. FRANGI, *Bilinear models for spatio-temporal point distribution analysis: Application to extrapolation of left ventricular, biventricular and whole heart cardiac dynamics*. International Journal of Computer Vision, 85 (2009), pp. 237–252.
- J3 **C. HOOGENDOORN**, R. SEBASTIÁN, J.F. RODRÍGUEZ, K. LEKADIR, AND A.F. FRANGI, *An Atlas and Data-Driven Approach to Initializing Cardiac Electrophysiology Simulations*. Under review.
- J4 A.R. PORRAS, G. PIELLA, A. BERRUEZO, **C. HOOGENDOORN**, D. ANDREU, M. SITGES, J. FERNANDEZ-ARMENTA, AND A.F. FRANGI, *Interventional endocardial motion estimation from electro-anatomical mapping data: Application to scar characterization*. IEEE Transactions on Biomedical Engineering, 60 (2013), pp. 1217–1224.
- J5 X. ALBÀ, R.M. FIGUERAS I VENTURA, K. LEKADIR, C. TOBÓN-GÓMEZ, **C. HOOGENDOORN**, AND A.F. FRANGI, *Automatic cardiac LV segmentation in MRI using modified graph cuts with smoothness and inter-slice constraints*. Magnetic Resonance in Medicine, Accepted for publication.
- J6 K. LEKADIR, **C. HOOGENDOORN**, M. PEREÁÑEZ, X. ALBÀ, A. PASHAEI, AND A.F. FRANGI, *Statistical personalization of ventricular fiber orientation using shape predictors*. IEEE Transactions on Medical Imaging, Accepted for publication.
- J7 C. BUTAKOFF, S. BALOCCO, F.M. SUKNO, **C. HOOGENDOORN**, C. TOBÓN-GÓMEZ, G. AVEGLIANO, AND A.F. FRANGI, *Left-ventricular epi- and endocardium extraction from 3D ultrasound images using an automatically constructed 3D ASM*. Under review.
- J8 K. LEKADIR, A. PASHAEI, **C. HOOGENDOORN**, M. PEREÁÑEZ, X. ALBÀ, AND A.F. FRANGI, *Effect of statistically derived fiber models on the estimation of cardiac electrical activation*. Under review.
- J9 M. PEREÁÑEZ, C. BUTAKOFF, **C. HOOGENDOORN**, K. LEKADIR, AND A.F. FRANGI, *Framework for the merging of pre-existing and correspondenceless 3D statistical shape models*. Under review.
- J10 V.A. ZIMMER, K. LEKADIR, G. PIELLA, **C. HOOGENDOORN**, G.-Z. YANG, AND A.F. FRANGI, *A framework for optimized kernel-based manifold embedding of medical image data*. Under review.

Book chapters

- B1 A. FRANGI, D. FRIBOULET, N. AYACHE, H. DELINGETTE, T. GLATARD, **C. HOOGENDOORN**, L. HUMBERT, K. LEKADIR, I. LARRABIDE, Y. MARTELLI, F. PEYRIN, X. PLANES, M. SERMESANT, M.-C. VILLA-URIOL, T. WHITMARSH, AND D. ATKINSON, *Image and signal-based modelling*, in Computational Biomedicine, P. Coveney, V. Diaz, P. Hunter, and M. Viceconti, eds., Oxford University Press. In press. Ch. 4.

Papers in peer-reviewed conferences

- C1 **C. HOOGENDOORN**, A. PASHAEI, R. SEBASTIÁN, F.M. SUKNO, O. CÁMARA, AND A.F. FRANGI, *Sensitivity analysis of mesh warping and subsampling strategies for generating large scale electrophysiological simulation data*. In Proc. Int. Conf. on Functional Imaging and Modeling of the Heart (FIMH), volume 6666 of Lecture Notes in Computer Science, 2011, pp. 418–426. Oral presentation.
- C2 **C. HOOGENDOORN**, A. PASHAEI, R. SEBASTIÁN, F.M. SUKNO, O. CÁMARA, AND A.F. FRANGI, *Influence of geometric variations on LV activation times: A study on an atlas-based virtual population*. In Proc. Workshop on Statistical Atlases and Computational Models of the Heart (STACOM), volume 6364 of Lecture Notes in Computer Science, 2010, pp. 242–251. Poster.
- C3 **C. HOOGENDOORN**, T. WHITMARSH, N. DUCHATEAU, F.M. SUKNO, M. DE CRAENE, AND A.F. FRANGI, *A groupwise mutual information metric for cost efficient selection of a suitable reference in cardiac computational atlas construction*. In Proc. SPIE Medical Imaging, volume 7623, 2010. Poster.
- C4 **C. HOOGENDOORN**, F.M. SUKNO, S. ORDÁS, AND A.F. FRANGI, *Bilinear models for spatiotemporal point distribution analysis: Application to extrapolation of whole heart cardiac dynamics*. In Proc. Int. Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA), 2007. Oral presentation.
- C5 K. LEKADIR, A. PASHAEI, **C. HOOGENDOORN**, M. PEREAÑEZ, X. ALBÀ, AND A.F. FRANGI, *Personalized modeling of cardiac electrophysiology using shape-based prediction of fiber orientation*. In Proc. Workshop on Statistical Atlases and Computational Models of the Heart (STACOM), 2013.
- C6 M. PEREAÑEZ, K. LEKADIR, C. BUTAKOFF, **C. HOOGENDOORN**, AND A.F. FRANGI, *Fusing correspondenceless 3D statistical shape models*. In Proc. Int. Conf. on Medical Image Computing and Computer Assisted Intervention (MICCAI), volume 8149 of Lecture Notes in Computer Science, 2013, pp. 251–258.
- C7 V.A.M. ZIMMER, R. FONOLLÀ, K. LEKADIR, G. PIELLA, **C. HOOGENDOORN**, AND A.F. FRANGI, *Patient-specific manifold embedding of multispectral images using kernel combination*. In Proc. Int. Workshop on Manifold Learning in Medical Imaging (MLMI), volume 8184 of Lecture Notes in Computer Science, 2013, pp. 82–89.
- C8 K. LEKADIR, **C. HOOGENDOORN**, N. DUCHATEAU, AND A.F. FRANGI, *The construction of a statistical atlas of the whole heart from a large 4D database*. In Proc. Computing in Cardiology, 2012, pp. 541–544.
- C9 A.R. PORRAS, G. PIELLA, **C. HOOGENDOORN**, D. ANDREU, A. BERRUEZO, AND A.F. FRANGI, *Endocardial motion estimation from electro-anatomical data*. In Proc. IEEE Int. Symp. on Biomedical Imaging (ISBI), 2012, pp. 282–285.
- C10 V.Y. WANG, **C. HOOGENDOORN**, G. ENGELBRECHT, A.F. FRANGI, A.A. YOUNG, P.J. HUNTER, AND M.P. NASH, *Unsupervised segmentation and personalised FE modelling of in vivo human myocardial mechanics based on an MRI atlas*. In Proc. IEEE Int. Symp. on Biomedical Imaging (ISBI), 2012, pp. 1360–1363.

- C11 V.Y. WANG, C. HOOGENDOORN, A.F. FRANGI, B.R. COWAN, P.J. HUNTER, AND M.P. NASH, *Automated personalised human left ventricular FE models to investigate heart failure mechanics*. In Proc. Workshop on Statistical Atlases and Computational Models of the Heart (STACOM), volume 7746 of Lecture Notes in Computer Science, 2012, pp. 307–316.
- C12 A. BURGER, B. DE BONO, P.J. HUNTER, J. BISBAL, A.F. FRANGI, C. HOOGENDOORN, D. DAVIDSON, X. GU, AND R. BALDOCK, *Integrating volumetric biomedical data in the virtual physiological human*. In Proc. Int. Workshop on Managing Interoperability and Complexity in Health Systems (MIX-HS), 2012, pp. 19–26.
- C13 K. LEKADIR, B. GHAFARYASL, E. MUÑOZ-MORENO, C. BUTAKOFF, C. HOOGENDOORN, AND A.F. FRANGI, *Predictive modeling of cardiac fiber orientation using the Knutsson mapping*. In Int. Conf. on Medical Image Computing and Computer Assisted Intervention (MICCAI), volume 6892 of Lecture Notes in Computer Science, 2011, pp. 50–57.
- C14 A. PASHAEI, C. HOOGENDOORN, R. SEBASTIÁN, D. ROMERO, O. CÁMARA, AND A.F. FRANGI, *Infarction and fast electrophysiology of human heart: In-silico study on atlas-based virtual population*. In Proc. Int. Conf. on Functional Imaging and Modeling of the Heart (FIMH), volume 6666 of Lecture Notes in Computer Science, 2011, pp. 427–436.
- C15 N. DUCHATEAU, M. DE CRAENE, G. PIELLA, C. HOOGENDOORN, E. SILVA, A. DOLTRA, L. MONT, M.Á. CASTEL, J. BRUGADA, M. SITGES, AND A.F. FRANGI, *Atlas-based quantification of myocardial motion abnormalities: Added-value for the understanding of CRT outcome?* In Proc. Workshop on Statistical Atlases and Computational Models of the Heart (STACOM), volume 6364 of Lecture Notes in Computer Science, 2010 pp. 65–74.
- C16 R.M. FIGUERAS I VENTURA, C. HOOGENDOORN, F.M. SUKNO, AND A.F. FRANGI, *Bilinear point distribution models for heart motion analysis*. In Proc. IEEE Int. Symp. on Biomedical Imaging (ISBI), 2010, pp. 476–479.

Media coverage

- *La Media del Corazón* (in Spanish), Quo, May 2013, page 22.
- *Un Corazón Español* (in Spanish) Quo online, May 2013.
- *La estadística dibuja un atlas del corazón humano* (in Spanish), Agencia SiNC, May 2013.
- *An atlas of the human heart is drawn using statistics* (through SiNC), EurekAlert and AlphaGalileo, May 2013.
- *L'estadística dibuixa un atlas del cor humà* (in Catalan), UPF.edu, May 2013.

ACKNOWLEDGMENTS

First and foremost, I would like to thank my supervisor Prof. Alejandro Frangi for giving me the opportunity to obtain the Ph.D. degree in an ambitious lab. I would also like to thank the day-to-day supervisors Drs. Federico Sukno and Karim Lekadir. Two very different people, who brought different views to adapt to.

I would like to thank the external collaborators that made some of the impossible things seem easy. Those are Prof. Damián Sánchez-Quintana of the University of Extremadura, for his contribution to Chapter Two, and Drs. Rafael Sebastián of the University of Valencia and José Rodríguez of the University of Zaragoza, for their contributions to Chapter Four. Thank you.

I would like to thank my colleagues in teaching Linear Algebra, Graph Theory, and Optical Engineering. Also I would like to thank the programme committees and the editors who have invited me to review for their conferences and journals. Both teaching and reviewing are skills that we are expected to acquire, but one needs the opportunity and the help to get started.

My collaborators in the RICORDO project deserve a special thanks, specifically Dr. Albert Burger, Prof. Peter Hunter, Dr. Vicky Wang and Prof. Martyn Nash.

Colleagues of CISTIB... Ph.D. students, postdocs, PMO, SSD. Thank you all for the environment you created. So many of us had nobody here but our colleagues, but like I, you left behind here many friends.

My old friends, Bart, Willjan, Oluseye, Reinhard... you don't live in Spain. I have been neglecting you but I will make up for it. Thank you for not holding it against me.

I believe the story of a Ph.D. starts when one begins to think about whether to do it or not. I would like to thank those who contributed to my decision: Erik Dam, Mads Nielsen, Marleen de Bruijne, Sune Keller, Bram van Ginneken, and Max Viergever. And Marco "if I could do it, you can definitely do it" Loog.

Finally, the family bond is still the strongest bond in life. My years here have seen my family grow; not only with two little pirates for nephews, but also with a beautiful wife, with parents-in-law, sister-in-law, brother-in-law and a few more additions to grandma's list of great-grandchildren. Thank you all for understanding that when I'm with you, it's only for a short while, but it's for being with you.

Mom, dad. You let me choose what to study, where to study. You supported me, sometimes silently, sometimes vocally. You pushed me to also support myself, and to live up to my potential while still allowing me to return to my comfort zone if I needed to. And you didn't ask me to explain all my work to you

yet you would still hear out my frustrations. It's all I could have asked for, if I had had the wits to do so. Thank you.

Finally I would like to thank Martha, who gave up so much to be by my side and even joined me to a small 14th-century castle to say "ja, ik wil". You picked me up when I was down, you pushed me onward when any doubts about continuing arose. You made my world a nicer place.

Te amo.

FUNDING

In this space I would like to acknowledge the various sources of funding of this research.

- The Department of Information and Communications Technologies of UPF for the departmental scholarship.
- European Commission's 7th Framework Programme, through project eu-Heart (FP7/ICT-2007-224495).
- European Commission's 7th Framework Programme, through project RICORDO (FP7/ICT-2009.5.4).
- Ministerio de Economía y Competitividad (MINECO), through project STIMATH (TIN2009-14536-C02-1).
- Centro para el Desarrollo Tecnológico Industrial (CDTI) of the MINECO's CENIT programme, through project cvRemod (CEN-20091044).
- Centro de Investigación Biomédica en Red – Bioingeniería, Biomateriales y Nanomedicina (CIBER-BBN), for a travel grant to participate in the IEEE EMBS International Summer School on Biomedical Imaging, Île de Berder, France, in 2008.



BIOSKETCH

Corné Hoogendoorn was born in 1981 in Werkendam, the Netherlands. After graduating high school from the Gymnasium Camphusianum in Gorinchem he proceeded to study at the University of Utrecht, the Netherlands. He obtained the B.Sc. degree in Computer Science and the M.Sc. degree in Biomedical Sciences, programme Biomedical Image Sciences, in 2004 and 2006, respectively. He wrote his M.Sc. thesis on the topic of point correspondence for 3D statistical shape modeling, during a research stay at the Image Group of the IT University Copenhagen, Denmark. In 2006, he commenced working towards the Ph.D. degree at CISTIB on the topic of cardiac shape modeling and atlas construction.



During his time at CISTIB, he taught Graph Theory and Linear Algebra to first-year students, and Optical Engineering to second-year students at UPF's Escola Superior Politècnica. In 2010 and 2012 he was a reviewer for the annual MICCAI conference, and he has reviewed manuscripts for the journals Computer Vision and Image Understanding, IEEE Transactions on Medical Imaging, and EuroPace.

Corné took part in the 2008 edition of the biennial International Summer School on Biomedical Imaging of IEEE's Engineering in Medicine and Biology Society, held on the island of Berder, France, and in the summer school on Non-linear Manifold Learning on the island of Hven, Sweden, in 2010, organized by the Danish Society for Pattern Recognition. In 2012, he was a volunteer at IEEE's International Symposium on Biomedical Imaging, held in Barcelona.

COLOR IMAGES

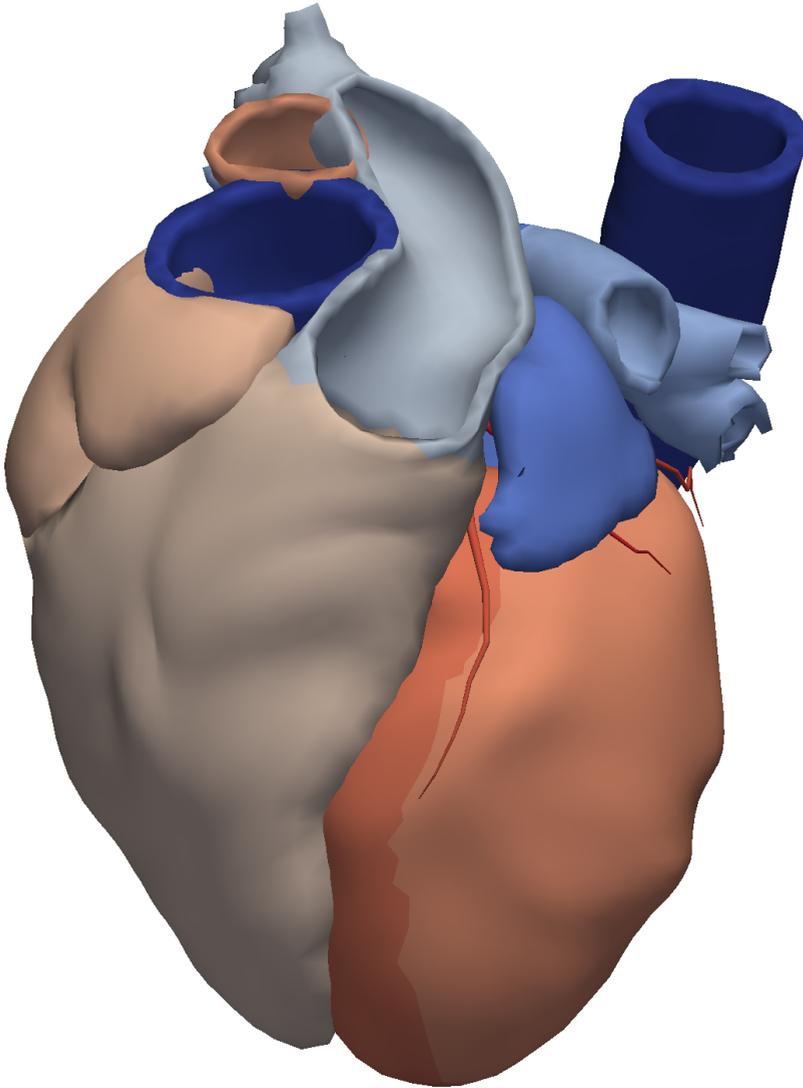


Figure 2.5: *The full atlas mesh, corresponding to the synthesized mean image. Different colors indicate different structures.*

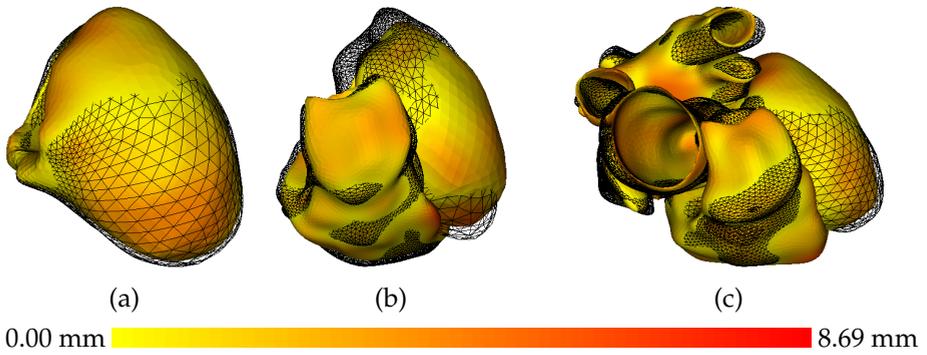


Figure 3.11: Reconstructions of the shapes with the median RMS P2S errors, constructed using the highest dimensional models. (a) Left ventricle (1.869 mm), (b) biventricular (2.212 mm) and (c) full heart (2.344 mm). The wireframe shows the ground truth for these shapes. In (b) and (c), the difference in point density between the left ventricle and the other structures is clearly visible.

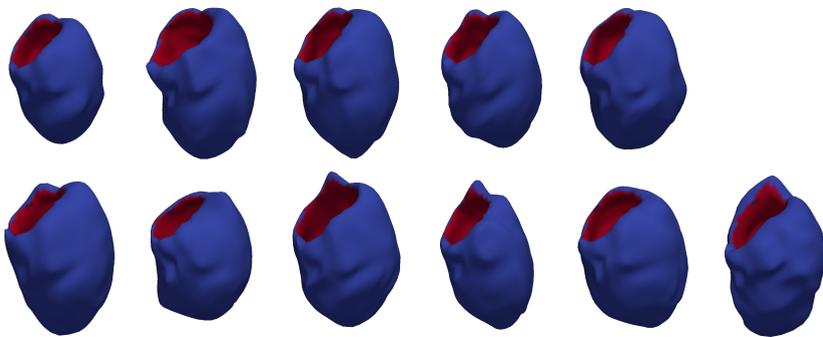


Figure 4.3: Surface renderings of the eleven geometries, with endocardium (red) and epicardium (blue).

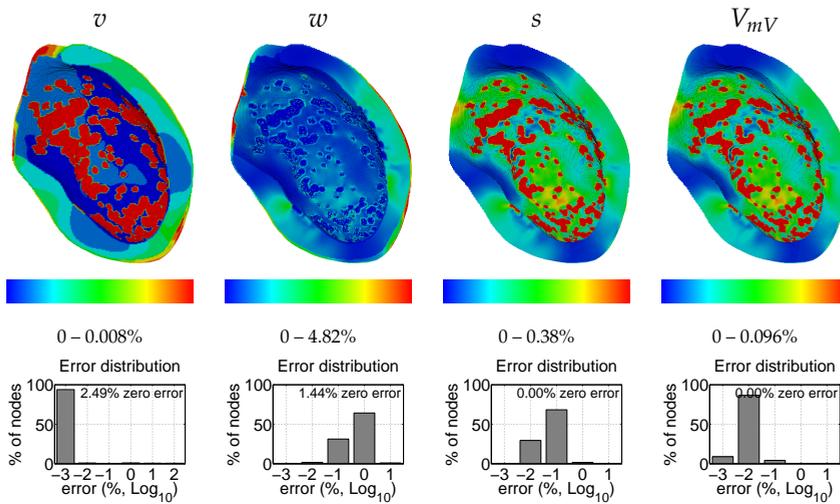


Figure 4.8: Mean absolute prediction errors at time $t = 5$ color coded on the surface of the mean myocardium. Each pair of a surface rendering and histogram corresponds to a model state variable. Color bar upper limits were set to the 95th percentile.

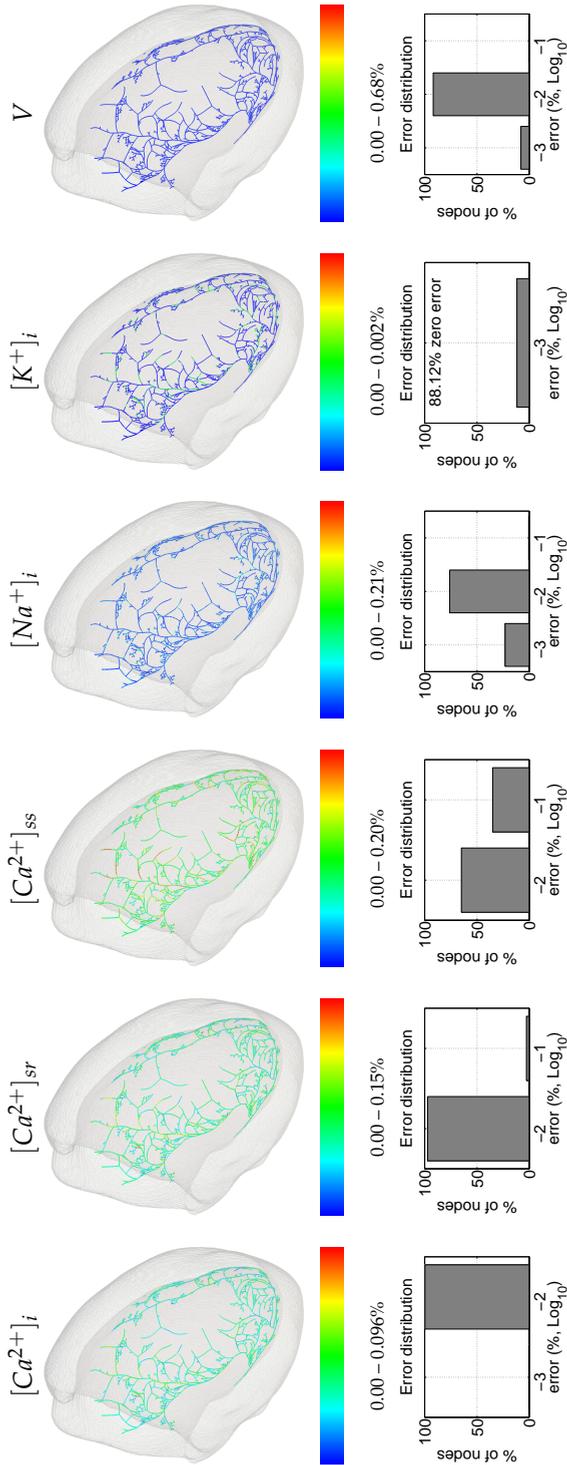


Figure 4.9: Mean absolute prediction errors at time $t = 5$ color coded on the Purkinje network corresponding to the mean geometry. Each pair of a tree rendering and histogram corresponds to a model state variable. Only the ionic concentrations and the voltage V are shown.

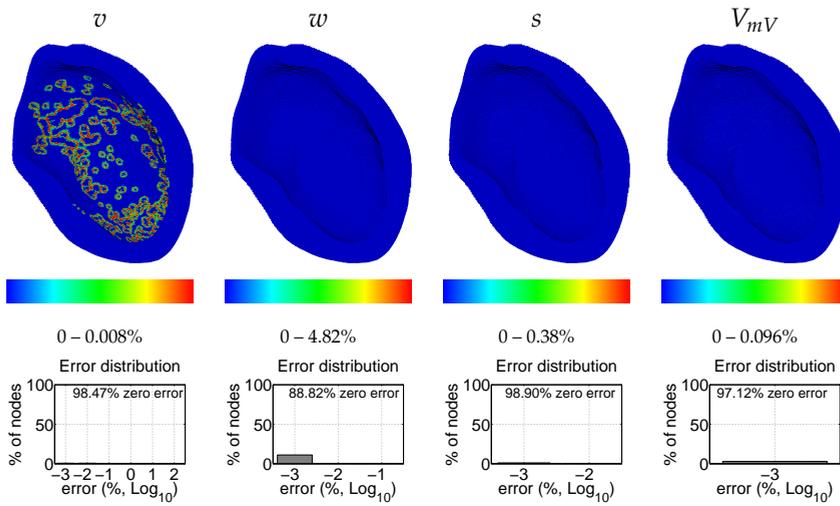


Figure 4.12: Mean absolute differences between the two approaches to reach time $t = 10$ color coded on the surface of the mean myocardium. Each pair of a surface rendering and histogram corresponds to a model state variable. Color bar ranges are the same as in Fig. 4.8.

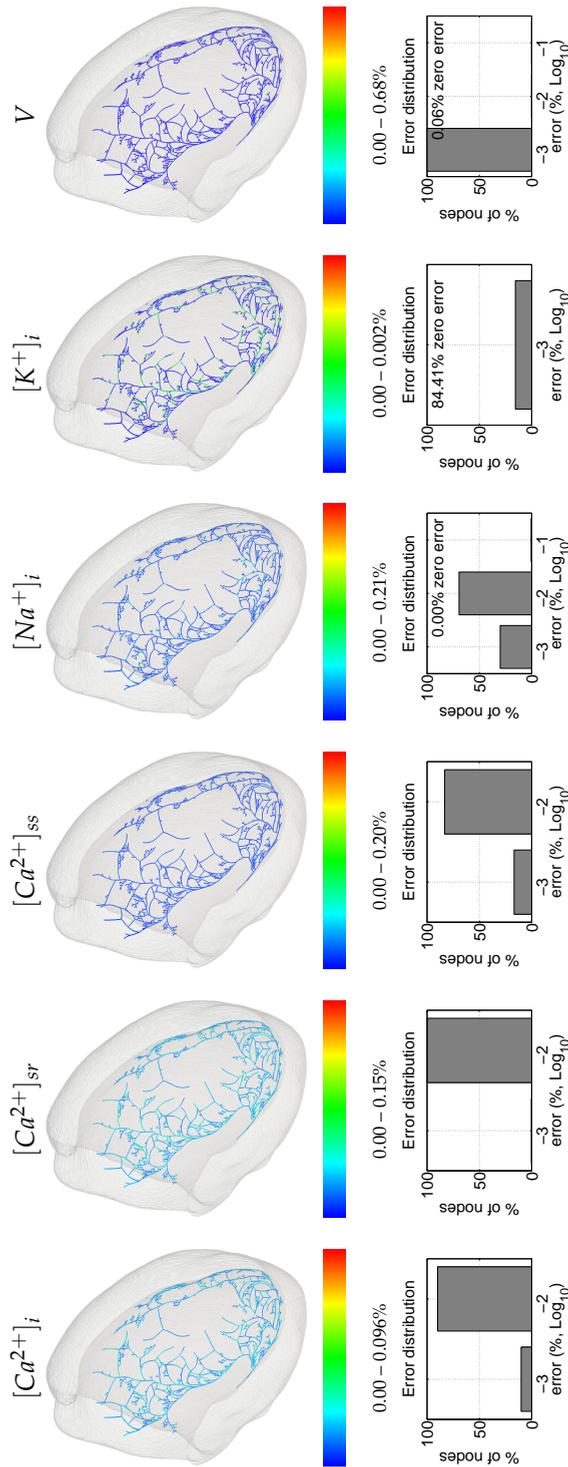


Figure 4.13: Mean absolute differences between the two different approaches to reach time $t = 10$ color coded on the Purkinje network corresponding to the mean geometry. Each pair of a tree rendering and histogram corresponds to a model state variable. Only the ionic concentrations and the voltage V are shown. Color bar ranges are the same as in Fig. 4.9.

