

Robust Personalization of Cardiac Computational Models

Robuste Personalisierung von
rechnergestützten Herzmodellen

Der Technischen Fakultät
der Friedrich-Alexander-Universität
Erlangen-Nürnberg

zur

Erlangung des Doktorgrades Dr.-Ing.

vorgelegt von

Dominik Neumann
aus
Amberg, Deutschland

Als Dissertation genehmigt
von der Technischen Fakultät
der Friedrich-Alexander-Universität Erlangen-Nürnberg

Tag der mündlichen Prüfung: 24.06.2019

Vorsitzender des Promotionsorgans: Prof. Dr.-Ing. Reinhard Lerch

Gutachter: Prof. Dr.-Ing. Joachim Hornegger

Prof. Dr. rer. nat. Bjoern H. Menze

Abstract

Heart failure (HF) is a major cause of morbidity and mortality in the Western world, yet early diagnosis and treatment remain a major challenge. As computational cardiac models are becoming more mature, they are slowly evolving into clinical tools to better stratify HF patients, predict risk and perform treatment planning. A critical prerequisite, however, is their ability to precisely capture an individual patient's physiology. The process of fitting a model to patient data is called personalization, which is the overarching topic of this thesis.

An image-based, multi-physics 3D whole-heart model is employed in this work. It consists of several components covering anatomy, electrophysiology, biomechanics and hemodynamics. Building upon state-of-the-art personalization techniques, the first goal was to develop an automated pipeline for personalizing all components of the model in a streamlined and reproducible fashion, based on routinely acquired clinical data. Evaluation was performed on a multi-clinic cohort consisting of 113 patients, the largest cohort in any comparable study to date. The goodness of fit between personalized models and ground-truth clinical data was mostly below clinical variability, while a full personalization was finalized within only few hours. This showcases the ability of the proposed pipeline to extract advanced biophysical parameters robustly and efficiently.

Designing such personalization algorithms is a tedious, model- and data-specific process. The second goal was to investigate whether artificial intelligence (AI) concepts can be used to learn this task, inspired by how humans manually perform it. A self-taught artificial agent based on reinforcement learning (RL) is proposed, which first learns how the model behaves, then computes an optimal strategy for personalization. The algorithm is model-independent; applying it to a new model requires only adjusting few hyper-parameters. The obtained results for two different models suggested that equivalent, if not better goodness of fit than standard methods could be achieved, while being more robust and with faster convergence rate. AI approaches could thus make personalization algorithms generalizable and self-adaptable to any patient and any model.

Due to limited data, uncertainty in the clinical measurements, parameter non-identifiability, and modeling assumptions, various combinations of parameter values may exist that yield the same quality of fit. The third goal of this work was uncertainty quantification (UQ) of the estimated parameters and to ascertain the uniqueness of the found solution. A stochastic method based on Bayesian inference and fast surrogate models is proposed, which estimates the posterior of the model, taking into account uncertainties due to measurement noise. Experiments on the biomechanics model showed that not only goodness of fits equivalent to the standard methods could be achieved, but also the non-uniqueness of the problem could be demonstrated and uncertainty estimates reported, crucial information for subsequent clinical assessments of the personalized models.

Kurzfassung

Herzfehler gehören zu den häufigsten Erkrankungen und Todesursachen in der westlichen Welt, dennoch stellen deren Früherkennung und Behandlung auch heute noch eine große Herausforderung dar. Computermodelle des Herzens werden immer ausgereifter und stehen kurz davor, Einzug in Kliniken zu erhalten, um dabei zu helfen, Patienten mit Herzfehlern besser zu stratifizieren, Risiken vorherzusagen und die Planung von Eingriffen zu unterstützen. Zuvor muss aber sichergestellt werden, dass die Modelle auch tatsächlich die Physiologie von individuellen Patienten abbilden können. Diese Dissertation befasst sich mit der Personalisierung solcher Modelle, das heißt, mit dem Prozess der Anpassung eines Modells an Patientendaten.

Hierfür kommt ein bildbasiertes Computermodell des Herzens zum Einsatz, welches aus Anatomie-, Elektrophysiologie-, Biomechanik- und Hämodynamikkomponenten besteht. Der erste Teil der Arbeit befasst sich mit der Entwicklung einer automatisierten Pipeline zur reproduzierbaren Personalisierung aller Komponenten, wobei nur Daten verwendet werden, die routinemäßig zur Verfügung stehen. Die Pipeline wurde auf 113 Patienten aus mehreren Kliniken evaluiert. Die erzielte Ähnlichkeit zwischen personalisierten Modellen und klinischen Daten war dabei größtenteils besser als die klinische Variabilität, und jede Personalisierung war innerhalb weniger Stunden abgeschlossen. Diese Ergebnisse lassen darauf schließen, dass biophysikalische Parameter robust und effizient aus den Daten extrahiert werden konnten.

Das Design solcher Personalisierungsalgorithmen ist oft langwierig und modell- und datenspezifisch. Im zweiten Teil der Arbeit wird untersucht, inwiefern künstliche Intelligenz dabei helfen kann, diese Aufgabe zu lernen, so wie es auch Menschen tun. Ein selbstlernender künstlicher Agent basierend auf "Reinforcement Learning" wird vorgestellt, der zunächst die Verhaltensweisen des Modells lernt, um im Anschluss daraus eine optimale Personalisierungsstrategie zu berechnen. Der Algorithmus funktioniert unabhängig vom verwendeten Modell, es müssen lediglich wenige "Hyperparameter" angepasst werden. Die Ergebnisse auf zwei verschiedenen Modellen zeigen, dass im Vergleich zu Standardmethoden äquivalente, oder sogar bessere Personalisierungsergebnisse erzielt werden konnten. Gleichzeitig war der künstliche Agent robuster und konvergierte schneller. Künstliche Intelligenz könnte also Personalisierungsalgorithmen generalisierbarer und adaptierbarer machen.

Aufgrund der limitierten Datenmenge, Unsicherheiten in klinischen Messwerten, Nicht-Identifizierbarkeit von Parametern und Modellannahmen könnten mehrere Kombinationen von Modellparametern existieren, die äquivalente Personalisierungsergebnisse erzeugen. Der dritte Teil der Arbeit befasst sich daher mit der Quantifizierung der Unsicherheit der geschätzten Parameter und der Frage der Eindeutigkeit der Ergebnisse. Es wird ein stochastischer Ansatz vorgestellt, der auf Bayes-Inferenz und schnellen Ersatzmodellen basiert, um die A-posteriori Wahrscheinlichkeit des Modells unter Berücksichtigung der Messunsicherheiten zu schätzen. Die Experimente zeigten nicht nur, dass damit vergleichbare Personalisierungsqualität wie bei deterministischen Standardmethoden erzielt werden konnte, sondern es wurde auch die Nicht-Eindeutigkeit des Personalisierungsproblems demonstriert. Die automatisch berechneten Konfidenzintervalle können außerdem zur klinischen Beurteilung der personalisierten Modelle verwendet werden.

Acknowledgments

This thesis would not have been possible without the support, encouragement, trust, inspiration and expertise of my supervisors, colleagues, collaboration partners, friends and family.

My first thank you goes to my doctoral advisor Prof. Dr.-Ing. Joachim Hornegger, who sparked my interest in pattern recognition and medical image processing, and gave me the opportunity and support necessary to work on this project.

I want to deeply thank Dr. Tommaso Mansi, the mentor and main driver behind my work. Without his exceptional guidance, his continuous support and push, his enthusiasm and passion, I would not be where I am today.

A special thanks to a number of key persons, both from the Pattern Recognition Lab and Siemens Healthineers, who further enabled and supported my work: Prof. Dr.-Ing. habil. Andreas Maier, Dr.-Ing. habil. Stefan Steidl, Dr. Dorin Comaniciu, Dr. Tobias Heimann, Dr. Sasa Grbic, Dr. Olivier Ecabert, Dr. Michael Sühling. I knew that I could always count on you.

There are innumerable colleagues and friends, who I was fortunate enough to meet throughout this phase of my life. I want to thank all of you for the joint trips to conferences, the often intense, but always fruitful discussions, the advice you shared, the long nights we spent together writing papers, and all the good times we've had: Felix Lugauer-Breunig, Martin Berger, Jian Wang, Matthias Ring, Chloé Audigier, Birgi Tamersoy, and many more. A big thanks goes also towards the students I met and worked with, especially Felix Meister and Philipp Seegerer.

A very special thank you goes to my family, to my parents Andrea and Siegfried, and to my siblings Sarah, Manuel, Sabrina and Ramona. Thank you for always being on my side, providing the necessary emotional support and safety net, and for making me the person I am today.

Last but certainly not least, I want to express my deepest gratitude to my wife Isabella: for her love, her understanding and support, her patience and encouragements.

Dominik Neumann

Contents

Chapter 1 Introduction	1
1.1 Motivation	1
1.1.1 Clinical Background	1
1.1.2 What If We Could Create a Virtual Heart of a Patient?	2
1.2 Scientific Contributions	3
1.2.1 Streamlined Multi-Physics Cardiac Model Personalization	3
1.2.2 Artificially Intelligent Personalization Agent	4
1.2.3 Uncertainty and Uniqueness of Parameter Estimates	5
1.2.4 Other Contributions Related to Cardiac Modeling	6
1.3 Organization of this Thesis	6
Chapter 2 The Human Heart	9
2.1 Heart Anatomy and Function	9
2.1.1 Cardiovascular System	9
2.1.2 Chambers and Valves	10
2.1.3 Myocardium	10
2.1.4 Cardiac Activity and Electrical Conduction System	11
2.2 Measuring Structure and Function of the Heart	11
2.2.1 Cardiac Imaging	11
2.2.2 Electrocardiography	12
2.2.3 Blood Pressure	13
2.2.4 Normal Values and Relations	14
I Patient-specific Virtual Hearts	17
Chapter 3 Image-Based, Multi-Physics Model of Heart Function	19
3.1 Anatomy	19
3.1.1 Geometry Segmentation from Imaging Data	19
3.1.2 Automatic Mesh Tagging	20
3.1.3 Rule-based Myocardial Fiber Architecture	22
3.2 Electrophysiology	23
3.2.1 First Model: LBM-EP	24
3.2.2 Second Model: Graph-EP	26
3.2.3 From Cardiac Potentials to Torso Potentials to ECG	28
3.3 Biomechanics	29
3.3.1 Passive Properties	29
3.3.2 Active Force	30
3.3.3 Finite Element Modeling	31

3.4	Boundary Conditions	32
3.4.1	Phase-Model of the Cardiac Cycle	32
3.4.2	Artery Model	33
3.4.3	Atrium Model	34
3.4.4	Spatial Constraints	35
Chapter 4 Whole-Heart Personalization Through Inverse Optimization		37
4.1	State-of-the-Art Model Personalization	38
4.2	Problem Formulation	40
4.3	Integrated Personalization Pipeline	40
4.3.1	Hemodynamics: Personalized Artery Boundary Conditions	42
4.3.2	Electrophysiology: Cascaded Inverse Optimization	43
4.3.3	Biomechanics: Global Volume and Pressure Fitting	45
4.3.4	Towards Regional Biomechanics Personalization	48
4.3.5	Implementation Details	50
4.4	Experiments and Results	50
4.4.1	Description of Study Cohort	50
4.4.2	Whole-Heart Personalization: Goodness of Fit	51
4.4.3	Preliminary Evaluation of Regional Mechanics Approach	57
4.5	Predictive Power of Personalized Model	61
4.6	Summary	62
II Advanced Personalization with Artificial Intelligence and Uncertainty Quantification		65
Chapter 5 Intelligent Agent for Computational Model Personalization		67
5.1	The Potential of Artificial Intelligence for Model Personalization	67
5.2	Learning How to Personalize a Model	69
5.2.1	Model-based Reinforcement Learning	70
5.2.2	Reformulation of the Model Personalization Problem	71
5.2.3	Learning Model Behavior through Model Exploration	72
5.2.4	From Computed Objectives to Representative MDP State	73
5.2.5	Transition Function as Probabilistic Model Representation	73
5.2.6	On-line Model Personalization	74
5.3	Experiments and Results	75
5.3.1	Synthetic Experiment: the Rosenbrock Function	75
5.3.2	Personalization of Cardiac Electrophysiology Model	77
5.3.3	Personalization of a Whole-Body Circulation Model	81
5.4	Summary	85
Chapter 6 Uncertainty Quantification in Personalized Modeling		87
6.1	Uncertainty Quantification	87
6.2	Robust Estimation of Model Parameters and Their Uncertainty	88
6.2.1	Overview of the Method	88
6.2.2	Bayesian Inference	89
6.2.3	Fast Surrogate of the Computational Model	90

6.2.4	Posterior Analysis under Uncertainty	91
6.3	Experiments and Results	93
6.3.1	Precision of PCE-based Surrogate Model	94
6.3.2	Estimation of Biomechanical Parameters	94
6.3.3	Comparison to Deterministic Personalization.	95
6.3.4	Empirical Evaluation of Estimated Uncertainty	95
6.4	Summary.	96
III	Conclusion	99
Chapter 7	Summary and Perspectives	101
7.1	Summary.	101
7.1.1	Patient-specific Virtual Hearts	101
7.1.2	Advanced Personalization with Artificial Intelligence and Uncertainty Quantification.	104
7.2	Perspectives.	106
7.2.1	From Global to Regional Personalization	106
7.2.2	Challenges of AI-based Personalization.	107
7.2.3	Challenges of Stochastic Personalization	108
7.2.4	Predictive Power of Personalized Models	109
7.2.5	What If We Could Create a Virtual Patient?.	109
Chapter A	Appendix	111
A.1	Software Tool for Interactive Personalization	111
A.1.1	VvTk Rapid Prototyping Platform.	111
A.1.2	HeartBuildR.	111
A.2	Supplementary Personalization Results	113
A.3	Additional Details of AI-based Personalization	114
A.3.1	Data-driven State Space Quantization.	114
A.3.2	Data-driven Initialization	115
	List of Acronyms	117
	List of Symbols	121
	List of Figures	127
	List of Tables	129
	Bibliography	131

Introduction

1.1 Motivation.....	1
1.2 Scientific Contributions.....	3
1.3 Organization of this Thesis.....	6

The first part of this chapter, Sec. 1.1.1, provides the clinical motivation of this work, followed by a discussion about the potential of patient-specific computational modeling in Sec. 1.1.2. The main scientific contributions are listed in Sec. 1.2, followed by an overview of the structure of the remainder of this thesis in Sec. 1.3.

1.1 Motivation

1.1.1 Clinical Background

Heart Failure

Heart failure (HF) is a cardiovascular condition characterized by the decline of cardiac function, which leads to insufficient blood delivery to the organs, thus causing organ dysfunction or failure [Lloy 10, McMu 12]. Increasing number of incidences and high mortality and morbidity rates make HF one of the leading causes of death in the Western world [Murr 97, Lope 06, Sant 13]. About 1–2% of the adult population is affected. The incidence rate among the elderly (70 years and older) is even above 10% [Poni 16]. HF accounts for a significant burden to the national healthcare systems. The wide variety of etiologies and pathophysiologies call for new strategies for early diagnosis and treatment planning.

Dilated Cardiomyopathy

Cardiomyopathies are a group of diseases affecting the heart muscle. Among the most common types, and one of the leading causes of HF, is dilated cardiomyopathy (DCM), which is characterized by an enlargement of the heart in one or more of its chambers [Kuma 17]. As the chambers dilate, the walls become thinner and contraction less powerful (Fig. 1.1). Patients suffering from DCM thus have impaired heart function, which can in turn affect other organs and systems of the body. DCM is often a genetic condition, but it can also have other non-genetic causes such as viral infections, auto-immune disease or exposure to toxins. Many DCM patients show no or only minor symptoms, which makes early diagnosis and treatment difficult. Others develop early symptoms such as shortness of breath, swelling of legs, fatigue,




Figure 1.1: Four-chamber views of healthy heart on the left compared to a DCM heart with severely dilated left ventricle on the right as captured by cardiac magnetic resonance imaging (MRI). First and third image show hearts during diastolic phase, where the chambers are filled with blood (white), versus systolic phase (heart is fully contracted) in images two and four. Images courtesy of Siemens Healthineers.

etc., which may become worse as severity of DCM increases, when chest pain, blood clots in the heart chambers or even sudden cardiac death can occur [Kuma 17]. If diagnosed early, progression of DCM can sometimes be slowed down using medication and changes in lifestyle only. At later stages of the disease, minimally-invasive or surgical procedures aimed towards improving pump function become necessary. Cardiac resynchronization therapy (CRT), implantable cardioverter-defibrillators (ICD), or heart transplants may be viable options depending on the severity of the disease and comorbidities. In light of these complex diseases, new avenues for improved patient management are being explored.

Cardiac Resynchronization Therapy

Cardiac resynchronization therapy (CRT) is an efficient and minimally invasive treatment for heart failure with cardiac dyssynchrony. CRT consists in pacing the cardiac myocardium with an advanced pacemaker to recover intra-ventricular and inter-ventricular synchronization. Despite progresses in therapy planning and patient selection, 30% to 50% of the patients do not respond to the therapy although they are within clinical guidelines [Isma 10]. Novel approaches are therefore necessary to compute CRT predictors from preoperative data for improved patient selection and therapy optimization.

1.1.2 What If We Could Create a Virtual Heart of a Patient?

If we could create a virtual heart, meaning that we can accurately simulate patient-specific cardiac function, we can better *describe* the state of the patient through non-invasive functional measurements. Second, a virtual heart can help to *diagnose* the patient's condition more holistically, it provides new ways for disease understanding and can help to improve patient stratification. Third, it could *predict* disease course and therapy outcome, and thus can be used to proactively *prescribe* an optimal intervention for the patient. In essence, a virtual heart could ultimately lead to improved clinical management of patients suffering from cardiomyopathies [Kayv 15, Amr 16].

Virtual hearts, i. e., computational cardiac models, attracted significant attention in cardiac research over the last decades [Fran 01, Nobl 02, Hunt 03, Kerc 08, Clay 11, Kuij 12, Kris 13, Nied 16]. However, their application in clinical practice faces many challenges. For example, model parameters need to be identified, the process of building the virtual heart needs to be robust and automated, and it needs to be intuitive to use to make model-enhanced diagnosis and prediction accessible. Furthermore, a framework needs to be set in place that can provide to the user the confidence of the model prediction, and that offers certain guarantees with respect to failure rates, etc., all of which may require large scale validation efforts. This thesis provides novel solutions to image-based model personalization and tackles some of these challenges by proposing robust algorithms for large-scale personalization and streamlined creation of image-based computational cardiac models. Furthermore, a novel **machine learning (ML)** approach to automatically learn optimal personalization strategies, and a stochastic approach to analyze model uncertainties, are proposed.

1.2 Scientific Contributions

In this section, the contributions of this work to the progress of research are reviewed briefly and the corresponding peer-reviewed scientific publications are listed.

1.2.1 Streamlined Multi-Physics Cardiac Model Personalization

A multitude of computational models of cardiac physiology have been proposed in the literature. One example is an image-based, multi-physics electromechanical 3D whole-heart model [Zett 13], which has been used and enhanced (Chap. 3) in the course of this work. In order to transform such models into tools that could become meaningful in clinical practice, they need to be personalized, i. e., fitted to an individual patient. The first major contribution is a robust and efficient pipeline for streamlined global personalization of this complex model integrating novel optimization-based inverse modeling algorithms, and its successful application to a very large population of 113 patients. The latter makes this work particularly distinctive as most publications in this field are limited to synthetic data and/or only few (typically less than ten) patients. Methods and results are detailed in Chap. 4 and early versions of the pipeline evaluated on smaller populations were presented at an international conference (accepted for oral presentation) and in two clinical publications:

- [Neum 14c] D. Neumann, T. Mansi, S. Grbic, I. Voigt, B. Georgescu, E. Kayvanpour, A. Amr, F. Sedaghat-Hamedani, J. Haas, H. Katus, B. Meder, J. Hornegger, A. Kamen, and D. Comaniciu. “Automatic Image-to-Model Framework for Patient-specific Electromechanical Modeling of the Heart”. In: *IEEE 11th International Symposium on Biomedical Imaging*, pp. 935–938, IEEE, Beijing, China, Apr. 2014.

- [Kayv 15] E. Kayvanpour, T. Mansi, F. Sedaghat-Hamedani, A. Amr, D. Neumann, B. Georgescu, P. Seegerer, A. Kamen, J. Haas, K. S. Frese, M. Irawati, E. Wirsz, V. King, S. Buss, D. Mereles, E. Zitron, A. Keller, H. A. Katus, D. Comaniciu, and B. Meder. “Towards Personalized Cardiology: Multi-Scale Modeling of the Failing Heart”. *PloS One*, Vol. 10, No. 7, pp. 1–18, July 2015.
- [Amr 16] A. Amr, E. Kayvanpour, F. Sedaghat-Hamedani, T. Passerini, V. Mihalef, A. Lai, D. Neumann, B. Georgescu, S. Buss, D. Mereles, E. Zitron, A. Posch, M. Würstle, T. Mansi, H. Katus, and B. Meder. “Personalized Computer Simulation of Diastolic Function in Heart Failure”. *Genomics, Proteomics & Bioinformatics*, Vol. 14, No. 4, pp. 244–252, Aug. 2016.

In the context of cardiac **electrophysiology (EP)** model personalization, additional research activities were carried out, which are, however, not the main focus of this thesis and therefore not explained in detail in this document. In particular, publications covering an **ML** approach for fast global **EP** personalization and a novel approach for regional personalization of electrical tissue properties for patients with localized cardiac conduction abnormalities, were co-authored:

- [Zett 14] O. Zettinig, T. Mansi, D. Neumann, B. Georgescu, S. Rapaka, P. Seegerer, E. Kayvanpour, F. Sedaghat-Hamedani, A. Amr, J. Haas, H. Steen, H. Katus, B. Meder, N. Navab, A. Kamen, and D. Comaniciu. “Data-Driven Estimation of Cardiac Electrical Diffusivity from 12-Lead ECG Signals”. *Medical Image Analysis*, Vol. 18, No. 8, pp. 1361–1376, Dec. 2014.
- [Seeg 15] P. Seegerer, T. Mansi, M.-P. Jolly, D. Neumann, B. Georgescu, A. Kamen, E. Kayvanpour, A. Amr, F. Sedaghat-Hamedani, J. Haas, H. Katus, B. Meder, and D. Comaniciu. “Estimation of Regional Electrical Properties of the Heart from 12-Lead ECG and Images”. In: *Statistical Atlases and Computational Models of the Heart*, pp. 204–212, Springer, Cham, Switzerland, Jan. 2015.

1.2.2 Artificially Intelligent Personalization Agent

The second major contribution is a novel paradigm for personalization based on **artificial intelligence (AI)** concepts. For the first time, personalization was formulated in terms of **reinforcement learning (RL)** to create an artificial agent with the ability to learn by itself how to personalize a model. Methods and results are described in Chap. 5 and were presented at an international conference, and published in a major international journal:

[Neum 15b] D. Neumann, T. Mansi, L. Itu, B. Georgescu, E. Kayvanpour, F. Sedaghat-Hamedani, J. Haas, H. Katus, B. Meder, S. Steidl, J. Hornegger, and D. Comaniciu. “Vito – A Generic Agent for Multi-physics Model Personalization: Application to Heart Modeling”. In: *Medical Image Computing and Computer-Assisted Intervention*, pp. 442–449, Springer International Publishing, Cham, Switzerland, Oct. 2015.

[Neum 16] D. Neumann, T. Mansi, L. Itu, B. Georgescu, E. Kayvanpour, F. Sedaghat-Hamedani, A. Amr, J. Haas, H. Katus, B. Meder, S. Steidl, J. Hornegger, and D. Comaniciu. “A Self-taught Artificial Agent for Multi-Physics Computational Model Personalization”. *Medical Image Analysis*, Vol. 34, pp. 52–64, Dec. 2016.

Based on very similar principles, a versatile artificial agent for fast and robust anatomical landmark detection in medical images was co-developed (not focus of this thesis):

[Ghes 16a] F. C. Ghesu, B. Georgescu, T. Mansi, D. Neumann, J. Hornegger, and D. Comaniciu. “An Artificial Agent for Anatomical Landmark Detection in Medical Images”. In: *Medical Image Computing and Computer-Assisted Intervention*, pp. 229–237, Springer International Publishing, Cham, Switzerland, Oct. 2016.

1.2.3 Uncertainty and Uniqueness of Parameter Estimates

The third area of scientific contributions towards the progress of research in image-based cardiac modeling is [uncertainty quantification \(UQ\)](#). A novel stochastic personalization method is proposed that enables not only the computation of the most likely model parameters of the biomechanics model, but also the estimation of uncertainty due to noise in the clinical data. Methods and results are described in Chap. 6 and were presented at an international conference (accepted for oral presentation and awarded the “MICCAI 2014 Student Travel Award”):

[Neum 14b] D. Neumann, T. Mansi, B. Georgescu, A. Kamen, E. Kayvanpour, A. Amr, F. Sedaghat-Hamedani, J. Haas, H. Katus, B. Meder, J. Hornegger, and D. Comaniciu. “Robust Image-Based Estimation of Cardiac Tissue Parameters and Their Uncertainty from Noisy Data”. In: *Medical Image Computing and Computer-Assisted Intervention*, pp. 9–16, Springer International Publishing, Cham, Switzerland, Sep. 2014.

In addition to this UQ-driven personalization framework, a case study analyzing the impact of uncertainty in myocardial fiber architecture on estimated model parameters was performed (not focus of this thesis):

- [Moll 15] R. Molléro, D. Neumann, M.-M. Rohé, M. Datar, H. Lombaert, N. Ayache, D. Comaniciu, O. Ecabert, M. Chinali, G. Rinelli, *et al.* “Propagation of Myocardial Fibre Architecture Uncertainty on Electromechanical Model Parameter Estimation: A Case Study”. In: *International Conference on Functional Imaging and Modeling of the Heart*, pp. 448–456, Springer, Cham, Switzerland, June 2015.

1.2.4 Other Contributions Related to Cardiac Modeling

Two other related research activities in the context of cardiac modeling were carried out (both not focus of this thesis). First, a framework for mitral valve segmentation and biomechanical closure computation based on *ex vivo* animal data:

- [Neum 14a] D. Neumann, S. Grbic, T. Mansi, I. Voigt, J. Rabbah, A. Siefert, N. Saikrishnan, A. Yoganathan, D. Yuh, and R. Ionasec. “Multi-modal Pipeline for Comprehensive Validation of Mitral Valve Geometry and Functional Computational Models”. In: *Statistical Atlases and Computational Models of the Heart*, pp. 188–195, Springer, Cham, Switzerland, Oct. 2014.

- [Grbi 17] S. Grbic, T. F. Easley, T. Mansi, C. H. Bloodworth, E. L. Pierce, I. Voigt, D. Neumann, J. Krebs, D. D. Yuh, M. O. Jensen, D. Comaniciu, and A. P. Yoganathan. “Personalized Mitral Valve Closure Computation and Uncertainty Analysis from 3D Echocardiography”. *Medical Image Analysis*, Vol. 35, pp. 238–249, Jan. 2017.

Finally, an ML-based probabilistic multi-modality cardiac image fusion approach that could facilitate integration of data from multiple imaging modalities into computational models was proposed:

- [Neum 15a] D. Neumann, S. Grbic, M. John, N. Navab, J. Hornegger, and R. Ionasec. “Probabilistic Sparse Matching for Robust 3D/3D Fusion in Minimally Invasive Surgery”. *IEEE Transactions on Medical Imaging*, Vol. 34, No. 1, pp. 49–60, Jan. 2015.

1.3 Organization of this Thesis

The remainder of this thesis is organized as follows. First, in Chap. 2, the clinical knowledge and terminology required to follow the ideas and concepts in this document are introduced. Anatomy and physiology of the human heart are described in Sec. 2.1 and ways to measure it using imaging and other devices in Sec. 2.2.

The core of this thesis divides into two main parts. Part one is concerned with comprehensive tools for building patient-specific virtual hearts. It starts from the

description of the multi-physics computational 3D whole-heart model that is used throughout this thesis in Chap. 3. Its four main components, anatomy, electrophysiology, biomechanics and hemodynamics boundary conditions are described in Sec. 3.1, Sec. 3.2, Sec. 3.3 and Sec. 3.4, respectively. Then in Chap. 4, the proposed integrated pipeline for comprehensive personalization of the whole-heart model from the previous chapter, which enables streamlined generation of reproducible, individualized heart models, is described. An overview of state-of-the-art model personalization approaches is given in Sec. 4.1 and a generic formulation of the personalization problem in Sec. 4.2. The main ideas behind each module of the personalization pipeline, the derived personalization strategies, and the employed inverse optimization methods and cost functions are then detailed in Sec. 4.3. Comprehensive evaluation of all aspects of the pipeline on a large patient population in Sec. 4.4 shows the good match between personalized models and patient data, with errors often below clinical variability of the measurements. The predictive power of the model personalized using the proposed pipeline is then tested in a pilot experiment reported in Sec. 4.5. A short summary of the chapter in Sec. 4.6 concludes the first main part.

While the first part is focused on streamlined processing, reproducibility and large patient cohorts, which are important steps towards clinical translation, the second main part is more futuristic and may suggest novel research directions. First, Chap. 5 explores the feasibility and opportunities (Sec. 5.1) of applying modern artificial intelligence concepts to computational model personalization. The core ideas and methods of this chapter are then detailed in Sec. 5.2. An artificially intelligent agent with the ability to learn by itself how to personalize a model from the available data, while being model- and data-independent, is introduced. This is the first time that reinforcement learning has been applied to computational biophysical model personalization. A large variety of experiments and their results are described in Sec. 5.3, which suggest that the proposed algorithm can indeed learn how to personalize models while being faster and more robust than hand-crafted methods in many cases. The chapter is briefly summarized in Sec. 5.4. The second topic of part two in Chap. 6 is concerned with the various uncertainties in computational modeling, in particular with inverse uncertainty quantification. To this end, the sources of uncertainty in computational modeling are identified in Sec. 6.1. Next, the proposed method for parameter uncertainty quantification is detailed in Sec. 6.2. In a nutshell, a stochastic framework based on Bayesian inference, polynomial chaos expansion, and probability density aggregation for estimating an approximation of the posterior of the model parameters given the clinical measurements is proposed. Experimental results on real clinical data in Sec. 6.3 showed that not only the approach is as effective as standard methods, but that it is also computationally efficient and can provide rich information about uncertainties in the model parameters due to noise in the data. Finally, the main ideas and findings from this chapter are summarized in Sec. 6.4.

Chap. 7 then concludes this thesis with a comprehensive summary in Sec. 7.1 and an outlook for future work in Sec. 7.2. Supplementary material is provided in the appendix, Chap. A.

The Human Heart

2.1 Heart Anatomy and Function	9
2.2 Measuring Structure and Function of the Heart	11

The first part of this chapter, Sec. 2.1, provides a brief overview of the morphology and physiology of the human heart. Sec. 2.2 introduces commonly used modalities to measure it.

2.1 Heart Anatomy and Function

The heart is a powerful muscular organ that weighs between 148 and 296 grams and 233 and 383 grams in normal, young women and men, respectively [Moli 15, Moli 12], and is about the size of a fist. It acts as a pump through rhythmic contractions of the heart muscle at an average rate of 70 to 75 beats per minute, and thus is the main driver of the cardiovascular system. Throughout a human lifespan, a heart can beat more than 3.5 billion times, each day pumping 7 500 liters of blood in more than 100 000 beats.

2.1.1 Cardiovascular System

The heart and the *circulatory system* define the human *cardiovascular system*, a closed loop network that transports blood to ensure continuous oxygen and nutrition supply to the entire human organism, to organs, tissues and individual cells, while removing the waste products produced by them. Continuous and flawless functioning of the cardiovascular system is essential for survival. It consists of two main, interconnected loops: the *pulmonary circulation* and the *systemic circulation*. The pulmonary circulation is responsible for the oxygenation of blood, while carbon dioxide is eliminated by exhalation. The heart pumps oxygen-deficient blood through the pulmonary artery to the lungs, where it is oxygenated and then flows back to the heart through the pulmonary veins. The systemic circulation transports oxygenated blood to the entire body starting from the aorta, then the other great arteries, arterioles and finally into the capillaries, where oxygen and nutrients are released. The blood then flows back to the heart through venules, veins and inferior and superior venae cavae.




Figure 2.1: Diagram of the human heart depicting chambers and valves. Drawing adapted, original version by Wapcaplet / CC BY-SA 3.0.

2.1.2 Chambers and Valves

The heart is divided into four chambers as shown in Fig. 2.1: the two upper chambers *left atrium (LA)* and *right atrium (RA)*, and the two lower chambers *left ventricle (LV)* and *right ventricle (RV)*. The two left (LA and LV) and the two right chambers (RA and RV) form the *left heart* and the *right heart*, respectively. The heart has four valves that regulate the flow of blood through the heart chambers and to the systemic and pulmonary circulation. All four valves are located on approximately one level, the valve plane. Atria and ventricles are separated by the *atrioventricular* valves: the *mitral valve* lets oxygenated blood from the lungs pass from the LA into the LV, and the *tricuspid valve* regulates the flow of deoxygenated blood from the RA into the RV. The *arterial valves*, i. e., the *aortic valve* and the *pulmonary valve* are the entry points to the systemic and pulmonary circulation, as they are located at the base of the aorta and pulmonary artery, respectively. Opening and closure of valves is determined by pressure differences in front of and behind the valves. Hereby, the healthy valves' structures ensure unidirectional blood flow from the ventricle to the artery for the arterial valves, and from the atrium to the ventricle for the atrioventricular valves. The *papillary muscles* and the *chordae tendineae*, which connect the papillary muscles to the atrioventricular valves, support the valves to remain sealed when the pressure in the ventricles is higher than in the atria. Certain types of cardiovascular disease can prevent the valves from closing properly and thus blood can flow backwards. This condition is called *regurgitation*. Chronic regurgitations can cause functional impairments and myocardial dilation [Beau 13].

2.1.3 Myocardium

The myocardium is a thick wall of muscle tissue, which is encapsulated and protected by two thin layers of endothelial tissue: the *endocardium* inside the cavities, and the

epicardium outside. It consists of *cardiomyocytes*, the cardiac muscle cells, which contain organelles consisting of long chains of flexible *sarcomeres* that enable shortening and stretching of the cell, thus enabling a contracting and relaxing heart. The left and right ventricle cavities are separated by the “inner” section of the myocardium: the *septum*. The outer myocardium is referred to as the *free wall*. The thickness of the myocardium varies between the chambers depending on the force it needs to generate. The atrial walls are rather thin (approximately 2 mm) as not much contraction force is required to push blood from atria to ventricles. The right ventricular myocardium is thicker (approximately 5 mm) to ensure that sufficient force can be generated to push blood through the lungs and pulmonary circulation. The left ventricular myocardium is the thickest (approximately 15 mm) as it must pump blood through the systemic circulation, which is much larger and exhibits greater vascular resistance than the pulmonary circulation. The heart is contained in a thin sac, the *pericardium*, consisting of a thin two-layered fibrous membrane. To minimize friction and thus to enable smooth cardiac motion, a thin fluid layer is located between the pericardium and the epicardium.

2.1.4 Cardiac Activity and Electrical Conduction System

Cardiac activity is controlled by the *sympathetic nervous system* [Silv09]. The *sinoatrial node* can be seen as the natural pacemaker as it periodically generates the electrical impulse that causes cardiac contraction. First the atria are stimulated. The impulse then travels to the *atrioventricular node*, where it is first delayed, then propagated to the *bundle of His*. From there it reaches *Purkinje’s system*, a network of fast conducting *Purkinje fibers* organized in bundles that rapidly transmit and distribute the electrical impulse to the endocardia of the ventricles. From there an electrical wave propagates through the entire ventricular myocardium until reaching the epicardia. The speed of electrical propagation through the myocardium is controlled by the tissue conductivity, which can vary spatially, especially in the presence of myocardial scars or cardiac fibrosis.

Electrical signals trigger the contraction of cardiomyocytes by depolarizing them, which causes the heart beat. The heart beat is the result of two biophysical phenomena. First, an electrical wave traveling through the heart’s electrical conduction system activates the cardiomyocytes, as described above. Second, on cell-level, the cardiomyocytes contract and later relax, which, on a macroscopic scale, translates to the periodic global cardiac motion, the beating of the heart. Myocyte contraction is mainly caused by two proteins: actin and myosin, also called *myofilaments*.

2.2 Measuring Structure and Function of the Heart

2.2.1 Cardiac Imaging

Cardiac anatomy and function can be assessed by a variety of imaging modalities such as *ultrasound (US)*, *magnetic resonance imaging (MRI)*, *computed tomography (CT)* and many others. In most cases, a cardiac US scan (echocardiography) is the first method of choice due to its low cost and fast acquisition speed. It can




Figure 2.2: Example **SSFP** cardiac cine **MRI**: one diastolic (maximum ventricular blood volume) and one systolic (contracted heart, minimum blood volume) frame among typically 25–30 frames are visualized. Images courtesy of Siemens Healthineers.

be used to measure the time-varying blood volume in the ventricles at the different heart phases and thus to compute important clinical indicators such **ejection fraction (EF)** and **stroke volume (SV)** based on the end-systolic and end-diastolic volume measurements. In addition, cardiac dynamics and motion can be assessed, however the imaging quality, resolution and signal-to-noise ratio are typically low, even with modern **US** devices, due to physical limitations. **CT** offers high quality anatomical information and fast acquisition speed at the cost of ionizing radiation, which can be harmful to the patient. **MRI** could be seen as a good compromise between **US** and **CT** as it can produce good quality time-resolved images without ionizing radiation, but at the cost of slow acquisition speeds of up to one hour depending on the imaging protocol. In this work, **steady-state free precession (SSFP)** cardiac cine **MRI** is the main source of (spatio-temporal) cardiac imaging information, see Fig. 2.2 for an example. Temporal images are acquired, which enable monitoring of cardiac motion and estimation of time-varying ventricular volumes.

2.2.2 Electrocardiography

The electrical activity of the heart can be recorded indirectly using electrodes on the surface of the body. The electrodes are designed to detect small electrical gradients caused by the depolarization and repolarization of muscle cells in the heart and propagated through the torso to the skin. These small changes are typically recorded continuously over a short period (e.g., one or more heart beats) in the form of an **electrocardiogram (ECG)** on a sheet of paper, or electronically. In today’s clinical cardiology practice, **ECG** is commonly used for diagnosis of heart failure due to its low cost, non-invasiveness, and high versatility. It can provide helpful information in a wide range of situations, for instance by diagnosing myocardial infarctions, identifying life-threatening cardiac arrhythmias, and many more [Thal 10]. In this work, standard 12-lead **ECG** is used, consisting in Einthoven and Goldberger limb leads (I, II, III, aVR, aVL, aVF), see Fig. 2.3, and Wilson precordial leads (V1, V2, . . . , V6). The Einthoven leads, i. e., the three basic bipolar limb leads (I, II, III) are computed from the electrodes on the extremities:

- I: left arm – right arm

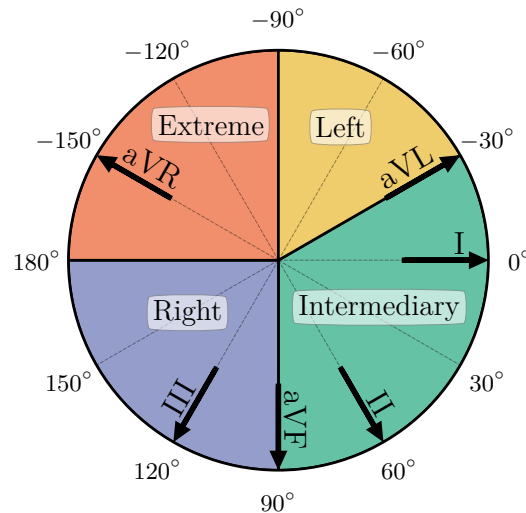


Figure 2.3: Einthoven and Goldberger leads and their correspondence to the heart’s electrical axis (EA). EA between -30° and 90° is considered normal. Variations can be classified as left, right or extreme axis deviation, represented by the different colors in the plot

- II: left leg – right arm
- III: left leg – left arm

The Goldberger leads, i. e., the three augmented limb leads (aVR, aVL, aVF) are derived from the same electrodes, however they enable a look at cardiac electrophysiology from different angles:

- aVR: right arm – (left arm + left leg) / 2
- aVL: left arm – (right arm + left leg) / 2
- aVF: left leg – (right arm + left arm) / 2

Finally, the Wilson leads, i. e., the six precordial leads (V1, V2, . . . , V6) are obtained by directly relating the six chest electrodes with Wilson’s central terminal: (right arm + left arm + left leg) / 3.

Among the most clinically relevant ECG features are QRS duration, i. e., the width of the QRS complex, QT duration, and electrical axis (EA) of the heart, which represents direction of the average electrical depolarization in the heart. EA can be classified into intermediary (normal) or deviated: left deviation, right deviation or extreme deviation, see Fig. 2.3.

2.2.3 Blood Pressure

Arterial pressure is commonly measured non-invasively using a *sphygmomanometer*, also known as blood pressure meter or “cuff” pressure. It is usually placed around the upper arm and consists of an inflatable cuff, which is inflated while the measurements

are being taken, and a measuring unit. Sphygmomanometers measure mean arterial pressure (MAP), diastolic arterial pressure (DBP) and systolic arterial pressure (SBP) within the brachial artery by oscillometric detection. When precise aortic or intra-cardiac pressure information such as ventricular or atrial pressures are required (see Fig. 2.4), *cardiac catheterization*, or “cath” pressure, is the method of choice, both for diagnostic and interventional purposes. In a minimally invasive procedure, a catheter is inserted into a chamber or vessel of the heart, accessed via, e.g., the femoral vein. Once the catheter reached its desired position, pressures are being monitored continuously over a certain time frame, yielding time-varying pressure curves covering one or more heart cycles.

2.2.4 Normal Values and Relations

Some key clinical measurements that can be extracted using the methods and modalities described above and their normal values according to [Mace 06a, Mace 06b, Klin 11] are listed in Tab. 2.1. Relationship between various time-varying measure-

Table 2.1: Measurements of heart function and corresponding normal values for the healthy population.

Measure	Modalities	Normal values
Heart rate	ECG	60–100 bpm
End-diastolic volume (EDV)	Imaging	142 ± 21 ml
End-systolic volume (ESV)	Imaging	47 ± 10 ml
Stroke volume (SV)	Imaging	95 ± 14 ml
Ejection fraction (EF)	Imaging	67 ± 5 %
Cardiac output (CO)	Imaging	4–8 L/min
QRS duration	ECG	0.06–0.1 s
QT duration	ECG	≈0.4 s
Electrical axis (EA)	ECG	−30°–105°
Diastolic arterial pressure (DBP)	Cath, cuff	60–90 mmHg
Systolic arterial pressure (SBP)	Cath, cuff	90–140 mmHg
Mean arterial pressure (MAP)	Cath, cuff	70–105 mmHg
Diastolic ventricular pressure	Cath	3–12 mmHg
Systolic ventricular pressure	Cath	100–140 mmHg

ments and certain cardiac events (e.g., valve opening and closing) during the course of the cardiac phases can be illustrated nicely in a so-called Wiggers diagram [Wigg 24] as shown in Fig. 2.4.




Figure 2.4: Wiggers diagram for two heart cycles of a healthy heart. Ventricular pressure and volume curves of a healthy left ventricle (LV) are depicted. ECG trace shows electrical waves (P, Q, R, S and T wave). Drawing adapted, original version by D. Chang / CC BY-SA 4.0.

PART I

Patient-specific Virtual Hearts

Image-Based, Multi-Physics Model of Heart Function

3.1 Anatomy	19
3.2 Electrophysiology	23
3.3 Biomechanics	29
3.4 Boundary Conditions	32

The heart is a highly complex organ, in which a variety of biophysical systems interact with each other at different scales. A cardiac model is a representation of the complex organ using a set of mathematical relationships. The model used in this work was developed by a team of researchers at Siemens Healthineers [Zett 13]. It consists of a set of coupled components, which can be computed independently, but interact with each other, see Fig. 3.1. Each of the four main components models an integral part of the heart morphology or physiology: anatomy, electrophysiology (EP), biomechanics and hemodynamics (HD) boundary conditions, as described in the following sections.

3.1 Anatomy

First, the patient-specific heart morphology and dynamics need to be modeled.

3.1.1 Geometry Segmentation from Imaging Data

In order to derive a patient-specific geometric representation of the cardiac anatomy, previously developed comprehensive geometric models of the heart [Zhen 07, Iona 10] are integrated. Physiological landmarks are explicitly encoded in the model and provide semantic associations to the underlying anatomy. The models are highly modular and can be customized depending on the application. For the purpose of this work, models of the left and right ventricle are selected. The models have explicit geometrical representations for the left ventricle endocardium, epicardium, mitral annulus, left ventricular outflow tract, ventricular regions, tricuspid and pulmonary valve locations. The models can be personalized from a variety of 3D imaging modalities, including computed tomography (CT), magnetic resonance imaging (MRI) or ultrasound (US), and the model parameters are automatically determined from imaging data using a database-guided machine learning (ML) framework [Geor 05, Wang 13]. The framework provides robust parameter estimation for complex cardiac models by




Figure 3.1: Schematic overview of the four main components of the computational heart model and their interactions (arrows).

exploiting expert knowledge from annotated medical imaging databases. Discriminative learning methods with semantic constraints have proven to be the solution of choice for estimation problems in high-dimensional spaces. While several other image segmentation algorithms to extract cardiac chambers exist [Fran 01, Peng 16], marginal space learning (MSL) [Zhen 08, Zhen 14] is among the most efficient ways to learn high-dimensional models and perform fast online search by operating on spaces of increasing dimensionality. The discriminative classifiers are trained with a probabilistic boosting tree [Tu 05], which accounts for patterns of large intra-class variability for complex data distributions. Using such methods that learn context and prior knowledge, even whole-body volume analysis can be performed efficiently [Seif 09]. Recently, the successor of MSL, marginal space deep learning (MSDL) [Ghes 16b], was introduced and showed potential towards further improving the performance of such algorithms by making use of the latest advancements in modern deep learning (DL). Motion manifold learning [Yang 08] is a state-of-the-art method for estimating temporal components and dynamics of the heart models derived explicitly from temporal patient scans such as steady-state free precession (SSFP) cine MRI.

Volumetric Representation

As shown in Fig. 3.2, after fitting the geometric model to the patient data at end-diastasis, the resulting left and right ventricular anatomies are fused into a single bi-ventricular volumetric tetrahedral mesh representation using the CGAL library [CGAL 17]. The vertices of the mesh are then tagged into surface zones according to the underlying anatomy. Myocardial scars or fibrosis identified in images like delayed-enhancement MRI can be mapped onto the anatomical model to simulate their specific properties.

3.1.2 Automatic Mesh Tagging

The anatomical model is enhanced with tags that represent certain types of tissues or tissue regions useful for modeling. Two such tags are highlighted here, fibrous tissue, and regional segments of the left ventricle (LV).

Fibrous Tissue Mapping

The aim of the fibrous tissue tags is to properly model the inactive electrical property of this tissue type (collagen tissue does not depolarize as it does not contain any

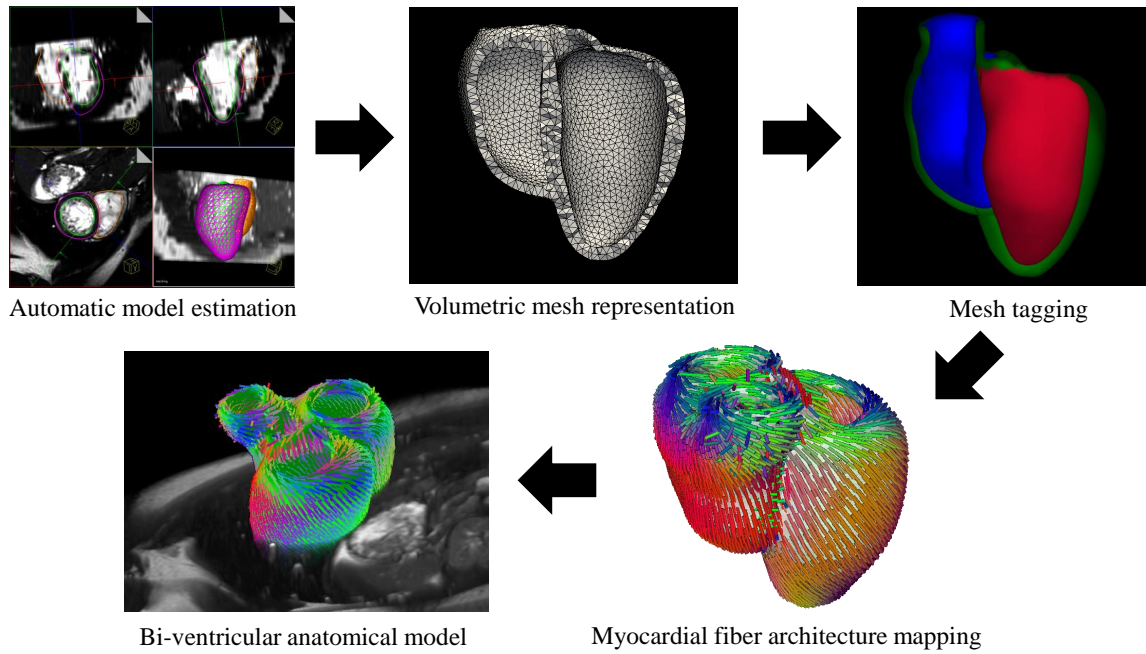


Figure 3.2: From images to patient-specific anatomical models.

myocyte-type cell). Based on communication with clinical collaborators, a rule-based classification of fibrous tissue is employed. It comprises the fibrous rings of the pulmonary and aortic valves as well as fibrous connections linking these rings to the atrioventricular valves. Therefore all mesh elements in the left ventricular outflow tract as well as all mesh elements in the right ventricular outflow tract and above the plane of the atrioventricular valves are tagged as fibrous tissue.

Automatic AHA Segmentation

In order to enable modeling regionally varying properties of the left-ventricular myocardium, such as regional alterations of conductivity (**EP**) in left-bundle-branch block patients, or regional alterations in active and passive force (biomechanics) due to myocardial scars, the **LV** is divided into 17 segments according to the widely-used standard established by the **American Heart Association (AHA)** [Cerq02], see Fig. 3.3, left panel.

The **AHA** segmentation procedure takes as input the tagged tetrahedral heart model and is performed as described below:

1. Compute the long axis of the heart defined by two points: the barycenter of the mitral valve and the **LV** barycenter.
2. Compute the endocardial apex, i. e., the point on the endocardium that is closest to the apex of the myocardium, by intersecting the **LV** endocardial surface with the long axis.
3. Divide **LV** endocardium into equal thirds perpendicular to the long axis (top: basal; mid: mid-cavity; bottom: apical). All myocardial mesh cells (tetrahedra)




Figure 3.3: Left: Bull’s eye representation of the 17 standardized American Heart Association (AHA) segments of the LV. The two outermost circles define basal (segments 1–6) and mid (segments 7–12) anterior, anteroseptal, inferoseptal, inferior, inferolateral and anterolateral segments, respectively. Segments 13–16 are apical segments: anterior, septal, inferior and lateral, and 17 is the apex segment. Right: AHA segmentation example. Different colors represent the 17 segments. Note that the right ventricle (RV) and connective tissue, visualized in dark blue color, are not part of the AHA segment model.

whose barycenters are below the apical region are classified as segment 17, the apex segment.

4. Further divide the three regions into circular segments by defining cutting planes along the long axis. The plane orientations are based on angles (short-axis plane) with respect to the vector connecting the LV and right ventricle (RV) barycenters. The left ventricular myocardial cells are then divided into:
 - (a) Basal and mid-cavity regions: six segments of 60° each.
 - (b) Apical region: four segments of 90° each.

The process described above runs fully automatically. An example result is illustrated in Fig. 3.3, right panel.

3.1.3 Rule-based Myocardial Fiber Architecture

Myocardium fibers play a crucial role in cardiac function. In particular, the myocardial tissue has orthotropic mechanical properties. It is therefore sensitive to fiber orientation and to the arrangement of fibers in sheets. Moreover, electrical conductivity is affected by fibers, because the electrical wave propagates faster along the local direction of fibers. Integrating fiber architecture into the anatomical model is therefore essential. The distribution of the fibers across the myocardium has been widely studied on *ex vivo* hearts using histology or diffusion tensor imaging (DTI) [Vett 98, Peyr 07, Pop 09]. Because measuring them *in vivo* is still an open challenge [Wu 09, Tous 13], atlases of fiber orientations are being computed from dog hearts [Peyr 07] or human hearts [Lomb 12]. However, it is not clear how well such atlases translate to pathological hearts.

While advanced *in vivo* DTI acquisition protocols are currently being developed by the community, the fiber architecture used in this work relies on a rule-based model [Serm 06a, Baye 12, Zett 13], if not stated otherwise. The model, which is derived from *ex vivo* studies [LeGr 95, Arts 05], covers fiber orientation and fiber sheets and is automatically mapped to the patient-specific anatomy. It can be adjusted by the user to adapt it to the investigated pathology.

The rule-based fibers are computed as described in [Zett 13]. Below the basal plane, the fiber elevation angle \angle_{fiber} , i. e., the angle with respect to the short-axis plane, varies linearly across the myocardium from $\angle_{\text{fiber}}^{\text{endo}} = -66^\circ$ (-80° for RV) on the epicardium to $\angle_{\text{fiber}}^{\text{epi}} = +66^\circ$ ($+80^\circ$ for RV) on the endocardium. Similarly, the sheet directions, defined by an angle \angle_{sheet} with respect to the outward transmural axis, vary transmurally from $\angle_{\text{sheet}}^{\text{endo}} = +45^\circ$ on the epicardium to $\angle_{\text{sheet}}^{\text{epi}} = -45^\circ$ on the endocardium. Note that these values could be personalized as well, if appropriate data such as DTI is available. Between the apex and the basal plane, for every vertex \mathbf{v} of the tetrahedral myocardium mesh, the elevation angle $\angle_{\text{fiber}}^{\mathbf{v}}$ is varied according to the geodesic distance of \mathbf{v} to the endocardium d_{endo} and epicardium d_{epi} :

$$\angle_{\text{fiber}}^{\mathbf{v}} = \frac{d_{\text{epi}} \cdot \angle_{\text{fiber}}^{\text{endo}} + d_{\text{endo}} \cdot \angle_{\text{fiber}}^{\text{epi}}}{d_{\text{epi}} + d_{\text{endo}}} . \quad (3.1)$$

The sheet angles $\angle_{\text{sheet}}^{\mathbf{v}}$ are varied accordingly.

Next, the fiber and sheet orientations around each valve are computed. Following [Moir 09], around the aortic valve the fibers are longitudinal, and tangential elsewhere. The sheet normals are oriented towards the barycenter of each valve. The local orthonormal basis is interpolated from the basal plane to the valve, first by following the myocardium surface and then transmurally.

All these operations are performed per vertex. As the employed computational model solvers require the coordinate system to be defined per tetrahedron, the tetrahedral fiber coordinate systems are computed using barycentric interpolation. To preserve orthonormality, the Log-Euclidean framework [Arsi 06] is used for all interpolation operations.

3.2 Electrophysiology

The first biophysical phenomenon that needs to be simulated is cardiac **electrophysiology (EP)**, which is one of the main determinants of cardiac function as it triggers myocardial contraction. Cardiac EP is solved on the patient-specific anatomical model (Sec. 3.1). Since the seminal work of [Hodg 52, Fitz 61, Nagu 62], a large variety of computational EP models have been proposed [Clay 08, Elsh 14] to deal with different biological scales and theoretical complexities. These models can be organized into three different categories: biophysical, phenomenological and Eikonal.

Biophysical Models

The most complex category of models aims to simulate the ionic interactions across the cell membrane and the biological phenomena underlying the ion channels [Nobl 62,

[Luo 91, Nobl 98, Ten 04]. These models are highly detailed and controlled by a large number of parameters. The cell models are then integrated at the tissue scale using semi-linear reaction-diffusion equations, where two main categories exist: mono-domain and bi-domain. The more computationally efficient mono-domain approach considers the myocardium as a single excitable tissue, while bi-domain separates the intra- and extra-cellular domains [Coud 06, Bour 09], which allows investigation of biological phenomena happening outside the cell.

Phenomenological Models

Historically, phenomenological models were the first to be developed, as they can be derived from experimental observations on nerves or cardiac tissue samples [Fitz 61, Nagu 62]. These models aim to reproduce the shape of the experimentally observed action potential (AP) and its variation under changing external conditions. While also integrated at the organ scale using mono-domain [Alie 96, Fent 98, Mitc 03, Corr 16] or bi-domain [Clay 08] equations, phenomenological models are simplifications of biophysical models as they do not compute the underlying ionic phenomena. This allows for a much smaller number of parameters (typically two or three), which are often directly related to the shape of the AP or electrocardiogram (ECG) measurements.

Eikonal Models

If the mathematical formulation is further simplified to the sole propagation of the electrical wave, one speaks of Eikonal models. In this category, the AP is not directly simulated. Instead, the propagation of the electrical wave is reproduced by computing its arrival time at each location on the myocardium [Fran 90, Serm 07]. The main advantages of Eikonal models are their high computational efficiency and simple parametrization using only one or two parameters. For example, solvers based on fast marching methods [Serm 07] enable almost real-time simulations, where a full heart depolarization can be simulated in seconds or less. Due to the drastic simplifications, however, many pathological phenomena cannot be simulated, including arrhythmias, fibrillations or tachycardia.

3.2.1 First Model: LBM-EP

For the purpose of this work it is sufficient to capture normal electrophysiology, cardiac asynchrony and minor to mild arrhythmias as observed in typical dilated cardiomyopathy (DCM) patients. While the modeling of ion channels may be useful for particular applications, a strong requirement for this work was a low number of parameters to overcome issues of identifiability and observability to be able to adjust the model to patients from limited data. Therefore, the cells are considered mono-domain spaces. The trans-membrane potential is computed directly. In-flow and out-flow channels are lumped in two variables to further limit the number of parameters of the system.

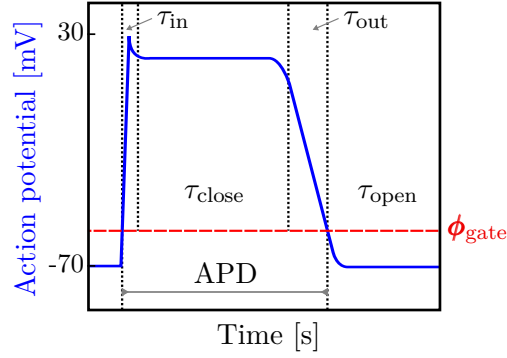


Figure 3.4: Qualitative plot of cardiac action potential (AP) voltage trace in a myocardial cell illustrating relationship between model parameters, action potential duration (APD) and shape of the AP.

Selected Cell Model

Phenomenological models are a good trade-off between modeling detail and computational complexity. Therefore, the Mitchell-Schaeffer mono-domain model proposed in [Mitc 03] is used. The time-varying trans-membrane potential field ϕ throughout the myocardium is described with the following equation:

$$\frac{\partial \phi}{\partial t} = J_{\text{in}} + J_{\text{out}} + J_{\text{stim}} + \nu \cdot \text{div}(\mathbf{D}\nabla\phi) , \quad (3.2)$$

scaled to the interval -70 mV to 30 mV. The model simplifies all ion channel interactions to only two currents. First, the inward gated current:

$$J_{\text{in}} = \frac{h\phi^2(1-\phi)}{\tau_{\text{in}}} , \quad (3.3)$$

which captures the fast acting ionic currents in the myocyte. The gating variable $h(t)$ models the state of the ion channels and is defined based on the change-over voltage ϕ_{gate} . It is computed as $\frac{dh}{dt} = (1-h)/\tau_{\text{open}}$ if $\phi < \phi_{\text{gate}}$ and $\frac{dh}{dt} = -h/\tau_{\text{close}}$ otherwise. Second, the ungated outward current:

$$J_{\text{out}} = \frac{-\phi}{\tau_{\text{out}}} , \quad (3.4)$$

which accounts for trans-membrane voltage decrease. A transient stimulus current J_{stim} can further be added to mimic electrical pacing. The time constants τ_{in} , τ_{out} , τ_{open} and τ_{close} are parameters of the model, which are related to the shape and duration of the action potential as illustrated in Fig. 3.4.

The last term in Eq. (3.2) models spatial propagation by diffusion, which is mainly controlled by diffusion coefficient ν and anisotropic diffusion tensor $\mathbf{D} = (1-\rho)\mathbf{I} + \rho\mathbf{o}\mathbf{o}^T$ along the fiber direction \mathbf{o} . ρ denotes the anisotropy ratio and \mathbf{I} the 3×3 identity matrix. The operators ∇ and div denote gradient, i. e., direction of steepest increase of ϕ , and divergence, a measure of increase in the given direction, respectively.

Fast Solver

Eq. (3.2) is solved using LBM-EP [Rapa 12], which is based on Lattice-Boltzmann method (LBM). It provides excellent scalability on modern computing architectures due to local computations that can be performed in parallel. LBM-EP models EP using fictive particles performing consecutive collision and propagation processes on a discrete Cartesian grid with 7-connectivity topology (six neighbors and central position) and Neumann boundary conditions. For each of the seven connections, a function $g_i(\mathbf{v}, t)$, $i = 1 \dots 7$, represents the probability of finding a particle traveling along the respective edge \mathbf{e}_i of node \mathbf{v} at time t . The trans-membrane potential $\phi(\mathbf{v}, t)$ is computed as the sum $\phi(\mathbf{v}, t) = \sum_{i=1}^7 g_i(\mathbf{v}, t)$.

The computation of g_i is decomposed into two phases: collision and streaming. First, the collision phase yields intermediate post-collision states $g'_i(\mathbf{v}, t)$, taking into account the anisotropic diffusion due to local fiber orientation and emphasizing the contribution from the central position by applying connection-dependent weighting factors. Due to node-wise computations, the collision phase can be implemented efficiently on a graphics processing unit (GPU). Second, the streaming phase propagates these functions along their corresponding edges: $g_i(\mathbf{v} + \mathbf{e}_i, t + \delta_t) = g'_i(\mathbf{v}, t)$, where δ_t is the time-step, and applies the boundary conditions. For more details see [Rapa 12].

After LBM-EP computations are finalized, the time-varying trans-membrane potentials are mapped back from the Cartesian domain to the tetrahedral volume mesh using tri-linear interpolation.

Purkinje Fibers

As described in Sec. 2.1, the electrical conduction system of the heart is composed of several elements including the Purkinje fibers (Purkinje's system). This macro-scale system can be implemented in the framework of EP solvers based on Cartesian computational grids, such as LBM-EP: the endocardial surface is rasterized on the computational grid and the grid cells corresponding to endocardium are assigned a high diffusion coefficient ν_{Purk} . A limitation of this approach is that the thickness of the layer of high-speed conducting tissue cannot be controlled, i. e., it is always in the order of one grid cell. Thus the Purkinje thickness changes if the grid resolution changes, which is undesirable. To overcome this limitation, a rule to classify the myocardial tissue as part of the high-speed bundles based on its distance from the endocardial septum is defined. Tissue diffusion coefficient ν is defined as a piece-wise constant field over the Cartesian grid, modeling each grid point as the center of a voxel. The value assigned to each grid point \mathbf{v} , $\nu(\mathbf{v})$, ranges from normal (ν_{normal}) to high (ν_{Purk}), based on the fraction φ of volume inside the voxel whose distance from the endocardium is smaller than a user-defined thickness threshold:

$$\nu(\mathbf{v}) = \varphi \cdot \nu_{\text{Purk}} + (1 - \varphi) \cdot \nu_{\text{normal}} \quad . \quad (3.5)$$

3.2.2 Second Model: Graph-EP

Although LBM-EP provides significant gain in computational efficiency compared to traditional methods (up to 45 times faster [Rapa 12]), the forward computation of one

heart beat is still in the order of seconds, depending on the available processing power, LBM grid resolution and size. This may not sound much, however, in the context of model personalization, where a forward model may have to run up to a few hundred times in order to converge to reasonable parameter estimates, these few seconds sum up quite rapidly. There is therefore a need for even faster EP solvers. Such solvers typically come at the price of reduced accuracy, which may however be acceptable in certain scenarios. A highly efficient graph-based method for computation of EP depolarization, inspired by [Wall 12], is called Graph-EP, an Eikonal EP model. The computational domain of Graph-EP is the tetrahedral mesh of the patient’s anatomical model, taking into account its per-vertex fiber orientations. Graph-EP computes the activation time at each node of the mesh, i. e., the time at which the electrical wave arrives, in the order of a few milliseconds.

Algorithm

The algorithm is based on a variant of Dijkstra’s shortest path algorithm [Dijk 59]. First, the nodes from where the electrical wave starts are added to a queue, and a value equal to their stimulation time is stored. In particular, the LV and RV septum nodes are added with an activation time corresponding to the atrioventricular delay. Any additional activation point (e. g., due to pacing devices) can be added to the queue analogously. The first node of the queue is then popped and all its neighbors in the graph are processed. Let \mathbf{v}_i be the node that is currently being processed, and \mathbf{v}_j one of its neighbors. A tentative activation time $t'_{\text{act}}(j)$ is computed as $t'_{\text{act}}(j) = t_{\text{act}}(i) + t_{\text{travel}}(i, j)$, where $t_{\text{travel}}(i, j)$ is the travel time from vertex i to j given by $t_{\text{travel}}(i, j) = \text{edge_cost}(i, j)/\nu_{ij}$, where ν_{ij} the conduction velocity of the tissue crossed by the edge $(\mathbf{v}_i, \mathbf{v}_j)$. If $t'_{\text{act}}(j)$ is smaller than the current activation time estimate $t_{\text{act}}(j)$, then $t_{\text{act}}(j)$ is updated, $t_{\text{act}}(j) = t'_{\text{act}}(j)$ and the node \mathbf{v}_j is added to the queue for later processing. This process is iterated until the queue is empty, i. e., no more nodes need to be processed and the graph (tetrahedral heart mesh) is fully processed. Finally, the potentials are mapped by assigning at any given time a value of -70 mV to the node that are still not activated, 30 mV otherwise.

Tissue Anisotropy

Anisotropy of myocardial tissue is modeled by modifying the edge cost to take into account fiber orientation:

$$\text{edge_cost}(i, j) = \frac{\sqrt{(\mathbf{v}_j - \mathbf{v}_i)^\top \mathbf{D} (\mathbf{v}_j - \mathbf{v}_i)}}{\|(\mathbf{v}_j - \mathbf{v}_i)\|}, \quad (3.6)$$

where \mathbf{D} is the anisotropy tensor defined by $\mathbf{D} = (1 - \rho)\mathbf{I} + \rho\mathbf{o}_{ij}\mathbf{o}_{ij}^\top$ and ρ denotes the anisotropy ratio, \mathbf{I} the identity matrix and \mathbf{o}_{ij} the edge-based fiber direction (cf. Sec. 3.1). Since fibers are defined node-wise, $\mathbf{o}_{ij} = (\mathbf{o}_i + \mathbf{o}_j)/2$.

Purkinje fibers

In Graph-EP, the Purkinje network is modeled by edge-based linear interpolation of the fast conduction velocities.




Figure 3.5: Overview of mapping procedure from cardiac extra-cellular potentials to torso potentials and ECG.

3.2.3 From Cardiac Potentials to Torso Potentials to ECG

The standard clinical measure of cardiac electrical activity is ECG, which is derived from electrical signals measured on the surface of the patient’s body (cf. Sec. 2.2.2). In order to compare measured activity with computed activity, one needs to propagate the potentials computed on the cardiac domain to the torso, and from there derive the ECG leads and features.

Torso Fitting

First, patient-specific torso geometry is either segmented automatically from the images or, when the full torso is not available, the following procedure is applied. First, the contours of the patient’s torso are outlined in coronal, sagittal and transverse slices of the survey MRI sequences and visualized together with the heart model. Then, an atlas of torso geometry, obtained from CT images, is manually registered to these contours (affine transformation).

Torso Potentials

Next, a three-step procedure is employed to compute potentials on the patient-specific torso mesh based on the previously computed cardiac EP as described in [Zett 14]. First, extra-cellular potentials are estimated from the trans-membrane potentials, which were computed using LBM-EP, by using an elliptic formulation [Chha 12]. In the case of Graph-EP, the computed potentials are used directly. Second, the extra-cellular potentials are mapped back to the epicardium surface mesh using trilinear interpolation. Third, the boundary element method (BEM) [Shou 09] is used to project the extra-cellular potentials onto the torso surface, from where ECG lead signals can be derived, see Fig. 3.5.

Electrocardiogram

Finally, based on the resulting torso potentials, the leads of the standard 12-lead ECG, namely the Einthoven and Goldberger limb leads (I, II, III, aVR, aVL, aVF) and the Wilson precordial leads (V1, V2, ..., V6), are computed, see Sec. 2.2 for details. ECG descriptors are derived from the computed leads as described in [Zett 14]. In particular, QRS duration Δ_{QRS} is detected by convolving the squared derivative of each limb lead with a sliding average kernel and a threshold operation [Kohl 02],

Electrical axis (EA), \angle_{EA} , is computed based on a trigonometric operation performed on the R and S peak amplitudes in leads I and II: $\angle_{EA} = \arctan(2h_{II} - h_I)/(\sqrt{3}h_I)$, where h_I and h_{II} are the sum of R and S peak amplitudes in the respective lead.

3.3 Biomechanics

The second biophysical phenomenon to be modeled is cardiac biomechanics [Tray 11, Kris 13, Lama 14, Wang 15]. The myocardium is an active, non-linear, anisotropic viscoelastic tissue. Its constitutive law, which describes its elastic behavior, comprises an active and a passive component. The active component is the active contraction of the myocytes controlled by the action potential (AP). The passive component relates to the elasticity of the tissue. Both components are linked together through the Hill-Maxwell framework [Fung 93]. In practice, the active contraction is viewed as a transient external force that makes the myocardium contract. The passive properties of the tissue are internal forces that ensure realistic motion [Hunt 03].

3.3.1 Passive Properties

Models of myocardium passive properties aim to reproduce how the myocardium deforms under given stress and boundary conditions. Mathematically, they solve the partial differential equations modeling constitutive laws of tissue behavior on the patient's anatomical model. A multitude of models have been proposed in the last decades to simulate myocardium passive properties, ranging from simple isotropic linear elasticity [Serm 06b] to more realistic orthotropic non-linear hyper-elasticity [Holz 09, Cans 15]. Linear elasticity assumes linear relationship between strain and stress. It is often implemented within the infinitesimal, linear strain theory (small deformation assumption) for computational efficiency, which, however, becomes inaccurate for large deformations.

Linear models are limited in their application since heart tissue is non-linear. Constitutive laws have been derived by stretching slabs of myocardium tissues in several directions using bi-axial machines to measure tissue strain under known load [Hunt 88, Hump 90a, Hump 90b, Gucc 93, Holz 09]. The most common laws are the pole-zero law [Hunt 88] and the more recent Costa law [Cost 01]. The latter considers both fiber orientation and their organization in sheets across the myocardium. Yet, improving model accuracy is achieved at the price of complexity, with increasing number of parameters. The Costa law, for instance, is governed by seven parameters, most of them difficult to estimate *in vivo*.

In this work, a simplified version of the constitutive model proposed by Holzapfel and Ogden [Holz 09] is used. Contrary to other, more phenomenological constitutive laws like Costa [Cost 01] or Guccione [Gucc 93], the Holzapfel-Ogden (HO) model is derived from considerations of the tissue structure and not by fitting exponential functions to stress-strain relationship observed experimentally. The model starts from the observation that cardiac tissue shows different behaviors whether it is stretched along the fiber, tangential to the sheet structure, perpendicular to the sheets or in the fibers-sheet direction. For all these conditions, the stress-strain relationship is exponential. In this work, terms related to sheets are neglected, because measuring

sheet orientation *in vivo* is difficult. Based on these considerations, the simplified HO stress-strain energy function HO comprises two exponentials:

$$\text{HO}(\beta) = \beta \cdot \frac{a}{2 \cdot b} \exp(b \cdot (I_1 - 3)) + \beta \cdot \frac{a_f}{2 \cdot b_f} \left\{ \exp(b_f \cdot (I_{4f} - 1)^2) - 1 \right\} , \quad (3.7)$$

where a , b , a_f and b_f are tissue parameters and I_1 and I_{4f} are invariants of the deformation tensor, see [Holz 09] for more details. To simplify the tuning of the model, a dimensionless factor $\beta > 0$ is introduced, which allows isotropic scaling of tissue stiffness by modifying a single parameter only. To this end, $a = 0.496$ kPa, $b = 7.209$, $a_f = 15.193$ kPa and $b_f = 20.417$ are set to their default values [Gokt 11], i. e., when $\beta = 1$, the model behaves like normal *ex vivo* myocardial tissue, if the value is larger, the tissue is stiffer, and softer if β is between zero and one.

3.3.2 Active Force

The second component of the biomechanical model mimics the active force (myocyte contraction) generated by every cell when it depolarizes, which fades out once the cell repolarizes. Three categories of active contraction models can be distinguished.

The first category is biophysical models that simulate the ion interactions and the actin-myosin bindings that generate cardiac motion [Hunt 88, Nied 08, Rice 08]. Most biophysical models are highly detailed and were derived from experimental studies on *ex vivo* animal hearts for hypothesis testing. In recent years, also clinically-focused studies are appearing [Tray 11, Nied 12, Mang 18]. The model proposed in [Hunt 98] and its extension [Nied 08] are commonly used for organ simulations as they stay computationally tractable, however, they rely on biophysical EP models. Today's standard for single cell simulations is [Rice 08], as it captures most of the sub-cellular and protein mechanisms involved in myofilament function, from troponin function and ion binding to force generation. Scaling up this model to the organ level is, however, challenging, as it comprises more than 40 differential equations per cell that need to be integrated at every time step.

The idea behind the second category, multi-scale phenomenological models, is to mathematically integrate the biological mechanisms spanning from the actin-myosin interaction to the organ [Chap 01, Best 01]. The transition from one scale to another (e. g., from the calcium concentration to actin-myosin binding) is achieved mathematically, ultimately resulting in a set of simplified equations that are controlled by fewer parameters (usually only four to five). In [Serm 06a], a simplified version of these models with analytical integration for model-based image analysis was proposed. More recently, a more comprehensive multi-scale model that also considers energy exchange during the heart cycle, in particular the relation between oxygen supply and energy consumption, was proposed [Chap 12].

The third category are lumped models, the most simple active contraction models consisting of analytical formulations of fiber contraction that do not consider spatial variability [Arts 91], thus they do not require meshes to be solved. Lumped models can be solved very efficiently but they cannot capture regional abnormalities of the myocardium, such as scars or localized fibrosis.

In this work, a simplified phenomenological model is used, assuming that active contraction is directly related to the action potential (AP) through a multi-scale, lin-




Figure 3.6: Qualitative plot showing variation of time-varying active force (blue curve), which is controlled by the the maximum asymptotic contraction force $\hat{\sigma}$, and the contraction and relaxation rates k_{ATP} and k_{RS} (red arrows).

ear law that relates trans-membrane voltage with the rates of adenosine triphosphate (ATP) binding and release.¹ Other sub-cellular mechanisms are not modeled. The model is a linearization of the phenomenological model proposed in [Best 01], derived in [Serm 06a]. It is computational efficient and has only few parameters, which are all clinically-related. The active force $\sigma(t)$ of a cell at a given time t , is controlled by depolarization and repolarization times, t_{dep} and t_{rep} , which are computed by the EP model, cf. Sec. 3.2.

$$\sigma(t) = \begin{cases} \hat{\sigma} \cdot [1 - \exp(k_{ATP} \cdot (t_{dep} - t))], & \text{if } t_{dep} \leq t \leq t_{rep} \\ \sigma(t_{rep}) \cdot \exp(-k_{RS} \cdot (t_{rep} - t)), & \text{if } t_{rep} < t < t_{dep} + t_{HP} \end{cases} \quad (3.8)$$

where t_{HP} is the duration of one heart cycle. The governing parameters of the model are contraction and relaxation rate, $k_{ATP} > 0$ and $k_{RS} < 0$, and the maximum asymptotic contraction force, $\hat{\sigma}$. An illustration of the time-varying active force is provided in Fig. 3.6.

3.3.3 Finite Element Modeling

The active and passive component of the biomechanical model are solved on the anatomical model using finite element method (FEM). The motion of the heart resulting from these laws is computed by solving the dynamic system for all the vertices of the mesh

$$\mathbf{M}\ddot{\mathbf{u}} + \mathbf{C}\dot{\mathbf{u}} + \mathbf{K}\mathbf{u} = \mathbf{h}_a + \mathbf{h}_p + \mathbf{h}_b, \quad (3.9)$$

where $\ddot{\mathbf{u}}$, $\dot{\mathbf{u}}$ and \mathbf{u} denote accelerations, velocities and displacements of the mesh nodes, and \mathbf{M} , \mathbf{C} and \mathbf{K} are the mass, Rayleigh damping and internal elastic stiffness matrix, respectively. Mass lumping is used with a mass density of 1.07 g/ml. \mathbf{h}_p captures the pressures applied to the endocardia during the cardiac phases and \mathbf{h}_b accounts for the external boundary conditions. \mathbf{h}_a is the active force generated by the depolarized cells. At a given time t , for tetrahedron i with fiber direction \mathbf{o}_i , the active stress $\sigma_i(t)$ results in a 3D anisotropic stress tensor $\mathbf{A} = \sigma_i(t)\mathbf{o}_i\mathbf{o}_i^\top$, from which the contraction force vector is obtained as follows:

$$\mathbf{h}_a = \int \mathbf{A}\mathbf{n} \, dS, \quad (3.10)$$

¹ATP is a molecule involved in intra-cellular energy transfer.

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

First, cardiac electrophysiology is solved and the time-varying potentials per vertex are stored. Then, for every iteration of the cardiac biomechanics solver, the potential value corresponding to the current simulated time point is used to compute the active force generated by the cell. The process is iterated to compute cardiac electromechanics (EM) over time. In this work, the non-linear total Lagrangian explicit dynamics (TLED) framework [Mill 07] is used to solve cardiac biomechanics due to its computational efficiency and support for GPU implementation [Coma 08].

3.4 Boundary Conditions

Heart function depends on external conditions mostly determined by neighboring organs, blood flow and the circulatory system. Neighboring organs play a pivotal role in cardiac mechanics. Heart ventricles are connected to the atrium and arteries, which create some additional stiffness at the valves. A common way to model these effects is to add additional stiffness forces in the base (valve) region using stiff springs [Serm 06a, Nied 08]. Blood flow can be modeled using fluid-structure interaction (FSI) methods. Such approaches are highly detailed but also complex and costly to solve, because coupled systems are governed by large sets of parameters, see for instance [McQu 00, Miha 09] and references therein. Alternatively, only the ventricular pressure is considered and applied as constraints to the biomechanical model. The constraint can be computed from lumped models of the circulatory system [Arts 05, Kerc 07], or directly input by the user [Serm 06b]. These approaches rely on clinically related parameters and are fast to compute, but they ignore the flow patterns in the ventricles, which may have long-term impact on the cardiac function or on the effects of a therapy.

Assuming that spatial variations in flow patterns inside the ventricle have only minor effects on global cardiac function, which we are interested in, the hemodynamics parameters are modeled using lumped parameter models (homogeneous intraventricular pressure) within each ventricle. The arterial pressure is modeled using lumped Windkessel (WK) models, one for each artery. The atrial pressure is modeled using an elastance model of atrial contraction. The cardiac phases are modeled as sequential states.

3.4.1 Phase-Model of the Cardiac Cycle

The model simulates the four cardiac phases (cf. Fig. 2.4): filling, isovolumetric contraction, ejection and isovolumetric relaxation, by alternating the boundary conditions of the model according to the following rules.

1. When the blood flow, computed as the derivative of ventricular volume, is positive and the ventricular and atrial pressure are equal, the ventricle is being filled. A pressure equal to the atrial pressure is applied to the endocardium.
2. As soon as myocardial cells start to contract, the blood flow is reverted and the atrioventricular valve closes. Yet, the pressure is still lower than the arterial pressure. The arterial valves stay closed. The ventricular volume thus




Figure 3.7: Electrical circuit analogy of the three-element **WK** model. See text for details.

stays constant. This phase is called isovolumetric contraction phase. A projection/prediction constraint [Zett 13] is employed to compute the ventricular pressure during this phase while ensuring constant ventricular volume.

3. When the ventricular pressure becomes higher than the arterial pressure, the arterial valve opens and blood is ejected. The pressure of the artery is applied to the endocardium to mimic the resistance of the vessels.
4. As soon as the ventricle starts to relax, the blood flow is reverted. The arterial valve closes and the isovolumetric relaxation phase starts. A projection/prediction constraint is employed to compute the ventricular pressure during this phase while ensuring constant ventricular volume.

During each phase, the pressure applied to the endocardial surfaces is computed for each vertex and gathered into the global pressure vector \mathbf{h}_p of the dynamic system, cf. Eq. (3.9). The cardiac phases are handled independently for **LV** and **RV** to enable capturing of asynchronous cardiac motion.

3.4.2 Artery Model

Arterial pressure is modeled using a three-element **Windkessel (WK)** model [West 71], which takes as input the arterial flow and returns the pressure within the artery at every time step of the simulation. The model is derived from electrical circuit analogies (Fig. 3.7), where the blood flow is the current and the arterial pressure is the voltage. The first element of the model is a peripheral resistance R_p , which accounts for the distal resistance of the circulatory system which is mainly due to the small vessels. The compliance C accounts for the elasticity of the arterial walls. The characteristic resistance R_c accounts for the blood mass and for the compliance of the artery proximal to the valves. Let $Q_{ar}(t)$ be the arterial flow (change in volume) at time t , defined as the inverse of the ventricular flow, $P_{ar}(t)$ be the arterial pressure at time t and P_{ref} be a constant low pressure of reference (pressure of the remote venous system). When during ejection blood flows from the ventricles into the arteries, $Q_{ar}(t) > 0$, the three-element **WK** model writes:

$$\frac{dP_{ar}(t)}{dt} = R_c \frac{dQ_{ar}(t)}{dt} + \left(1 + \frac{R_c}{R_p}\right) \frac{Q_{ar}(t)}{C} - \frac{P_{ar}(t) - P_{ref}}{R_p \cdot C} \quad (3.11)$$

When the valves are closed, the blood flow is stopped, $Q_{\text{ar}}(t) = 0$, and the model writes:

$$\frac{dP_{\text{ar}}(t)}{dt} = -\frac{P_{\text{ar}}(t) - P_{\text{ref}}}{R_p \cdot C} \quad (3.12)$$

Two independent [WK](#) models are used for the aorta and the pulmonary artery.

3.4.3 Atrium Model

Atrium contraction, which happens just after diastasis and before systole, optimizes ventricular filling. Because simulating atrial contraction explicitly in 3D may be computationally prohibitive, a common approach is to rely on lumped models that mimic the raise of ventricular pressure due to atrial contraction. While some simplified models consider atrial pressure constant, the most common strategy consists in using phenomenological models of atrial pressure based on sigmoid functions, e. g., [[Serm 06b](#)]. More predictive elastance models have also been proposed to capture the interactions between atrial volume, pressure, tissue stiffness and the circulatory system [[Kerc 07](#)].

The approach used in this work relies on a simplification of the circulatory model proposed by [[Kerc 07](#)]. In particular, the arteries are decoupled from the atria: pulmonary and systemic circulations are neglected. Atrial contraction is modeled using a lumped time-varying elastance model. Minimum and maximum elastance parameters enable to set the peak systolic and diastolic stiffness, which then controls atrial pressure based on the current volume. A simple model of atrial activation synchronized with ventricular electrophysiology through a time-shift parameter corresponding to the duration of the PR interval Δ_{PR} enables controlling atrial volume. Vena cava pressure (P_{LA}) and pulmonary vein pressure (P_{RA}), are kept constant throughout the cardiac cycle. They can be manually adjusted to control the baseline pressure of the atrium.

Atrial pressure $P_{\text{atr}}(t)$ is computed according to the equation:

$$P_{\text{atr}}(t) = E(t) \cdot (V_{\text{atr}}(t) - V_{\text{atr}}^{\text{rest}}(t)) \quad , \quad (3.13)$$

where elastance E and the rest volume $V_{\text{atr}}^{\text{rest}}$ are defined as:

$$E(t) = \phi(t) \cdot (E^{\text{max}} - E^{\text{min}}) + E^{\text{min}} \quad , \quad (3.14)$$

and

$$V_{\text{atr}}^{\text{rest}} = (1 - \phi(t)) \cdot (V_{\text{atr}}^{\text{dia}} - V_{\text{atr}}^{\text{sys}}) + V_{\text{atr}}^{\text{sys}} \quad . \quad (3.15)$$

In these equations, four model parameters appear: E^{max} and E^{min} denote the maximum and minimum elastance, respectively, and $V_{\text{atr}}^{\text{sys}}$ and $V_{\text{atr}}^{\text{dia}}$ the diastolic and systolic atrial volumes at zero pressure, respectively. The activation function ϕ is defined as:

$$\phi(t) = \begin{cases} -12 \cos\left(2\pi \frac{t_{\text{atr}}(t)}{t_{\text{twitch}}}\right) + 0.5, & \text{if } t_{\text{atr}}(t) < t_{\text{twitch}} \\ 0, & \text{otherwise} \end{cases} \quad (3.16)$$

where t_{twitch} is the duration of ventricular contraction and t_{atr} :

$$t_{\text{atr}}(t) = \begin{cases} \text{mod}(t - t_{\text{active}} + \Delta_{\text{PR}}, t_{\text{HP}}), & \text{if } t \geq t_{\text{active}} - \Delta_{\text{PR}} \\ 0, & \text{otherwise} \end{cases} \quad (3.17)$$

where Δ_{PR} denotes the ECG PR interval and t_{HP} the heart period (duration of one cycle). Finally, the atrial volume is given by integrating:

$$\frac{dV_{\text{atr}}(t)}{dt} = Q_{\text{vein}}(t) - Q_{\text{mitral}}(t) , \quad (3.18)$$

where Q_{mitral} is the blood flow from the atrium through the atrioventricular valve into the ventricle, and Q_{vein} is the flow from the vein to atrium. The latter is given by:

$$Q_{\text{vein}}(t) = \frac{P_{\text{vein}} - P_{\text{atr}}(t)}{R_{\text{vein}}} , \quad (3.19)$$

where P_{vein} is the vein pressure and R_{vein} the resistance of the pulmonary vein.

Two independent models are used for **left atrium (LA)** and **right atrium (RA)**. In the **LA** model, Q_{vein} is the pulmonary vein pressure and in the **RA** model the vena cava pressure, denoted P_{LA} and P_{RA} , respectively.

3.4.4 Spatial Constraints

Base Stiffness

The effect of arteries and atria on the ventricular motion is simulated by connecting the vertices of the valve plane to springs whose stiffness is K_{base} . The fixed extremity of the springs corresponds to the rest position of the nodes, taken at mid diastasis, when the heart is at rest. The spring stiffness K_{base} is anisotropic. A strong stiffness, K_{base}^1 is applied along the long axis of the heart, while a lower stiffness, $K_{\text{base}}^{\text{s1}}$ and $K_{\text{base}}^{\text{s2}}$, is employed in the short-axis plane to allow free radial motion. Under these definitions, the base stiffness force writes:

$$\mathbf{h}_{\text{base}}(\mathbf{v}) = \mathbf{T} \begin{bmatrix} K_{\text{base}}^1 & 0 & 0 \\ 0 & K_{\text{base}}^{\text{s1}} & 0 \\ 0 & 0 & K_{\text{base}}^{\text{s2}} \end{bmatrix} \mathbf{T}^{-1}(\mathbf{v} - \mathbf{v}_0) , \quad (3.20)$$

where \mathbf{T} is the transformation matrix from world coordinates to the coordinate system defined by the **LV** long axis and the short-axis plane, and \mathbf{v} and \mathbf{v}_0 are the position of a mesh vertex at a given time step and its initial position, respectively. The computed forces \mathbf{h}_{base} of all myocardial nodes are added to the boundary condition force vector \mathbf{h}_b from Eq. (3.9).

Pericardium Constraint

Many computational heart models rely on strong base stiffness to fix the heart in 3D space. However, as shown in [Nied09, Mans10], computed cardiac motion highly depends on the strength of this boundary condition and also on the positions of the spring attachments. In reality, the effect of the great arteries and of the atrium on the bi-ventricular myocardium do not dominate overall cardiac contraction. Indeed, the heart contracts longitudinally, the base moving downwards towards the apex. Anatomically, the heart and the root of the great vessels are protected by a stiff sac, the pericardium (cf. Sec. 2.1.3). The pericardium fixes the heart to the diaphragm

and the sternum. Heart motion is also constrained by the neighboring lungs and liver. On the other hand, the myocardium is free to move inside the pericardium bag. A thin layer of pericardial fluid separates the epicardium to the pericardium, allowing friction-free sliding of the myocardium.

In [Mans10], a contact-based model of the pericardium has been proposed to mimic the effects of the neighboring organs and of the pericardium on the cardiac motion, while enabling fine-tuning the stiffness of the great vessels and atria attachments to get realistic basal motion. Promising results were obtained in patients with dilated right ventricles due to a tetralogy of Fallot repair. This constraint is also integrated into the model used in this thesis.

The idea consists in limiting the motion of the epicardial nodes inside an authorized area while preserving friction-free sliding. To this end, a signed distance map is created from the epicardial boundary at end-diastole to identify the interior and exterior of the pericardium. Let $D(\mathbf{v})$ be that distance map, where \mathbf{v} denotes the spatial coordinate. D is negative inside the pericardium sac, positive outside. Next, a gradient map is generated, which allows to compute for every point \mathbf{v} a normalized direction vector $\mathbf{g}(\mathbf{v})$ that points towards the closest point on the epicardial surface. Based on this spatial information, motion of the epicardial nodes is limited to a restricted zone by applying a force to every node that goes outside the pericardial sac (to bring it back into the authorized area). The pericardium force, $\mathbf{h}_{\text{peric}}$, is computed as follows:

$$\mathbf{h}_{\text{peric}}(\mathbf{v}) = \begin{cases} k \frac{D(\mathbf{v})^2}{D(\mathbf{v})^2 + m^2} \mathbf{g}(\mathbf{v}), & \text{if } D(\mathbf{v}) < d_{\text{in}} \\ -k \frac{D(\mathbf{v})^2}{D(\mathbf{v})^2 + m^2} \mathbf{g}(\mathbf{v}), & \text{if } D(\mathbf{v}) > d_{\text{out}} \\ \mathbf{0}, & \text{otherwise} \end{cases} \quad (3.21)$$

where k is the maximum strength of the force, m is the rate at which the force increases, and d_{in} and d_{out} are parameters that enable to control the extent of the authorized area around the pericardial region. With this definition, the epicardial nodes can slide along the pericardium cavity, but their radial motion is limited to emulate neighboring organs and the stiff pericardial sac. The computed forces $\mathbf{h}_{\text{peric}}$ of the epicardial nodes are added to \mathbf{h}_b , the full boundary condition force vector from Eq. (3.9).

Whole-Heart Personalization Through Inverse Optimization

4.1 State-of-the-Art Model Personalization.....	38
4.2 Problem Formulation	40
4.3 Integrated Personalization Pipeline.....	40
4.4 Experiments and Results.....	50
4.5 Predictive Power of Personalized Model.....	61
4.6 Summary.....	62

Calculating the output of a computational model given its input parameters is called *forward modeling*. Estimating the model parameters such that the output of the forward model matches a set of observations is the inverse problem of computational modeling, also known as *inverse modeling*, or *model personalization*. In this work, the term *inverse optimization* denotes a class of approaches that aim to solve inverse modeling by using optimization techniques. In this chapter, a comprehensive, modular pipeline for personalizing the coupled **electromechanics (EM)** whole-heart model (cf. Chap. 3) using advanced inverse optimization is presented.

First, the state-of-the-art in cardiac model personalization is described in Sec. 4.1. Then, in Sec. 4.2 the main nomenclature is introduced and the model personalization problem is formulated mathematically. An overview of the automated personalization pipeline developed during this thesis is presented in Sec. 4.3. Sec. 4.3.1 describes the estimation of patient-specific **hemodynamics (HD)** boundary conditions from time-varying ventricular volume and arterial pressure data. In Sec. 4.3.2, an automatic, cascaded algorithm for global **electrophysiology (EP)** personalization from 12-lead **electrocardiogram (ECG)** is presented. Sec. 4.3.3 then describes two algorithms for global and regional **electromechanics (EM)** model personalization from ventricular (regional) volume and pressure data. Next, quantitative and qualitative personalization results based on the full personalization pipeline on a large and heterogeneous patient population are presented in Sec. 4.4. Finally, Sec. 4.5 reports the results of a pilot study concerning the predictive capabilities of the model personalized using the proposed methodology. A short summary in Sec. 4.6 concludes the chapter.

4.1 State-of-the-Art Model Personalization

Over the last decades, computational cardiac modeling has attracted significant interest and great progress towards more realistic simulations and increased efficiency of the models has been achieved by the community (Chap. 3). As models mature, a critical step that is currently being undertaken is their translation into clinical applications and clinical practice in general [Nied 16], which requires the transition from synthetic simulations to precisely personalized patient-specific models. A wide variety of manual and (semi-)automatic image-based model personalization approaches using inverse problem techniques have been explored for the different components of cardiac models. For instance, an overview of common arterial Windkessel (WK) personalization approaches can be found in [Ster 95]. Personalization of complex 3D electrophysiology (EP) and biomechanics models is, however, more challenging due to the significantly higher computational cost of the forward simulations. Therefore, such algorithms are still an area of active research. A number of different classes of techniques can be identified in the literature [Talb 14, Mang 18], including gradient-based and gradient-free inverse optimization methods, data assimilation driven methods, methods based on machine learning (ML), and stochastic methods.

Inverse Optimization

In [Webe 10], an approach to personalize an atrial EP model using gradient-based inverse optimization was proposed. Most authors of 3D model personalization approaches, however, tend to choose gradient-free over gradient-based methods, as gradients of the involved, rather complex cost functions, are often difficult to obtain analytically, and finite difference approximations are computationally expensive and can introduce numerical difficulties. For instance, [Wong 15] proposed a gradient-free inverse optimization approach for personalization of biomechanical contractility parameters using a velocity-based objective function. In [Le F 13], mathematical currents represent information about heart shape and motion, which was exploited to evaluate the goodness of fit between the biomechanical simulation and the real heart. A multi-step optimization procedure for the estimation of left ventricular passive myocardial properties was proposed in [Gao 15]. Recently, a combination of gradient-free optimization and one-dimensional parameter sweeps for myocardial passive material parameter estimation was presented in [Hadj 17], but like the above approach, it was only applied to single-ventricle models. The authors of [Xi 13] suggested a similar multi-step idea, where parameter sweeps are performed iteratively on pre-defined two- and one-dimensional sample grids to obtain the most reasonable biomechanical parameter estimates among the tested ones. Generally, while methods based on parameter sweeps (exhaustive search) quickly become intractable, especially for large numbers of parameters, they enable the exploration of the full landscape of the objective function, but only at the pre-defined resolution. Optimization-based methods for cardiac EP further include the work by [Chin 08, Serm 09], a multi-level estimation approach for regional EP parameters, as well as the two-step approach proposed by [Rela 11], where, first, Purkinje activity is estimated, then global myocardial conductivity, in order to mitigate problems of parameter identifiability. [Seeg 15]

proposed a method for iterative regional personalization of cardiac EP tailored for patients with left-bundle-branch block.

Data Assimilation

In approaches based on data assimilation, the personalization problem is framed in a way that aims at identifying unknown variables using observations of a dynamical system. In each iteration, a forecast of the computational model and the assimilated observations are analyzed to estimate the current state regarding uncertainty in the state and in the observations. Two types of data assimilation are commonly used: variational [Deli 12] and filtering approaches. The latter includes methods for global [Chab 12] and regional [Marc 13] active contractility estimation, or passive material parameter estimation [Xi 11], all based on unscented Kalman filters.

Machine Learning

Recently, data-driven machine learning (ML) methods started to appear, such as the one-shot global EP personalization method proposed in [Zett 14], where a polynomial regression model is trained to directly map observed ECG features to conductivity parameters of the model. In another work [Prak 14], the non-linear relationship between heart motion descriptors derived from temporal images and parameters of electrical propagation was learned to regress electrical parameters from observed motion using a support vector machine. These examples demonstrated that by using ML, non-linear mappings from observations to cardiac model parameters can indeed be learned. Even when only limited data was available, by carefully generating synthetic training databases, the regression models could still produce valid predictions for real data in the reported experiments. Another interesting idea with its root in ML is to use genetic algorithms, which was successfully tested for estimation of regional Purkinje system parameters in [Cama 10]. More recently, genetic algorithms were combined with a multi-fidelity modeling approach, where the full 3D model is approximated by an efficient 0D surrogate model to speed up of the overall estimation process of active and passive tissue parameters [Moll 18].

Stochastic Methods

The last class of methods presented in this section are stochastic methods, which aim to not only estimate a single “optimal” parameter set, but also provide insights regarding confidence of the estimation and solution uniqueness taking into consideration the variety of uncertainties involved in the estimation process [Kou 11, Wall 14, Neum 14b]. More details are provided in Chap. 6 of this thesis.

Streamlined Whole-Heart Personalization

Most of the methods described above focus only on isolated sub-components of much more complex cardiac models. Only few researchers tackled the holistic problem of comprehensive and fast 3D whole-heart personalization by proposing automated and robust personalization pipelines. [Serm 12] presented one such approach, however it involves significant manual steps and its robustness to different patients, pathologies




Figure 4.1: A computational model f is a dynamic system that maps model input parameters \mathbf{x} to model state (output) variables \mathbf{y} . The goal of personalization is to tune \mathbf{x} such that the objectives \mathbf{c} , defined as the misfit between \mathbf{y} and the corresponding measured data \mathbf{z} of a given patient, are optimized (the misfit is minimized).

and data quality remains an open question, because it was tested only on a very limited number of patients (two). To our knowledge, no comprehensive, integrated pipeline has been presented to personalize full 3D whole-heart models, including anatomy, electrophysiology, biomechanics and hemodynamics boundary conditions, in a streamlined, consistent and automatic fashion on a large number of patients.

4.2 Problem Formulation

As illustrated in Fig. 4.1, any computational model f is governed by a set of parameters $\mathbf{x} = (x_1, \dots, x_{n_x})^\top$, bounded within a physiologically plausible domain Ω , and characterized by (observable) state variables $\mathbf{y} = (y_1, \dots, y_{n_y})^\top$. The state variables can be used to estimate \mathbf{x} . Note that some parameters may be pre-estimated or assigned fixed values. The goal of personalization is to optimize a set of objectives $\mathbf{c} = (c_1, \dots, c_{n_c})^\top$. The objectives are scalars defined as $c_i = \varepsilon(y_i, z_i)$, where ε is a measure of misfit, and z_i denotes the patient’s measured data (\mathbf{z}) corresponding to y_i . For instance, $\varepsilon(y_i, z_i) = y_i - z_i$ or $\varepsilon(y_i, z_i) = (y_i - z_i)^2$. Personalization is considered successful if all user-defined convergence criteria $\boldsymbol{\psi} = (\psi_1, \dots, \psi_{n_c})^\top$ are met. The criteria are usually defined in terms of maximum acceptable misfit per objective: $\forall i \in \{1, \dots, n_c\} : |c_i| < \psi_i$.

4.3 Integrated Personalization Pipeline

A robust, modular pipeline for comprehensive personalization of the coupled **electromechanics (EM)** whole-heart model has been developed. It enables generation of reproducible, patient-specific models at a large scale by estimating model parameters efficiently with minimal user interaction from routinely acquired clinical data.

An overview of the personalization modules, the input data, and their interactions is illustrated in Fig. 4.2. First, the “Anatomy” module creates the detailed bi-ventricular anatomical model as described in Sec. 3.1 from expert-guided segmentations of the patient’s **magnetic resonance imaging (MRI)** data. Next, the endocardia are tracked over time in the images to compute the time-varying ventricular volume curves. Together with the pressure curves from cardiac catheterization, these are input to the “Hemodynamics” module, which estimates the parameters of the **WK** models for both arteries, see Sec. 4.3.1. These lumped models serve as boundary constraints for the biomechanics model in the final step. Next, the “Electrophysiology”

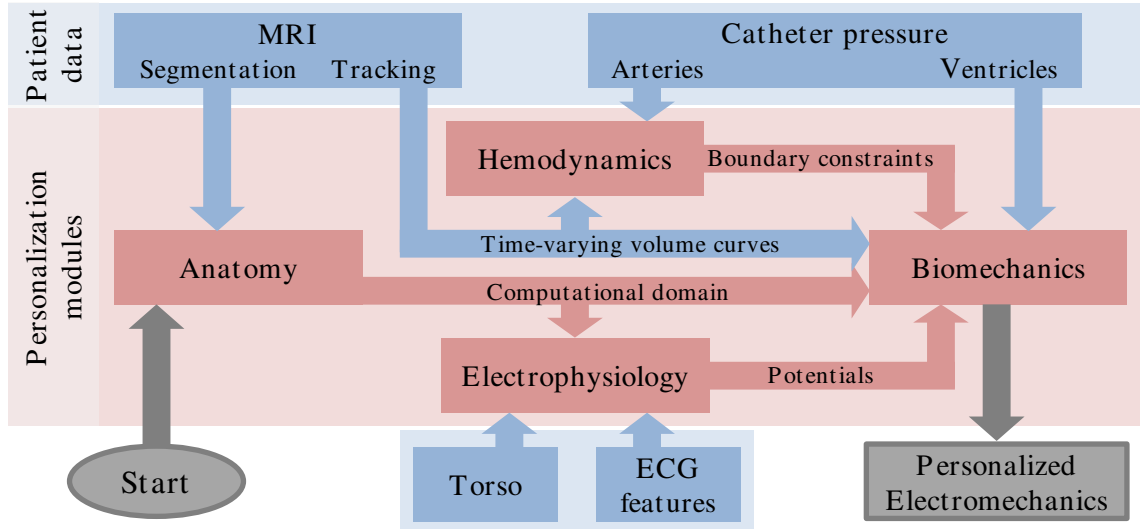


Figure 4.2: Flowchart of the integrated personalization pipeline. Blue represents data and flow of data, and red represents personalization modules.

module described in Sec. 4.3.2 uses 12-lead ECG and a torso atlas registered to the patient to compute the personalized electrical potentials on the anatomical model domain. The potentials trigger myocyte contraction in the biomechanics model. Finally, the active and passive mechanical tissue properties are estimated in the “Biomechanics” module such that the output of the model matches well the time-varying ventricular volume and ventricular pressure curves, as well as other important features derived from the data, see Sec. 4.3.3.

The output of the pipeline is the personalized whole-heart EM model defined by the patient’s anatomy and a set of personalized model parameters. In total, 23 model parameters are estimated:

- Twelve parameters related to hemodynamics (HD) boundary conditions:
 - Five WK parameters (C , R_c , R_p , P_{ref} , P_{init}) each for both arteries (aorta and pulmonary artery).
 - One parameter each for the pressure of the pulmonary vein (P_{LA}) and vena cava (P_{RA}).
- Four EP parameters:
 - Three diffusivity values (ν_{Myo} , ν_{LV} , ν_{RV}).
 - Time during which the ion channels are closed (τ_{close}).
- Seven active and passive parameters of the biomechanics model:
 - Maximum contractility ($\hat{\sigma}$) for both ventricles.
 - Contraction rates (k_{ATP}) for both ventricles.
 - Relaxation rates (k_{RS}) for both ventricles.
 - Global stiffness factor (β) of the Holzapfel-Ogden (HO) model.

Note that in the case of regional biomechanics personalization (Sec. 4.3.4), up to one contractility parameter per each of the 17 American Heart Association (AHA) segments of the left ventricle (LV) is estimated, and one for the right ventricle (RV), thus increasing the total count to up to 39 estimated parameters. In the following sections, the individual modules of the proposed personalization pipeline are presented.

4.3.1 Hemodynamics: Personalized Artery Boundary Conditions

The hemodynamics personalization consists in estimating the WK parameters of both arteries.

Estimated Parameters

As described in Sec. 3.4, the parameters of the three-element WK models, f_{WK} , are: artery compliance C , characteristic resistance R_c , peripheral resistance R_p and remote reference pressure P_{ref} . In addition, the initial pressure P_{init} is estimated to ensure that the first computed cycle during a 3D whole-heart EM simulation is already close to the steady state.

Personalization Approach

Some manual pre-processing steps are required to extract and synchronize the time-varying pressure curves from the catheter data and the volume curve derived from MRI tracking of the endocardium. First, the user interactively selects a cardiac cycle in the pressure curve and low-pass filters the artery and ventricular pressures, $\overline{P_{\text{ar}}}$ and $\overline{P_{\text{ven}}}$. Next, because the heart rate of a patient at the time of catheterization and at the time of MRI acquisition can be different, the pressure curve is automatically adjusted to match the heart rate at the MRI acquisition. As a simple temporal scaling would not be physiologically coherent, the following algorithm is applied. First, the systolic portion of the pressure curve is stretched such that the ejection time (ET) observed in the pressure measurement (time during which ventricular pressure is higher or equal than arterial pressure) matches the ET measured in the volume curve (time during which the ventricular flow is negative). Then, the pressure curves are shifted interactively such that they are synchronized with the volume curve.

Although pre-processing requires user interaction, the parameter estimation procedure itself is fully automatic. Let \mathbf{x}_{WK} be a vector of all parameters of the WK model except the initial pressure P_{init} . \mathbf{x}_{WK} is estimated automatically using inverse optimization:

$$\begin{aligned} & \underset{\mathbf{x}_{\text{WK}}}{\text{minimize}} && \zeta(\mathbf{x}_{\text{WK}}) \\ & \text{subject to} && \mathbf{x}_{\text{WK}} \in \Omega_{\text{WK}} , \end{aligned} \tag{4.1}$$




Figure 4.3: The different steps involved in the estimation of the **WK** parameters of the arteries for an example case.

where Ω_{WK} is the domain of physiologically plausible **WK** parameters (upper and lower bounds for each parameter). The cost function ζ consists of three terms:

$$\begin{aligned} \zeta(\mathbf{x}_{\text{WK}}) = & \frac{1}{N} \sum_{t=1}^N \varepsilon(\overline{P_{\text{ar}}}[t], P_{\text{ar}}(f_{\text{WK}}(\mathbf{x}_{\text{WK}}))[t]) + \\ & \varepsilon(\min \overline{P_{\text{ar}}}, \min P_{\text{ar}}(f_{\text{WK}}(\mathbf{x}_{\text{WK}}))) + \\ & \varepsilon(\max \overline{P_{\text{ar}}}, \max P_{\text{ar}}(f_{\text{WK}}(\mathbf{x}_{\text{WK}}))) \end{aligned} \quad (4.2)$$

$P_{\text{ar}}(f_{\text{WK}}(\mathbf{x}_{\text{WK}}))$ is the time-series of artery pressure computed by the model f_{WK} using the current parameters \mathbf{x}_{WK} and $\overline{P_{\text{ar}}}$ is the time-series of measured artery pressure. N is the number of measured samples and $\varepsilon(\cdot, \cdot)$ is the misfit defined as the squared difference between the two inputs. The first term in Eq. (4.2) ensures good overall fit between both curves. It computes the normalized sum of squared differences between the measured and computed curves. The second and third terms are penalty terms that ensure accurate fit of the arterial pressure at **end-diastole (ED)** and **end-systole (ES)**, where the pressures reach their minimum and maximum, respectively. Solving Eq. (4.1) amounts to estimating all the **WK** parameters but the initial pressure P_{init} . The latter is obtained automatically from the computed pressure curve over several cycles such that the first computed cycle during a 3D whole-heart simulation is close to the steady state.

Fig. 4.3 illustrates the full **WK** personalization procedure described above. For each patient with complete catheter pressure data, the procedure is performed twice: once for the aorta and once for the pulmonary artery.

4.3.2 Electrophysiology: Cascaded Inverse Optimization

In Sec. 3.2, two cardiac **EP** models were presented. The first option is LBM-EP, a **Lattice-Boltzmann method (LBM)** approach that solves the mono-domain Mitchell-Schaeffer cell model on a grid. The model parameters are closely related to the shape of the **action potential (AP)**. The second option is Graph-EP, a very fast graph-based solver that uses a simplified **AP** formulation to compute repolarization. In this section, the personalization strategy is presented in detail for LBM-EP, but the approach works for both models.

Estimated Parameters

Four parameters need to be personalized for the EP model f_{EP} . The first parameter is related to the timing of the AP: τ_{close} , the time during which the ion channels are closed. The other three parameters control the speed of electrical wave propagation for three different regions of the computational domain: the left and right endocardium with fast electrical diffusivity to mimic the Purkinje network, ν_{LV} and ν_{RV} , and the remaining myocardium, ν_{Myo} . A major goal of this thesis was to develop robust personalization strategies that do not require specialized electrical data such as contact mapping catheters. The proposed approach is therefore solely based on routinely acquired 12-lead ECG data.

Personalization Approach

The personalization algorithm, called CASCADEEP, is outlined in Alg. 1.

Algorithm 1 Iterative, cascaded EP personalization

Require: Anatomical model, ECG features $\overline{\Delta_{QRS}}$, $\overline{\Delta_{QT}}$, $\overline{\angle_{EA}}$

- 1: **procedure** CASCADEEP
 - 2: Initialize $\tau_{close}^0, \nu_{Myo}^0, \nu_{LV}^0, \nu_{RV}^0$ ▷ Default values
 - 3: **for** $i \in \{1, \dots, n_{iter}\}$ **do**
 - 4: $\nu_{Myo}^i, \nu_{LV}^i, \nu_{RV}^i \leftarrow \text{FITQRS}(\tau_{close}^{i-1}, \nu_{Myo}^{i-1}, \nu_{LV}^{i-1}, \nu_{RV}^{i-1})$ ▷ Stage 1
 - 5: $\nu_{LV}^i, \nu_{RV}^i \leftarrow \text{FITEA}(\tau_{close}^{i-1}, \nu_{Myo}^i, \nu_{LV}^i, \nu_{RV}^i)$ ▷ Stage 2
 - 6: $\tau_{close}^i \leftarrow \text{FITAPD}(\tau_{close}^{i-1}, \nu_{Myo}^i, \nu_{LV}^i, \nu_{RV}^i)$ ▷ Stage 3
 - 7: **return** $\tau_{close}^{n_{iter}}, \nu_{Myo}^{n_{iter}}, \nu_{LV}^{n_{iter}}, \nu_{RV}^{n_{iter}}$ ▷ Final personalized parameters
-

Subroutines

- 8: **function** FITQRS($\tau_{close}, \nu_{Myo}, \nu_{LV}, \nu_{RV}$) ▷ Jointly scale diffusivities
 - 9: $\hat{\kappa} \leftarrow \text{argmin}_{\kappa} \varepsilon(\overline{\Delta_{QRS}}, \Delta_{QRS}(f_{EP}(\tau_{close}, \kappa\nu_{Myo}, \kappa\nu_{LV}, \kappa\nu_{RV})))$
 - 10: $\nu_{Myo}^*, \nu_{LV}^*, \nu_{RV}^* \leftarrow \hat{\kappa}\nu_{Myo}, \hat{\kappa}\nu_{LV}, \hat{\kappa}\nu_{RV}$
 - 11: **return** $\nu_{Myo}^*, \nu_{LV}^*, \nu_{RV}^*$
 - 12: **function** FITEA($\tau_{close}, \nu_{Myo}, \nu_{LV}, \nu_{RV}$) ▷ Tune Purkinje diffusivities
 - 13: $\nu_{LV}^*, \nu_{RV}^* \leftarrow \text{argmin}_{\nu_{LV}, \nu_{RV}} \varepsilon(\overline{\angle_{EA}}, \angle_{EA}(f_{EP}(\tau_{close}, \nu_{Myo}, \nu_{LV}, \nu_{RV})))$
 - 14: **return** ν_{LV}^*, ν_{RV}^*
 - 15: **function** FITAPD($\tau_{close}, \nu_{Myo}, \nu_{LV}, \nu_{RV}$) ▷ Rectify time constant
 - 16: $\tau_{close}^* \leftarrow \tau_{close} + \overline{\Delta_{QT}} - \Delta_{QT}(f_{EP}(\tau_{close}, \nu_{Myo}, \nu_{LV}, \nu_{RV}))$
 - 17: **return** τ_{close}^*
-

First, the model parameters are initialized to default values, which were empirically set to $\tau_{close}^0 = 150$ ms, $\nu_{Myo}^0 = 500$ mm²/s, $\nu_{LV}^0 = 2000$ mm²/s, and $\nu_{RV}^0 =$

2000 mm²/s in all experiments. Next, the cascade strategy splits the fitting of different personalization objectives into three different personalization stages (lines 4–6). The first and second stage are geared towards ventricular depolarization, i. e., the QRS complex, whereas the third stage aims at [action potential duration \(APD\)](#) and thus covers also ventricular repolarization.

The purpose of the first stage, FITQRS, is to match measured $\overline{\Delta_{\text{QRS}}}$ and computed Δ_{QRS} , the QRS duration derived from the model output f_{EP} . This is achieved by jointly scaling all three diffusivity parameters using a common scaling factor $\hat{\kappa}$, which is determined by optimization. The second stage, FITEA, reduces the [electrical axis \(EA\)](#) misfit between measured $\overline{\angle_{\text{EA}}}$ and computed \angle_{EA} by tuning the fast endocardial diffusivities ν_{LV} and ν_{RV} . In this stage ν_{Myo} remains fixed in order to minimize the effect on Δ_{QRS} . The third personalization stage, FITAPD, rectifies the time constant τ_{close} by adding the difference between $\overline{\Delta_{\text{QT}}}$ and Δ_{QT} to its previous value. The optimization calls (lines 9 and 13) are performed using [bound optimization by quadratic approximation \(BOBYQA\)](#) [Powe09], a robust and fast, bound-constrained, gradient-free optimizer. In each stage, BOBYQA is initialized with the most recent parameter values. The objective function $\zeta_{\text{EP}} = \varepsilon(\cdot, \cdot)$ in lines 9 and 13 returns the squared distance between the two scalar inputs. The bound-constraints for the optimization are determined on-the-fly based on two requirements. First, myocardial diffusivity has to be slower than endocardial (Purkinje) diffusivity, $\nu_{\text{Myo}} \leq \nu_{\text{LV}}$, $\nu_{\text{Myo}} \leq \nu_{\text{RV}}$. Second, at any time, all parameters have to remain within the physiologically plausible parameter domain Ω_{EP} of the EP model.

The multi-stage cascade procedure described above is repeated $n_{\text{iter}} = 3$ times, which was set empirically to facilitate escaping from local optima and to yield more accurate and robust results, even for highly pathological cases.

4.3.3 Biomechanics: Global Volume and Pressure Fitting

Global biomechanics personalization of the cardiac EM model described in Sec. 3.3 consists in estimating the governing parameters of the active cell contraction model and of the passive HO tissue model. In addition, HD parameters of the atrial models are personalized during this process.

Estimated Parameters

Seven tissue parameters of the EM model, f_{EM} , gathered in a parameter vector $\mathbf{x}_{\text{tissue}}$, are personalized, one for the passive component of the model, the global stiffness factor β , and three parameters each for LV tissue and RV tissue for the active force model. The active parameters per ventricle comprise maximum asymptotic contraction force $\hat{\sigma}$, contraction rate k_{ATP} , and relaxation rate k_{RS} . The choice of granularity of these parameters (per ventricle) is based on considerations of computational tractability and parameter identifiability, given the amount and quality of available clinical measurements. In particular, the EM personalization approach has to rely on ventricular volume measurements derived from cardiac cine MRI; and ventricular pressures obtained from cardiac catheterization (cf. Sec. 4.3.1). The pressures are further used to personalize two additional parameters of the left atrium (LA) and

right atrium (RA) models: pulmonary vein pressure P_{LA} , and vena cava pressure P_{RA} , combined in \mathbf{x}_{veins} .

Personalization Approach

The proposed EM personalization algorithm, GLOBALEM, is outlined in Alg. 2.

Algorithm 2 Global EM personalization

Require: Anatomical model, ventricular volume and pressure measurements \mathbf{z}_{EM}

- 1: **procedure** GLOBALEM
 - 2: Initialize $\mathbf{x}_{tissue}^0, \mathbf{x}_{veins}^0$ ▷ Default values
 - 3: $\mathbf{x}_{veins}^* \leftarrow \text{FITDP}(\mathbf{x}_{tissue}^0, \mathbf{x}_{veins}^0)$ ▷ Step 1
 - 4: $\mathbf{x}_{tissue}^* \leftarrow \text{FITPV}(\mathbf{x}_{tissue}^0, \mathbf{x}_{veins}^*)$ ▷ Step 2
 - 5: **return** $\mathbf{x}_{tissue}^*, \mathbf{x}_{veins}^*$ ▷ Final personalized parameters
-

Subroutines

- 6: **function** FITDP($\mathbf{x}_{tissue}, \mathbf{x}_{veins} = (P_{LA}, P_{RA})$) ▷ Adjust vein pressures
 - 7: $\mathbf{y}_{EM} = f_{EM}(\mathbf{x}_{tissue}, \mathbf{x}_{veins})$
 - 8: $P_{LA}^* \leftarrow P_{LA} + \min P_{ven}(\mathbf{z}_{EM}(LV)) - \min P_{ven}(\mathbf{y}_{EM}(LV))$
 - 9: $P_{RA}^* \leftarrow P_{RA} + \min P_{ven}(\mathbf{z}_{EM}(RV)) - \min P_{ven}(\mathbf{y}_{EM}(RV))$
 - 10: **return** (P_{LA}^*, P_{RA}^*)
 - 11: **function** FITPV($\mathbf{x}_{tissue}, \mathbf{x}_{veins}$) ▷ Optimize tissue parameters
 - 12: $\mathbf{x}_{tissue}^* \leftarrow \text{argmin}_{\mathbf{x}} \zeta_{EM}(\mathbf{z}_{EM}, f_{EM}(\mathbf{x}, \mathbf{x}_{veins}))$
 - 13: **return** \mathbf{x}_{tissue}^*
-

After initializing all parameters (\mathbf{x}_{tissue} and \mathbf{x}_{veins}) with their default values, the vein pressure parameters \mathbf{x}_{veins} of both atrial models are calibrated in FITDP. The goal of this first step is to fit the the diastolic pressure of the ventricles to the measurements. To this end, a forward model run of the EM model f_{EM} is performed, yielding the model output \mathbf{y}_{EM} . In lines 8–9, the parameters are adjusted by adding the difference between measured and computed diastolic pressures (minimum ventricular pressure approximately corresponds to the diastolic pressure) to the initial parameter values. This fitting strategy is motivated by the fact that during diastole, when the atrioventricular valves are open, pressure in the ventricles is dominated by pressure in the atria, which in turn is dominated by pressure originating from the veins.

The purpose of FITPV (step 2) is then to jointly optimize the biomechanical tissue parameters \mathbf{x}_{tissue} by fitting of pressure and volume curves using BOBYQA optimization: a carefully designed objective function ζ_{EM} determines the similarity between

measured (\mathbf{z}_{EM}) and computed (\mathbf{y}_{EM}) pressure and volume curves by computing a weighted sum of objectives:

$$\zeta_{\text{EM}}(\mathbf{z}_{\text{EM}}, \mathbf{y}_{\text{EM}}) = \sum_{i=1}^{12} w_i \varepsilon_i \left(\frac{z_i - m_{z_i}}{\text{SD}_{z_i}}, \frac{y_i - m_{z_i}}{\text{SD}_{z_i}} \right), \quad (4.3)$$

where w_i is the weight of the i th objective, $z_i \in \mathbf{z}_{\text{EM}}$ and $y_i \in \mathbf{y}_{\text{EM}}$ are individual measurements and model outcomes, respectively, and $\varepsilon_i(\cdot, \cdot)$ is the misfit function (cf. Sec. 4.2). Incorporating the means and standard deviations (computed off-line from all available patients) per measurement, m_{z_i} and SD_{z_i} , balances contributions of the individual objectives to ζ_{EM} (prior to weighting), as they cancel out discrepancies caused by different units and value ranges of the individual measurements.

The objectives and their manually tuned weights are listed in Tab. 4.1. All ob-

Table 4.1: Objectives of FITPV.

Measurement of objective c_i	Misfit function ε_i	Weight w_i
Time-varying LV volume	SSD	22.2%
Time-varying LV pressure	SSD	13.3%
LV ejection fraction (EF)	L_2 -norm	8.9%
LV stroke volume (SV)	L_2 -norm	8.9%
LV end-systolic volume (ESV)	L_2 -norm	8.9%
LV end-diastolic volume (EDV)	L_2 -norm	8.9%
LV end-systolic pressure (ESP)	L_2 -norm	8.9%
LV end-diastolic pressure (EDP)	L_2 -norm	4.4%
RV EDV	L_2 -norm	4.4%
RV ESP	L_2 -norm	4.4%
RV SV (computed) vs. LV SV (measured)	L_2 -norm	4.4%
RV EDP	L_2 -norm	2.4%

jectives can be computed from the measured and computed time-varying volume and pressure curves. The misfit function **SSD** (first two objectives in Tab. 4.1) denotes the sum of squared differences between two time-series (e. g., measured versus computed ventricular volume V_{ven}).

Remarks

As one may notice from Tab. 4.1, the combined weights of all **RV**-related objectives are lower than those of the **LV**. This is an attempt to minimize the impact of the often lower quality **RV** segmentation and tracking performance owed to the thin **RV** myocardium, which is difficult to outline and track accurately in routine **MRI**. Furthermore, the second-last objective in Tab. 4.1, the misfit between **SV** of **LV** and **RV**, drives the model personalization towards equal **SV** for both ventricles. This is

needed because measured ejected blood volume (SV) of RV might differ from LV due to measurement errors. In a real heart, the amounts are precisely equal for both ventricles [Fran 62]. Finally, in order to minimize transient effects, whenever the forward model f_{EM} is computed, two heart cycles are simulated and the computed pressures and volumes in \mathbf{y}_{EM} are based on the second cycle only.

Parameter Default Values

Default values of the personalized atrium model parameters are based on normal values for healthy hearts, $P_{LA} = 1.6$ kPa and $P_{RA} = 0.7$ kPa, respectively. For LV and RV maximum contraction force, the default value is 100 kPa and must remain within $[0; 500]$ kPa during personalization. Contraction and relaxation rates for both ventricles are initialized with $k_{ATP} = 20$ s⁻¹, $k_{RS} = -20$ s⁻¹ and bounded by $[5; 100]$ s⁻¹ and $[-100; -5]$ s⁻¹, respectively, and the stiffness parameter is initialized with $\beta = 1$ and bounded by $[1/5, 5]$.

4.3.4 Towards Regional Biomechanics Personalization

While the global biomechanics personalization approach GLOBALEM described above can accurately adapt the model to fit global indicators of heart physiology such as EF , SV , and pressure, it does not take into account potential localized functional variations that are occurring only in specific regions of the myocardium, which could be induced by scars or other (non-)pathological alterations. This could potentially lead to inconsistencies between cardiac dynamics observed in the images and the computed motion.

Estimated Parameters

An enhancement of the global pressure and volume fitting procedure FITPV from the previous section is proposed, which introduces a set of regional active force parameters based on the AHA segmentation (cf. Sec. 3.1) of the LV part of the anatomical model. More precisely, instead of having only one global contractility parameter $\hat{\sigma}_{LV}$ for the entire LV , one value for each of the 17 AHA segments is defined: $\hat{\sigma}_{AHA1} \dots \hat{\sigma}_{AHA17}$ (the one for RV , $\hat{\sigma}_{RV}$, remains).

Regional Volume Measurements

In order to identify regional alterations in cardiac function, localized observations are extracted from the clinical data and compared against corresponding model output. The proposed method relies on “regional volumes” [Marc 13] computed from the LV endocardium meshes tracked from cine MRI , and from the simulated dynamic meshes of the EM model f_{EM} . Regional volume curves represent portions of the time-varying LV cavity volume. These portions are computed from regional volume meshes, which are formed by the endocardial surface patch of a given AHA segment and the LV barycenter, cf. Fig. 4.4. Evaluation and comparison of regional volumes could help identifying local alterations of heart function, as supported by clinical literature [Baxl 71].




Figure 4.4: Visualization of regional volumes meshes of three example AHA segments derived from a real patient’s anatomical model. From left to right: basal anterolateral, mid anterolateral, and apical anterior (cf. Fig. 3.3). White dots represent the LV barycenter.

Joint Estimation of Regional Parameters

The global objective function from Eq. (4.3) is adapted (denoted ζ_{EM}^{reg}) to incorporate regional volume measurements. In particular, the time-varying LV volume objective from the global objectives listed in Tab. 4.1, is replaced by 17 equally-weighted regional volume objectives, where the misfits are SSDs per time-varying regional volume curve (one per AHA segment). The same is applied to all other objectives computed from the LV volume curve. For example, 17 “regional EFs” replace the global LV EF objective. All EM parameters including regional contractility values can be estimated jointly by replacing the global cost function in the pressure and volume fitting procedure FITPV by ζ_{EM}^{reg} in Alg. 2. It should be noted, however, that this approach may not yield good results due to parameter identifiability issues, and it is computationally demanding and may quickly fall victim to the curse of dimensionality [Bell57].

Hierarchical Coarse-to-Fine Estimation Approach

To overcome these issues, a faster and more robust approach is proposed, which follows a multi-step coarse-to-fine strategy inspired by recent work in lung model estimation [Fuer15]. The idea is to address the multi-parameter optimization problem in a hierarchical way. First, the global LV myocardium active force value $\hat{\sigma}_{LV}$ is estimated such that global EF, SV, etc., are fitted as described in Sec. 4.3.3. Then, the regional volume for each AHA segment is analyzed to identify potential deviations between model and measurements. The LV is finally divided into two regions, one with the largest deviation, and the rest, and the active force is estimated separately for both regions. The process could be iterated if needed. Such an approach is motivated by the proven good performance record of hierarchical, coarse-to-fine optimization strategies in terms of optimization results (less prone to local minima) and speed (require less iterations overall), as demonstrated in machine learning (ML) [Zhen14, Ghes16b], model personalization [Chin08, Fuer15, Dham16, Dham16, Phei17, Bala17], and image registration [Verc09].

4.3.5 Implementation Details

The entire personalization pipeline is implemented in an integrated environment based on C++ with an interactive [user interface \(UI\)](#) for live-testing (see appendix A.1 for more details), but also with bindings to the Python scripting language [Van 11] to enable rapid experimentation and batch processing, which was used for all experiments in this thesis. The framework is controlled through a per-patient [XML](#) study file, which gathers the settings of the optimizations and other personalization steps to ensure reproducibility, as well as the resulting estimated parameters. Python scripts control the estimation sequence: first anatomy (cf. Sec. 3.1), then [HD](#) boundary conditions (cf. Sec. 4.3.1), then [EP](#) (cf. Sec. 4.3.2), and finally biomechanics (cf. Sec. 4.3.3 or Sec. 4.3.4). At each stage, the estimated parameters can be saved to the [XML](#) study file such that the process can be restarted or refined from any accomplished stage. A generic [BOBYQA](#) optimizer [Powe 09] has also been integrated to handle the [EP](#) and biomechanics personalization seamlessly, where the model is seen as a black box that is called by the optimizer during each iteration.

4.4 Experiments and Results

4.4.1 Description of Study Cohort

A total of $n_{\text{datasets}} = 113$ patients, many suffering from [dilated cardiomyopathy \(DCM\)](#), see Sec. 1.1.1, were included in this study. Patient data was acquired at four clinical sites:

- Department of Medicine III, University of Heidelberg
- Ospedale Pediatrico Bambino Gesù, Rome
- Great Ormond Street Hospital, University College London
- Department of Congenital Heart Diseases, German Heart Center Berlin

The latter three sites and the collaboration partner of this thesis, Siemens Healthineers, were involved in the [MD-Paedigree](#) project funded by the European Commission under grant agreement 600932 (FP7-ICT-2011-9), which focused on pediatric cardiology. Therefore, a significant portion of the patients, about one third (33), are children. The youngest patient was one year at the time of data acquisition, while the oldest was 81 years (mean: 44.5 years). In some of the following analyses, the cohort is divided into pediatric (< 18 years) and adult (≥ 18 years) population. The two groups can be distinguished by their respective coloring in figures: yellow for pediatric and purple for adult.

Required Data

A complete patient dataset (for the purpose of this work) consists in a [steady-state free precession \(SSFP\)](#) cardiac cine [MRI](#) acquisition, a 12-lead [ECG](#) reading, and time-varying pressure curves from catheterization of both ventricles and both arteries.

The different personalization modules (Sections 4.3.1 to 4.3.3) require different inputs. If parts of the data were missing for a patient, some personalization steps could either not be performed at all or performed only in a limited fashion:

- If **MRI** was incomplete or missing, no segmentation, no time-varying ventricular volume curves and no **anatomical model** could be generated. Those constitute essential inputs for the different models and personalization modules. Therefore, such datasets were never included. The total number of successfully generated patient-specific anatomical models is 113.
- If artery pressures were missing, the respective **WK models** could not be personalized and their parameters were set to default values. The number of available aorta and pulmonary artery pressure data were 67 and 57, respectively: a total of 124 **WK models** were personalized.
- If 12-lead **ECG** was missing, the **EP model** could not be personalized and its parameters were set to default values. Only five **ECG** were missing, thus the total number of personalized **EP models** was 108.
- If ventricle pressures were missing, the affected terms in the cost function for personalizing the **biomechanics model** were disabled and thus the fitting was based on volume curves only. For 55 patients, all pressures were available.

Although the number of *complete* patient datasets was only 53, the final output, i. e., the whole-heart **EM models**, could be generated for all 113 patients that were included in this study.

4.4.2 Whole-Heart Personalization: Goodness of Fit

Earlier versions of the proposed personalization pipeline and personalization outcomes on a subset (only one clinical center and reduced number of patients) of the cohort used in this thesis were reported at a biomedical imaging conference [Neum 14c] and in two clinical journal publications [Kayv 15, Amr 16]. The results presented below are—for the first time—based on the latest model and personalization pipeline advancements, also featuring the largest multi-clinic patient cohort yet.

Anatomical models

First, the 3D anatomical models for all cases were generated from cine **MRI** under expert guidance and tracked over time to compute the time-varying volume curves for both ventricles, see Sec. 3.1 for more details. Some results are visualized in appendix A.2. The myocardium mass according to the myocardium volume at **ED** and assuming muscle tissue density of 1.06 g/ml [Urba 01] ranged from 78.6 g to 537.4 g with mean and SD of $274.2 \text{ g} \pm 82.4 \text{ g}$. The **LV EF** and **SV** means were measured as $36.1 \pm 12.4\%$ and $79.5 \pm 28.2 \text{ ml}$, and ranged from 10.2 to 63.2% and 18.0 to 146.7 ml, respectively. The differences in terms of **EF** were only minor between children and adults, but for myocardium mass and **SV** a shift towards lower values was observed for the pediatric group. For instance, the mean \overline{SV}_{LV} was about 30 ml less. This was expected since young hearts are smaller and thus can have much lower **EDV**.




Figure 4.5: Three example results of measured and fitted arterial pressure curves computed using the personalized **WK** parameters and the flow derived from the ventricular volume curves.

The severity of cardiac disease is often characterized in terms of left ventricular ejection fraction \overline{EF}_{LV} . In this cohort, eight patients were identified as normal with \overline{EF}_{LV} between 55% and 70%, 36 patients with moderate \overline{EF}_{LV} reduction between 40% and 55%, and the majority, 69, with severe **EF** reduction ($\overline{EF}_{LV} < 40\%$).

Estimation of Artery Parameters

Next, the **Windkessel (WK)** model parameters for both the aorta and pulmonary artery were personalized for all cases with available arterial pressure measurements. The semi-automatic pre-processing steps described in Sec. 4.3.1 could be performed within minutes, after which the automatic personalization computed the five **WK** parameters per artery within seconds. Personalized artery pressures matched qualitatively well for almost all patients. Example results are shown in Fig. 4.5. In some cases, the trends were more difficult to capture. This could be explained by inaccurate flow data, especially for **RV** flow, where segmentation is difficult. However, in all cases, diastolic and peak pressures were still captured, which are the most important features for the subsequent biomechanical computations, as the flow will then be directly computed by the model (cf. Fig. 4.10). Cohort-specific analysis was not performed as artery pressures were not acquired for pediatric patients.

Estimation of Electrophysiology Parameters

To evaluate the performance of the **EP** personalization module, first, the fit between measured and computed **ECG** features after **EP** model personalization was analyzed for all 108 patients with proper **ECG** data. Measured **ECG** statistics based on all patients are shown in Tab. 4.2.

Fig. 4.6 shows the agreement between measured and computed electrical features after **EP** personalization. Absolute errors are listed in Tab. 4.3. As one can see, the QRS and QT fit were excellent. The mean errors were well below 1% of the respective cohort-average values (cf. Tab. 4.2), and even the maximum observed errors could be considered minor, which leads to the conclusion that the personalization pipeline could match successfully both objectives for all cases. The **EA** is a more complex measurement that was expected to be much harder to fit. Still, the mean absolute error over the entire cohort was below 5% of the maximum possible error (180°), and

Table 4.2: Functional characteristics of the cohort derived from ECG.

Feature	Mean	SD	Min	Max	Unit
$\overline{\Delta_{QRS}}$	110.0	27.6	68.0	185.0	ms
$\overline{\Delta_{QT}}$	398.6	42.6	302.0	520.0	ms
$\overline{\angle_{EA}}$	12.0	54.9	-174.0	120.0	deg




Figure 4.6: Bland-Altman plots showing the good agreement between measured and computed EP features: QRS and QT durations are shown in the left and mid panel, respectively. The blue line is the mean of the differences and the red lines delineate the 95% confidence interval (CI). The purple and yellow dots represent adults and children, respectively. As one can see in the right panel, the goodness of fit in terms of EA was good in the majority of patients, however for some patients the fit was poor and even some larger outliers were present. The visualization uses polar coordinates, where the center represents a difference of $\angle_{EA} - \overline{\angle_{EA}} = -90^\circ$ and the thick, black outer circle $+90^\circ$. The dashed inner circles are located at 45° steps: the circle closest to the mean (blue) represents zero difference (desired, perfect fit), and the intermediate circles $\pm 45^\circ$.


almost 80% of the cases yielded acceptable errors of less than 15° . One potential reason for larger errors is solution non-existence due to possibly invalid assumptions of the employed EP model for patients with complex pathologies.

One might expect that discrepancies are larger for patients with pathological heart axis. However, as one can see from Fig. 4.6, right panel, even for the few patients with extreme axis deviation (between -180° and -90° , cf. Fig. 2.3), some models could be fitted almost perfectly. Average errors in the “right-intermediary” and right-deviated segments (between 30° and 120°) were lower compared to the “left-intermediary” and left-deviated segments (between -60° and 30°). However, no general rules could be derived since outliers (positive or negative) existed throughout the full EA spectrum. These observations highlight the complexity and high uncertainties involved in the inverse EP problem.

Next, the distributions of estimated EP parameters were analyzed, see Fig. 4.7. The pediatric and the adult patient group were considered separately. Both $\tau_{\text{close-}}$

Table 4.3: Goodness of fit in terms of absolute differences between ECG features.

Absolute error	Mean	SD	Min	50%	80%	95%	Max	Unit
$ \Delta_{\text{QRS}} - \overline{\Delta_{\text{QRS}}} $	0.3	0.5	0.0	0.1	0.7	1.6	2.2	ms
$ \Delta_{\text{QT}} - \overline{\Delta_{\text{QT}}} $	3.1	2.3	0.0	2.7	4.7	6.7	11.3	ms
$ \angle_{\text{EA}} - \overline{\angle_{\text{EA}}} $	8.9	17.1	0.0	0.5	16.2	43.9	86.3	deg

**Figure 4.7:** Histograms showing the distributions of estimated EP parameters. Purple and yellow colors represent adult and children group, respectively.

distributions were rather mono-modal, except for few outliers. However, the peak was shifted slightly towards lower values for the pediatric group. This can be attributed to the fact that especially in young children, heart rates are typically much higher compared to adults. Because an increased heart rate can reduce [action potential duration \(APD\)](#) and τ_{close} is closely related to the shape of the AP, such a difference could be expected. The ν_{Myo} -distributions looked qualitatively similar, yet the mean myocardium diffusivity was more than 13% faster in the pediatric patient group. This is in accordance with findings from animal experiments, suggesting a negative correlation between age and conduction velocity [Jones04].

Run-time performance was good considering the iterative nature of the cascaded EP personalization algorithm (Alg. 1). The average number of forward model runs until convergence was 108.7 ± 25.2 . This translates to about five to 15 minutes on a workstation computer with mid-range [graphics processing unit \(GPU\)](#). To reach this level of performance, the internal parameters of BOBYQA (initial step lengths, convergence criteria, ...) were tuned carefully on a subset of the patients. The goal was to find the best trade-off between accuracy and minimum number of required (costly) LBM-EP forward runs.

Estimation of Biomechanics Parameters

Next, the biomechanics personalization module was evaluated in terms of its ability to capture global indicators of cardiac function, focusing on LV function. From the results in Tab. 4.4 and Fig. 4.8 one can see that the personalization algorithm enabled to fit the model such that computed EF and SV were accurate within $1.3 \pm 1.3\%$ and




Figure 4.8: Bland-Altman plots showing the good agreement between measured and computed global indicators of LV cardiac function. The blue line is the mean of the differences and the red lines delineate the 95% CI. The purple and yellow dots represent adults and children, respectively.

4.1 ± 3.2 ml of the corresponding measurements, which is below clinical variability of 6% [Mood 15] (for EF measured from SSFP cardiac MRI, the imaging protocol used in this work). Errors between computed RV SV and measured LV SV were

Table 4.4: Goodness of fit in terms of absolute differences between measured and computed global LV function indicators.

Absolute error	Mean	SD	Min	50%	80%	95%	Max	Unit
$ \overline{EF}_{LV} - \overline{EF}_{LV} $	1.3	1.3	0.0	1.0	1.9	4.3	5.8	%
$ \overline{SV}_{LV} - \overline{SV}_{LV} $	4.1	3.2	0.0	3.5	5.5	11.2	14.8	ml
$ \overline{SV}_{RV} - \overline{SV}_{LV} $	9.0	9.7	0.1	5.8	15.1	26.6	56.0	ml

larger on average, albeit this was expected and is acceptable given the setup of the the personalization procedure (GLOBALEM), and in particular the employed objective function, which was designed to respect the difficulties in RV segmentation and tracking (cf. Sec. 4.3.4). Overall, the personalized model could capture quantitatively well the main global features of cardiac function for a large patient population.

Also qualitatively, the goodness of fit was evaluated. For instance, the dynamic MRI sequences were compared to EM-simulated motion (Fig. 4.9), and the fit between measured and computed time-varying ventricular pressure and volume curves were analyzed. For most patients, very good agreement was observed by visual inspection, see Fig. 4.10. Notable misfits occurred mostly towards the end of ventricular diastole, when the atria contract (atrial systole) and rapidly push blood through the atrioventricular valves into the ventricles, which can cause sudden increase in ventricular volume and pressure. Since in the employed whole-heart model, the atrium models are highly simplified, and because the current personalization pipeline covers only rudimentary atrium model personalization, such mismatches were expected for certain patients.




Figure 4.9: Qualitative goodness of fit between cine MRI acquisition and motion computed by the model. Contours overlaid on top of the short-axis image slice represent cut through the computed myocardium surface at three different time points, from diastole (left) to peak-systole (right). Computed motion and images fit qualitatively well throughout the heart cycle. Colors encode active cell force (red is strongest).




Figure 4.10: Goodness of fit between measured (black) and computed (red) LV pressure (top) and volume (bottom) curves for three example patients (left to right) after whole-heart EM personalization. Thin curves in pressure plots represent arterial (aortic) pressures.

Distributions of estimated passive and active biomechanics parameters for all patients are visualized in Fig. 4.11. Estimates follow a fairly mono-modal distribution with highest frequency near the parameter default values (cf. last paragraph of Sec. 4.3.3). No major differences between both patient groups were found, although the passive stiffness parameter, β , appears to have more “outliers” towards larger values (stiffer tissue) for the adult group, which may be explained to some extent by fibrotic cardiac remodeling of the aging heart as suggested in literature [Bier 11].

Finally, run-time performance was evaluated. The average number of forward model runs until convergence of the biomechanics personalization module was 58.5 ± 21.7 . Considering that two heart cycles were computed in each iteration (measurements taken from second cycle only) and the computation time of each cycle is typically two to five minutes (depending on mesh size, heart rate, etc.), the full procedure takes between four and twelve hours on standard hardware. It is therefore not suitable for specific clinical applications where immediate answers are required. Nevertheless, for many conceivable use-cases, especially in diagnostics and treatment planning, such




Figure 4.11: Histograms showing the distributions of estimated LV active ($\hat{\sigma}$, k_{ATP} , k_{RS}) and passive (β) biomechanics parameters. Purple and yellow colors represent adult and children group, respectively.

wait times are acceptable (e. g., running the personalization over night and analyze the results the next day).

4.4.3 Preliminary Evaluation of Regional Mechanics Approach

Next, the regional biomechanics personalization was evaluated.

Experiments on Synthetic Data

The performance of the regional parameter estimation method was first analyzed assuming an idealized setup where perfect fit can be achieved between measurements and simulation. To this end, instead of relying on noisy and potentially erroneous regional volume curves derived from the MRI-based LV tracking, synthetic measurements were generated directly from the EM model of the selected patient. This implies that the obtained measurements are free from measurement noise and thus they can theoretically be matched perfectly during the personalization process. More importantly, the ground-truth regional contractility parameters, denoted $\hat{\sigma}_{AHA_1} \dots \hat{\sigma}_{AHA_{17}}$, were known, enabling quantitative evaluation of the estimated parameters.

Two patient datasets were selected for experimentation based on their evident asynchronous cardiac motion. A coarse biomechanical model calibration was performed based on measurements derived from global left and right ventricular volume curves to define the baseline model, see Sec. 4.3.3. The seven estimated model parameters include contractility values of LV and RV ($\hat{\sigma}_{LV}^*$, $\hat{\sigma}_{RV}^*$). Since the regional personalization method focuses on contractility only, all other parameters remained fixed from now on. Next, all regional LV values were initialized to the estimated global LV contractility: $\hat{\sigma}_{AHA_i} = \hat{\sigma}_{LV}^* \forall i \in \{1 \dots 17\}$. The to be tuned regional parameters $\hat{\sigma}_{AHA_i}$ and the RV contractility $\hat{\sigma}_{RV}$ are gathered in a parameter vector denoted \mathbf{x}_{EM}^{reg} .

Sensitivity Analysis. Before looking into evaluation of the regional parameter estimation process itself, the idea was to make sure that the parameters in \mathbf{x}_{EM}^{reg} can be identified from the regional volumes. Ideally, one would observe a one-to-one mapping between parameters and measurements: modifying the contractility of a single AHA

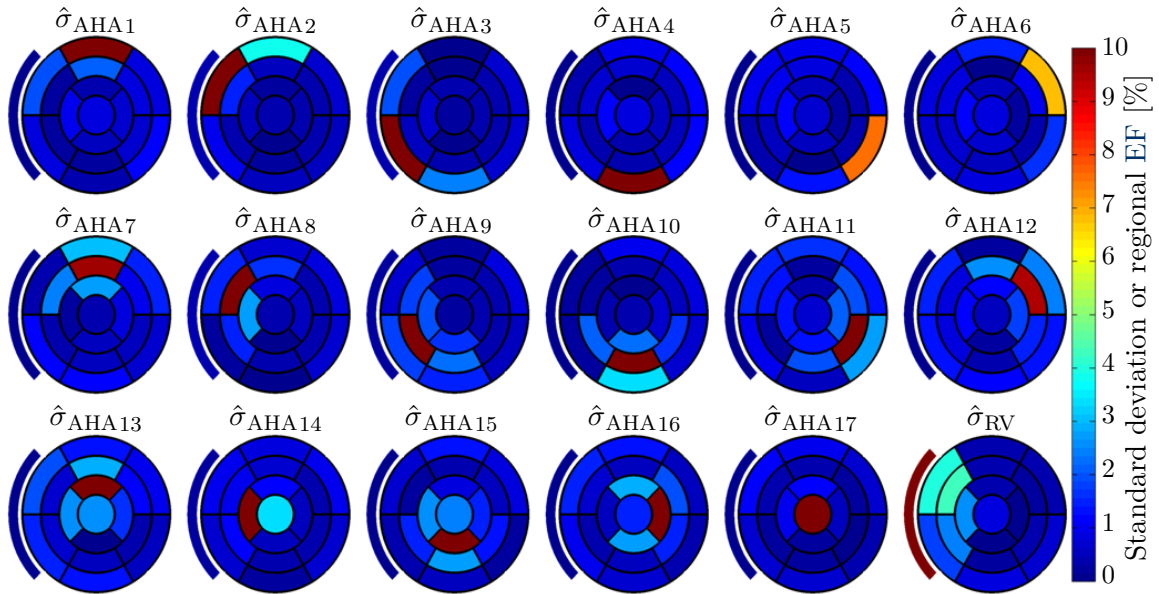


Figure 4.12: Regional parameter sensitivity for the first evaluated patient. Wedge on the left to each AHA “Bull’s eye” plot represents the RV. The title of each plot lists the parameter that has been modified to generate that plot.

segment would affect only the regional volume corresponding to that segment (or a small subset of AHA segments in close neighborhood). To this end, a sensitivity analysis was carried out. The forward EM model was computed using four different settings per parameter, where the value of each parameter $\hat{\sigma}_{AHA_i} \in \mathbf{x}_{EM}^{\text{reg}}$ was modified individually by ± 50 kPa and ± 100 kPa. Then, the regional volume variability was quantified by computing the standard deviation of the regional volumes among the four simulations per parameter.

Results for the variability in “regional EF” for one patient are illustrated in Fig. 4.12. As one can see, there is a clear correlation between model parameter and the region of measurement variability for all parameters. This is true for both datasets (plots for the second evaluated patient are not shown but were qualitatively similar), suggesting generalizability of these findings. Although a perfect one-to-one mapping was not observed (due to the intrinsic inter-dependency between the different regions in the ventricle), these findings suggest that the selected measurements and parameters still provide a good basis for regional personalization of active contraction force.

Regional Personalization in a Simplified Setup. Building upon the results of the sensitivity analysis, a simple experiment was designed to evaluate the identifiability of the regional model parameters from the regional volume measurements in the framework of an actual (but highly simplified) parameter estimation process. In particular, a series of experiments was conducted, where for each experiment a single parameter $\hat{\sigma}_{AHA_i} \in \mathbf{x}_{EM}^{\text{reg}}$ was modified by a given value. This mimics a scenario where the LV is contracting normally everywhere, except in the AHA segment i , which may be caused by a myocardial scar in that segment. Then the forward model was computed and the resulting simulated regional volume curves were treated as mea-

measurements in the subsequent estimation process. The goal was to estimate $\hat{\sigma}_{\text{AHA}_i}$, while all other parameters remained fixed at their ground-truth values. This one-parameter estimation problem was solved using **BOBYQA** optimization. Initialized at the global **LV** contractility value, $\hat{\sigma}_{\text{AHA}_i} = \hat{\sigma}_{\text{LV}}$, the optimizer iteratively minimized the overall regional cost $\zeta_{\text{EM}}^{\text{reg}}$ from Sec. 4.3.4, yielding $\hat{\sigma}_{\text{AHA}_i}^* = \text{argmin} \zeta_{\text{EM}}^{\text{reg}}$.

Quantitative evaluation based on 20 experiments (two patients \times five selected regional contractility parameters ($\hat{\sigma}_{\text{AHA}_i}, i \in \{1, 4, 8, 16, 17\}$) \times four parameter modifications (± 100 kPa)) showed that the estimation was robust and yielded low misfit $\varepsilon_i = |\hat{\sigma}_{\text{AHA}_i}^* - \overline{\hat{\sigma}_{\text{AHA}_i}}|$ between estimated ($\hat{\sigma}_{\text{AHA}_i}^*$) and ground-truth ($\overline{\hat{\sigma}_{\text{AHA}_i}}$) parameter values. On average, the absolute misfit was $\varepsilon = 0.3 \pm 0.4$ kPa, i. e., less than 1% of the parameter modification (± 100 kPa).

Towards Realistic Regional Personalization Scenarios. Finally, a more complex estimation scenario, much closer to a real case setup, was mimicked. In particular, the experiment design from the previous section was partially re-used, but instead of estimating only the parameter of the modified region, $\hat{\sigma}_{\text{AHA}_i}$, and assuming the values for all other regions were known, the ideas of the hierarchical multi-step estimation procedure described at the end of Sec. 4.3.4 were followed. The goal was to evaluate whether the regional personalization approach can be used to identify automatically the **AHA** region whose contractility was altered, by analyzing residual errors. The contractility value $\hat{\sigma}_{\text{AHA}_i}$ for the modified region and the value for all other regions were estimated as follows:

1. Synthetic data setup (same as in previous experiment)
 - (a) Select one parameter $\hat{\sigma}_{\text{AHA}_i} \in \mathbf{x}_{\text{EM}}^{\text{reg}}$ and alter its value by ± 100 kPa.
 - (b) Compute (synthetic) regional volume measurements using forward **EM** model.
2. Global personalization
 - (a) Estimate global **LV** contractility $\hat{\sigma}_{\text{LV}}^*$.
 - (b) Initialize regional **LV** contractility values with the estimated global **LV** value: $\hat{\sigma}_{\text{AHA}_i} = \hat{\sigma}_{\text{LV}}^* \forall i \in \{1 \dots 17\}$.
3. Regional refinement
 - (a) Analyze residual errors for each **AHA** regional volume curve.
 - (b) Identify the modified parameter, $\hat{\sigma}_{\text{AHA}_i}$, as the one corresponding to the **AHA** region with the largest error.
 - (c) Estimate one contractility value for $\hat{\sigma}_{\text{AHA}_i}$ and one for all other regions, $\mathbf{x}_{\text{EM}}^{\text{reg}} \setminus \hat{\sigma}_{\text{AHA}_i}$, using **BOBYQA** and the regional cost function $\zeta_{\text{EM}}^{\text{reg}}$.

The modified parameter could be identified correctly in all 20 experiments as selecting the largest residual regional volume error after the global personalization step. Furthermore, low errors after regional refinement were observed for both estimated parameters: 2.2 ± 2.6 kPa and 0.0 ± 0.1 kPa for the contractility values of the identified




Figure 4.13: Left: residual regional EF misfit (color-coded) after initial global LV contractility personalization. Right: Manual scar annotation based on review of non-routine MRI protocol geared towards myocardial scar localization. A pronounced transmural scar in the basal inferior (4), basal inferolateral (5), mid inferior (10), mid inferolateral (11), and apical inferior (15) AHA segments is present according to the annotation.

region and the value for all other regions, respectively. An explanation why the mean error of the former is notably larger is the equal weighting of all regions in the cost function, resulting in a $16\times$ larger weight for the other regions. Adapted weights in the cost function of the refinement step could possibly further improve the results.

Preliminary Test on Real Data

Finally, the regional personalization approach was evaluated in a real-case scenario. The goal of this preliminary experiment was to determine whether the proposed strategy can identify pathological regions in the LV myocardium after the initial coarse personalization, which is the foundation for the subsequent (finer) personalization steps. To this end, first, global contractility estimation was performed. Then, the residual regional volume errors per AHA segment were computed and the spatial error distribution was compared qualitatively against manual clinical scar annotations as shown in Fig. 4.13. A desirable outcome would be a clear overlap between scarred segments in the clinical annotation (Fig. 4.13), and segments with large residual errors (greenish and red colors in the plot on the left). Although not the entire extent of the annotated transmural scar was identified, the good agreement in segments 11 and 15 suggests that the obtained results are promising. More work is needed in the future to make sure this conclusion holds for larger populations.

Next, still blinded to the clinical annotation, the contractility was estimated in two regions: the abnormal regions exhibiting the largest errors (here comprising AHA segments 11 and 15) and the healthy region (all other AHA segments combined). BOBYQA personalization using the regional cost function (cf. Sec. 4.3.4) estimated contractility values of 52 kPa for the region defined as abnormal and 102 kPa for the healthy region, respectively. The much lower contractility in the abnormal region could indeed be seen as an indicator of potential regional pathologies such as scars.

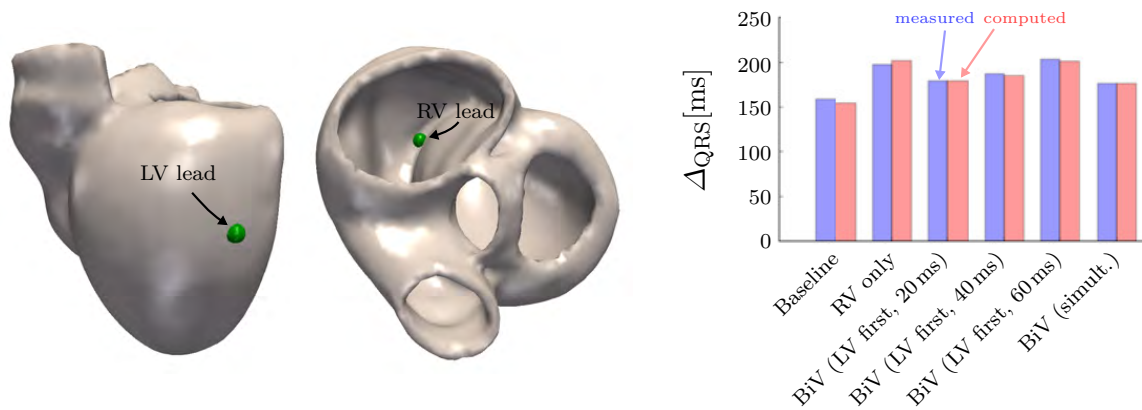


Figure 4.14: Left: Interactive positioning of CRT leads to mimic placement observed in interventional imaging data. Right: Comparison of measured and predicted QRS duration for the tested pacing protocols. “RV only” means that only the RV was stimulated, “BiV” stands for bi-ventricular pacing. In brackets: the CRT lead that was first stimulated and the delay between the first and second lead are noted.

4.5 Predictive Power of Personalized Model

The results presented in the previous section demonstrate that the proposed personalization pipeline can indeed fit the model to the available patient data. However, good fit does not necessarily guarantee good predictive power of the model when it comes to prediction of disease course or therapy outcome. As a first step towards a large-scale and in-depth analysis of the model’s predictive power, a preliminary virtual cardiac resynchronization therapy (CRT) pilot experiment was conducted. The hypothesis was that the 3D whole-heart model personalized using the proposed pipeline can provide predictors of acute CRT outcomes (cf. Sec. 1.1.1).

One patient from the cohort with wide QRS complex ($\Delta_{QRS} = 158$ ms) and LV EF of 45%, who underwent CRT, was selected for this experiment. Once the model was fully personalized (baseline), virtual CRT was performed according to the following protocol. First, the positions of the CRT leads (cf. Fig. 4.14) were selected interactively on the 3D anatomical model (Sec. 3.1) at approximately the same position as during the real procedure (based on interventional X-ray images). CRT stimulation was then modeled by adding stimulation currents in the EP model at the identified lead positions (cf. Sec. 3.2.1). Duration, timing and frequency of the pacing stimulation were adjusted according to the electrophysiologist’s report. A forward simulation of the 3D whole-heart model (with altered EP due to external stimulation) was then performed per tested pacing protocol and acute CRT predictors were computed from the model response. These predictors were then validated against the measurements (Δ_{QRS} from ECG plots) acquired by the electrophysiologist during device programming. Fig. 4.14, right panel, shows measured versus predicted Δ_{QRS} at baseline and for the five different pacing protocols. No protocol managed to improve Δ_{QRS} of that patient (ideally, Δ_{QRS} should be below 120 ms), nevertheless, the model predicted these outcomes accurately.

Next, acute hemodynamics effects after virtual CRT were analyzed based on the varying model responses of the different pacing protocols in terms of EF and maximum rate of ventricular pressure increase ($\max(dP_{\text{ven}}/dt)$). The changes in hemodynamics responses, despite being small in absolute value, were consistent with the changes in Δ_{QRS} in terms of the relative performance of the different CRT protocols. Interestingly, all tested CRT protocols except “RV only” produced an improvement in the hemodynamics parameters (e.g., 11% increase in $\max(dP_{\text{ven}}/dt)$ for simultaneous bi-ventricular pacing), although Δ_{QRS} values worsened. This suggests that cardiac function may be more sensitive to the spatial pattern of electrical activation than to the overall width of the QRS complex.

These promising results and findings need to be further analyzed and validated. Hemodynamics predictions need to be compared against true measurements of acute hemodynamics response directly after CRT programming (e.g., based on interventional echocardiography images), and the study needs to be extended to a larger number of patients.

4.6 Summary

This chapter presented a streamlined pipeline for comprehensive personalization of the multi-physics 3D whole-heart electromechanics (EM) model described in Chap. 3 using only routinely acquired clinical imaging and 12-lead ECG data, as well as intracardiac pressures from cardiac catheterization, if available. Based on the anatomical model described in Sec. 3.1, first, an automatic approach for computation of patient-specific hemodynamics (HD) boundary conditions from the available pressure and volume data was presented. It relies on some simple semi-automated data pre-processing steps, in particular for pressure selection and smoothing, and for synchronization of pressure and volume curves, to automatically compute five parameters each for both arterial Windkessel (WK) models (representing pulmonary artery and aorta) using bound-constrained optimization and a specifically designed cost function. Next, a fully-automated personalization strategy for four patient-specific cardiac electrophysiology (EP) model parameters was proposed. It was designed to match model-simulated ECG to the clinical data by fitting various ECG descriptors using a carefully designed multi-stage cascaded inverse optimization procedure, CASCADEEP. Next, an approach for comprehensive biomechanics personalization by global volume and pressure fitting, GLOBALEM, was presented. It adjusts seven biomechanical parameters, both active and passive, and two additional HD parameters of the atrial models to the patient measurements in a fully automatic fashion. Furthermore, two methods towards a more regional personalization of active biomechanics parameters were proposed, where instead of one contractility parameter for the entire left ventricle (LV), up to 17 parameters, one for each LV AHA segment, are estimated.

The full pipeline was then evaluated on a heterogeneous cohort of 113 patients, which to our knowledge is the largest patient number in any comparable study to date. About one third were children suffering from various cardiomyopathies, and the remaining ones were adults with dilated cardiomyopathy (DCM). The main objective was to assess the ability of the pipeline to accurately capture patients’ measured

cardiac physiology in the model. To this end, once personalization was performed, the model outputs were compared to various measured clinical indicators of cardiac function for all personalized model components. Throughout the entire spectrum of models and patients, in the vast majority, good to excellent agreement could be achieved: most errors were below the clinical variability of the measurements. Finally, regional biomechanics personalization was analyzed. Results from partially synthetic or idealized personalization setups were promising, and the test on real data even suggested that the approach could potentially help to detect and localize scars in the myocardium without using specialized imaging protocols. However, more work will be needed in the future to verify this preliminary finding. Finally, a pilot study to analyze the predictive capabilities of the model personalized with the proposed pipeline was conducted in a virtual CRT prediction scenario. Computed acute predictors of cardiac EP matched well the true measurements acquired by the electrophysiologist during CRT device implantation.

PART II

Advanced Personalization with Artificial Intelligence and Uncertainty Quantification

Intelligent Agent for Computational Model Personalization

5.1	The Potential of Artificial Intelligence for Model Personalization	67
5.2	Learning How to Personalize a Model	69
5.3	Experiments and Results	75
5.4	Summary	85

The design and implementation of robust model personalization algorithms like the ones described in the previous chapter is often a tedious and time-consuming task, and the resulting algorithms are highly model- and data-specific. Many “lessons learned” from the design of one algorithm for one particular model do not necessarily apply to a new algorithm for a new model. However, for specific clinical questions, specialized models are required, therefore, novel personalization algorithms continue to be designed frequently.

This chapter is based on [Neum 15b, Neum 16]. It revolves around the question whether **artificial intelligence (AI)** can potentially support the designers of new models by automatically designing efficient personalization methods that are tailored to *their model* and are robust to *their data*. In the first section, Sec. 5.1, the potential of AI for computational modeling, in particular model personalization, is discussed. Sec. 5.2 presents the proposed application of AI concepts to computational biophysical model personalization: an intelligent agent that learns by itself how to estimate model parameters from clinical data while being model-independent. In Sec. 5.3, the results of both synthetic and real-data experiments are presented. The chapter concludes with a short summary in Sec. 5.4.

5.1 The Potential of Artificial Intelligence for Model Personalization

As elaborated in Sec. 4.1, over the last decades, a wide variety of computational model personalization approaches have been proposed. Most methods aim to iteratively reduce the misfit between model output and measurements using inverse optimization algorithms. Applied blindly, those techniques could easily fail on unseen data, if not supervised, due to parameter ambiguity, data noise and local min-

ima [Komu 11, Neum 14b, Wall 14]. Therefore, complex algorithms, including CASCADEEP from Sec. 4.3.2, have been designed combining cascades of optimizers in a very specific way to achieve high levels of robustness, even on larger populations [Kayv 15, Seeg 15, Neum 14c]. However, those methods are often designed from tedious, trial-and-error-driven manual tuning, they are model-specific rather than generic, and their generalization to varying data quality cannot be guaranteed. On the contrary, if the personalization task is assigned to an experienced human, given enough time, he or she almost always succeeds in manually personalizing a model for any subject (although solution uniqueness is not guaranteed, see Chap. 6, but this is inherent to the problem).

There are several potential reasons why a human expert is often superior to standard automatic methods in terms of personalization accuracy and success rates. First, an expert is likely to have an intuition of the model’s behavior from his prior knowledge of the physiology of the modeled organ. Second, knowledge about model design and assumptions, and model limitations and implementation details certainly provide useful hints on the “mechanics” of the model. Third, past personalization of other datasets allows the expert to build up experience. The combination of prior knowledge, intuition and experience enables to solve the personalization task more effectively, even on unseen data.

Inspired by humans and contrary to previous works, the personalization problem is addressed from an AI perspective in this chapter. In particular, reinforcement learning (RL) methods [Sutt 98] developed by the AI community are applied to solve the parameter estimation task for computational physiological models. With its roots in control theory on the one hand, and neuroscience theories of learning on the other hand, RL encompasses a set of approaches to make an artificial agent learn from experience generated by interacting with its environment. Contrary to standard (supervised) machine learning (ML) [Bish 06], where the objective is to compute a direct mapping from input features to a classification label or regression output, RL aims to learn *how to* perform tasks. The goal of RL is to compute an optimal problem-solving strategy (agent behavior), e. g., a strategy to play the game “tic-tac-toe” successfully. In the AI field, such a behavior is often represented as a policy, a mapping from states, describing the current “situation” the agent finds itself in (such as the current locations of all “X” and “O” on the tic-tac-toe grid), to actions, which allow the agent to interact with the environment (place “X” on an empty cell) and thus influence that situation. The key underlying principle of RL is that of reward [Kael 96], which provides an objective means for the agent to judge the outcome of its actions. In tic-tac-toe, the agent may receive a high, positive reward if the latest action led to a horizontal, vertical or diagonal row full of “X” marks (winning), and a negative reward (punishment) if the latest action would allow the opponent to win in his next move. Based on such rewards, the artificial agent learns an optimal winning policy through trial-and-error interactions with the environment.

RL was first applied to game [Tesa 95] or simple control tasks. However, the past few years saw tremendous breakthroughs in RL for more complex, real-world problems [Nguy 11, Kvet 12]. Some noteworthy examples include [Mull 13], where the control entity of a robot arm learned to select appropriate motor primitives to play table tennis, and [Mnih 15], where the authors combine RL with deep learning to




Figure 5.1: Overview of VITO: a self-taught artificial agent for computational model personalization, inspired by how human operators approach the personalization problem.

train an agent to play 49 Atari games, yielding better performance than an expert in the majority of them. RL now also starts to appear more and more in the medical imaging domain [Litj 17]. For instance, in [Ghes 16a], an agent learns how to detect landmarks in 3D medical images, while being faster, more accurate and more generic than previous algorithms. RL applied to medical image registration is explored in [Liao 17].

Motivated by these recent successes, an RL-based personalization approach, henceforth called VITO, is proposed. The goal is to design a framework that can learn by itself how to estimate model parameters from clinical data while being model-independent. As illustrated in Fig. 5.1, first, like a human expert, VITO assimilates the behavior of the physiological model under consideration in an off-line, one-time only, data-driven exploration phase. From this knowledge, VITO learns the optimal strategy using RL [Sutt 98]. The goal of VITO during the on-line personalization phase is then to sequentially choose actions that maximize future rewards, and therefore bring VITO to the state representing the solution of the personalization problem. To setup the algorithm, the user needs to define what observations need to be matched, the allowed actions, and a hyper-parameter related to the desired granularity of the state space. Then everything is learned automatically. The algorithm does not depend on the underlying biophysical model.

VITO was evaluated on three different tasks. First, in a synthetic environment, convergence properties of the algorithm were analyzed. Then, two tasks involving real clinical data were evaluated: the inverse problem of cardiac electrophysiology (cf. Sec. 3.2) and the personalization of a lumped-parameter model of whole-body circulation (cf. Sec. 3.4). The obtained results suggested that VITO can achieve equivalent (or better) goodness of fit as standard optimization methods, increased robustness and faster convergence rates.

5.2 Learning How to Personalize a Model

This section presents the RL framework for computational model personalization. Sec. 5.2.1 introduces Markov decision process (MDP). Sec. 5.2.2 shows how the personalization problem can be reformulated in terms of an MDP. Sec. 5.2.3 describes how the artificial agent, VITO, learns how the model behaves. Next, Sec. 5.2.4 provides details about state space quantization, and Sec. 5.2.5 describes how the model knowledge is encoded in the form of transition probabilities. All steps mentioned so

far are performed in an off-line training phase. Finally, Sec. 5.2.6 explains how the learned knowledge is applied on-line to personalize unseen data.

5.2.1 Model-based Reinforcement Learning

MDP Definition

A crucial prerequisite for applying RL is that the problem of interest, here personalization, can be modeled as a [Markov decision process \(MDP\)](#). An MDP is a mathematical framework for modeling decision making when the decision outcome is partly random and partly controlled by a decision maker [[Sutt 98](#)]. Formally, an MDP is a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \gamma)$, where:

- \mathcal{S} is the finite set of $n_{\mathcal{S}}$ states that describe the agent’s environment, and $s_t \in \mathcal{S}$ is the state at time t .
- \mathcal{A} is the finite set of $n_{\mathcal{A}}$ actions, which allow the agent to interact with the environment, and $a_t \in \mathcal{A}$ denotes the action performed at time t .
- $\mathcal{T} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0; 1]$ is the stochastic transition function, where $\mathcal{T}(s_t, a_t, s_{t+1})$ describes the probability of arriving in state s_{t+1} after the agent performed action a_t in state s_t .
- $\mathcal{R} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ is the scalar reward function, where $r_{t+1} = \mathcal{R}(s_t, a_t, s_{t+1})$ is the immediate reward the agent receives at time $t + 1$ after performing action a_t in state s_t resulting in state s_{t+1} .
- $\gamma \in [0; 1]$ is the discount factor that controls the importance of future versus immediate rewards.

Value Iteration

The value of a state, $\mathcal{V}^*(s)$, is the expected discounted reward the agent accumulates when it starts in state s and acts optimally in each step:

$$\mathcal{V}^*(s) = E \left\{ \sum_{k=0}^{\infty} \gamma^k r_{t+k+1} \middle| s_t = s \right\} , \quad (5.1)$$

where $E\{\cdot\}$ denotes the expected value given the agent always selects the optimal action, and t is any time step. \mathcal{V}^* can be computed using value iteration [[Sutt 98](#)], an iterative algorithm based on dynamic programming. In the first iteration $i = 0$, let $\mathcal{V}_i : \mathcal{S} \rightarrow \mathbb{R}$ denote an “initial guess” for the value function that maps states to arbitrary values. Further, let $\mathcal{Q}_i : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ denote the i^{th} “state-action value function”-guess, which is computed as:

$$\mathcal{Q}_i(s, a) = \sum_{s' \in \mathcal{S}} \mathcal{T}(s, a, s') [\mathcal{R}(s, a, s') + \gamma \mathcal{V}_i(s')] . \quad (5.2)$$

Value iteration iteratively updates \mathcal{V}_{i+1} from the previous \mathcal{Q}_i :

$$\forall s \in \mathcal{S} : \quad \mathcal{V}_{i+1}(s) = \max_{a \in \mathcal{A}} \mathcal{Q}_i(s, a) , \quad (5.3)$$

until the left- and right-hand side of Eq. (5.3) are equal for all $s \in \mathcal{S}$; then $\mathcal{V}^* \leftarrow \mathcal{V}_{i+1}$ and $\mathcal{Q}^* \leftarrow \mathcal{Q}_{i+1}$. From this equality relation, also known as the Bellman equation [Bell57], one can obtain an optimal problem-solving strategy for the problem described by the MDP (assuming that all components of the MDP are known precisely). It is encoded in terms of a deterministic optimal policy $\pi^* : \mathcal{S} \rightarrow \mathcal{A}$:

$$\pi^*(s) = \operatorname{argmax}_{a \in \mathcal{A}} \mathcal{Q}^*(s, a) , \quad (5.4)$$

i. e., a mapping that tells the agent in each state the optimal action to take.

Stochastic Policy

In this work, not all components of the MDP are known precisely, instead some are approximated from training data. Value iteration, however, assumes an exact MDP to guarantee optimality of the computed policy. Therefore, instead of relying on the deterministic policy π^* (Eq. (5.4)), a generalization to stochastic policies $\tilde{\pi}^*$ is proposed here to mitigate potential issues due to approximations. Contrary to Eq. (5.4), where for each state only the one action with maximum \mathcal{Q}^* -value is considered, a stochastic policy stores several candidate actions with similar high \mathcal{Q}^* -value and returns one of them through a random process each time it is queried. To this end, the $\mathcal{Q}^*(s, \cdot)$ -values for a given state s are first normalized:

$$\tilde{\mathcal{Q}}_s^*(a) = \frac{\mathcal{Q}^*(s, a) - \min_{a' \in \mathcal{A}}[\mathcal{Q}^*(s, a')]}{\max_{a' \in \mathcal{A}}[\mathcal{Q}^*(s, a')] - \min_{a' \in \mathcal{A}}[\mathcal{Q}^*(s, a')]} . \quad (5.5)$$

All actions whose normalized $\tilde{\mathcal{Q}}_s^*$ -value is below a threshold of $\epsilon = \frac{4}{5}$ (set empirically) are discarded, while actions with large values are stored as potential candidates. Each time the stochastic policy is queried, $a = \tilde{\pi}_\epsilon^*(s)$, it returns one of the candidate actions, a , selected randomly with probability proportional to its $\tilde{\mathcal{Q}}_s^*$ -value: $\tilde{\mathcal{Q}}_s^*(a) / \sum_{a'} \tilde{\mathcal{Q}}_s^*(a')$; the sum is over all candidate actions a' .

5.2.2 Reformulation of the Model Personalization Problem

The model personalization problem described in Sec. 4.2 is mapped to an MDP as follows:

- **States:** An MDP state encodes the misfit between the computed model state (outcome of forward model run) and the patient’s measurements. Thus, MDP states carry the same type of information as objective vectors \mathbf{c} , yet the number of MDP states has to be finite (Sec. 5.2.1), while there are an infinite number of different objective vectors due to their continuous nature. Therefore, the space of objective vectors in \mathbb{R}^{n_c} is reduced to a finite set of *representative states*: the MDP states \mathcal{S} , each $s \in \mathcal{S}$ covering a small region in that space. One of those states, $\hat{s} \in \mathcal{S}$, encodes personalization success as it is designed such that it covers exactly the region where all convergence criteria are satisfied. The goal of VITO is to learn how to reach that state.

- **Actions:** VITO’s actions modify the parameters \mathbf{x} to fulfill the objectives \mathbf{c} . An action $a \in \mathcal{A}$ consists in either in- or decrementing one parameter x_i by $1\times$, $10\times$ or $100\times$ a user-specified reference value δ_i with $\boldsymbol{\delta} = (\delta_1, \dots, \delta_{n_x})^\top$. This empirically defined quantization of the intrinsically continuous action space yielded good results for the problems considered in this work.
- **Transition function:** \mathcal{T} encodes the agent’s knowledge about the computational model f and is learned automatically as described in Sec. 5.2.5.
- **Rewards:** Inspired by the “mountain car” benchmark [Sutt 98], the rewards are defined as always being equal to $\mathcal{R}(s, a, s') = -1$ (punishment), except when the agent performs an action resulting in personalization success, i. e., when $s' = \hat{s}$. In that case, $\mathcal{R}(\cdot, \cdot, \hat{s}) = 0$ (no punishment).
- **Discount factor:** The large discount factor $\gamma = 0.99$ encourages finding a policy that favors future over immediate rewards, as the latter could prefer finding local over global optima.

5.2.3 Learning Model Behavior through Model Exploration

Like a human operator, VITO first learns how the model “behaves” by experimenting with it. This is done through a “self-guided sensitivity analysis”. A batch of sample transitions is collected through model exploration episodes $\mathcal{E}^p = \{e_1^p, e_2^p, \dots\}$. An episode e_i^p is a sequence of $n_{e\text{-steps}}$ consecutive transitions generated from the model f and the patient p for whom the target measurements \mathbf{z}^p are known. An episode is initiated at time $t = 0$ by generating random initial model parameters \mathbf{x}_t within the physiologically plausible domain Ω . From the outputs of a forward model run $\mathbf{y}_t = f(\mathbf{x}_t)$, the misfits to the patient’s corresponding measurements are computed, yielding the objectives vector $\mathbf{c}_t = \varepsilon(\mathbf{y}_t, \mathbf{z}^p)$. Next, a random exploration policy π_{rand} that selects an action according to a discrete uniform probability distribution over the set of actions is employed. The obtained $a_t \in \mathcal{A}$ is then applied to the current parameter vector, yielding modified parameter values $\mathbf{x}_{t+1} = a_t(\mathbf{x}_t)$. From the output of the forward model run $\mathbf{y}_{t+1} = f(\mathbf{x}_{t+1})$ the next objectives \mathbf{c}_{t+1} are computed. The next action a_{t+1} is then selected according to π_{rand} , and this process is repeated $n_{e\text{-steps}} - 1$ times. Hence, each episode can be seen as a set of consecutive tuples:

$$e = \{(\mathbf{x}_t, \mathbf{y}_t, \mathbf{c}_t, a_t, \mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{c}_{t+1}), \quad t = 0, \dots, n_{e\text{-steps}} - 1\} . \quad (5.6)$$

In this work, $n_{e\text{-steps}} = 100$ transitions are created in each episode as a trade-off between sufficient length of an episode to cover a real personalization scenario and sufficient exploration of the parameter space.

The model is explored with many different training patients and the resulting episodes are combined into one large training episode set $\mathcal{E} = \bigcup_p \mathcal{E}^p$. The underlying hypothesis (verified in experiments) is that the combined \mathcal{E} allows to cancel out peculiarities of individual patients, i. e., to abstract from patient-specific to model-specific knowledge.

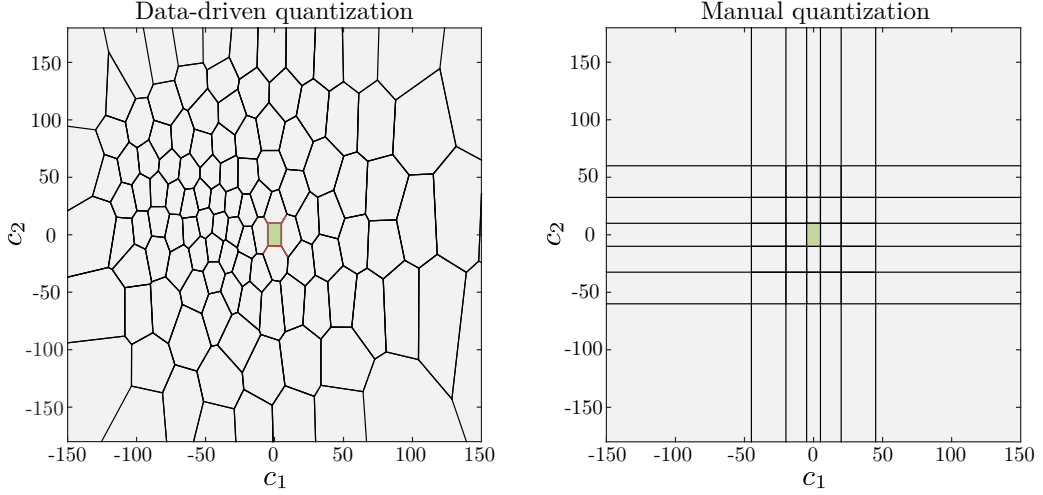


Figure 5.2: Left: Example data-driven quantization of a two-dimensional state space into $n_{\mathcal{S}} = 120$ representative states. The states are distributed according to the observed objective vectors \mathbf{c} in one of the experiments in Sec. 5.3.2, where the objectives were QRS duration [ms] (c_1) and electrical axis [deg] (c_2). The center rectangle (green region) denotes the success state \hat{s} . Right: Manual quantization as suggested initially in [Neum 15b].

5.2.4 From Computed Objectives to Representative MDP State

As mentioned above, the continuous space of objective vectors is quantized into a finite set of representative MDP states \mathcal{S} . A data-driven approach is proposed. First, all objective vectors observed during training are clustered according to their distance to each other. To ensure equal contribution among all objectives (to cancel out different units, etc.), the distance between two objective vectors ($\mathbf{c}_1, \mathbf{c}_2$) is defined relative to the inverse of the convergence criteria $\boldsymbol{\psi}$:

$$\|\mathbf{c}_1 - \mathbf{c}_2\|_{\boldsymbol{\psi}} = \sqrt{(\mathbf{c}_1 - \mathbf{c}_2)^\top \text{diag}(\boldsymbol{\psi})^{-1} (\mathbf{c}_1 - \mathbf{c}_2)}, \quad (5.7)$$

where $\text{diag}(\boldsymbol{\psi})^{-1}$ denotes a diagonal matrix with $(\frac{1}{\psi_1}, \frac{1}{\psi_2}, \dots)$ along its diagonal. The centroid of a cluster is the centroid of a representative state. In addition, a special “success state” representing personalization success, denoted \hat{s} , is created, which covers the region in state space where all objectives are met: $\forall i : |c_i| < \psi_i$. The full algorithm is described in appendix A.3.1.

Finally, an operator $\phi : \mathbb{R}^{n_c} \rightarrow \mathcal{S}$ that maps continuous objective vectors \mathbf{c} to representative MDP states is introduced:

$$\phi(\mathbf{c}) = \underset{s \in \mathcal{S}}{\text{argmin}} \|\mathbf{c} - \boldsymbol{\xi}_s\|_{\boldsymbol{\psi}} \quad (5.8)$$

where $\boldsymbol{\xi}_s$ denotes the centroid corresponding to state s . For an example state space quantization see Fig. 5.2.

5.2.5 Transition Function as Probabilistic Model Representation

In this work, the stochastic MDP transition function \mathcal{T} encodes the agent’s knowledge about the computational model f . It is learned from the training data \mathcal{E} . First, the

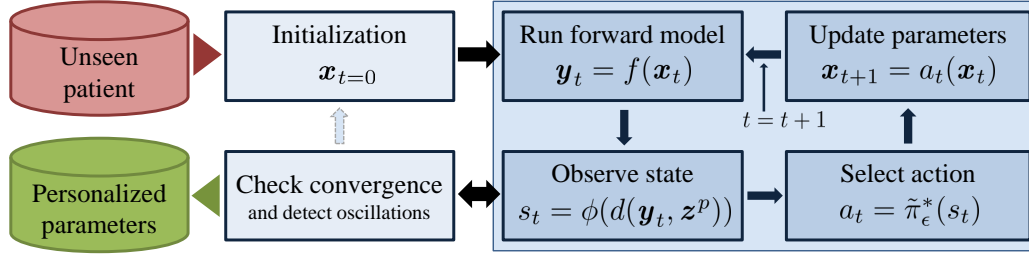


Figure 5.3: VITO's probabilistic on-line personalization phase. See text for details.

individual samples $(\mathbf{x}_t, \mathbf{y}_t, \mathbf{c}_t, a_t, \mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{c}_{t+1})$ are converted to state-action-state transition tuples $\hat{\mathcal{E}} = \{(s, a, s')\}$, where $s = \phi(\mathbf{c}_t)$, $a = a_t$ and $s' = \phi(\mathbf{c}_{t+1})$. Then, \mathcal{T} is approximated from statistics over the observed transition samples:

$$\mathcal{T}(s, a, s') = \frac{|\{(s, a, s') \in \hat{\mathcal{E}}\}|}{\sum_{s'' \in \mathcal{S}} |\{(s, a, s'') \in \hat{\mathcal{E}}\}|}, \quad (5.9)$$

where $|\{\cdot\}|$ denotes the cardinality of a set. If $n_{\mathcal{S}}$ and $n_{\mathcal{A}}$ are large compared to the total number of samples it may occur that some state-action combinations are not observed: $|\{(s, a, \cdot) \in \hat{\mathcal{E}}\}| = 0$. In that case uniformity is assumed: $\forall s'' \in \mathcal{S} : \mathcal{T}(s, a, s'') = 1/n_{\mathcal{S}}$.

\mathcal{M} is now fully defined. Value iteration (Sec. 5.2.1) is applied and the stochastic policy $\tilde{\pi}_\epsilon^*$ is computed, which completes the off-line phase.

5.2.6 On-line Model Personalization

On-line personalization, as illustrated in Fig. 5.3, can be seen as a two-step procedure. First, VITO initializes the personalization of unseen patients from training data. Second, VITO relies on the computed policy $\tilde{\pi}_\epsilon^*$ to guide the personalization process.

Data-driven Initialization

Good initialization can be decisive for a successful personalization. VITO's strategy is to search for forward model runs in the training database \mathcal{E} for which the model state $f(\mathbf{x}) = \mathbf{y} \approx \mathbf{z}^p$ is similar to the patient's measurements. To this end, VITO examines all parameters $\Xi = \{\mathbf{x} \in \mathcal{E} \mid f(\mathbf{x}) \approx \mathbf{z}^p\}$ that yielded model states similar to the patient's measurements. Due to ambiguities induced by the different training patients, data noise and model assumptions, Ξ could contain significantly dissimilar parameters. Hence, picking a single $\mathbf{x} \in \Xi$ might not yield the best initialization. Analyzing Ξ probabilistically instead helps VITO to find likely initialization candidates. The details of the initialization procedure are described in appendix A.3.2. Given the patient's measurements \mathbf{z}^p , the procedure outputs a set of initialization candidates $\{\mathbf{x}_0\}$, sorted by likelihood.

Probabilistic Personalization

As illustrated in Fig. 5.3, from the most likely initial parameters $\mathbf{x}_0 \in \{\mathbf{x}_0\}$, VITO computes the forward model $\mathbf{y}_0 = f(\mathbf{x}_0)$ and the misfit between the model output

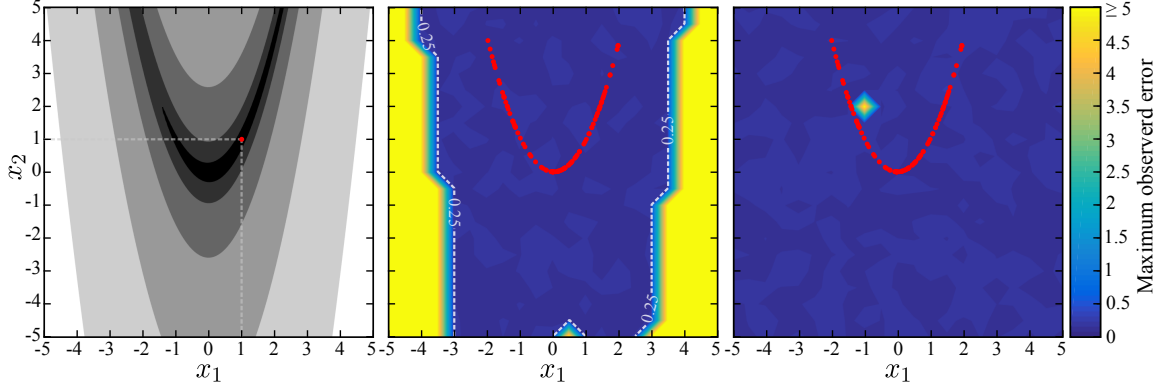


Figure 5.4: Left: Contour plot of the Rosenbrock function $f^{\alpha=1}$ with global minimum at $\mathbf{x} = (1, 1)^\top$ (red dot). The color scale is logarithmic for visualization purposes: the darker, the lower the function value. Mid: Maximum L_2 -error in parameter space after personalization over all functions for varying initial parameter values. See text for details. Yellow represents errors ≥ 5 (maximum observed error ≈ 110). Right: Same as mid panel, except the extended action set was used. The red dots are the 100 ground-truth parameters $\mathbf{x} = (\alpha, \alpha^2)^\top$ generated for random α .

and the patient’s measurements $\mathbf{c}_0 = \varepsilon(\mathbf{y}_0, \mathbf{z}^p)$ to derive the first state $s_0 = \phi(\mathbf{c}_0)$. Given s_0 , VITO decides from its policy the first action to take $a_0 = \tilde{\pi}_\epsilon^*(s_0)$, and walks through state-action-state sequences to personalize the computational model f by iteratively updating the model parameters through MDP actions. Bad initialization could lead to oscillations between states as reported in literature RL works [Kvet 12, Neum 15b]. Therefore, upon detection of an oscillation, which is done by monitoring the parameter traces to detect recurring sets of parameter values, the personalization is re-initialized at the second-most-likely $\mathbf{x}_0 \in \{\mathbf{x}_0\}$, etc. If all $|\{\mathbf{x}_0\}|$ initialization candidates have been tested, a potential re-initialization will default to fully random within the physiologically plausible parameter domain Ω . The process terminates once VITO reaches state \hat{s} (success), or when a pre-defined maximum number of iterations is reached (failure).

5.3 Experiments and Results

VITO was applied to a synthetic parameter estimation problem and to two challenging problems involving real clinical data: personalization of a cardiac electrophysiology (EP) model (cf. Sec. 3.2), and a whole-body circulation (WBC) model. All experiments were conducted using leave-one-out cross-validation.

5.3.1 Synthetic Experiment: the Rosenbrock Function

First, VITO was employed in a synthetic scenario, where the ground-truth model parameters were known. The goals were to test the ability of VITO to optimize cost functions generically, and to directly evaluate the performance in parameter space.

Forward Model Description

The Rosenbrock function [Rose 60], see Fig. 5.4, left panel, is a non-convex function that is often used to benchmark optimization algorithms. It was treated as the forward model in this experiment:

$$f^\alpha(x_1, x_2) = (\alpha - x_1)^2 + 100 \cdot (x_2 - x_1^2)^2, \quad (5.10)$$

where $\mathbf{x} = (x_1, x_2)^\top$ were the model parameters to estimate for any α , and $f^\alpha : \Omega \rightarrow \mathbb{R}$. As described in Sec. 5.2, each of VITO’s actions $a \in \mathcal{A}$ in- or decrements a parameter value by multiples ($1\times, 10\times, 100\times$) of parameter-specific reference values. The reference values were set to $\boldsymbol{\delta} = (0.01, 0.01)^\top$, determined as 0.1% of the defined admissible parameter space per dimension, $\Omega = [-5; 5]^2$. The parameter $\alpha \in \mathbb{R}$ defines a family of functions $\{f^\alpha\}$. The goal was to find generically $\operatorname{argmin}_{x_1, x_2} f^\alpha(x_1, x_2)$.

The Rosenbrock function has a unique global minimum at $\mathbf{x} = (\alpha, \alpha^2)^\top$, where both terms $T_1 = (\alpha - x_1)$ and $T_2 = (x_2 - x_1^2)$ evaluate to 0. The personalization objectives were therefore defined as $\mathbf{c} = (|T_1 - 0|, |T_2 - 0|)^\top$, with the measured data $\mathbf{z} = (0, 0)^\top$ were zero for both objectives and the computed data $\mathbf{y} = (T_1, T_2)^\top$. The convergence criteria were set empirically to $\boldsymbol{\psi} = (0.05, 0.05)^\top$.

Evaluation

VITO was evaluated on $n_{\text{datasets}} = 100$ functions f^α with randomly generated $\alpha \in [-2, 2]$. In the off-line phase, for each function, $n_{\text{samples}} = 10 \cdot n_{e\text{-steps}} = 1000$ samples, i. e., ten training episodes, each consisting in $n_{e\text{-steps}} = 100$ transitions (Sec. 5.2.3), were generated to learn the policy. The number of representative states was set to $n_{\mathcal{S}} = 100$. To focus on VITO’s on-line personalization capabilities, both the data-driven initialization and the re-initialization on oscillation (Sec. 5.2.6) were disabled. In total, 441 experiments with different initializations (sampled on a 21×21 uniform grid spanned in Ω) were conducted. For each experiment all 100 functions were personalized using leave-one-family-function-out cross validation, and the error value from the function exhibiting the maximum L_2 -error (worst-case scenario) between ground-truth (α, α^2) and estimated parameters was plotted. As one can see from the large blue region in Fig. 5.4, mid panel, for the majority of initial parameter values VITO always converged to the solution (maximum L_2 -error < 0.25 ; the maximum achievable accuracy depended on the specified convergence criteria $\boldsymbol{\psi}$ and on the reference values $\boldsymbol{\delta}$, which “discretized” the parameter space). However, especially for initializations far from the ground-truth (near border regions of Ω), VITO was unable to personalize some functions properly, which was likely due to the high similarity of the Rosenbrock function shape in these regions.

To investigate this issue, the experiment was repeated after additional larger parameter steps were added to the set of available actions: $\mathcal{A}' = \mathcal{A} \cup \{\pm 500\delta_1; \pm 500\delta_2\}$. As shown in Fig. 5.4, right panel, VITO could now personalize successfully starting from any point in Ω . The single spot with larger maximum error (bright spot at approximately $\mathbf{x} = (-1, 2)^\top$) can be explained by VITO’s stochastic behavior: VITO may have become *unlucky* if it selected many unfavorable actions in sequence due to the randomness introduced by the stochastic policy. Enabling re-initialization on

oscillation solved this issue entirely. In conclusion, this experiment showed that VITO can learn how to minimize a cost function generically.

5.3.2 Personalization of Cardiac Electrophysiology Model

VITO was then tested in a cardiac electrophysiology scenario, in particular for personalization of the Graph-EP model coupled with 12-lead ECG as described in Sec. 3.2. Personalization was performed for real patients from actual clinical data. A total of $n_{\text{datasets}} = 83$ patients from the study cohort described in Sec. 4.4.1 were used for experimentation. As outlined in Sec. 3.1, for each patient, the end-diastolic bi-ventricular anatomy was segmented from short-axis steady-state free precession (SSFP) cine magnetic resonance imaging (MRI), a tetrahedral anatomical model including myofibers was estimated and a torso atlas affinely registered to the patient based on MRI scout images.

Forward Model Description

The Graph-EP algorithm itself, the propagation of computed myocardium potentials to torso potentials, and the derivation of electrocardiogram (ECG) measurements thereof is described in detail in Sec. 3.2. The model is controlled by the conduction velocities (in m/s) of myocardial tissue and left and right Purkinje network: $\mathbf{x} = (\nu_{\text{Myo}}, \nu_{\text{LV}}, \nu_{\text{RV}})^{\top}$. The admissible parameter space Ω was set to $[200; 1000]$ mm/s for ν_{Myo} and $[500; 5000]$ mm/s for both ν_{LV} and ν_{RV} . Reference increment values to build the action set \mathcal{A} were set to $\boldsymbol{\delta} = (5, 5, 5)^{\top}$ mm/s for the three model parameters. The goal of personalization was to estimate \mathbf{x} from the measured QRS duration and electrical axis (EA), denoted Δ_{QRS} and \angle_{EA} . Accounting for uncertainty in the measurements and errors in the model, a patient was considered personalized if Δ_{QRS} and \angle_{EA} misfits were below $\boldsymbol{\psi} = (5 \text{ ms}, 10^{\circ})^{\top}$, respectively.

Number of Representative States

In contrast to [Neum 15b], where state space quantization required manual tuning of various threshold values, the latest version of VITO relies on a single hyperparameter only: $n_{\mathcal{S}}$, the number of representative states (Sec. 5.2.4). To specify $n_{\mathcal{S}}$, eight patients were selected for scouting. Exhaustive search was performed for $n_{\mathcal{S}} \in \{10, 20, \dots, 490, 500\}$ representative states. The goodness of a given configuration was evaluated based on the success rate (relative number of successfully personalized cases according to convergence criteria $\boldsymbol{\psi}$) over five independent, consecutive, leave-one-patient-out cross-validated personalization runs of the eight patients. Furthermore, the average number of required forward model runs was considered. To this end, 100 training episodes ($100 \cdot n_{e\text{-steps}} = 10^4$ transition samples) per patient were generated for each personalization run as described in Sec. 5.2.3. As one can see from Fig. 5.5, good performance was achieved from 50 to 300 representative states. The large range of well performing $n_{\mathcal{S}}$ indicates a certain level of robustness with respect to that hyperparameter. A slight performance peak at 120 representative states was observed. Therefore, $n_{\mathcal{S}} = 120$ was selected for further experimentation

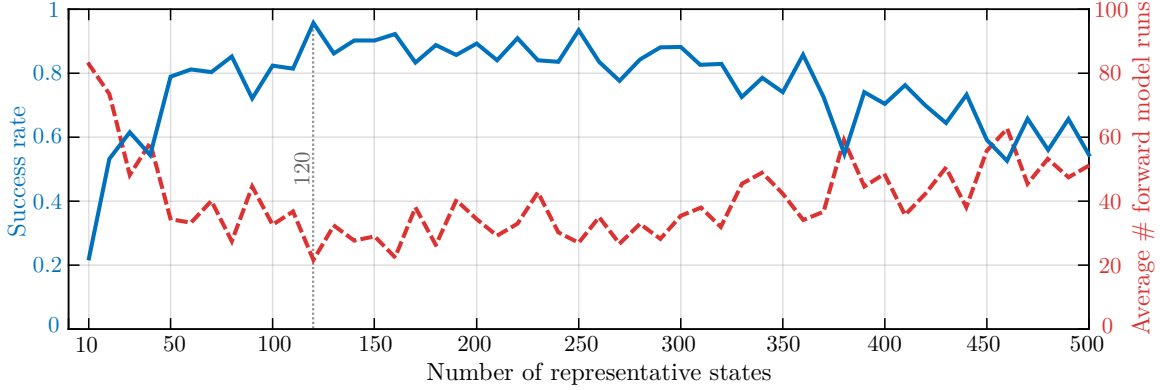


Figure 5.5: Hyper-parameter scouting: VITO’s performance for varying number of representative states n_S . The plot shows aggregated results over five consecutive personalization runs on the eight scouting datasets, each run based on different training data.

as compromise between maintaining a low number of states and sufficient state granularity. An example quantization with $n_S = 120$ is visualized in Fig. 5.2. The eight scouting datasets were discarded for the following experiments to avoid bias in the analysis.

Reference Methods

VITO’s results were compared to two optimization-based personalization methods using BOBYQA [Powe09]. The first approach, called SIMPLEGEP, mimicked the most basic estimation setup, where only the minimum level of model and problem knowledge were assumed. The objective function was the sum of absolute Δ_{QRS} and \angle_{EA} errors: $\sum_{i=1}^{n_c} |c_i|$. It was minimized in a single optimizer run where all three parameters in \mathbf{x} were tuned simultaneously. The algorithm terminated once all convergence criteria ψ were satisfied (success) or if the number of forward model evaluations exceeded 100 (failure). The second approach, CASCADEGEP, implemented an advanced estimator with strong focus on robustness, which computed the optimum parameters in a multi-step iterative fashion. In particular, Alg. 1 from Sec. 4.3.2 was adapted for Graph-EP, simply by replacing the LBM-EP forward model by Graph-EP and removing FITAPD (step 3) from the personalization procedure, as this experiment focuses on depolarization only.

To remove bias towards the choice of initial parameter values, for each of the two methods all datasets were personalized 100 times with different random initializations within the range of physiologically plausible values Ω . The differences in performance were striking: only by changing the initialization, the number of successfully personalized cases varied from 13 to 37 for SIMPLEGEP, and from 31 to 51 for CASCADEGEP (variability of more than 25% of the total number of patients). These results highlight the non-convexity of the cost function to minimize.

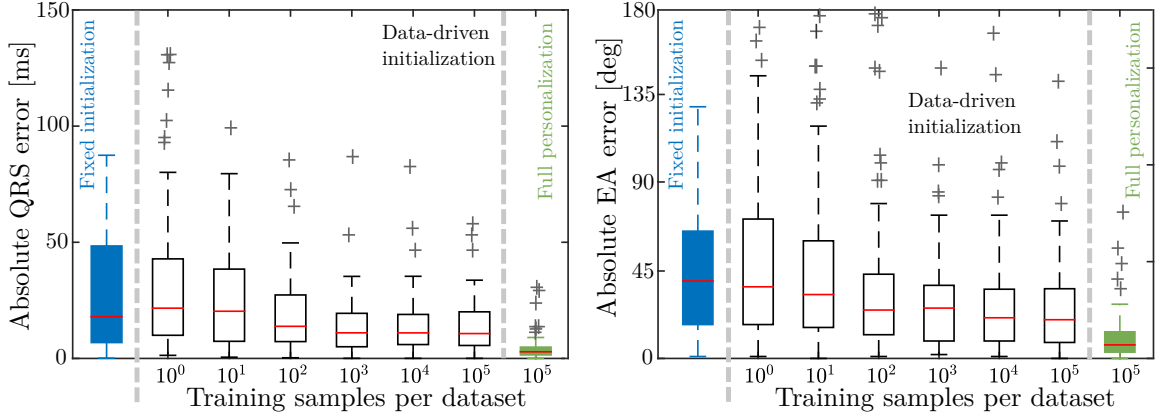


Figure 5.6: Absolute errors over all patients after initialization with fixed parameter values (blue), after data-driven initialization for increasing amount of training data (white), and after full personalization with VITO (green). Data-driven initialization yielded significantly reduced errors if sufficient training data were available ($> 10^2$) compared to initialization with fixed values. Full personalization further reduced the errors significantly. The red bar and the box edges indicate the median absolute error, and the 25 and 75 percentiles, respectively. Left: Δ_{QRS} errors. Right: \angle_{EA} errors.

Full Personalization Performance

First, VITO’s overall performance was evaluated. The full personalization pipeline consisting in off-line learning, initialization, and on-line personalization was run on all patients with leave-one-patient-out cross-validation using 1 000 training episodes ($n_{\text{samples}} = 1\,000 \cdot n_{e\text{-steps}} = 10^5$ transition samples) per patient. The maximum number of iterations was set to 100. The green box plots in the two panels of Fig. 5.6 summarize the results. The mean absolute errors were 4.1 ± 5.6 ms and $12.4 \pm 13.3^\circ$ in terms of Δ_{QRS} and \angle_{EA} , respectively, a significant improvement over the residual error after initialization. In comparison to the reference methods, the best SIMPLEGEP run yielded absolute errors of 4.4 ± 10.8 ms Δ_{QRS} and $15.5 \pm 18.6^\circ$ \angle_{EA} on average, and the best CASCADEGEP run 0.1 ± 0.2 ms Δ_{QRS} and $11.2 \pm 15.8^\circ$ \angle_{EA} , respectively. Thus, in terms of \angle_{EA} error all three methods yielded comparable performance, and while SIMPLEGEP and VITO performed similarly in terms of Δ_{QRS} CASCADEGEP outperformed both in this regard. However, considering success rates, i. e., successfully personalized patients according to the defined convergence criteria (ψ) divided by total number of patients, both the performance of VITO (67%) and CASCADEGEP (68%) were equivalent, while SIMPLEGEP reached only 49% or less. In terms of run-time, i. e., average number of forward model runs until convergence, VITO (31.8) almost reached the high efficiency of SIMPLEGEP (best: 20.1 iterations) and clearly outperformed CASCADEGEP (best: 86.6 iterations), which means VITO was $\approx 2.5\times$ faster.

Residual Error after Initialization

A major advantage over standard methods such as the two BOBYQA approaches is VITO’s automated, data-driven initialization method (Sec. 5.2.6), which eliminates the need for user-provided initial parameter values. To evaluate the utility of this step, personalization using VITO was stopped directly after initialization (the most likely \mathbf{x}_0 was used) and the errors in terms of Δ_{QRS} and \angle_{EA} resulting from a forward model run f with the computed initial parameter values were quantified. This experiment was repeated for increasing number of transition samples per dataset: $n_{\text{samples}} = 10^0 \dots 10^5$, and the results were compared to the error after initialization when fixed initial values were used (the initialization of the best performing BOBYQA experiment was used). As one can see from Fig. 5.6, with increasing amount of training data both errors decreased notably. As few as 10^2 transitions per dataset already provided more accurate initialization than the best tested fixed initial values. Thus, not only does this procedure simplify the setup of VITO for new problems (no user-defined initialization needed), this experiment showed that it can reduce initial errors by a large margin, even when only few training transitions were available. It should be noted that VITO further improves the model fit in its normal operating mode (continue personalization after initialization), as shown in the previous experiment.

Convergence Analysis

An important question in any RL application relates to the amount of training needed until convergence of the artificial agent’s behavior. For VITO in particular, this translates to the amount of transition samples required to accurately estimate the MDP transition function \mathcal{T} to compute a solid policy on the one hand, and to have enough training data for reliable parameter initialization on the other hand. To this end, VITO’s overall performance (off-line learning, initialization, personalization) was evaluated for varying number of training transition samples per dataset. As one can see from the results in Fig. 5.7, with increasing amount of training data the performance increased, suggesting that the learning process was working properly. Even with relatively limited training data of only $n_{\text{samples}} = 10^2$ samples per patient, VITO outperformed the best version of SIMPLEGEP (49% success rate). Starting from $n_{\text{samples}} \approx 3000$, a plateau at $\approx 66\%$ success rate was reached, which remained approximately constant until the maximum tested number of samples. This was almost on par with the top CASCADEGEP performance (68% success rate). Also the run-time performance increased with more training data. For instance, VITO’s average number of iterations was 36.2 at 10^3 samples, 31.5 at 10^4 samples, or 31.8 at 10^5 samples.

These results suggested that not only VITO can achieve similar performance as an advanced, manually engineered method, but also the number of required training samples was not excessive. In fact, a rather limited and thus well manageable amount of data, which can be computed in a reasonable time-frame, sufficed.




Figure 5.7: EP model personalization results: personalization success rate in blue and average number of iterations in red. Left: VITO’s performance for increasing number of training transition samples per dataset. Each dot represents results from one experiment (cross-validated personalization of all 75 datasets), solid/dashed line is low-pass filtered mean, shaded areas represent $0.5 \times \text{SD}$ and $1 \times \text{SD}$ (standard deviation). Right: Performance of both reference methods. Each shade represents 10% of the results, sorted by performance.

5.3.3 Personalization of a Whole-Body Circulation Model

Next, VITO was asked to personalize a lumped-parameter **WBC** model from pressure catheterization (cf. Sec. 2.2) and volume data. A subset of $n_{\text{datasets}} = 56$ patients from the previous experiments were used for experimentation. The discrepancy was due to missing catheterization data for some patients, which was required for **WBC** personalization only.

Forward Model Description

The **WBC** model to personalize was the one proposed in [Itu 14]. It contained a heart model (consisting of left ventricle (LV) and left atrium (LA), right ventricle (RV) and right atrium (RA), and valves), the systemic circulation (arteries, capillaries, veins) and the pulmonary circulation (arteries, capillaries, veins). Time-varying elastance models were used for all four chambers of the heart. The valves were modeled through a resistance and an inertance. A three-element Windkessel (WK) model was used for the systemic and pulmonary arterial circulation, while a two-element WK model was used for the systemic and pulmonary venous circulation. The full model and its components are described in detail in [Itu 14, Neum 15b]. Personalization was performed with respect to the patient’s heart rate as measured during catheterization.

The goal of this experiment was to compare VITO’s personalization performance for the systemic part of the model in setups with increasing number of parameters to tune and objectives to match. To this end, VITO was employed on setups with two to six parameters (2p, 3p, 5p, 6p): initial blood volume, LV maximum elastance, time until maximum elastance is reached, total aortic resistance and compliance, and LV dead volume. The reference values δ to define VITO’s allowed actions \mathcal{A} were set to 0.5% of the admissible parameter range Ω for each individual parameter, see Tab. 5.1 for details. The personalization objectives were LV MRI-derived end-diastolic volume (EDV), LV end-systolic volume (ESV), ejection time (time duration during which

Table 5.1: WBC parameters \boldsymbol{x} , their default values and domain Ω . The last column denotes the experiment setups in which a parameter was personalized (e. g., “5”: parameter was among the estimated parameters in “WBC 5p” experiment). Default values were used in experiments where the respective parameters were not personalized.

\boldsymbol{x}	Default value	Ω	Setups
Initial volume	400 ml	[200; 1000] ml	6, 5, 3, 2
LV max. elastance	2.4 mmHg/ml	[0.2; 5] mmHg/ml	6, 5, 3, 2
Aortic resistance	1100 g / (cm ⁴ s)	[500; 2500] g / (cm ⁴ s)	6, 5, 3
Aortic compliance	1.4 · 10 ⁹ cm ⁴ s ² /g	[0.5; 6] · 10 ⁹ cm ⁴ s ² /g	6, 5
LV dead volume	10 ml	[−50; 500] ml	6, 5
Time to E_{\max}	300 ms	[100; 600] ms	6

Table 5.2: WBC objectives \boldsymbol{c} , their convergence criteria $\boldsymbol{\psi}$ and range of measured values in the patient population used for experimentation.

\boldsymbol{c}	$\boldsymbol{\psi}$	Measured range	Setups
End-diastolic LV volume	20 ml	[129; 647] ml	6, 5, 3, 2
End-systolic LV volume	20 ml	[63; 529] ml	6, 5, 3, 2
Mean aortic pressure	10 mmHg	[68; 121] mmHg	6, 5, 3
Peak-systolic aortic pressure	10 mmHg	[83; 182] mmHg	6, 5
End-diastolic aortic pressure	10 mmHg	[48; 99] mmHg	6, 5
Ejection time	50 ms	[115; 514] ms	6

the aortic valve is open and blood is ejected), and peak-systolic, end-diastolic, and mean aortic blood pressures as measured during cardiac catheterization, see Fig. 5.8. To account for measurement noise, personalization was considered successful if the misfits per objective were below acceptable threshold values $\boldsymbol{\psi}$ as listed in Tab. 5.2.

Number of Representative States

Along the same lines as Sec. 5.3.2, the hyper-parameter for state space quantization was tuned based on the eight scouting patients. The larger the dimensionality of the state space, the more representative states were needed to yield good performance. In particular, for the different WBC setups, the numbers of representative states (n_S) yielding the best scouting performance were 70, 150, 400 and 600 for the 2p, 3p, 5p and 6p setup, respectively. The scouting datasets were discarded for the following experiments.

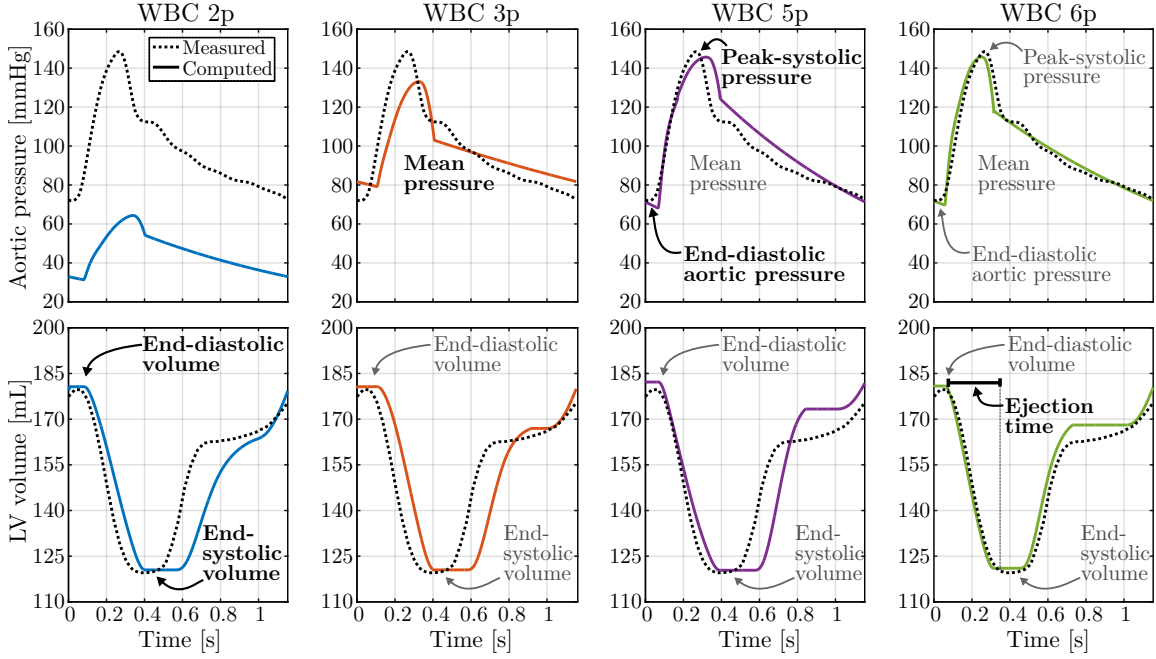


Figure 5.8: Goodness of fit in terms of time-varying *LV* volume and aortic pressure for VITO personalizing an example patient based on the different *WBC* setups. The added objectives per setup are highlighted in the respective column. With increasing number of parameters and objectives VITO manages to improve the fit between model and measurements.

Reference Method

A gradient-free optimizer [Laga 98] based on the simplex method was used to benchmark VITO. The objective function was the sum of squared differences between computed and measured values, weighted by the inverse of the convergence criteria to respect the different ranges of objective values: $\|\mathbf{c}\|_{\psi}$, cf. Eq. (5.7). Personalization was terminated once all convergence criteria were satisfied (success), or when the maximum number of iterations was reached (failure). To account for the increasing complexity of optimization with increasing number of parameters n_x , the maximum number of iterations was set to $50 \cdot n_x$ for the different setups.

As one can see from Fig. 5.9, right panels, with increasing number of parameters to be estimated, the performance in terms of success rate and number of forward model runs decreased slightly. This is expected as the problem becomes harder. To suppress bias originating from (potentially poor) initialization, the reference method was run 100 times per setup (as in EP experiments), each time with a different, randomly generated set of initial parameter values. The individual performances varied significantly for all setups.

Convergence Analysis

For each *WBC* setup the full VITO personalization pipeline was evaluated for increasing training data ($n_{\text{samples}} = 10^0 \dots 10^5$) using leave-one-patient-out cross-validation. The same iteration limits as for the reference method were used. The results are

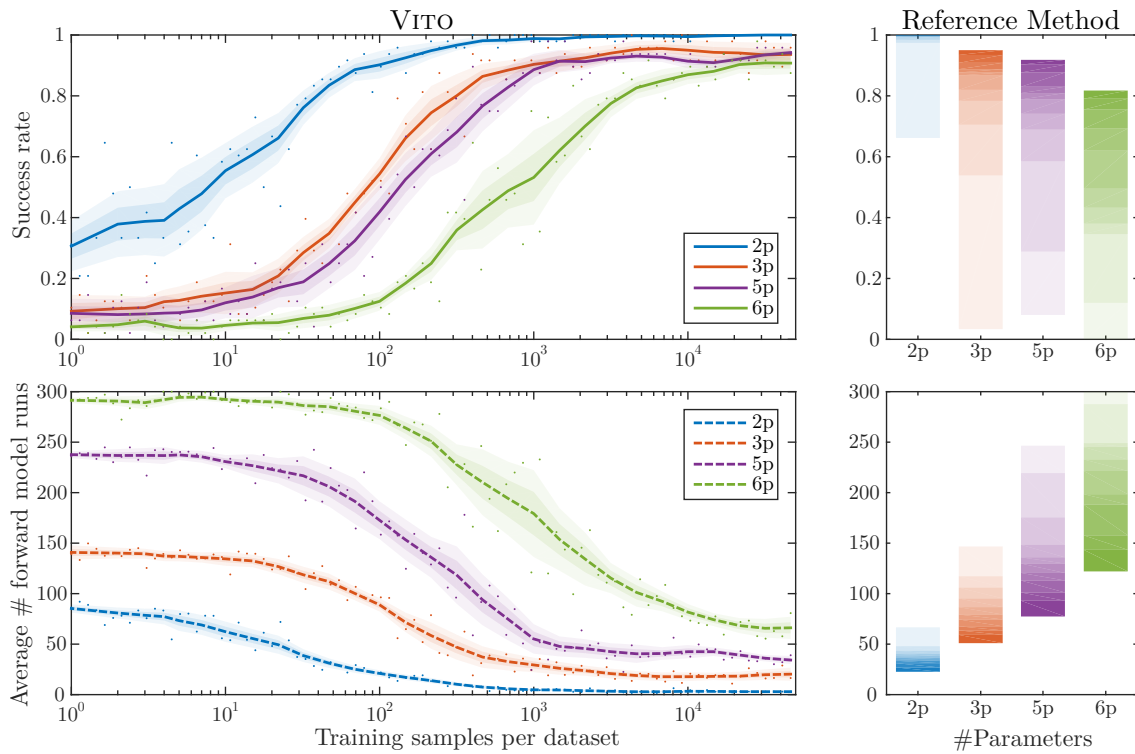


Figure 5.9: WBC model personalization results (top: success rate, bottom: average number of forward model runs until convergence) for various estimation setups (different colors), see text for details. Left: VITO’s performance for increasing number of training transition samples per dataset. Each dot represents results from one experiment (cross-validated personalization of all 48 datasets), solid/dashed lines are low-pass filtered mean, shaded areas represent $0.5 \times \text{SD}$ and $1 \times \text{SD}$. Right: Performance of reference method. Each shade represents 10% of the results, sorted by performance; darkest shade: best 10%.

presented in Fig. 5.9, left panels. With increasing data, VITO’s performance, both in terms of success rate and run-time (iterations until convergence), increased steadily until reaching a plateau. As one would expect, the more complex the problem, i. e., the more parameters and objectives involved in the personalization, the more training data was needed to reach the same level of performance. For instance, VITO reached 80% success rate with less than $n_{\text{samples}} = 50$ training samples per dataset in the 2p setup, whereas almost $90\times$ as many samples were required to achieve the same performance in the 6p setup.

Compared to the reference method, given enough training data, VITO reached equivalent or better success rates (e. g., up to 11% higher success rate for 6p) while significantly outperforming the reference method in terms of run-time performance. In the most basic setup (2p), if $n_{\text{samples}} \geq 10^3$, VITO converged after 3.0 iterations on average, while the best reference method run required 22.6 iterations on average, i. e., VITO was seven times faster. For the more complex setups (3p, 5p, 6p), the speed-up was not as drastic. Yet, in all cases VITO outperformed even the best run of the reference method by a factor of 1.8 or larger.

5.4 Summary

In this chapter, a novel personalization approach called VITO has been presented. To our knowledge, it is the first time that cardiac model personalization is addressed using [artificial intelligence \(AI\)](#) concepts. Inspired by how humans approach the personalization problem, VITO first learns the characteristics of the computational model under consideration using a data-driven approach. This knowledge is then utilized to learn how to personalize the model using [reinforcement learning \(RL\)](#). VITO is generic in the sense that it requires only minimal and intuitive user input (parameter ranges, authorized actions, number of representative states) to learn by itself how to personalize a model.

VITO was applied to a synthetic scenario and to two challenging personalization tasks in cardiac computational modeling. The problem setups and hyper-parameter configurations are listed in Tab. 5.3. In most setups the majority of hyper-parameters were identical and only few ($n_{\mathcal{S}}$) required manual tuning, suggesting good generalization properties of VITO. Another key result was that VITO was up to 11% more robust (higher success rates) compared to standard personalization methods. VITO’s ability to generalize the knowledge obtained from a set of training patients to personalize unseen patients was shown as all experiments were based on cross-validation. Furthermore, VITO’s robustness against training patients for whom we could not find a solution was tested. In particular, for about 20% of the patients, in none of the electrophysiology experiments in Sec. 5.3.2 any personalization (neither VITO nor the reference methods) could produce a result that satisfied all convergence criteria. Hence, for some patients no solution may exist under the given electrophysiology model configuration, possibly due to invalid assumptions of Graph-EP for patients with complex pathologies. Still, all patients were used to train VITO, and surprisingly VITO was able to achieve equivalent success rates as the manually engineered personalization approach for cardiac [EP](#).

Table 5.3: Applications considered in this chapter described in terms of the number of parameters (n_x), objectives (n_c) and datasets (n_{datasets}) used for experimentation (in brackets: excluding scouting patients, if applicable); and VITO’s hyper-parameters: the number of representative MDP states (n_S) and the number of actions per parameter (n_A/n_x). The last column (n_{plateau}) denotes the approximate number of samples needed to reach the performance “plateau” (see convergence analyses in Sec. 5.3.2 and Sec. 5.3.3).

Application	n_x	n_c	n_{datasets}	n_S	n_A/n_x	n_{plateau}
Rosenbrock	2	2	100	100	6	n/a
Rosenbrock ext.	2	2	100	100	8	n/a
EP	3	2	83 (75)	120	6	3 000
WBC 2p	2	2	56 (48)	70	6	450
WBC 3p	3	3	56 (48)	150	6	2 000
WBC 5p	5	5	56 (48)	400	6	3 500
WBC 6p	6	6	56 (48)	600	6	20 000

Generating training data could be considered VITO’s computational bottleneck. However, on the one hand, training is performed off-line and one-time only, and on the other hand, it is independent for each training episode and each patient. Therefore, large compute clusters could be employed to perform rapid training by parallelizing this phase. On-line personalization, on the contrary, is not parallelizable in its current form: the parameters for each forward model run depend on the outcome of the previous iteration. Since the forward computations are the same for every “standard” personalization method (not including surrogate-based approaches), the number of forward model runs until convergence was used for benchmarking: VITO was up to seven times faster compared to the reference methods. The on-line overhead introduced by VITO (convert data into an MDP state, then query policy) is negligible.

As such, VITO could become a unified framework for personalization of any computational physiological model, potentially eliminating the need for an expert operator with in-depth knowledge to design and engineer complex optimization procedures.

Uncertainty Quantification in Personalized Modeling

6.1	Uncertainty Quantification	87
6.2	Robust Estimation of Model Parameters and Their Uncertainty	88
6.3	Experiments and Results	93
6.4	Summary	96

The heart is a highly complex system, computational heart models are based on assumptions, and the data acquired to measure the heart and to drive personalization are sparse and noisy. Still, for the vast majority of cardiac computational model personalization approaches, including those described in the previous chapters, personalization equals estimation of a single, “most likely” solution (parameter estimate). Considering the many sources of uncertainties involved, the clinical value of a single estimate can, however, be questioned, if no information about the shape of the solution space is offered: solution uniqueness cannot be guaranteed.

This chapter is based on [Neum 14b]. First, sources of uncertainties in the context of cardiac modeling and personalization are listed and the **uncertainty quantification (UQ)** problem is introduced in Sec. 6.1. Then the proposed approach that shows how personalization with inverse UQ could be tackled is described in Sec. 6.2 and evaluated in Sec. 6.3. Finally, the main ideas and results are summarized in Sec. 6.4.

6.1 Uncertainty Quantification

Uncertainty quantification (UQ) encompasses a set of methods that aim to characterize and quantify uncertainties. A comprehensive overview of the various sources of uncertainties that are relevant for computational modeling applications can be found in [Kenn 01, Mira 16]. These include:

- **Parameter uncertainty** due to model input parameters whose exact values are not known.
- **Experimental uncertainty** due to the variability of experimental measurements (e. g., noise).
- **Structural uncertainty** due to modeling assumptions.
- **Algorithmic uncertainty** due to numerical errors and approximations.

Two main types of UQ problems exist: forward UQ and inverse UQ (also called parameter UQ). Just like forward modeling computes deterministic model output given a set of deterministic model parameters, forward UQ computes the variability in model output due to uncertain model parameters. An example for forward UQ is the following. Consider a patient-specific computational cardiac model that computes ejection fraction (EF) as output based on a number of input parameters, some of which are known as they can be measured directly, some are personalized from measurements, and others are not identifiable. For the latter, default values are used as the model requires all parameters to be set in order to evaluate its mathematical equations to be able to compute the EF value. If instead of default values one would provide a distribution of input values, e. g., from population studies, and propagate this uncertainty through the model to its output, then the output would change from a single EF value to a distribution, which can then be used to analyze the likelihood of certain EF values and their variability due to the uncertain input. Such information can be decisive if e. g., the output is used in guidelines for clinical decision making. The more difficult problem is inverse UQ, where uncertainties in the measurements used to personalize certain parameters of the model need to be propagated backwards through the model parameter estimation process in order to compute a distribution of input parameters that may have lead to that observation.

6.2 Robust Estimation of Model Parameters and Their Uncertainty

Uncertainty in clinical measurements and model assumptions are known to increase the non-identifiability of parameters. In particular, solution uniqueness is not guaranteed and multiple solutions or entire manifolds of solutions with equal level of confidence may exist. As a result, the clinical value of one single estimate can be questioned. Yet, to the best of our knowledge, only little work exists in the cardiac modeling community which addresses these challenges. Parameter UQ due to noise in the data has been investigated in [Wall 14, Konu 11] for the cardiac electrophysiology (EP) problem. However, the authors do not take into account this knowledge to estimate a robust optimum for the model parameters under consideration as their main focus was on UQ. In [Konu 11], the mean of the posterior density was selected as parameter value, which can become inaccurate for skewed or multi-modal distributions. In [Wall 14], the maximum posterior was used and evaluated on synthetic data only. It is not clear though whether that choice would be robust under uncertain data noise level.

6.2.1 Overview of the Method

This section presents a stochastic method for the robust estimation of biomechanical parameters of the myocardium and their uncertainty due to noisy data. A general overview of the main steps of the method with inputs and outputs is shown in Fig. 6.1. First, a fast surrogate model of the image-based cardiac electromechanics (EM) model of the heart from Chap. 3 is estimated by using polynomial chaos ex-

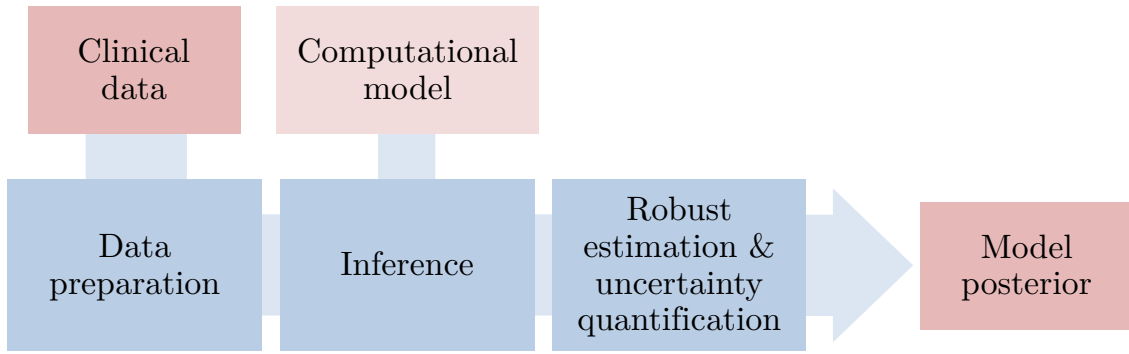


Figure 6.1: Proposed robust estimation framework. From imaging and clinical data, the posterior of the biomechanical computational model parameters, i. e., their most likely value and their uncertainty, are inferred through robust estimation and UQ. See text for details.

pansion (PCE). The surrogate model is then used in a Bayesian inference framework to estimate posterior probabilities of model parameters. Then, mean-shift is applied on the posteriors to find the optimal parameter value by integrating the space of measurement uncertainties. Experiments on eight dilated cardiomyopathy (DCM) patients from the cohort described in Sec. 4.4.1 showed that the proposed approach yielded goodness of fit equivalent to a deterministic method, while being computationally efficient and providing additional information such as confidence intervals (CIs). More importantly, we show that the manifold of possible solutions is much larger, including multi-modal posteriors, which are automatically identified, therefore enabling quantitative assessment of the clinical utility of estimated parameters.

For the purpose of presenting the stochastic personalization and UQ approach, the computational model to be personalized is reformulated as a statistical problem, denoted $f(\mathbf{x}) = \mathbf{y}$, where \mathbf{x} and \mathbf{y} are random input (model parameters) and output (model responses) variables. The approach is designed to be suitable even for the highly complex and computationally demanding EM model from Sec. 3.3.

6.2.2 Bayesian Inference

Inverse UQ is solved within the framework of Bayesian calibration [Kenn 01]: the parameter values \mathbf{x} are inferred and their uncertainty due to noisy measurements \mathbf{z} is quantified. The goal is to compute the probability density function (PDF) $p(\mathbf{x}|\mathbf{z})$, i. e., the *posterior* of the model parameters given the patient’s measurements \mathbf{z} . This is achieved by forward propagation of uncertainty [Marz 07] through the model.

Following Bayes’ rule:

$$p(\mathbf{x}|\mathbf{z}) \propto p(\mathbf{z}|\mathbf{x}) \cdot p(\mathbf{x}) . \quad (6.1)$$

The first term on the right-hand side, $p(\mathbf{z}|\mathbf{x})$, is called *likelihood*. It describes how well each set of parameters is supported by the measured data. As in [Wall 14], the

likelihood is approximated by expressing it in terms of the misfit between responses and measurements:

$$p(\mathbf{z}|\mathbf{x}) = \exp\left(-\frac{1}{2}\varepsilon(\mathbf{y}, \mathbf{z})^\top \mathbf{S}^{-1}\varepsilon(\mathbf{y}, \mathbf{z})\right), \quad (6.2)$$

where $\varepsilon(\mathbf{y}, \mathbf{z}) = \mathbf{y} - \mathbf{z}$. The misfit is modeled as a normal distribution $\mathcal{N}(\mathbf{0}, \mathbf{S})$ with zero mean and covariance matrix \mathbf{S} . In this work, all sources of error are aggregated under ε . The second term in Eq. (6.1), $p(\mathbf{x})$, is the *prior*:

$$p(\mathbf{x}) \sim \mathcal{U}_{\mathbf{x}}. \quad (6.3)$$

It represents the knowledge on the parameters independently from the measurements, and is modeled in this work using a uniform distribution $\mathcal{U}_{\mathbf{x}}$.

Now that the posterior $p(\mathbf{x}|\mathbf{z})$ is defined analytically, i. e., it can be evaluated for any given set of input parameters, the next step is to estimate its density. To this end, a **Markov chain Monte Carlo (MCMC)** method is used. In brief, **MCMC** samples from the statistical distribution of input parameters \mathbf{x} to iteratively build up the distribution of the output (the model posterior $p(\mathbf{x}|\mathbf{z})$). The problem with **MCMC** is that it requires a large number of evaluations of the likelihood to estimate even simple properties of the distribution, let alone its entire shape. Because evaluating the likelihood (cf. Eq. (6.2)) requires the model output for the given parameter, $\mathbf{y} = f(\mathbf{x})$, each step involves computation of the forward model, which is by itself computationally demanding. Various extensions and improvements of **MCMC** aimed towards reducing the computational burden are available, such as “delayed rejection adaptive Metropolis” [Haar 06], which is used in this work. However, still thousands of forward model evaluations are required, thus rendering **MCMC** intractable for the models employed in this work.

6.2.3 Fast Surrogate of the Computational Model

To make estimation of the posterior density using **MCMC** sampling computationally tractable, a very fast forward model is required. As the full **EM** model is too computationally prohibitive to be used directly in **MCMC**, the involvement of a surrogate model is proposed. The surrogate model, denoted \tilde{f} , should reasonably approximate the full model f for a specified range of input parameters at vastly reduced computational cost.

To this end, an approach based on finite **polynomial chaos expansion (PCE)** is employed [Adam 13]. **PCE** provides an efficient functional mapping from model input \mathbf{x} to individual responses $y \in \mathbf{y} = f(\mathbf{x})$ through multidimensional orthogonal polynomial approximations of f . More specifically, for each individual response y , a dedicated **PCE** surrogate $\tilde{f}_y(\mathbf{x})$ is estimated. $\tilde{f}_y(\mathbf{x})$ corresponds to expressing y in terms of a linear combination of polynomials:

$$\tilde{f}_y(\mathbf{x}) = \sum_{i=0}^{n_{\text{coeff}}-1} \alpha_i \cdot \Psi_i(\mathbf{x}) \quad (6.4)$$

where the multivariate polynomial basis $\Psi_i(\mathbf{x})$ is defined as combination of univariate 1D basis Legendre polynomials, and n_{coeff} denotes the total number of coefficients.

Let n_x denote the number of parameters to estimate, the total number of coefficients is determined as

$$n_{\text{coeff}} = (n_{\text{poly}} + 1)^{n_x} , \tag{6.5}$$

where n_{poly} is the maximum polynomial order as specified by the user. The higher n_{poly} , the more fidel is the approximation, but also more coefficients need to be estimated.

The PCE coefficients α_i are obtained in a training phase using the spectral projection approach [Eldr 09]. The main idea is that model responses are projected against each basis function Ψ_i using inner products, then polynomial orthogonality properties are employed to extract the coefficients. The reader is referred to [Marz 07, Eldr 09, Adam 13] for additional details. A total of n_{coeff} forward model runs of the full model are required as “training data”.

6.2.4 Posterior Analysis under Uncertainty

From the surrogate model and the Bayesian calibration described above, the posterior PDF $p(\mathbf{x}|\mathbf{z})$ of the model parameters can be estimated under the assumption that the level of noise in the data is known. Coming back to the main objective, i. e., model personalization, the next step is to derive an estimate of the most likely model parameters, denoted \mathbf{x}^* , and their uncertainty. In particular, the main goal is to estimate parameter values that are most robust to varying level of noise in the data, because the exact noise level is often difficult to precisely identify. To that end, an algorithm to aggregate posteriors $p(\mathbf{x}|\mathbf{z})$ calculated for different levels of noise into one PDF is proposed. The most likely parameter values are then estimated from the aggregated PDF. A visual illustration for a simplified estimation setup is shown in Fig. 6.2.

Mixture Model Fitting

Let $p_i(\mathbf{x}|\mathbf{z})$ denote the i th estimated posterior, which was calculated for a certain noise level \mathbf{S}_i (cf. Eq. (6.2)) according to the method described in Sec. 6.2.2. First, the number of modes, k_i , in $p_i(\mathbf{x}|\mathbf{z})$ is estimated using the mean-shift algorithm [Coma 02]. Next, a GMM \mathcal{G}_i with k_i components is fitted to $p_i(\mathbf{x}|\mathbf{z})$ using the expectation-maximization algorithm. These steps are repeated n_S times with distinct levels of measurement noise uncertainty by varying the error variances of the individual responses in the covariance matrix $\mathbf{S} = \mathbf{S}_i$ of the likelihood, cf. Eq. (6.2). At the end of this process, a set of n_S mixture models $\mathcal{G} = \{\mathcal{G}_1 \dots \mathcal{G}_{n_S}\}$ is generated. The variability of \mathbf{S}_i and n_S are specified by the user, as these parameters are application- and data-dependent.

Extraction of Most Likely Parameters

Next, the \mathcal{G}_i are aggregated to get the final estimate. Consider one particular mixture model $\mathcal{G}_i \in \mathcal{G}$. For each of the $j = 1 \dots k_i$ means \mathbf{m}_{ij} of the components of \mathcal{G}_i , its




Figure 6.2: Illustration of Gaussian mixture model (GMM) aggregation for a highly simplified demonstrative estimation example (only one parameter x to be estimated). Left: Colored curves are mixture models \mathcal{G}_i fitted to the posterior samples of $n_S = 3$ different noise levels \mathcal{S}_i . \mathcal{G}_2 is mono-modal, while others are multi-modal. The modes m_{ij} and their weights are visualized on horizontal axis. Right: Final aggregated mixture model $\hat{\mathcal{G}}$ based on clustering of weighted modes and combination of covariance matrices per cluster. Two clusters \diamond_1 and \diamond_2 are identified, thus $\hat{\mathcal{G}}$ has two modes. Most likely parameter estimate is the mode of the cluster with highest combined weights, $\mathbf{x}^* = \hat{\mathbf{m}}_l$, visualized as red-yellow dot. Black-white dot is *alternative* solution.

“support” is computed by summation of the log-probabilities of \mathbf{m}_{ij} in each of the other \mathcal{G}_t :

$$w_{ij} = \sum_{t \neq i} \log \mathcal{G}_t(\mathbf{m}_{ij}) . \quad (6.6)$$

The supports w_{ij} are normalized such that all values are mapped between 0 and 1, the latter representing the highest support. Next, the \mathbf{m}_{ij} are separated into k^* clusters $\diamond = \diamond_1 \dots \diamond_{k^*}$ using w_{ij} -weighted k -means clustering. The number of clusters k^* is determined by voting among all k_i in \mathcal{G} . Finally, the centroid $\hat{\mathbf{m}}_l$ of the \mathbf{m}_{ij} in the cluster with the highest combined support \hat{w}_l :

$$\hat{w}_l = \sum_{ij \in \diamond_l} w_{ij} , \quad (6.7)$$

is selected as the final estimate $\mathbf{x}^* = \hat{\mathbf{m}}_l$. This way, the robustness in the estimate is increased even without explicitly knowing the level of noise in the measurements, while still being able to capture multi-modal PDFs.

Explicit Representation of Uncertainty

Uncertainty in estimated parameters is mainly described by confidence regions or intervals in this work. Let \mathbf{Y}_{ij} be the covariance matrix of the j th mixture component in the i th mixture model \mathcal{G}_i . It is assumed that in each cluster \diamond_l , all contained \mathbf{m}_{ij} with corresponding \mathbf{Y}_{ij} are distorted (through noise) manifestations of the same

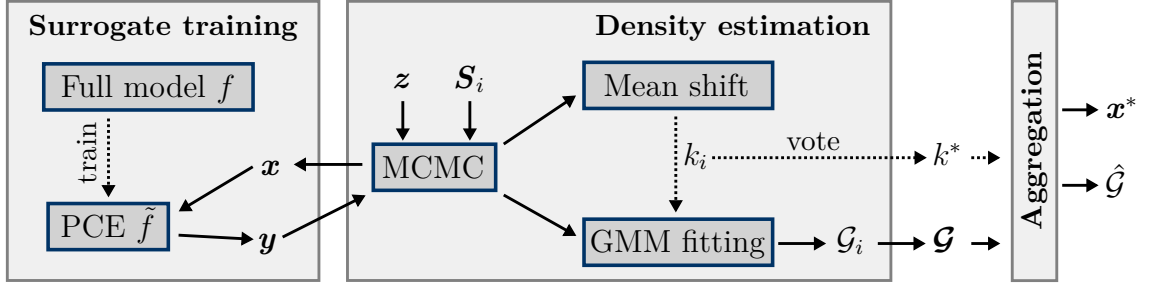


Figure 6.3: Illustration of the main methods and variables involved in computing the personalized model parameters and uncertainty. First the surrogate PCE model \tilde{f} is trained from the full model f . MCMC then samples the model posterior given the measurements \mathbf{z} , assuming a specific noise level \mathbf{S}_i . Sampling involves evaluation of the surrogate model for generated parameters \mathbf{x} to predict model outputs \mathbf{y} . Mean shift is applied on a large set of samples to identify the number of modes k_i in the posterior, which is input to a GMM fitting procedure to estimate the i th posterior density \mathcal{G}_i . This process is repeated for $i = 1, \dots, n_S$ with varying noise levels \mathbf{S}_i . All generated mixture models \mathcal{G} are then aggregated into a robust representation of uncertainty $\hat{\mathcal{G}}$ (see text for details) and the most likely parameters \mathbf{x}^* are extracted.

normal PDF, which is centered at the centroid $\hat{\mathbf{m}}_l$ with unknown covariance $\hat{\mathbf{\Upsilon}}_l^*$. $\hat{\mathbf{\Upsilon}}_l^*$ is approximated by linear combination of the covariance matrices:

$$\hat{\mathbf{\Upsilon}}_l = \hat{w}_l^{-1} \sum_{ij \in \diamond_l} w_{ij} \mathbf{\Upsilon}_{ij} . \quad (6.8)$$

Finally, all information from the k^* clusters is merged into one “robust GMM”:

$$\hat{\mathcal{G}} = \sum_{l=1}^{k^*} \hat{w}_l \cdot \mathcal{N}(\hat{\mathbf{m}}_l, \hat{\mathbf{\Upsilon}}_l) , \quad (6.9)$$

thus forming an explicit representation of uncertainty. The main steps of the entire procedure are illustrated in Fig. 6.3.

6.3 Experiments and Results

The proposed stochastic parameter estimation and inverse UQ method was evaluated for personalization of the patient-specific EM model from Sec. 3.3. The focus of the evaluation was on two parameters representing the main determinants of active and passive cardiac biomechanics. In particular, the maximum active contraction force of the left ventricle (LV), and the global stiffness factor of the passive model were selected: $\mathbf{x} = (\hat{\sigma}_{LV}, \beta)$. Heart geometry, EP, hemodynamics (HD), and remaining biomechanics parameters were assumed to be known for each patient. As this approach is still in a phase of active research and methodological improvements (see also discussions in Sec. 7.2.3), the results and conclusions of the following experiments should be considered preliminary, in particular because only a small group of eight patients was involved. However, care was taken to select heterogeneous subjects with

large variety of disease severity (EF ranging from 19–47%) among the DCM patients from the cohort described in Sec. 4.4.1.

The main objective was to investigate uniqueness and CIs of the biomechanical parameter estimates given the clinically measured data, which was affected by unknown levels of noise. In particular, the measurements (\mathbf{z}) chosen for evaluating the method were based on the time-varying LV volume and pressure curves: end-diastolic volume (EDV), end-systolic volume (ESV) and mean ventricular volume; and end-diastolic, peak-systolic and mean ventricular pressure.

The PCE surrogate models were computed with maximum polynomial order $n_{\text{poly}} = 4$, if not stated otherwise. The number of required forward model runs of the full model to train the PCE depends on the maximum polynomial order, cf. Eq. (6.5). In total, the full model needed to be computed 25 times per patient on a 5×5 isotropic grid spanned by physiological parameter ranges of $\hat{\sigma}_{\text{LV}} \in [50; 300]$ kPa and $\beta \in [0.2, 1.5]$.

The mean-shift based posterior analysis described in Sec. 6.2.4 was carried out using $n_s = 15$ random noise levels. Each intermediate GMM (\mathcal{G}_i , one per noise level) was estimated based on 50 000 MCMC samples. The response-specific level of noise was modeled as the standard deviation (SD) of the assumed measurement error, individually drawn from a uniform distribution $\text{SD}_y \sim \mathcal{U}(l_y, u_y)$. Lower and upper bounds $l_y = 1 \text{ units}$ and $u_y = 3 \text{ units}$, where *units* denotes ml or kPa for volume and pressure responses, respectively, were chosen heuristically to model plausible levels of data noise. The error covariance matrix \mathbf{S} (cf. Eq. (6.2)) is defined as diagonal matrix: $\mathbf{S} = \text{diag}(\text{SD}_{y_1}^2, \text{SD}_{y_2}^2, \dots)$.

6.3.1 Precision of PCE-based Surrogate Model

The goal of the first experiment was to verify that the selected PCE maximum polynomial order $n_{\text{poly}} = 4$ was enough to model the posterior of the forward model. To this end, the responses computed by the PCE model of order 4, \tilde{f} , were compared with those obtained using a high-fidelity PCE model, \tilde{f}_{10} , of order $n_{\text{poly}} = 10$. Note that following the formula in Eq. (6.5), calculation of the high-fidelity model \tilde{f}_{10} required almost five times (121) as many forward runs of the full model compared to \tilde{f} (25).

A total of 1 000 parameters were sampled randomly per patient and all model responses were calculated using both \tilde{f} and \tilde{f}_{10} . Averaged over all patients, the errors over volume and pressure responses were 1.6 ± 1.1 ml and 0.6 ± 0.4 kPa, respectively, two orders of magnitude less than their absolute values (cf. Tab. 2.1). Furthermore, posterior PDFs estimated from the two surrogates compared qualitatively well. These results suggested that \tilde{f} was enough to reliably estimate the posterior of the model.

6.3.2 Estimation of Biomechanical Parameters

Next, the biomechanical model parameters $\mathbf{x} = (\hat{\sigma}_{\text{LV}}, \beta)$ were estimated for all cases using the proposed stochastic approach. The estimated parameters yielded high goodness of fit (errors were below clinical variability), with a mean EF error of $2.3 \pm 1.3\%$ and stroke volume (SV) error of 8.6 ± 3.6 ml, respectively. In Fig. 6.4, the

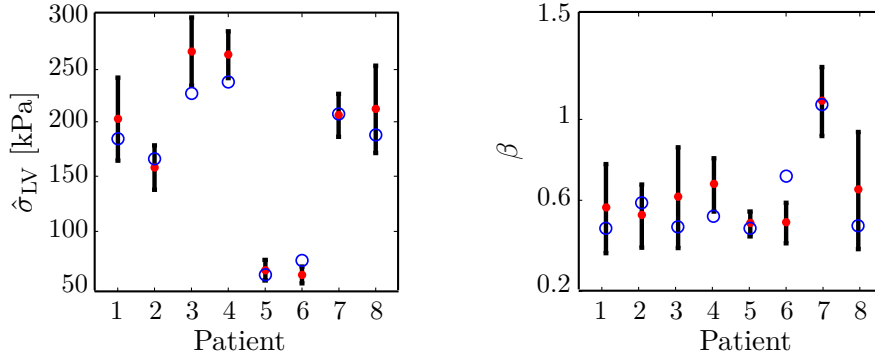


Figure 6.4: Visualization of 95% confidence interval (CI) of GMM mode with highest combined support estimated by the proposed stochastic approach. The red dots and blue circles represent most likely parameter values computed by the stochastic approach and by the BOBYQA reference method, respectively. All BOBYQA results lie within (or near) the CIs.

values of the estimated parameters and their 95% CI extracted from the main modes of the model posteriors are visualized. As one can see, the noise in the data can have significant impact on the confidence of the estimated values, with in some cases intervals of up to ± 40 kPa and 0.3 for $\hat{\sigma}$ and β , respectively. Such estimates could be employed as indicators of model fitting quality to suggest to the operator to acquire more or higher-quality data in order to reduce these uncertainties. The method can further capture multi-model PDFs as shown in Fig. 6.5, left panel.

6.3.3 Comparison to Deterministic Personalization

Next, the personalization results of the proposed method were compared to those obtained using BOBYQA optimization [Powe 09]. The employed objective function writes $\zeta(\mathbf{y}) = \|\mathbf{y} - \mathbf{z}\|_2^2$. The BOBYQA-based EF errors after personalization of the same patients were approximately equivalent to those resulting from the stochastic method: on average, EF and SV errors were $2.0 \pm 0.9\%$, and 7.6 ± 2.6 ml, respectively.

In terms of run-time, BOBYQA required 13.3 ± 1.9 evaluations of the full EM forward model to converge to the final solution. In comparison, the proposed stochastic approach always required exactly 25 forward runs in order to estimate the surrogate model as described earlier. The stochastic method is therefore slightly less efficient, but it provides valuable additional outputs in terms of confidence intervals and other properties of the model posterior.

6.3.4 Empirical Evaluation of Estimated Uncertainty

Finally, in an empirical experiment on one dataset with complex posterior, the “meaningfulness” of the computed uncertainty was investigated by analyzing how the goodness of fit of the model responses \mathbf{z} varied between regions of high confidence and regions of low confidence of the posterior. To this end, the full forward model was computed based on several sets of parameters \mathbf{x} along two one-dimensional sub-spaces of the full parameters space as shown in Fig. 6.5. As one can see from the error profile

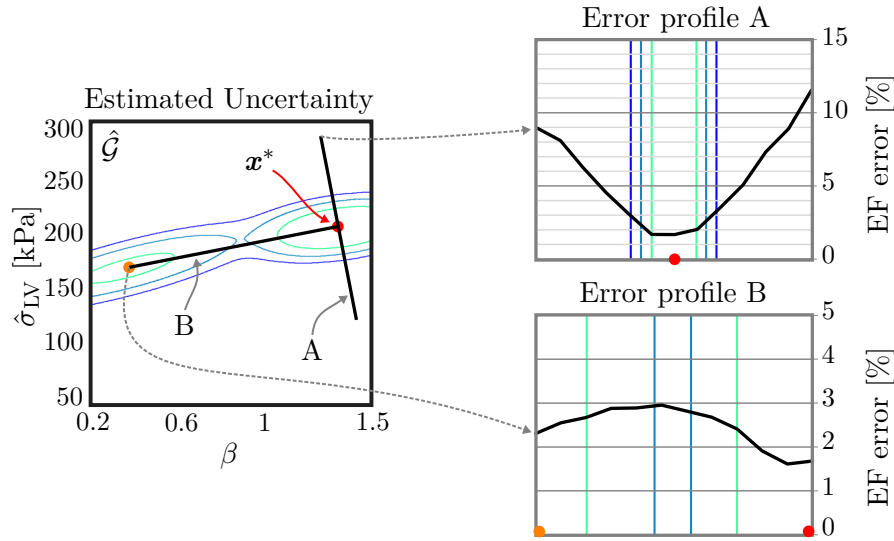


Figure 6.5: Empirical evaluation of estimated uncertainty. Left: Automatically estimated robust model posterior $\hat{\mathcal{G}}$ visualized by 80%, 95% and 99% confidence regions. The overlaid black lines depict the one-dimensional subspaces A and B that were analyzed in this experiment (see text for details). The red and orange dots depict the most likely solution and an “alternative solution” as computed by the stochastic algorithm, respectively. Middle and right: LV EF error profiles of A and B, overlaid by vertical lines representing the intersections of the 1D subspaces with the 80% (light green), 95% and 99% (blue) confidence regions of the estimated posterior.

“A”, the absolute errors in terms of EF are significantly lower for parameters within the marked high confidence regions, while outside the 99% confidence region (blue contours), the errors increase suddenly and drastically. The further away from these regions, the larger the errors became. The same trends were observed for other errors such as SV and systolic and diastolic ventricular pressures. In the second profile “B”, the underlying one-dimensional sub-space connects the two modes of the bi-modal posterior $\hat{\mathcal{G}}$. All points on the resulting line are inside regions of high confidence. The obtained errors along this line are consistently low, while the two modes (solution and alternative solution) mark points that are approximately local minima of the error profile. These empirical findings suggest that the proposed aggregation method is meaningful. More comprehensive and theoretical evaluation will be needed in the future.

6.4 Summary

In this chapter, the impact of data noise on the estimated biomechanical parameters was explored. The proposed parameter uncertainty quantification (UQ) framework relies on stochastic parameter estimation and aggregates the probabilities estimated under different noise levels to derive a robust parameter estimate without explicitly knowing the level of noise in the data. Experimental results on real clinical data showed that not only the approach is as effective as well-established deterministic, optimization-based methods, but that it is also computationally efficient while provid-

ing rich information regarding the most likely parameter and measures of confidence and uncertainty. Furthermore, using this approach, the non-uniqueness of the inverse problem of cardiac biomechanics could be demonstrated by reporting different solution spaces, which can be automatically identified through the estimated posterior PDFs. Such an approach could therefore provide precious insights when analyzing the clinical relevance of estimated parameters and personalized model predictions. In addition to that, it could constitute a criterion to select and refine data acquisition protocols used for model personalization.

PART III

Conclusion

Summary and Perspectives

7.1 Summary	101
7.2 Perspectives	106

What if we could create a virtual heart of a patient? We could better *describe* the state of the patient, better *diagnose* the patient's condition, *predict* therapy outcome and proactively *prescribe* optimal interventions. The work performed during the course of this thesis tackles some of the most challenging aspects in creating virtual hearts and in their translation into clinical practice: from streamlined personalization of a multi-physics computational **electromechanics (EM)** heart model, to intelligent agents aiding the design of parameter estimation methods, to a stochastic framework for robust parameter **uncertainty quantification (UQ)**. A comprehensive summary in Sec. 7.1 and an outlook for future work in Sec. 7.2 conclude this thesis.

7.1 Summary

Dilated cardiomyopathy (DCM) is one of the most common causes of heart failure (HF) with significant mortality and morbidity rates (Chap. 1). However, clinical management of these patients is challenged by the wide variety of disease causes and therapies. There is therefore an important need for new tools to better stratify patients. Computational models of heart function are being explored to improve patient stratification, risk prediction and therapy planning. They can not only be used to estimate hidden information from clinical data, but also to predict therapy outcome and disease course due to their generative nature. However, in order to use such models in clinical practice, they need to be precisely personalized to the patient. The multi-physics 3D computational whole-heart model that has been used and enhanced in the course of this thesis aims to capture the most relevant aspects of real cardiac physiology, which are briefly introduced in Chap. 2 together with the most common tools and devices to measure heart morphology and function.

7.1.1 Patient-specific Virtual Hearts

The model itself, which was developed by a team of researchers at Siemens Healthineers, is then described in Chap. 3, which marks the beginning of the first main part of this thesis, which concerns the path from routine clinical data to patient-specific virtual hearts. The whole-heart model consists of a set of coupled components, which can be computed independently, but interact with each other. Each component models

an integral part of the heart's morphology or physiology: anatomy, electrophysiology (EP), biomechanics or boundary conditions.

The first component, anatomical model, captures the geometric and structural properties of the heart. To this end, bi-ventricular heart anatomy is efficiently segmented under expert guidance from magnetic resonance imaging (MRI) using machine learning (ML) techniques and converted to a volumetric mesh representation. A rule-based model of myocardial fiber architecture is fitted to the patient geometry, and the mesh is automatically tagged with certain spatially-varying information required by the other components.

Next, cardiac EP models the propagation of the electrical signal in the myocardium. Two interchangeable EP models are available in the framework. The first one, LBM-EP, is based on the Mitchell-Schaeffer cellular model, and the second one, Graph-EP, is a highly efficient graph-based method, which is more simplified than LBM-EP, but still able to capture normal electrical activity and certain pathologies. Both approaches support anisotropic tissue properties by incorporating local fiber orientation into the computations. Models of the fast Purkinje network are also available for both approaches. The EP model further connects to an electrocardiogram (ECG) model by propagating electrical potentials on the heart to the torso.

The third component models cardiac biomechanics triggered by the potentials of the EP model. It is based on intertwined active and passive forces: a simplified model of active myocyte contraction and an orthotropic passive tissue constitutive law, which incorporates fiber orientation. Biomechanics is solved using a fast GPU-accelerated finite element method (FEM) implementation.

Finally, several boundary conditions are added to the model to mimic external conditions such as neighboring organs, blood flow and the circulatory system. First, a phase model of the cardiac cycle simulates the four cardiac phases by altering ventricular boundary conditions. Second, two three-element Windkessel (WK) models simulate arterial pressures, one for aorta, one for pulmonary artery. Third, the effect of the atria is simulated using two independent lumped elastance models. Fourth, two spatial constraints are introduced: base stiffness to mimic the effect of arteries and atria on ventricular motion, and pericardium constraint, a contact-based model of the pericardium to mimic the effects of the neighboring organs and of the pericardium on the cardiac motion.

The model described above is used for the design and evaluation of the methods developed in the course of this work, which should be considered the main contributions of this thesis. In Chap. 4, the first of three major contributions is presented: an integrated pipeline for streamlined personalization of the full cardiac model. A critical prerequisite for applying computational models in clinical practice is their ability to precisely capture an individual patient's physiology. The process of fitting a model to patient data by estimating model parameters is called personalization. First, the state-of-the-art in cardiac model personalization techniques is reviewed. Then each module of the proposed personalization pipeline is described. It should be noted that a major requirement was to build tools that work with routinely acquired data and do not rely specialized acquisition protocols. Based on the anatomical model segmented from MRI, first, an automated approach for personalizing hemodynamics (HD) boundary conditions from time-varying pressure and volume curves is presented.

In a semi-automated pre-processing step, one heart cycle needs to be selected from the typically multi-cycle pressure data, which is then smoothed and synchronized to the volume curves. Based on that data, five parameters of each arterial **WK** model are then personalized fully automatically by minimizing a specifically designed cost function in order to fit computed arterial pressure to the measurements. Next, the **EP** personalization module is presented. It personalizes four parameters of the **EP** model, three related to depolarization, one to repolarization, in order to compute simulated **ECG** traces that match the shape of the measured ones. To this end, a carefully designed multi-stage cascaded inverse optimization procedure is proposed to counter the ill-posed nature of the inverse **EP** problem. Next, the biomechanics personalization procedure is described, which adjusts seven biomechanical parameters, both passive and active, and two additional parameters of the atrial boundary conditions. As with **EP**, biomechanics personalization is fully automatized and based on inverse optimization involving hand-crafted cost functions. Finally, a hierarchical coarse-to-fine strategy for regional personalization of active biomechanical properties is proposed.

All modules of the pipeline were then evaluated comprehensively on a very large and heterogeneous cohort of more than one hundred patients, where about one third were children or adolescents, and two thirds were adults. The goal was to assess the ability of the personalization pipeline to capture cardiac physiology of the patients, thereby integrating the available clinical data into the model. To this end, the full pipeline was applied to all patients and the outputs of each step were compared to measured clinical indicators of cardiac function, i. e., by evaluating the goodness of fit between patient-specific model output and measurements.

The personalized artery pressures after **WK** personalization matched well for almost all patients, except few ones, where the pressure trends were difficult to capture due to the limited quality of available imaging data. However, in all cases, at least the diastolic pressure and the peak pressure were matched well, which are the most relevant features for the biomechanical computations. Next, **EP** results were analyzed. QRS and QT duration of the computed **ECG** matched excellently to the measured data for all patients. However, while analyzing **electrical axis (EA)**, a measurement that describes the shape of the QRS complex in the **ECG** and can depend on regional variations in electrical conductivity of the myocardial tissue, some outliers were observed. The outliers could be the results of potentially invalid assumptions of the **EP** model for certain pathologies. For the majority of patients (80%), however, only very low or acceptable errors within clinical variability were observed. Next, the biomechanics personalization was evaluated. The goodness of fit between global indicators of cardiac function were analyzed. Again, for the majority of patients, simulated clinical features such as **ejection fraction (EF)** and **stroke volume (SV)** matched almost perfectly to the data, which is further verified by qualitative comparison of heart dynamics as observed in **MRI** and simulated motion, as well as comparison between simulated and measured time-varying pressure and volume curves. To summarize these results: the proposed personalization pipeline can indeed capture physiology of patients as demonstrated on this large and heterogeneous cohort. Furthermore, an analysis of the distribution of personalized **EP** and biomechanics parameters for both patient groups revealed some minor differences between adults and children,

which can be explained well based on physiological considerations and findings in literature. Runtime for the full pipeline is typically between four to twelve hours per patient depending on various factors.

Next, regional mechanics personalization was tested in several experiments on both synthetic and real data. The first results were promising, and the preliminary experiment on real data even suggested that the approach could potentially identify and localize regional abnormalities in the myocardium. However, more work and larger test sets will be needed to validate these observations.

Finally, a virtual **cardiac resynchronization therapy (CRT)** pilot experiment to test the predictive power of the model was conducted. First, the model was personalized with respect to the state of a patient's heart at baseline, i. e., prior to implantation. Next, virtual **CRT** was performed *in silico* by adding electrical stimulation to the model at the location where the leads of the real **CRT** device were attached to mimic the device, and the model was then simulated under these modified conditions. Finally, quantitative predictors of **CRT** outcome, such as change in QRS duration, were computed from the model and compared to the real clinical outcome as observed by the electrophysiologist during device programming. The results suggested that the model could indeed predict acute electrophysiological changes for this patient, and predicted acute hemodynamics changes were consistent.

7.1.2 Advanced Personalization with Artificial Intelligence and Uncertainty Quantification

Chap. 5 opens part two of this thesis, which covers both the second and the third major contribution of this work. Starting from the observation that designing robust personalization algorithms as the ones presented in the previous chapter is a tedious, time-consuming, model- and data-specific process, the theme of the second major contribution was to explore whether **artificial intelligence (AI)** concepts can be used to learn this task. A novel personalization algorithm called VITO is presented, which is inspired by how humans approach personalization and after some time excel at the task: by enduring tedious trial-and-error phases, thereby gaining experience, and through intuition. Based on a set of training patients, VITO first learns the characteristics of the computational model to be personalized using a data-driven model exploration approach, i. e., it learns how the model behaves under change of parameters. Then this knowledge is transferred into the framework of **reinforcement learning (RL)** to automatically compute an optimal personalization strategy, which can then be applied to robustly personalize unseen patients. The algorithm is generic and requires only minimal user input to learn by itself how to personalize a model.

VITO was first tested in a synthetic scenario in order to verify its ability to minimize cost functions generically. Then it was evaluated extensively on two challenging cardiac models: the Graph-EP electrophysiology model from Sec. 3.2 and a lumped-parameter model of **whole-body circulation (WBC)**. One key observation was that it reached equivalent success rates as a carefully engineered hand-crafted approach, and it was substantially more robust than naïve cost function minimization. At the same time, as shown in detail for several configurations of the **WBC** model, personalization of unseen cases converged consistently faster than standard methods: by a factor of

up to seven. The hyper-parameters to be set by the user were only few and mostly intuitive to the designers of computational models. Those where more in-depth knowledge of the algorithm is required (e. g., the number of representative states for the decision process model) could in part be tuned automatically using a set of “scouting” patients. These positive results are evidence of good generalization properties and suggest that VITO was indeed able to generalize the knowledge obtained from a set of training patients to personalize unseen patients. Such an approach could therefore evolve into a unified framework for personalization of any computational physiological model and support designers of new models by automatically designing tailored and efficient personalization methods.

The third major contribution concerns the sometimes neglected, yet highly important topic of uncertainties in computational modeling. Chap. 6 starts off with an overview of the various sources of uncertainties in modeling, then the focus shifts towards inverse **uncertainty quantification (UQ)**, which in this context means the problem of propagating uncertainties backwards through the model personalization process. In particular, a stochastic personalization approach is proposed, with the ability to quantify uncertainties in estimated model parameters, while taking into account uncertainties due to sparse and limited data and noise in the clinical measurements. It can further identify solution non-uniqueness. The method is based on Bayesian inference and **Markov chain Monte Carlo (MCMC)** sampling is employed to generate samples of the model posterior. To make the sampling practical, an efficient surrogate model based on **polynomial chaos expansion (PCE)** is utilized, which approximates the computationally expensive full forward model to ascertain computational tractability of the problem. The full approach can be summarized as follows. In a first step, the surrogate **PCE** model is trained from the full model. Then for various levels of data noise a set of posterior densities are estimated, which are then aggregated to generate the final model posterior. The posterior is then analyzed to derive the most likely parameters given the unknown levels of noise in the data, and to identify high confidence regions.

This stochastic method was then evaluated on a number of heterogeneous **DCM** patients and the cardiac biomechanics model, where the goal was to estimate two of the most critical parameters related to cardiac dynamics: active contraction force and passive myocardial tissue stiffness. In a first experiment, the required fidelity of the **PCE** surrogate model was evaluated. Next, the proposed method was used to automatically compute most likely model parameters and confidence intervals. The fit between the output of the full model computed with the identified parameters and the real clinical measurements was below clinical variability and equivalent to results computed using standard deterministic personalization. Furthermore, deterministically estimated model parameters were within the automatically estimated 95% confidence intervals of the stochastic method. Computed uncertainties (as captured by the automatically estimated model posteriors) were evaluated empirically and the results suggested that the proposed inference and aggregation method can indeed provide meaningful uncertainty outputs. Finally, it could be shown that non-uniqueness of the problem can be identified automatically. Such an approach could therefore provide precious insights when analyzing the clinical relevance of estimated parameters

and personalized model predictions. In addition to that, it could constitute a criterion to select and refine data acquisition protocols used for model personalization.

7.2 Perspectives

The novel concepts and algorithms proposed in this thesis pushed forward the state-of-the-art in efficient and robust image-based computational biophysical model personalization, but at the same time they are opening up new avenues for future research opportunities. In this section, potential future directions as well as some of the most pressing open challenges towards clinical translation are discussed.

7.2.1 From Global to Regional Personalization

Most of the proposed methods in this thesis focus on global personalization with the goal to capture some of the most important clinical indicators of cardiac function, which are mostly of global nature (e. g., *EF* and *SV* capture blood output of an entire ventricle). For instance, one of the parameters that is adjusted in the *EP* model in Sec. 4.3.2 represents electrical conductivity of myocardial tissue. This parameter is estimated once for the entire myocardium, thus assuming homogeneity of the myocardial tissue (regionally varying fiber orientation is considered). Perfectly homogeneous tissue is obviously not encountered in many patients, if any at all. Disease-related variability due to cardiac infarction causing myocardial scars or fibrotic remodeling, but also non-pathological variability due to natural variations in tissue properties are not fully taken into account. There are several reasons for the selected choice of parameter granularity in each of the models and personalization algorithms presented in this thesis.

The first one concerns theoretical limitations in parameter identifiability. One major goal of this work was to develop methods that work with routinely acquired clinical data, which is inherently sparse. Coming back to the *EP* example above, the only standard measurement of cardiac electrical activity available in clinical routine is 12-lead *ECG*, which is not even measured directly at the heart, but rather indirectly on the torso. *ECG* therefore cannot provide detailed regional information and localized alterations cannot always be identified without additional information. Therefore, if more detailed parametrization was used, the solution space would explode quickly since more and more combinations of parameters could be identified that lead to the same *ECG* output. Parameters are more identifiable when less parameters are used.

The second consideration for the selected choice of parameter granularity is computational cost. The complexity of any personalization method heavily depends on the dimensionality of the parameter space. The best trade-off between computational burden and model flexibility needed to be found.

Nevertheless, novel methods are being developed to overcome these limitations to some extent, such as [Seeg 15], where additional information in the form of mechanical strain maps is incorporated into the *EP* personalization procedure to derive spatially varying parameters. In Sec. 4.3.4 of this thesis, an approach for regional biomechanical model personalization is proposed, where instead of personalizing a single contractility parameter for the entire left ventricle (*LV*), the ventricle is divided into

17 segments and one parameter is assigned to each segment. Based on “regional volumes”, a set of advanced descriptors of regional cardiac motion derived from the tracked endocardium meshes, regional information is extracted and incorporated in the personalization cost function to increase identifiability of regional alterations. First results were promising, however, to fully capture local variations in heart physiology for all patients including those with rare pathologies, additional methods may have to be developed to extract all relevant information from the available data. As a next step, time-varying spatial distributions of landmarks of the endocardial surfaces could be incorporated to further enhance identifiability of spatially varying parameters of the cardiac model.

7.2.2 Challenges of AI-based Personalization

Open questions and challenges related to VITO, the self-taught artificial personalization agent presented in Chap. 5, still remain. In its current form, the number of actions available to VITO to adjust intermediate parameter estimates is finite and their nature is discrete, i. e., only discrete increments or decrements to the parameters can be made, one parameter at a time. This may be limiting in terms of effectiveness and efficiency of on-line personalization. The incorporation of continuous action spaces [Van 12] could further improve its performance and versatility. However, care needs to be taken in order to keep computational costs manageable. The same holds for the state space. In this thesis, a data-driven state space quantization strategy is proposed. Contrary to [Neum 15b], where a threshold-based state-quantization involving several manually tuned threshold values (Fig. 5.2) was employed, the new method is based on a single hyper-parameter only: the number of representative states. Although it simplifies the setup of VITO, this quantization strategy may still not be optimal, especially if only little training data is available. Therefore, advanced approaches for continuous RL with value function approximation [Sutt 98, Mnih 15] could be integrated to fully circumvent quantization issues.

At the same time, such methods could improve VITO’s scalability towards high-dimensional estimation tasks. In this work it was shown that VITO can be applied to typical problems emerging in cardiac modeling, which could be described as medium-scale problems with moderate number of parameters to personalize and objectives to match. In unreported experiments involving >10 parameters, however, VITO could no longer reach satisfactory performance, which is likely due to the steeply increasing number of transition samples needed to sample the continuous state space of increasing dimensionality sufficiently during training. The trends in Sec. 5.3.3 confirm the need for more data. In the future, experience replay [Lin 93, Adam 12] or similar techniques could be employed to increase training data efficiency. Furthermore, massively parallel approaches [Nair 15] are starting to emerge, opening up new avenues for large-scale RL.

Although the employed RL techniques guarantee convergence to an optimal policy, the computed personalization strategy may not be optimal for the model under consideration as the environment is only partially observable and the personalization problem ill-posed: there is no guarantee for solution existence or uniqueness. Yet, it could be shown that VITO can solve personalization more robustly and more ef-

fectively than standard methods under the same conditions. However, a theoretical analysis in terms of convergence guarantees and general stability of the method would be desirable, in particular with regards to the proposed re-initialization strategy. As a first step towards this goal, in preliminary (unreported) experiments on the EP and the WBC model, it could be observed that the number of patients which do not require re-initialization (due to oscillation) to converge to a successful personalization consistently increased with increasing training data.

Beyond these challenges, VITO showed promising performance and versatility, making it a first step towards an automated, self-taught model personalization agent.

7.2.3 Challenges of Stochastic Personalization

Among the methods proposed in this work, the inverse UQ approach for stochastic model personalization described in Chap. 6 is the one with the largest potential for additional research. Theories and methodologies for forward UQ are relatively well established, but in inverse UQ approaches as the one described here, several difficulties remain unsolved in the UQ community. Challenges due to high dimensionality, which may become a limiting factor even in “standard” (deterministic) personalization methods, play a much larger role in inverse UQ. Computational cost increases much more rapidly with the dimensionality of the problem. In the approach presented in this thesis, thousands of samples of the posterior are needed even for low-dimensional parameter spaces.

On top of sampling the posterior, the complex model to be personalized needs to be sampled in an initial training step to determine coefficients of the PCE surrogate model. The complexity of PCE training is in fact exponential with respect to the dimensionality of the parameter space: at least as many forward runs of the full model are required as PCE coefficients need to be determined (cf. Eq. (6.5)). With increasing desired PCE model fidelity, the computational burden grows even worse. One potential direction for speeding up training of the PCE model is to employ advanced sampling techniques such as the recently proposed “adaptive sparse sampling” approach [Quic 16]. In this approach, each polynomial is analyzed with respect to its contribution to the quality of the surrogate model, and low-contribution polynomials are later excluded to reduce the overall number of coefficients to be estimated. However, even if such an approach improves training speed by a certain factor, still, the complexity class remains exponential and high-dimensional problems will remain out of reach. Finally, also other types of surrogate models exist such as Gaussian process emulators [Rasm 03] or reduced-order physiology-based models [Moll 16], whose advantages and limitations should be evaluated carefully for further development of the method. To explore the limits of the proposed stochastic personalization approach in terms of parameter dimensionality, as a next step, the number of model parameters to be estimated could be increased steadily, one by one, e. g., by including additional biomechanics parameters or parameters of other parts of the model such as the arterial WK models or EP model.

Despite all these challenges, on the path to translation of patient-specific computational models into clinical applications, uncertainties in data and modeling assumptions need to be considered carefully and extensively in order to prove the credibility

of the computational models. More work on UQ-driven approaches such as the one proposed in this work is therefore necessary. Fortunately, such approaches seem to gain popularity in the computational modeling research community, as very similar methods already started to appear in other related areas such as brain tumor growth modeling [Le 15]. This further suggests good generalization properties of such methods.

7.2.4 Predictive Power of Personalized Models

The generative power of the computational whole-heart model described in Chap. 3, namely its ability to compute cardiac function under various conditions, makes it ideal for computing predictors of therapy outcome, provided it is personalized accurately. Therefore, one important next step should be to further investigate the predictive power of the patient-specific model after personalization using the proposed methodology on a large number of cases. The positive virtual CRT results presented in Sec. 4.5, where virtual CRT treatment was performed on the personalized heart of one real CRT patient, should be considered only a starting point. Similar studies for model-based patient-specific prediction on small patient cohorts and different types of models have also been carried out [Serm 09, Seeg 15, Croz 16, Okad 17]. However, only in a large-scale validation effort one can truly evaluate the predictive power of the model when personalized using the algorithms proposed in this thesis.

7.2.5 What If We Could Create a Virtual Patient?

With the recent advances and the continuous rapid progress in imaging and personalized computational modeling, one can envision a future, in which new drugs are discovered by modeling diseases in a computer, and clinical trials are carried out on cohorts of virtual patients, before even enrolling any real patient into the trial [Vodo 14]. The implications of such a vision becoming reality are manifold with huge benefits towards the progress of research and medicine through reduction of costs and patient risk, and thereby acceleration of drug discovery and clinical trials, while at the same time, the necessity for animal experiments may decrease. Furthermore, new medical devices (advanced stents, pacemakers, and others) can be tested *in silico*, with huge potential towards more efficient regulatory approval and certification [Gott 17].

However, in order to accomplish this vision, society needs to commit to large interdisciplinary efforts that will allow the evolution of the current models, which already cover multiple scales—from cells, to tissues, to organs—into virtual models of entire patients. This undertaking will require long-standing and close collaboration between academia, translational researchers and industry. Until the necessary technological advancements, which may include fast simulation hardware, innovative mathematical formulations for coupling various multi-scale models, and additional medical insights, become available, the current organ-scale models, like the patient-specific virtual heart employed throughout this work, can and should be further extended, improved and evaluated by clinically-driven hypothesis testing and validated for patient-specific therapy outcome prediction.

Towards this goal, a potential next area of research could focus on gradually extending the current model to, e. g., a four-chamber version by adding atria, to increase the level of detail of the model by explicitly including other structures such as papillary muscles, or to replace the currently employed lumped models of the circulatory system by higher-dimensional ones. All of these changes may have significant impact on overall cardiac function and may allow to better capture certain pathologies. Progress on the modeling side will also entail the investigation of extended segmentation and personalization methods for these additional structures.

A.1 Software Tool for Interactive Personalization

A.1.1 VvTk Rapid Prototyping Platform

During the course of this thesis, significant contributions to Vitruvio toolkit (VvTk) have been made in terms of software engineering and algorithm development. VvTk is a modular C++ software platform for rapid prototyping and for delivering advanced clinical solutions, developed as proprietary software by the “Digital Technology and Innovation” group of the “Digital Services” department in Siemens Healthineers, the industrial collaboration partner of this thesis. The toolkit enables integration of algorithms and realization of modular workflows, while supporting large-scale batch processing and evaluation.

A.1.2 HeartBuildR

Fig. A.1 shows screenshots of the interactive computational heart model personalization software tool called “HeartBuildR”, which is based on VvTk and has been enhanced and extended continuously throughout this work. First, a patient browser lists all datasets available in a local database and the user selects one of them or creates a new entry by importing cardiac MRI data. Next, as described in Sec. 3.1, cardiac anatomy is automatically segmented from the imaging data. The user can then review and modify the resulting 3D surface meshes, if needed. Once the segmentation is accepted by the user, the meshes are fused into a single, watertight bi-ventricular mesh and a volumetric (tetrahedral) mesh is computed automatically from it. The volumetric mesh is then input to the anatomical model generation procedure, which adds myocardial fiber architecture, tagging, etc., to the model. Furthermore, a torso model for EP / ECG modeling is aligned semi-automatically to the anatomical model and MRI scout sequences, if available. Finally, based on the anatomical model and torso geometry, cardiac EP and biomechanics are personalized according to the algorithms described in Sec. 4.3.2 (or Sec. 5.2) and Sec. 4.3.3. As this process will take some time, there is an option to perform personalization off-line (script-based batch-processing environment), then load the results into HeartBuildR once again for inspection as soon as processing has completed.

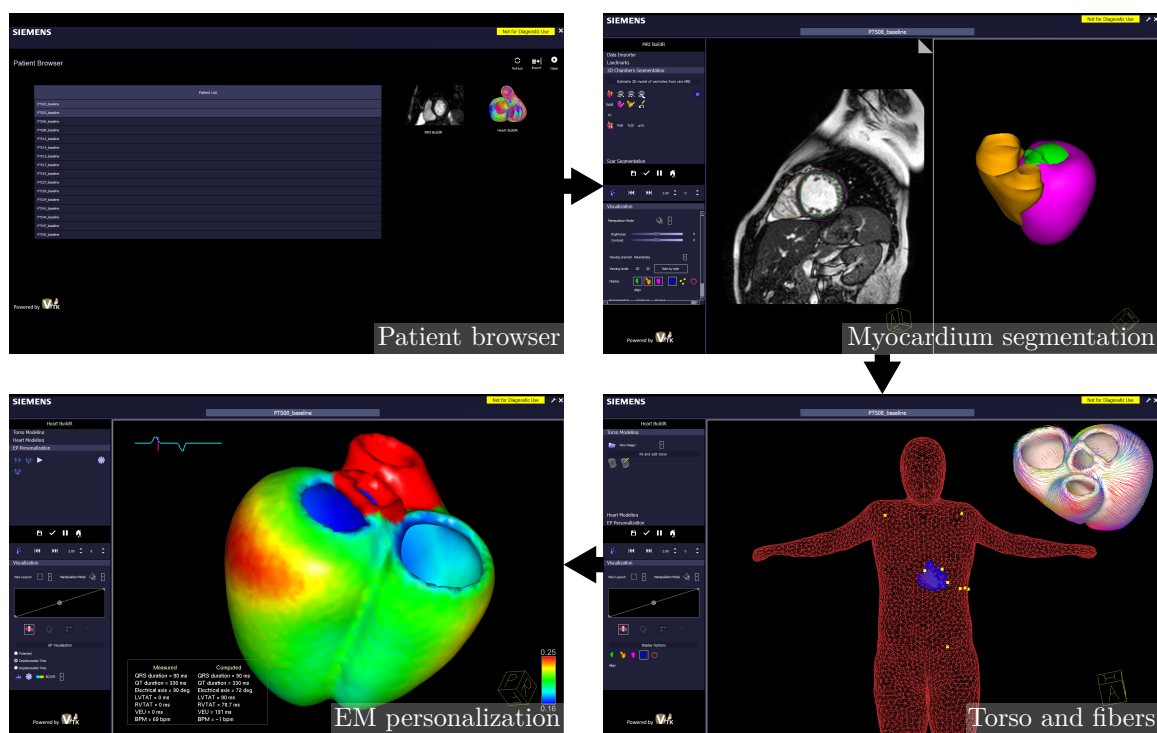


Figure A.1: Overview of the developed user interface (UI) for comprehensive personalization of the multi-physics whole-heart model.

A.2 Supplementary Personalization Results

This section presents additional results of the integrated personalization pipeline (Sec. 4.4) in terms of anatomical models of several patients from the study cohort, segmented from MRIs, illustrating the large variability in shape.




Figure A.2: Example anatomical models based on segmentation from MRI. The colored overlay represents myocardial fiber architecture (color-coded by fiber orientation) generated using the rule-based fiber model described in Sec. 3.1.

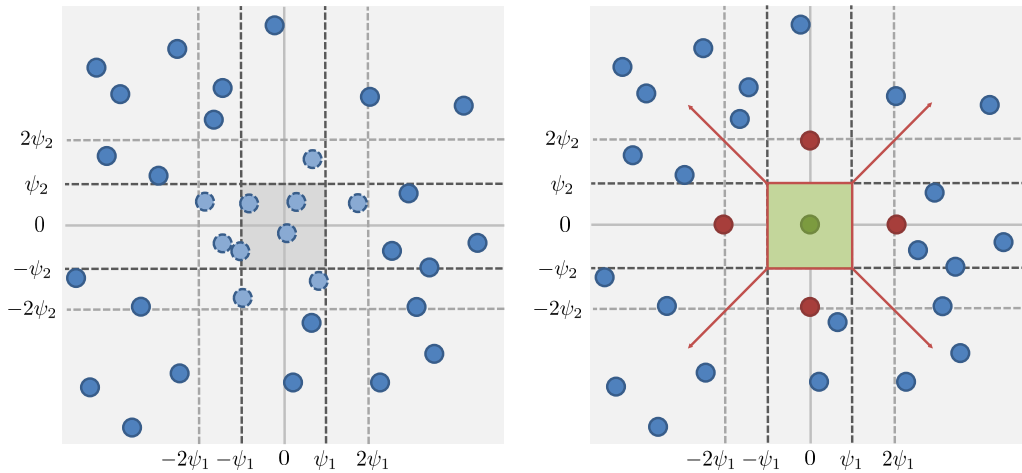


Figure A.3: Pre-processing of k -means input data to enforce the success state \hat{s} . Left: Continuous state space with observed objective vectors \mathbf{c} (blue points). The points with dashed outline will be canceled out. Right: Delineation of \hat{s} in green, enforced by inserted vectors (green / red points) with large weights. See text for details.

A.3 Additional Details of AI-based Personalization

A.3.1 Data-driven State Space Quantization

This section describes the details of the proposed data-driven quantization approach to define the set of representative MDP states \mathcal{S} (see Sec. 5.2.4). It is based on clustering, in particular on the weighted k -means algorithm described in [Arth07]. To this end, all objective vectors $\mathcal{C} = \{\mathbf{c} \in \mathcal{E}\}$ are extracted from the training data (Sec. 5.2.3). $\mathcal{C} \subset \mathbb{R}^{n_c}$ represents all observed “continuous states”. The goal is to convert \mathcal{C} into the finite set of representative MDP states \mathcal{S} while taking into account that VITO relies on a special “success state” \hat{s} encoding personalization success.

The success state \hat{s} does not depend on the data, but on the maximum acceptable misfit ψ . In particular, since personalization success implies that all objectives are met, \hat{s} should approximate a hyper-rectangle centered at $\mathbf{0}$ and bounded at $\pm\psi$, i. e., a small region in \mathbb{R}^{n_c} where $\forall i : |c_i| < \psi_i$. To enforce \hat{s} , the input to weighted k -means is pre-processed as follows.

First, the $\mathbf{0}$ -vector is inserted into \mathcal{C} , along with two vectors per dimension i , where all components are zero, except the i^{th} component, which is set to $\pm 2\psi_i$. These $2n_c + 1$ inserted vectors are later converted into centroids of representative states to delineate the desired hyper-rectangle for \hat{s} as illustrated in Fig. A.3. Furthermore, to avoid malformation of \hat{s} , no other representative state should emerge within that region. Therefore, all vectors $\mathbf{c} \in \mathcal{C}$, where $\forall i : |c_i| < 2\psi_i$ (except for the inserted vectors) are canceled out by assigning zero weight, while the inserted vectors are assigned large weights $\rightarrow \infty$ and all remaining vectors weights of 1.

Next, k -means is initialized by placing a subset of the initial centroids at the locations of the inserted states, and the remaining $n_S - 2n_c - 1$ centroids at random vectors in \mathcal{C} . Both the large weight and the custom initialization enforce the algorithm

to converge to a solution where one cluster centroid is located at each inserted vector, while the other centroids are distributed according to the training data. To ensure equal contribution of all objectives (cancel out different units, etc.), similarity is defined relative to the inverse of the user-defined convergence criteria (Eq. (5.7)).

Finally, after k -means converged, the resulting centroids, denoted ξ_s , are used to delineate the region in \mathbb{R}^{n_c} assigned to a representative state s .

A.3.2 Data-driven Initialization

This section describes the details of the proposed data-driven initialization approach to compute a set of candidate initialization parameters $\{\mathbf{x}_0\} \in \Omega$ for a new patient p based on the patient's measurements \mathbf{z}^p and the training database \mathcal{E} (see Sec. 5.2.6).

First, all model states are extracted from the training database: $\mathcal{Y} = \{\mathbf{y} \in \mathcal{E}\}$. Next, \mathcal{Y} is fed to a clustering algorithm (e.g. k -means). As in appendix A.3.1, the distance measure is defined relative to the inverse of the convergence criteria (Eq. (5.7)). The output is a set of centroids (for simplicity, in this work the number of centroids was set to $n_{\mathcal{S}}$), and each vector is assigned to one cluster based on its closest centroid. Let $\mathcal{Y}_p \subseteq \mathcal{Y}$ denote the members of the cluster whose centroid is closest to \mathbf{z}_p and

$$\Xi_p = \{\mathbf{x} \in \mathcal{E} \mid f(\mathbf{x}) \in \mathcal{Y}_p\} \quad (\text{A.1})$$

the set of corresponding model parameters. For each cluster, an approximation of the likelihood over the generating parameters is computed in terms of a [probability density function \(PDF\)](#). In this work a [Gaussian mixture model \(GMM\)](#) \mathcal{G} is assumed:

$$\mathcal{G}_p(\mathbf{x}) = \sum_{i=1}^K w_i \cdot \mathcal{N}(\mathbf{x}; \mathbf{m}_i, \Sigma_i) . \quad (\text{A.2})$$

The parameter vectors in Ξ_p are treated as random samples drawn from \mathcal{G}_p . Its properties, namely the number of mixture components K , their weights w_i , and their means \mathbf{m}_i and covariance matrices Σ_i , are estimated from these samples using a multivariate kernel density estimator with automated kernel bandwidth estimation, see [\[Kris 11\]](#) for more details. Finally, the M estimated means are selected as initialization candidates $\{\mathbf{x}_0\} = \{\mathbf{m}_i\}$. The candidates are sorted according to their w_i -value to prioritize more likely initializations.

List of Acronyms

AHA	American Heart Association 19, 20, 40, 46, 47, 55–58, 60, 117
AI	artificial intelligence iii, 4, 65, 66, 83, 100
AP	action potential 22, 23, 27, 28, 41, 42, 52
APD	action potential duration 23, 43, 52
ATP	adenosine triphosphate 28
BEM	boundary element method 26
BOBYQA	bound optimization by quadratic approximation 43, 44, 48, 52, 56–58, 75, 77, 92, 93
CI	confidence interval 51, 53, 87, 91–93
CO	cardiac output 12
CRT	cardiac resynchronization therapy 2, 59–61, 100, 105
CT	computed tomography 9, 10, 17, 26
DBP	diastolic arterial pressure 12
DCM	dilated cardiomyopathy 1, 2, 22, 48, 60, 87, 91, 97, 101
DL	deep learning 18
DTI	diffusion tensor imaging 20, 21
EA	electrical axis 11, 12, 26, 43, 50, 51, 75, 99
ECG	electrocardiogram 10–13, 22, 25, 26, 32, 35, 37, 39, 42, 48–51, 59, 60, 75, 98, 99, 102, 107, 113, 115, 116
ED	end-diastole 41, 49
EDP	end-diastolic pressure 45
EDV	end-diastolic volume 12, 45, 49, 79, 91
EF	ejection fraction 10, 12, 45–47, 49, 52, 56, 58, 59, 86, 91, 93, 94, 99, 102, 113, 114
EM	electromechanics 30, 35, 38–40, 43, 44, 46, 47, 49, 53–55, 57, 60, 86–88, 91, 93, 97, 114
EP	electrophysiology 3, 17, 19, 21, 23–26, 28, 29, 35–37, 39, 41–43, 48–52, 59–61, 73, 83, 84, 86, 91, 98, 99, 102, 104, 107, 114, 117
ES	end-systole 41
ESP	end-systolic pressure 45

ESV	end-systolic volume 12, 45, 79, 91
ET	ejection time 40
FEM	finite element method 29, 98
FSI	fluid-structure interaction 30
GMM	Gaussian mixture model 89–92, 111, 114
GPU	graphics processing unit 24, 30, 52, 98
HD	hemodynamics 17, 35, 39, 43, 48, 60, 91, 98
HF	heart failure iii, 1, 97
HO	Holzapfel-Ogden 27, 39, 43, 113
ICD	implantable cardioverter-defibrillators 2
LA	left atrium 8, 33, 43, 79
LBM	Lattice-Boltzmann method 23, 24, 41
LV	left ventricle 8, 13, 18–20, 25, 31, 33, 40, 43, 45–47, 49, 52–60, 79, 81, 91, 94, 102, 113, 115, 117
MAP	mean arterial pressure 12
MCMC	Markov chain Monte Carlo 88, 91, 92, 101
MDP	Markov decision process 67–72, 78, 83, 84, 113–117
ML	machine learning 3, 5, 17, 36, 37, 47, 66, 98
MRI	magnetic resonance imaging 2, 9, 10, 17, 18, 26, 38, 40, 43, 45, 46, 48, 49, 52–55, 58, 75, 79, 98, 99, 107, 109
MSDL	marginal space deep learning 18
MSL	marginal space learning 18
PCE	polynomial chaos expansion 86, 88, 89, 91, 92, 101, 104, 114, 115
PDF	probability density function 87, 89, 90, 92–94, 111, 114, 116, 118
RA	right atrium 8, 33, 43, 79
RL	reinforcement learning iii, 4, 66–68, 72, 83, 100, 103
RV	right ventricle 8, 20, 21, 25, 31, 40, 43, 45, 46, 50, 52, 55, 56, 59, 79, 113–115, 117
SBP	systolic arterial pressure 12
SSD	sum of squared differences 45, 46
SSFP	steady-state free precession 10, 18, 48, 52, 75
SV	stroke volume 10, 12, 45–47, 49, 52, 93, 99, 102, 117
TLED	total Lagrangian explicit dynamics 30
UI	user interface 47, 108

- UQ uncertainty quantification iii, 4, 85–87, 91, 94, 97, 101, 104
- US ultrasound 9, 10, 17
- VvTk Vitruvio toolkit 107
- WBC whole-body circulation 73, 79–83, 100, 104
- WK Windkessel 30, 31, 36, 38–41, 49, 50, 60, 79, 98, 99, 104, 113, 114, 116
- XML extensible markup language 47, 48

List of Symbols

a	MDP action
\mathcal{A}	set of MDP actions
\angle_{fiber}	fiber angle
\angle_{sheet}	sheet angle
β	dimensionless factor to isotropically scale tissue stiffness (HO model)
C	compliance: accounts for the elasticity of the arterial walls (WK model)
\mathbf{C}	Rayleigh damping matrix
c	personalization objective
\mathcal{C}	set of personalization objective vectors \mathbf{c}
\mathbf{c}	vector of personalization objectives c
\diamond	cluster of points
\mathbf{D}	anisotropy tensor
D	distance map
δ	reference value
$\boldsymbol{\delta}$	vector of reference values
e	model exploration episode
\mathcal{E}	set of model exploration episodes e
ε	measure of misfit between two variables or two vectors
\angle_{EA}	computed electrical axis (ECG)
$\overline{\angle_{\text{EA}}}$	measured electrical axis (ECG)
EF	computed ejection fraction (EF)
EF _{LV}	computed LV EF
$\overline{\text{EF}}_{\text{LV}}$	measured LV EF
EF _{RV}	computed RV EF

\overline{EF}_{RV}	measured RV EF
f	computational model
\tilde{f}	surrogate model based on PCE
\tilde{f}_{10}	high-fidelity surrogate model based on PCE
f_{EM}	computational EM (biomechanics) model
f_{EP}	computational EP model
f_{WK}	computational WK model
\mathbf{g}	normalized gradient vector
\mathcal{G}	Gaussian mixture model (GMM)
\mathcal{G}	vector of GMMs \mathcal{G}
γ	MDP discount factor
\mathbf{h}_a	active stress force vector
\mathbf{h}_b	boundary conditions force vector
\mathbf{h}_{base}	base stiffness force vector
HO	Holzapfel-Ogden energy function
\mathbf{h}_p	ventricular pressure force vector
\mathbf{h}_{peric}	pericardium constraint force vector
\mathbf{I}	identity matrix
\mathbf{K}	internal elastic stiffness matrix
k_{ATP}	contraction rate (active force model)
k_{RS}	relaxation rate (active force model)
K_{base}	base stiffness
\mathbf{A}	anisotropic stress tensor
\mathbf{M}	mass matrix
\mathcal{M}	Markov decision process (MDP)
m	mean
\mathcal{N}	normal PDF
\mathbf{n}	normal vector (e. g., surface normal)

$n_{\mathcal{A}}$	number of MDP actions
n_c	number of personalization objectives
n_{coeff}	number of PCE coefficients
n_{datasets}	number of datasets
$n_{\mathcal{E}}$	number of model exploration episodes
$n_{e\text{-steps}}$	number of consecutive transitions in a model exploration episode
n_{iter}	maximum number of iterations
n_{poly}	maximum polynomial order of PCE model
n_{ψ}	number of convergence criteria
$n_{\mathcal{S}}$	number of MDP states
n_{samples}	number of training samples
$n_{\mathcal{S}}$	number of noise levels
ν	electrical conductivity (Graph-EP) or diffusivity (LBM-EP)
ν_{LV}	electrical conductivity (Graph-EP) or diffusivity (LBM-EP) of the LV endocardium domain
ν_{Myo}	electrical conductivity (Graph-EP) or diffusivity (LBM-EP) of the myocardium domain
ν_{Purk}	electrical conductivity (Graph-EP) or diffusivity (LBM-EP) of the Purkinje fibers
ν_{RV}	electrical conductivity (Graph-EP) or diffusivity (LBM-EP) of the RV endocardium domain
ν_{torso}	electrical conductivity of the torso (ECG model)
$n_{\mathbf{x}}$	number of model parameters
$n_{\mathbf{y}}$	number of model state variables
n_z	number of patient measurements
Ω	domain of physiologically plausible model parameters
\mathbf{o}	fiber orientation vector
P_{ar}	time-varying artery pressure
P_{atr}	time-varying atrium pressure
$\overline{P_{\text{ar}}}$	measured time-varying artery pressure

p	PDF
ϕ	voltage / electrical potential
π	deterministic MDP policy
π^*	optimal deterministic MDP policy for a given problem
$\tilde{\pi}$	stochastic MDP policy
$\tilde{\pi}^*$	optimal stochastic MDP policy for a given problem
P_{init}	initial arterial pressure (WK model)
P_{LA}	pulmonary vein pressure (left atrium model)
Δ_{PR}	computed PR duration (ECG)
P_{ref}	constant low pressure of reference (WK model)
P_{RA}	vena cava pressure (right atrium model)
$\overline{\Delta_{\text{PR}}}$	measured PR duration (ECG)
ψ	convergence criterion
$\boldsymbol{\psi}$	vector of convergence criteria ψ
Ψ	multivariate polynomial basis
P_{ven}	time-varying ventricle pressure
$\overline{P_{\text{ven}}}$	measured time-varying ventricle pressure
Q_{ar}	time-varying artery flow
Δ_{QRS}	computed QRS duration (ECG)
$\overline{\Delta_{\text{QRS}}}$	measured QRS duration (ECG)
Δ_{QT}	computed QT duration (ECG)
$\overline{\Delta_{\text{QT}}}$	measured QT duration (ECG)
\mathcal{R}	scalar MDP reward function
\mathbb{R}	set of real numbers
R_c	characteristic resistance: accounts for the blood mass and the compliance of the artery proximal to the valves (WK model)
R_p	peripheral resistance: accounts for the distal resistance of the circulatory system (WK model)
σ	time-varying active contraction force (biomechanics model)

$\hat{\sigma}$	maximum asymptotic active contraction force
$\hat{\sigma}_{\text{AHA}}$	maximum asymptotic active contraction force in tissue of specific AHA segment
$\hat{\sigma}_{\text{LV}}$	maximum asymptotic active contraction force in LV tissue
$\hat{\sigma}_{\text{RV}}$	maximum asymptotic active contraction force in RV tissue
\mathbf{S}	noise level covariance matrix
s	MDP state
\hat{s}	MDP state representing the solution of a given personalization problem
\mathcal{S}	set of MDP states
SD	standard deviation
SV	computed stroke volume (SV)
SV_{LV}	computed LV SV
$\overline{\text{SV}_{\text{LV}}}$	measured LV SV
SV_{RV}	computed RV SV
$\overline{\text{SV}_{\text{RV}}}$	measured RV SV
\mathcal{T}	stochastic MDP transition function
t_{act}	time of electrical activation
τ_{close}	EP model parameter related to the time during which the ion channels are closed
τ_{in}	EP model parameter related to the timing of the ion channels
τ_{open}	EP model parameter related to the timing of the ion channels
τ_{out}	EP model parameter related to the timing of the ion channels
t_{dep}	depolarization time
t_{HP}	heart period (duration of one heart cycle)
t_{rep}	repolarization time
t_{travel}	travel time of the electrical signal
\mathbf{u}	mesh node displacement
$\dot{\mathbf{u}}$	mesh node velocity
$\ddot{\mathbf{u}}$	mesh node acceleration

\mathcal{U}	uniform PDF
\mathbf{v}	vertex of a mesh
V_{atr}	time-varying atrium volume
V_{ven}	time-varying ventricle volume
$\overline{V_{\text{ven}}}$	measured time-varying ventricle volume
w	weighting factor
ξ	centroid
$\boldsymbol{\xi}$	vector of centroids
x	model parameter
\mathbf{x}	vector of model parameters x
$\boldsymbol{\Upsilon}$	estimated covariance matrix
y	observable model state variable, also called model response
\mathbf{y}	vector of observable model state variables y
z	patient measurement (observed data)
\mathbf{z}	vector of patient measurements z
ζ	cost Function

List of Figures

1.1	Healthy heart versus dilated cardiomyopathy	2
2.1	Diagram of the human heart	10
2.2	Example cardiac MRI	12
2.3	Electrical axis	13
2.4	Wiggers diagram	15
3.1	Overview of computational heart model	20
3.2	From images to anatomical model	21
3.3	AHA segments bull’s eye view with example	22
3.4	Action potential	25
3.5	Forward model of cardiac electrocardiography	28
3.6	Active force model	31
3.7	Three element WK model: electrical circuit analogy	33
4.1	Computational model personalization	40
4.2	Integrated EM personalization pipeline	41
4.3	WK personalization procedure	43
4.4	Regional volumes	49
4.5	Personalized Windkessel models	52
4.6	Goodness of fit after whole-heart personalization: electrophysiology	53
4.7	Distribution of estimated EP parameters	54
4.8	Goodness of fit after EM personalization: biomechanics	55
4.9	Goodness of fit: dynamic images versus simulated motion	56
4.10	Goodness of fit: ventricular pressure and volume curves	56
4.11	Distribution of estimated biomechanics parameters	57
4.12	Regional parameter sensitivity (patient 1)	58
4.13	Comparison of residual regional volume errors	60
4.14	CRT experiment results	61
5.1	Self-taught personalization agent overview	69
5.2	Data-driven state space quantization example	73
5.3	Probabilistic on-line personalization	74
5.4	Rosenbrock convergence plots	75
5.5	Hyper-parameter scouting for Graph-EP personalization	78
5.6	Misfit after data-driven initialization	79
5.7	EP personalization convergence plots	81
5.8	WBC personalization setups	83
5.9	WBC personalization results	84
6.1	Proposed robust estimation framework	89

6.2	Illustration of GMM aggregation steps	92
6.3	From surrogate model to parameters and uncertainty	93
6.4	Estimated confidence intervals	95
6.5	Empirical evaluation of estimated uncertainty	96
A.1	Interactive user interface of model personalization pipeline	112
A.2	Anatomical models	113
A.3	Pre-processing for automatic state space quantization	114

List of Tables

2.1	Measurements of heart function	14
4.1	Objectives for pressure and volume fitting	47
4.2	Cohort characteristics	53
4.3	Goodness of fit: measured vs. computed ECG	54
4.4	Goodness of fit: measured vs. computed biomechanics	55
5.1	WBC parameters	82
5.2	WBC objectives	82
5.3	Overview of tested applications	86

Bibliography

- [Adam 12] S. Adam, L. Busoniu, and R. Babuska. “Experience Replay for Real-time Reinforcement Learning Control”. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, Vol. 42, No. 2, pp. 201–212, March 2012.
- [Adam 13] B. M. Adams, K. R. Dalbey, M. S. Eldred, D. M. Gay, and L. P. Swiler. “DAKOTA, a Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 5.4 User’s Manual”. Tech. Rep., Sandia National Laboratories, Livermore, CA, USA, Dec. 2013.
- [Alie 96] R. R. Aliev and A. V. Panfilov. “A Simple Two-Variable Model of Cardiac Excitation”. *Chaos, Solitons and Fractals*, Vol. 7, No. 3, pp. 293–301, March 1996.
- [Amr 16] A. Amr, E. Kayvanpour, F. Sedaghat-Hamedani, T. Passerini, V. Mihalef, A. Lai, D. Neumann, B. Georgescu, S. Buss, D. Mereles, E. Zitron, A. Posch, M. Würstle, T. Mansi, H. Katus, and B. Meder. “Personalized Computer Simulation of Diastolic Function in Heart Failure”. *Genomics, Proteomics & Bioinformatics*, Vol. 14, No. 4, pp. 244–252, Aug. 2016.
- [Arsi 06] V. Arsigny, P. Fillard, X. Pennec, and N. Ayache. “Log-Euclidean Metrics for Fast and Simple Calculus on Diffusion Tensors”. *Magnetic Resonance in Medicine*, Vol. 56, No. 2, pp. 411–421, Aug. 2006.
- [Arth 07] D. Arthur and S. Vassilvitskii. “k-means++: The Advantages of Careful Seeding”. In: *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1027–1035, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, Jan. 2007.
- [Arts 05] T. Arts, T. Delhaas, P. Bovendeerd, X. Verbeek, and F. W. Prinzen. “Adaptation to Mechanical Load Determines Shape and Properties of Heart and Circulation: the CircAdapt Model”. *American Journal of Physiology – Heart and Circulatory Physiology*, Vol. 288, No. 4, pp. 1943–1954, Apr. 2005.
- [Arts 91] T. Arts, P. Bovendeerd, F. W. Prinzen, and R. S. Reneman. “Relation between Left Ventricular Cavity Pressure and Volume and Systolic Fiber Stress and Strain in the Wall”. *Biophysical Journal*, Vol. 59, No. 1, pp. 93–102, Jan. 1991.
- [Bala 17] G. Balaban, H. Finsberg, H. H. Odland, M. Rognes, S. Ross, J. Sundnes, and S. Wall. “High Resolution Data Assimilation of Cardiac Mechanics Applied to a Dyssynchronous Ventricle”. *International Journal for Numerical Methods in Biomedical Engineering*, Vol. 33, No. 11, pp. 1–17, Dec. 2017.
- [Baxl 71] W. A. Baxley and T. J. Reeves. “Abnormal Regional Myocardial Performance in Coronary Artery Disease”. *Progress in Cardiovascular Diseases*, Vol. 13, No. 5, pp. 405–421, March 1971.

- [Baye 12] J. Bayer, R. Blake, G. Plank, and N. Trayanova. “A Novel Rule-based Algorithm for Assigning Myocardial Fiber Orientation to Computational Heart Models”. *Annals of Biomedical Engineering*, Vol. 40, No. 10, pp. 2243–2254, Oct. 2012.
- [Beau 13] J. Beaudoin, M. D. Handschumacher, X. Zeng, J. Hung, E. L. Morris, R. A. Levine, and E. Schwammenthal. “Mitral Valve Enlargement in Chronic Aortic Regurgitation as a Compensatory Mechanism to Prevent Functional Mitral Regurgitation in the Dilated Left Ventricle”. *Journal of the American College of Cardiology*, Vol. 61, No. 17, pp. 1809–1816, Apr. 2013.
- [Bell 57] R. Bellman. *Dynamic Programming*. Princeton University Press, Princeton, NJ, USA, 1st Ed., 1957.
- [Best 01] J. Bestel, F. Clément, and M. Sorine. “A Biomechanical Model of Muscle Contraction”. In: *Medical Image Computing and Computer-Assisted Intervention*, pp. 1159–1161, Springer International Publishing, Cham, Switzerland, Oct. 2001.
- [Bier 11] A. Biernacka and N. G. Frangogiannis. “Aging and Cardiac Fibrosis”. *Aging and Disease*, Vol. 2, No. 2, pp. 158–173, Apr. 2011.
- [Bish 06] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer Science+Business Media, New York, NY, USA, 4th Ed., 2006.
- [Bour 09] Y. Bourgault, Y. Coudiere, and C. Pierre. “Existence and Uniqueness of the Solution for the Bidomain Model used in Cardiac Electrophysiology”. *Nonlinear Analysis: Real World Applications*, Vol. 10, No. 1, pp. 458–482, Feb. 2009.
- [Cama 10] O. Camara, A. Pashaei, R. Sebastian, and A. F. Frangi. “Personalization of Fast Conduction Purkinje System in Eikonal-Based Electrophysiological Models with Optical Mapping Data”. *Statistical Atlases and Computational Models of the Heart*, Vol. 6364, pp. 281–290, Sep. 2010.
- [Cans 15] F. B. C. Cansız, H. Dal, and M. Kaliske. “An Orthotropic Viscoelastic Material Model for Passive Myocardium: Theory and Algorithmic Treatment”. *Computer Methods in Biomechanics and Biomedical Engineering*, Vol. 18, No. 11, pp. 1160–1172, Aug. 2015.
- [Cerq 02] M. D. Cerqueira, N. J. Weissman, V. Dilsizian, A. K. Jacobs, S. Kaul, W. K. Laskey, D. J. Pennell, J. A. Rumberger, T. Ryan, M. S. Verani, *et al.* “Standardized Myocardial Segmentation and Nomenclature for Tomographic Imaging of the Heart: a Statement for Healthcare Professionals from the Cardiac Imaging Committee of the Council on Clinical Cardiology of the American Heart Association”. *Circulation*, Vol. 105, No. 4, pp. 539–542, Jan. 2002.
- [CGAL 17] “CGAL User and Reference Manual”. <http://doc.cgal.org/4.8.1/Manual/packages.html>, 2017. Accessed: Nov 25, 2017.
- [Chab 12] R. Chabiniok, P. Moireau, P.-F. Lesault, A. Rahmouni, J.-F. Deux, and D. Chapelle. “Estimation of Tissue Contractility from Cardiac Cine-MRI using a Biomechanical Heart Model”. *Biomechanics and Modeling Mechanics*, Vol. 11, No. 5, pp. 609–630, May 2012.

- [Chap 01] D. Chapelle, F. Clément, F. Génot, P. Le Tallec, M. Sorine, and J. M. Urquiza. “A Physiologically-Based Model for the Active Cardiac Muscle Contraction”. In: *International Workshop on Functional Imaging and Modeling of the Heart*, pp. 128–133, Springer, Berlin, Heidelberg, Germany, Oct. 2001.
- [Chap 12] D. Chapelle, P. Le Tallec, P. Moireau, and M. Sorine. “Energy-preserving Muscle Tissue Model: Formulation and Compatible Discretizations”. *International Journal for Multiscale Computational Engineering*, Vol. 10, No. 2, pp. 189–211, 2012.
- [Chha 12] M. Chhay, Y. Coudière, and R. Turpault. *How to Compute the Extracellular Potential in Electrocardiology from an Extended Monodomain Model*. PhD thesis, École Nationale Supérieure des Mines de Paris, Paris, France, March 2012.
- [Chin 08] P. Chinchapatnam, K. S. Rhode, M. Ginks, C. A. Rinaldi, P. Lambiase, R. Razavi, S. Arridge, and M. Sermesant. “Model-based Imaging of Cardiac Apparent Conductivity and Local Conduction Velocity for Diagnosis and Planning of Therapy”. *IEEE Transactions on Medical Imaging*, Vol. 27, No. 11, pp. 1631–1642, Nov. 2008.
- [Clay 08] R. Clayton and A. Panfilov. “A Guide to Modelling Cardiac Electrical Activity in Anatomically Detailed Ventricles”. *Progress in Biophysics and Molecular Biology*, Vol. 96, No. 1, pp. 19–43, Jan. 2008.
- [Clay 11] R. Clayton, O. Bernus, E. Cherry, H. Dierckx, F. Fenton, L. Mirabella, A. Panfilov, F. Sachse, G. Seemann, and H. Zhang. “Models of Cardiac Tissue Electrophysiology: Progress, Challenges and Open Questions”. *Progress in Biophysics and Molecular Biology*, Vol. 104, No. 1, pp. 22–48, Jan. 2011.
- [Coma 02] D. Comaniciu and P. Meer. “Mean Shift: A Robust Approach toward Feature Space Analysis”. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 24, No. 5, pp. 603–619, May 2002.
- [Coma 08] O. Comas, Z. A. Taylor, J. Allard, S. Ourselin, S. Cotin, and J. Passenger. “Efficient Nonlinear FEM for Soft Tissue Modelling and its GPU Implementation within the Open Source Framework SOFA”. In: *Biomedical Simulation*, pp. 28–39, Springer, Berlin, Heidelberg, Germany, July 2008.
- [Corr 16] C. Corrado and S. A. Niederer. “A Two-Variable Model Robust to Pacemaker Behaviour for the Dynamics of the Cardiac Action Potential”. *Mathematical Biosciences*, Vol. 281, pp. 46–54, Nov. 2016.
- [Cost 01] K. D. Costa, J. W. Holmes, and A. D. McCulloch. “Modelling Cardiac Mechanical Properties in Three Dimensions”. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, Vol. 359, No. 1783, pp. 1233–1250, June 2001.
- [Coud 06] Y. Coudière and C. Pierre. “Stability and Convergence of a Finite Volume Method for Two Systems of Reaction-Diffusion Equations in Electro-Cardiology”. *Nonlinear Analysis: Real World Applications*, Vol. 7, No. 4, pp. 916–935, Sep. 2006.

- [Croz 16] A. Crozier, B. Blazevic, P. Lamata, G. Plank, M. Ginks, S. Duckett, M. Sohal, A. Shetty, C. A. Rinaldi, R. Razavi, S. A. Niederer, and N. P. Smith. “Analysis of Lead Placement Optimization Metrics in Cardiac Resynchronization Therapy with Computational Modelling”. *Europace*, Vol. 18, No. 4, pp. 113–120, Dec. 2016.
- [Deli 12] H. Delingette, F. Billet, K. C. L. Wong, M. Sermesant, K. Rhode, M. Ginks, C. A. Rinaldi, R. Razavi, and N. Ayache. “Personalization of Cardiac Motion and Contractility From Images Using Variational Data Assimilation”. *IEEE Transactions on Biomedical Engineering*, Vol. 59, No. 1, pp. 20–24, Jan. 2012.
- [Dham 16] J. Dhamala, J. L. Sapp, M. Horacek, and L. Wang. “Spatially-Adaptive Multi-scale Optimization for Local Parameter Estimation: Application in Cardiac Electrophysiological Models”. In: *Medical Image Computing and Computer-Assisted Intervention*, pp. 282–290, Springer International Publishing, Cham, Switzerland, Oct. 2016.
- [Dijk 59] E. W. Dijkstra. “A Note on Two Problems in Connexion With Graphs”. *Numerische Mathematik*, Vol. 1, No. 1, pp. 269–271, Dec. 1959.
- [Eldr 09] M. S. Eldred and J. Burkardt. “Comparison of Non-intrusive Polynomial Chaos and Stochastic Collocation Methods for Uncertainty Quantification”. *AIAA Aerospace Sciences Meeting*, Vol. 976, No. 2009, pp. 1–20, Jan. 2009.
- [Elsh 14] M. M. Elshrif and E. M. Cherry. “A Quantitative Comparison of the Behavior of Human Ventricular Cardiac Electrophysiology Models in Tissue”. *PloS One*, Vol. 9, No. 1, pp. 1–13, Jan. 2014.
- [Fent 98] F. Fenton and A. Karma. “Vortex Dynamics in Three-dimensional Continuous Myocardium with Fiber Rotation: Filament Instability and Fibrillation”. *Chaos: an Interdisciplinary Journal of Nonlinear Science*, Vol. 8, No. 1, pp. 20–47, 1998.
- [Fitz 61] R. FitzHugh. “Impulses and Physiological States in Theoretical Models of Nerve Membrane”. *Biophysical Journal*, Vol. 1, No. 6, pp. 445–466, July 1961.
- [Fran 01] A. F. Frangi, W. J. Niessen, and M. Viergever. “Three-dimensional Modeling for Functional Analysis of Cardiac Images, a Review”. *IEEE Transactions on Medical Imaging*, Vol. 20, No. 1, pp. 2–5, Jan. 2001.
- [Fran 62] D. L. Franklin, R. L. Van Citters, and R. F. Rushmer. “Balance Between Right and Left Ventricular Output”. *Circulation Research*, Vol. 10, No. 1, pp. 17–26, Jan. 1962.
- [Fran 90] P. C. Franzone, L. Guerri, and S. Rovida. “Wavefront Propagation in an Activation Model of the Anisotropic Cardiac Tissue: Asymptotic Analysis and Numerical Simulations”. *Journal of Mathematical Biology*, Vol. 28, No. 2, pp. 121–176, Feb. 1990.
- [Fuer 15] B. Fuerst, T. Mansi, F. Carnis, M. Sälzle, J. Zhang, J. Declerck, T. Boettger, J. Bayouth, N. Navab, and A. Kamen. “Patient-specific Biomechanical Model for the Prediction of Lung Motion from 4-D CT Images”. *IEEE Transactions on Medical Imaging*, Vol. 34, No. 2, pp. 599–607, Feb. 2015.

- [Fung 93] Y.-C. Fung. *Biomechanics: Mechanical Properties of Living Tissues*. Springer Science+Business Media, New York, NY, USA, 2nd Ed., 1993.
- [Gao 15] H. Gao, W. Li, L. Cai, C. Berry, and X. Luo. “Parameter Estimation in a Holzapfel-Ogden Law for Healthy Myocardium”. *Journal of Engineering Mathematics*, Vol. 95, No. 1, pp. 231–248, Jan. 2015.
- [Geor 05] B. Georgescu, X. S. Zhou, D. Comaniciu, and A. Gupta. “Database-guided Segmentation of Anatomical Structures with Complex Appearance”. In: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 429–436, IEEE, San Diego, CA, USA, July 2005.
- [Ghes 16a] F. C. Ghesu, B. Georgescu, T. Mansi, D. Neumann, J. Hornegger, and D. Comaniciu. “An Artificial Agent for Anatomical Landmark Detection in Medical Images”. In: *Medical Image Computing and Computer-Assisted Intervention*, pp. 229–237, Springer International Publishing, Cham, Switzerland, Oct. 2016.
- [Ghes 16b] F. C. Ghesu, E. Krubasik, B. Georgescu, V. Singh, Y. Zheng, J. Hornegger, and D. Comaniciu. “Marginal Space Deep Learning: Efficient Architecture for Volumetric Image Parsing”. *IEEE Transactions on Medical Imaging*, Vol. 35, No. 5, pp. 1217–1228, March 2016.
- [Gokt 11] S. Göktepe, S. Acharya, J. Wong, and E. Kuhl. “Computational Modeling of Passive Myocardium”. *International Journal for Numerical Methods in Biomedical Engineering*, Vol. 27, No. 1, pp. 1–12, Jan. 2011.
- [Gott 17] S. Gottlieb. “How FDA Plans to Help Consumers Capitalize on Advances in Science”. <https://blogs.fda.gov/fdavoices/index.php/tag/in-silico-tools/>, 2017. Accessed: Nov 28, 2017.
- [Grbi 17] S. Grbic, T. F. Easley, T. Mansi, C. H. Bloodworth, E. L. Pierce, I. Voigt, D. Neumann, J. Krebs, D. D. Yuh, M. O. Jensen, D. Comaniciu, and A. P. Yoganathan. “Personalized Mitral Valve Closure Computation and Uncertainty Analysis from 3D Echocardiography”. *Medical Image Analysis*, Vol. 35, pp. 238–249, Jan. 2017.
- [Gucc 93] J. Guccione, L. Waldman, and A. McCulloch. “Mechanics of Active Contraction in Cardiac Muscle: Part II – Cylindrical models of the Systolic Left Ventricle”. *Journal of Biomechanical Engineering*, Vol. 115, No. 1, pp. 82–90, Feb. 1993.
- [Haar 06] H. Haario, M. Laine, A. Mira, and E. Saksman. “DRAM: Efficient Adaptive MCMC”. *Statistics and Computing*, Vol. 16, No. 4, pp. 339–354, Dec. 2006.
- [Hadj 17] M. Hadjicharalambous, L. Asner, R. Chabiniok, E. Sammut, J. Wong, D. Peressutti, E. Kerfoot, A. King, J. Lee, R. Razavi, *et al.* “Non-invasive Model-based Assessment of Passive Left-Ventricular Myocardial Stiffness in Healthy Subjects and in Patients with non-ischemic Dilated Cardiomyopathy”. *Annals of Biomedical Engineering*, Vol. 45, No. 3, pp. 605–618, March 2017.
- [Hodg 52] A. L. Hodgkin and A. F. Huxley. “A Quantitative Description of Membrane Current and its Application to Conduction and Excitation in Nerve”. *The Journal of Physiology*, Vol. 117, No. 4, pp. 500–544, Aug. 1952.

- [Holz 09] G. A. Holzapfel and R. W. Ogden. “Constitutive Modelling of Passive Myocardium: a Structurally based Framework for Material Characterization”. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, Vol. 367, No. 1902, pp. 3445–3475, Sep. 2009.
- [Hump 90a] J. Humphrey, R. Strumpf, and F. Yin. “Determination of a Constitutive Relation for Passive Myocardium: I. A New Functional Form”. *Journal of Biomechanical Engineering*, Vol. 112, No. 3, pp. 333–339, Aug. 1990.
- [Hump 90b] J. Humphrey, R. Strumpf, and F. Yin. “Determination of a Constitutive Relation for Passive Myocardium: II. Parameter Estimation”. *Journal of Biomechanical Engineering*, Vol. 112, No. 3, pp. 340–346, Aug. 1990.
- [Hunt 03] P. J. Hunter and T. K. Borg. “Integration from Proteins to Organs: the Physiome Project”. *Nature Reviews Molecular Cell Biology*, Vol. 4, No. 3, pp. 237–243, March 2003.
- [Hunt 88] P. J. Hunter and B. H. Smaill. “The Analysis of Cardiac Function: a Continuum Approach”. *Progress in Biophysics and Molecular Biology*, Vol. 52, No. 2, pp. 101–164, 1988.
- [Hunt 98] P. Hunter, A. McCulloch, and H. Ter Keurs. “Modelling the Mechanical Properties of Cardiac Muscle”. *Progress in Biophysics and Molecular Biology*, Vol. 69, No. 2, pp. 289–331, March 1998.
- [Iona 10] R. I. Ionasec, I. Voigt, B. Georgescu, Y. Wang, H. Houle, F. Vega-Higuera, N. Navab, and D. Comaniciu. “Patient-specific Modeling and Quantification of the Aortic and Mitral valves from 4-D Cardiac CT and TEE”. *IEEE Transactions on Medical Imaging*, Vol. 29, No. 9, pp. 1636–1651, Sep. 2010.
- [Isma 10] H. Ismail and A. N. Makaryus. “Predictors of Response to Cardiac Resynchronization Therapy: the Holy Grail of Electrophysiology”. *The International Journal of Cardiovascular Imaging*, Vol. 26, No. 2, pp. 197–198, Feb. 2010.
- [Itu 14] L. Itu, P. Sharma, B. Georgescu, A. Kamen, C. Suci, and D. Comaniciu. “Model-based Non-invasive Estimation of PV Loop from Echocardiography”. In: *36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 6774–6777, IEEE, Jan. 2014.
- [Jone 04] S. A. Jones, M. K. Lancaster, and M. R. Boyett. “Ageing-related Changes of Connexins and Conduction within the Sinoatrial Node”. *The Journal of Physiology*, Vol. 560, No. 2, pp. 429–437, Oct. 2004.
- [Kael 96] L. P. Kaelbling, M. L. Littman, and A. W. Moore. “Reinforcement Learning: a Survey”. *Journal of Artificial Intelligence Research*, Vol. 4, pp. 237–285, May 1996.
- [Kayv 15] E. Kayvanpour, T. Mansi, F. Sedaghat-Hamedani, A. Amr, D. Neumann, B. Georgescu, P. Seegerer, A. Kamen, J. Haas, K. S. Frese, M. Irawati, E. Wirsz, V. King, S. Buss, D. Mereles, E. Zitron, A. Keller, H. A. Katus, D. Comaniciu, and B. Meder. “Towards Personalized Cardiology: Multi-Scale Modeling of the Failing Heart”. *PloS One*, Vol. 10, No. 7, pp. 1–18, July 2015.

- [Kenn 01] M. C. Kennedy and A. O'Hagan. "Bayesian Calibration of Computer Models". *Journal of the Royal Statistical Society: Series B – Statistical Methodology*, Vol. 63, No. 3, pp. 425–464, Jan. 2001.
- [Kerc 07] R. C. Kerckhoffs, M. L. Neal, Q. Gu, J. B. Bassingthwaite, J. H. Omens, and A. D. McCulloch. "Coupling of a 3D Finite Element Model of Cardiac Ventricular Mechanics to Lumped Systems Models of the Systemic and Pulmonic Circulation". *Annals of Biomedical Engineering*, Vol. 35, No. 1, pp. 1–18, Jan. 2007.
- [Kerc 08] R. C. P. Kerckhoffs, J. Lumens, K. Vernoooy, J. Omens, L. Mulligan, T. Delhaas, T. Arts, A. McCulloch, and F. Prinzen. "Cardiac Resynchronization: Insight from Experimental and Computational Models". *Progress in Biophysics and Molecular Biology*, Vol. 97, No. 2, pp. 543–561, June 2008.
- [Klin 11] M. E. Klingensmith, A. Aziz, A. Bharat, A. C. Fox, and M. R. Porembka. *The Washington Manual of Surgery*. Lippincott Williams & Wilkins, Philadelphia, PA, USA, 6th Ed., Dec. 2011.
- [Kohl 02] B.-U. Kohler, C. Hennig, and R. Orglmeister. "The Principles of Software QRS Detection". *IEEE Engineering in Medicine and Biology Magazine*, Vol. 21, No. 1, pp. 42–57, Aug. 2002.
- [Konu 11] E. Konukoglu, J. Relan, U. Cilingir, B. H. Menze, P. Chinchapatnam, A. Jadidi, H. Cochet, M. Hocini, H. Delingette, P. Jaïs, M. Haïssaguerre, N. Ayache, and M. Sermesant. "Efficient Probabilistic Model Personalization Integrating Uncertainty on Data and Parameters: Application to Eikonal-Diffusion Models in Cardiac Electrophysiology". *Progress in Biophysics and Molecular Biology*, Vol. 107, No. 1, pp. 134–146, Oct. 2011.
- [Kris 11] M. Kristan, A. Leonardis, and D. Skočaj. "Multivariate Online Kernel Density Estimation with Gaussian Kernels". *Pattern Recognition*, Vol. 44, No. 10, pp. 2630–2642, Oct. 2011.
- [Kris 13] A. Krishnamurthy, C. T. Villongco, J. Chuang, L. R. Frank, V. Nigam, E. Belezzuoli, P. Stark, D. E. Krummen, S. Narayan, J. H. Omens, A. D. McCulloch, and R. C. P. Kerckhoffs. "Patient-specific Models of Cardiac Biomechanics". *Journal of Computational Physics*, Vol. 244, pp. 4–21, July 2013.
- [Kuij 12] N. H. Kuijpers, E. Hermeling, P. H. Bovendeerd, T. Delhaas, and F. W. Prinzen. "Modeling Cardiac Electromechanics and Mechanoelectrical Coupling in Dyssynchronous and Failing Hearts". *Journal of Cardiovascular Translational Research*, Vol. 5, No. 2, pp. 159–169, Apr. 2012.
- [Kuma 17] V. Kumar, A. K. Abbas, and J. C. Aster. *Robbins Basic Pathology E-Book*. Elsevier Health Sciences, 10th Ed., March 2017.
- [Kvet 12] B. Kveton and G. Theodorou. "Kernel-Based Reinforcement Learning on Representative States". In: *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, pp. 977–983, July 2012.
- [Laga 98] J. C. Lagarias, J. A. Reeds, M. H. Wright, and P. E. Wright. "Convergence Properties of the Nelder-Mead Simplex Method in Low Dimensions". *SIAM Journal on Optimization*, Vol. 9, No. 1, pp. 112–147, 1998.

- [Lama 14] P. Lamata, R. Casero, V. Carapella, S. A. Niederer, M. J. Bishop, J. E. Schneider, P. Kohl, and V. Grau. “Images as Drivers of Progress in Cardiac Computational Modelling”. *Progress in Biophysics and Molecular Biology*, Vol. 115, No. 2, pp. 198–212, Aug. 2014.
- [Le F 13] L. Le Folgoc, H. Delingette, A. Criminisi, and N. Ayache. “Current-based 4D Shape Analysis for the Mechanical Personalization of Heart Models”. In: *International MICCAI Workshop on Medical Computer Vision*, pp. 283–292, Springer, Berlin, Heidelberg, Germany, Oct. 2013.
- [Le 15] M. Lê, H. Delingette, J. Kalpathy-Cramer, E. R. Gerstner, T. Batchelor, J. Unkelbach, and N. Ayache. “Bayesian Personalization of Brain Tumor Growth Model”. In: *Medical Image Computing and Computer-Assisted Intervention*, pp. 424–432, Springer International Publishing, Cham, Switzerland, Oct. 2015.
- [LeGr 95] I. J. LeGrice, B. Smaill, L. Chai, S. Edgar, J. Gavin, and P. J. Hunter. “Laminar Structure of the Heart: Ventricular Myocyte Arrangement and Connective Tissue Architecture in the Dog”. *American Journal of Physiology – Heart and Circulatory Physiology*, Vol. 269, No. 2, pp. 571–582, Aug. 1995.
- [Liao 17] R. Liao, S. Miao, P. de Tournemire, S. Grbic, A. Kamen, T. Mansi, and D. Comaniciu. “An Artificial Agent for Robust Image Registration”. In: *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, pp. 4168–4175, AAAI Press, San Francisco, CA, USA, Feb. 2017.
- [Lin 93] L.-J. Lin. *Reinforcement Learning for Robots Using Neural Networks*. PhD thesis, Carnegie-Mellon University, School of Computer Science, Pittsburgh, PA, USA, Jan. 1993.
- [Litj 17] G. Litjens, T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghafoorian, J. A. van der Laak, B. van Ginneken, and C. I. Sánchez. “A Survey on Deep Learning in Medical Image Analysis”. *Medical Image Analysis*, Vol. 42, pp. 60–88, Dec. 2017.
- [Lloy 10] D. Lloyd-Jones, R. J. Adams, T. M. Brown, M. Carnethon, S. Dai, G. De Simone, T. B. Ferguson, E. Ford, K. Furie, C. Gillespie, *et al.* “Heart Disease and Stroke Statistics – 2010 Update: A Report from the American Heart Association”. *Circulation*, Vol. 121, No. 7, pp. 46–215, Feb. 2010.
- [Lomb 12] H. Lombaert, J.-M. Peyrat, P. Croisille, S. Rapacchi, L. Fanton, F. Chriet, 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, pp. 1436–1447, July 2012.
- [Lope 06] A. D. Lopez, C. D. Mathers, M. Ezzati, D. T. Jamison, and C. J. Murray. “Global and Regional Burden of Disease and Risk Factors, 2001: Systematic Analysis of Population Health Data”. *The Lancet*, Vol. 367, No. 9524, pp. 1747–1757, June 2006.
- [Luo 91] C.-h. Luo and Y. Rudy. “A Model of the Ventricular Cardiac Action Potential. Depolarization, Repolarization, and their Interaction”. *Circulation Research*, Vol. 68, No. 6, pp. 1501–1526, June 1991.

- [Mace06a] A. M. Maceira, S. K. Prasad, M. Khan, and D. J. Pennell. “Reference Right Ventricular Systolic and Diastolic Function Normalized to Age, Gender and Body Surface Area from Steady-State Free Precession Cardiovascular magnetic resonance”. *European Heart Journal*, Vol. 27, No. 23, pp. 2879–2888, Dec. 2006.
- [Mace06b] A. Maceira, S. Prasad, M. Khan, and D. Pennell. “Normalized Left Ventricular Systolic and Diastolic Function by Steady State Free Precession Cardiovascular Magnetic Resonance”. *Journal of Cardiovascular Magnetic Resonance*, Vol. 8, No. 3, pp. 417–426, Feb. 2006.
- [Mang18] K. Mangion, H. Gao, D. Husmeier, X. Luo, and C. Berry. “Advances in Computational Modelling for Personalised Medicine after Myocardial Infarction”. *Heart*, Vol. 104, No. 107, pp. 550–557, Apr. 2018.
- [Mans10] T. Mansi. *Image-based Physiological and Statistical Models of the Heart: Application to Tetralogy of Fallot*. PhD thesis, École Nationale Supérieure des Mines de Paris, Paris, France, Sep. 2010.
- [Marc13] S. Marchesseau, H. Delingette, M. Sermesant, R. Cabrera-Lozoya, C. Tobon-Gomez, P. Moireau, R. M. Figueras i Ventura, K. Lekadir, A. Hernandez, M. Garreau, E. Donal, C. Leclercq, S. G. Duckett, K. Rhode, C. A. Rinaldi, A. F. Frangi, R. Razavi, D. Chapelle, and N. Ayache. “Personalization of a Cardiac Electromechanical Model using Reduced Order Unscented Kalman Filtering From Regional Volumes”. *Medical Image Analysis*, Vol. 17, No. 7, pp. 816–829, Oct. 2013.
- [Marz07] Y. M. Marzouk, H. N. Najm, and L. A. Rahn. “Stochastic Spectral Methods for Efficient Bayesian Solution of Inverse Problems”. *Journal of Computational Physics*, Vol. 224, No. 2, pp. 560–586, June 2007.
- [McMu12] J. McMurray, S. Adamopoulos, S. Anker, A. Auricchio, K. Dickstein, V. Falk, G. Filippatos, C. Fonseca, and M. Gomez-Sanchez. “ESC Guidelines for the Diagnosis and Treatment of Acute and Chronic Heart Failure”. *European Heart Journal*, Vol. 33, No. 14, pp. 1787–1847, Aug. 2012.
- [McQu00] D. M. McQueen and C. S. Peskin. “A Three-dimensional Computer Model of the Human Heart for Studying Cardiac Fluid Dynamics”. *ACM SIGGRAPH Computer Graphics*, Vol. 34, No. 1, pp. 56–60, Feb. 2000.
- [Miha09] V. Mihalef, D. Metaxas, M. Sussman, V. Hurmusiadis, and L. Axel. “Atrioventricular Blood Flow Simulation Based on Patient-specific Data”. In: *International Conference on Functional Imaging and Modeling of the Heart*, pp. 386–395, Springer, Berlin, Heidelberg, Germany, June 2009.
- [Mill07] K. Miller, G. Joldes, D. Lance, and A. Wittek. “Total Lagrangian Explicit Dynamics Finite Element Algorithm for Computing Soft Tissue Deformation”. *Communications in Numerical Methods in Engineering*, Vol. 23, No. 2, pp. 121–134, Aug. 2007.
- [Mira16] G. R. Mirams, P. Pathmanathan, R. A. Gray, P. Challenor, and R. H. Clayton. “Uncertainty and Variability in Computational and Mathematical Models of Cardiac Physiology”. *The Journal of Physiology*, pp. 6833–6847, March 2016.

- [Mitc 03] C. C. Mitchell and D. G. Schaeffer. “A Two-Current Model for the Dynamics of Cardiac membrane”. *Bulletin of Mathematical Biology*, Vol. 65, No. 5, pp. 767–793, Sep. 2003.
- [Mnih 15] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumara, D. Wierstra, S. Legg, and D. Hassabis. “Human-level Control through Deep Reinforcement Learning”. *Nature*, Vol. 518, No. 7540, pp. 529–533, Feb. 2015.
- [Moir 09] P. Moireau. *Assimilation de Données par Filtrage pour les Systèmes Hyperboliques du Second Ordre. Applications à la Mécanique Cardiaque*. PhD thesis, École Nationale Supérieure des Mines de Paris, Paris, France, Feb. 2009.
- [Moli 12] D. Molina and V. DiMaio. “Normal Organ Weights in Men: Part I – The Heart”. *The American Journal of Forensic Medicine and Pathology*, pp. 362–367, Dec. 2012.
- [Moli 15] D. Molina and V. DiMaio. “Normal Organ Weights in Women: Part I – The Heart”. *The American Journal of Forensic Medicine and Pathology*, pp. 176–181, Sep. 2015.
- [Moll 15] R. Molléro, D. Neumann, M.-M. Rohé, M. Datar, H. Lombaert, N. Ayache, D. Comaniciu, O. Ecabert, M. Chinali, G. Rinelli, *et al.* “Propagation of Myocardial Fibre Architecture Uncertainty on Electromechanical Model Parameter Estimation: A Case Study”. In: *International Conference on Functional Imaging and Modeling of the Heart*, pp. 448–456, Springer, Cham, Switzerland, June 2015.
- [Moll 16] R. Molléro, X. Pennec, H. Delingette, N. Ayache, and M. Sermesant. “A Multiscale Cardiac Model for Fast Personalisation and Exploitation”. In: *Medical Image Computing and Computer-Assisted Intervention*, pp. 448–456, Springer, Cham, Switzerland, Oct. 2016.
- [Moll 18] R. Molléro, X. Pennec, H. Delingette, A. Garny, N. Ayache, and M. Sermesant. “Multifidelity-CMA: a Multifidelity Approach for Efficient Personalisation of 3D Cardiac Electromechanical Models”. *Biomechanics and Modeling in Mechanobiology*, Vol. 17, No. 1, pp. 285–300, Feb. 2018.
- [Mood 15] W. Moody, N. Edwards, C. Chue, R. Taylor, C. Ferro, J. Townsend, and R. Steeds. “Variability in Cardiac MR Measurement of Left Ventricular Ejection Fraction, Volumes and Mass in Healthy Adults: Defining a Significant Change at 1 Year”. *The British Journal of Radiology*, Vol. 88, No. 1049, pp. 1–4, May 2015.
- [Mull 13] K. Mülling, J. Kober, O. Kroemer, and J. Peters. “Learning to Select and Generalize Striking Movements in Robot Table Tennis”. *The International Journal of Robotics Research*, Vol. 32, No. 3, pp. 263–279, March 2013.
- [Murr 97] C. J. Murray and A. D. Lopez. “Mortality by Cause for Eight Regions of the World: Global Burden of Disease Study”. *The Lancet*, Vol. 349, No. 9061, pp. 1269–1276, May 1997.

- [Nagu 62] J. Nagumo, S. Arimoto, and S. Yoshizawa. “An Active Pulse Transmission Line Simulating Nerve Axon”. *Proceedings of the IRE*, Vol. 50, No. 10, pp. 2061–2070, Oct. 1962.
- [Nair 15] A. Nair, P. Srinivasan, S. Blackwell, C. Alcicek, R. Fearon, A. De Maria, V. Panneershelvam, M. Suleyman, C. Beattie, S. Petersen, S. Legg, V. Mnih, K. Kavukcuoglu, and D. Silver. “Massively Parallel Methods for Deep Reinforcement Learning”. *arXiv:1507.04296*, 2015.
- [Neum 14a] D. Neumann, S. Grbic, T. Mansi, I. Voigt, J. Rabbah, A. Siefert, N. Saikrishnan, A. Yoganathan, D. Yuh, and R. Ionasec. “Multi-modal Pipeline for Comprehensive Validation of Mitral Valve Geometry and Functional Computational Models”. In: *Statistical Atlases and Computational Models of the Heart*, pp. 188–195, Springer, Cham, Switzerland, Oct. 2014.
- [Neum 14b] D. Neumann, T. Mansi, B. Georgescu, A. Kamen, E. Kayvanpour, A. Amr, F. Sedaghat-Hamedani, J. Haas, H. Katus, B. Meder, J. Hornegger, and D. Comaniciu. “Robust Image-Based Estimation of Cardiac Tissue Parameters and Their Uncertainty from Noisy Data”. In: *Medical Image Computing and Computer-Assisted Intervention*, pp. 9–16, Springer International Publishing, Cham, Switzerland, Sep. 2014.
- [Neum 14c] D. Neumann, T. Mansi, S. Grbic, I. Voigt, B. Georgescu, E. Kayvanpour, A. Amr, F. Sedaghat-Hamedani, J. Haas, H. Katus, B. Meder, J. Hornegger, A. Kamen, and D. Comaniciu. “Automatic Image-to-Model Framework for Patient-specific Electromechanical Modeling of the Heart”. In: *IEEE 11th International Symposium on Biomedical Imaging*, pp. 935–938, IEEE, Beijing, China, Apr. 2014.
- [Neum 15a] D. Neumann, S. Grbic, M. John, N. Navab, J. Hornegger, and R. Ionasec. “Probabilistic Sparse Matching for Robust 3D/3D Fusion in Minimally Invasive Surgery”. *IEEE Transactions on Medical Imaging*, Vol. 34, No. 1, pp. 49–60, Jan. 2015.
- [Neum 15b] D. Neumann, T. Mansi, L. Itu, B. Georgescu, E. Kayvanpour, F. Sedaghat-Hamedani, J. Haas, H. Katus, B. Meder, S. Steidl, J. Hornegger, and D. Comaniciu. “Vito – A Generic Agent for Multi-physics Model Personalization: Application to Heart Modeling”. In: *Medical Image Computing and Computer-Assisted Intervention*, pp. 442–449, Springer International Publishing, Cham, Switzerland, Oct. 2015.
- [Neum 16] D. Neumann, T. Mansi, L. Itu, B. Georgescu, E. Kayvanpour, F. Sedaghat-Hamedani, A. Amr, J. Haas, H. Katus, B. Meder, S. Steidl, J. Hornegger, and D. Comaniciu. “A Self-taught Artificial Agent for Multi-Physics Computational Model Personalization”. *Medical Image Analysis*, Vol. 34, pp. 52–64, Dec. 2016.
- [Nguy 11] D. Nguyen-Tuong and J. Peters. “Model Learning for Robot Control: a Survey”. *Cognitive Processing*, Vol. 12, No. 4, pp. 319–340, Nov. 2011.
- [Nied 08] S. A. Niederer and N. P. Smith. “An Improved Numerical Method for Strong Coupling of Excitation and Contraction Models in the Heart”. *Progress in Biophysics and Molecular Biology*, Vol. 96, No. 1, pp. 90–111, Jan. 2008.

- [Nied 09] S. Niederer, K. Rhode, R. Razavi, and N. Smith. “The Importance of Model Parameters and Boundary Conditions in Whole Organ Models of Cardiac Contraction”. In: *International Conference on Functional Imaging and Modeling of the Heart*, pp. 348–356, Springer, Berlin, Heidelberg, Germany, June 2009.
- [Nied 12] S. A. Niederer, A. Shetty, G. Plank, J. Bostock, R. Razavi, N. Smith, and C. Rinaldi. “Biophysical Modeling to Simulate the Response to Multisite Left Ventricular Stimulation using a Quadripolar Pacing Lead”. *Pacing and Clinical Electrophysiology*, Vol. 35, No. 2, pp. 204–214, Feb. 2012.
- [Nied 16] S. A. Niederer and N. P. Smith. “Using Physiological Based Models for Clinical Translation, Predictive Modelling, Data Interpretation or Something In-between?”. *The Journal of Physiology*, pp. 6849–6863, Dec. 2016.
- [Nobl 02] D. Noble. “Modeling the Heart – from Genes to Cells to the Whole Organ”. *Science*, Vol. 295, No. 5560, pp. 1678–1682, March 2002.
- [Nobl 62] 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, pp. 317–352, Feb. 1962.
- [Nobl 98] D. Noble, A. Varghese, P. Kohl, and P. Noble. “Improved Guinea-Pig Ventricular Cell Model incorporating a Diadic Space, IKr and IKs, and Length- and Tension-dependent Processes”. *Canadian Journal of Cardiology*, Vol. 14, No. 1, pp. 123–134, Jan. 1998.
- [Okad 17] J.-i. Okada, T. Washio, M. Nakagawa, M. Watanabe, Y. Kadooka, T. Kariya, H. Yamashita, Y. Yamada, S.-i. Momomura, R. Nagai, *et al.* “Multi-scale, Tailor-made Heart Simulation can Predict the Effect of Cardiac Resynchronization Therapy”. *Journal of Molecular and Cellular Cardiology*, Vol. 108, pp. 17–23, July 2017.
- [Peng 16] P. Peng, K. Lekadir, A. Gooya, L. Shao, S. E. Petersen, and A. F. Frangi. “A Review of Heart Chamber Segmentation for Structural and Functional Analysis Using Cardiac Magnetic Resonance Imaging”. *Magnetic Resonance Materials in Physics, Biology and Medicine*, Vol. 29, No. 2, pp. 155–195, Apr. 2016.
- [Peyr 07] 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, pp. 1500–1514, Nov. 2007.
- [Phei 17] T. Pheiffer, D. Soto-Iglesias, Y. Nikulin, T. Passerini, J. Krebs, M. Sitges, A. Berruezo, O. Camara, and T. Mansi. “Estimation of Local Conduction Velocity from Myocardium Activation Time: Application to Cardiac Resynchronization Therapy”. In: *International Conference on Functional Imaging and Modeling of the Heart*, pp. 239–248, Springer, Berlin, Heidelberg, Germany, June 2017.
- [Poni 16] P. Ponikowski, A. A. Voors, S. D. Anker, H. Bueno, J. G. Cleland, A. J. Coats, V. Falk, J. R. González-Juanatey, V.-P. Harjola, E. A. Jankowska, *et al.* “2016 ESC Guidelines for the Diagnosis and Treatment of Acute and Chronic Heart Failure”. *European Heart Journal*, pp. 2129–2200, Aug. 2016.

- [Pop 09] M. Pop, M. Sermesant, T. Mansi, E. Crystal, J. Detsky, Y. Yang, P. Fefer, E. R. McVeigh, A. Dick, N. Ayache, *et al.* “Characterization of Post-Infarct Scars in a Porcine Model – a Combined Experimental and Theoretical Study”. In: *International Conference on Functional Imaging and Modeling of the Heart*, pp. 1–10, Springer, Berlin, Heidelberg, Germany, June 2009.
- [Powe 09] M. J. Powell. “The BOBYQA Algorithm for Bound Constrained Optimization without Derivatives”. Tech. Rep., University of Cambridge, Cambridge, UK, June 2009.
- [Prak 14] A. Prakosa, M. Sermesant, P. Allain, N. Villain, C. Rinaldi, K. Rhode, R. Razavi, H. Delingette, and N. Ayache. “Cardiac Electrophysiological Activation Pattern Estimation from Images using a Patient-Specific Database of Synthetic Image Sequences”. *IEEE Transactions on Biomedical Engineering*, Vol. 61, No. 2, pp. 235–245, Feb. 2014.
- [Quic 16] S. Quicken, W. P. Donders, E. M. van Disseldorp, K. Gashi, B. M. Mees, F. van de Vosse, R. G. Lopata, T. Delhaas, and W. Huberts. “Application of an Adaptive Polynomial Chaos Expansion on Computationally Expensive Three-dimensional Cardiovascular Models for Uncertainty Quantification and Sensitivity Analysis”. *Journal of Miomechanical Engineering*, Vol. 138, No. 12, pp. 1–11, Nov. 2016.
- [Rapa 12] S. Rapaka, T. Mansi, B. Georgescu, M. Pop, G. A. Wright, A. Kamen, and D. Comaniciu. “LBM-EP: Lattice-Boltzmann Method for Fast Cardiac Electrophysiology Simulation from 3D Images”. In: *Medical Image Computing and Computer-Assisted Intervention*, pp. 33–40, Springer, Berlin, Heidelberg, Germany, Sep. 2012.
- [Rasm 03] C. E. Rasmussen, J. Bernardo, M. Bayarri, J. Berger, A. Dawid, D. Heckerman, A. Smith, and M. West. “Gaussian Processes to Speed Up Hybrid Monte Carlo for Expensive Bayesian Integrals”. In: *Bayesian Statistics 7*, pp. 651–659, Oxford University Press, Oxford, UK, Jan. 2003.
- [Rela 11] J. Relan, P. Chinchapatnam, M. Sermesant, K. Rhode, M. Ginks, H. Delingette, C. Rinaldi, R. Razavi, and N. Ayache. “Coupled Personalization of Cardiac Electrophysiology Models for Prediction of Ischaemic Ventricular Tachycardia”. *Interface Focus*, Vol. 1, No. 3, pp. 396–407, March 2011.
- [Rice 08] J. J. Rice, F. Wang, D. M. Bers, and P. P. De Tombe. “Approximate Model of Cooperative Activation and Crossbridge Cycling in Cardiac Muscle using Ordinary Differential Equations”. *Biophysical Journal*, Vol. 95, No. 5, pp. 2368–2390, Sep. 2008.
- [Rose 60] H. H. Rosenbrock. “An Automatic Method for Finding the Greatest or Least Value of a Function”. *The Computer Journal*, Vol. 3, No. 3, pp. 175–184, Jan. 1960.
- [Sant 13] G. Santulli. “Epidemiology of Cardiovascular Disease in the 21st Century: Updated Numbers and Updated Facts”. *Journal of Cardiovascular Disease*, Vol. 1, No. 1, pp. 1–2, July 2013.
- [Seeg 15] P. Seegerer, T. Mansi, M.-P. Jolly, D. Neumann, B. Georgescu, A. Kamen, E. Kayvanpour, A. Amr, F. Sedaghat-Hamedani, J. Haas, H. Katus, B. Meder, and D. Comaniciu. “Estimation of Regional Electrical Properties of the Heart from 12-Lead ECG and Images”. In: *Statistical*

Atlases and Computational Models of the Heart, pp. 204–212, Springer, Cham, Switzerland, Jan. 2015.

- [Seif 09] S. Seifert, A. Barbu, S. K. Zhou, D. Liu, J. Feulner, M. Huber, M. Suehling, A. Cavallaro, and D. Comaniciu. “Hierarchical Parsing and Semantic Navigation of Full Body CT Data”. In: *SPIE Medical Imaging*, pp. 1–4, International Society for Optics and Photonics, March 2009.
- [Serm 06a] M. Sermesant, H. Delingette, and N. Ayache. “An Electromechanical Model of the Heart for Image Analysis and Simulation”. *IEEE Transactions on Medical Imaging*, Vol. 25, No. 5, pp. 612–625, May 2006.
- [Serm 06b] M. Sermesant, P. Moireau, O. Camara, J. Sainte-Marie, R. Andriantsimiavona, R. Cimrman, D. L. Hill, D. Chapelle, and R. Razavi. “Cardiac Function Estimation from MRI using a Heart Model and Data Assimilation: Advances and Difficulties”. *Medical Image Analysis*, Vol. 10, No. 4, pp. 642–656, June 2006.
- [Serm 07] M. Sermesant, E. Konukoglu, H. Delingette, Y. Coudière, P. Chinchapatnam, K. S. Rhode, R. Razavi, and N. Ayache. “An Anisotropic Multi-Front Fast Marching Method for Real-time Simulation of Cardiac Electrophysiology”. In: *International Conference on Functional Imaging and Modeling of the Heart*, pp. 160–169, Springer, Berlin, Heidelberg, Germany, June 2007.
- [Serm 09] M. Sermesant, F. Billet, R. Chabiniok, T. Mansi, P. Chinchapatnam, P. Moireau, J.-M. Peyrat, K. Rhode, M. Ginks, P. Lambiase, S. Arridge, H. Delingette, M. Sorine, C. A. Rinaldi, D. Chapelle, R. Razavi, and N. Ayache. “Personalised Electromechanical Model of the Heart for the Prediction of the Acute Effects of Cardiac Resynchronisation Therapy”. In: *International Conference on Functional Imaging and Modeling of the Heart*, pp. 239–248, Springer, Berlin, Heidelberg, Germany, June 2009.
- [Serm 12] M. Sermesant, R. Chabiniok, P. Chinchapatnam, T. Mansi, F. Billet, P. Moireau, J. Peyrat, K. Wong, J. Relan, and K. Rhode. “Patient-specific Electromechanical Models of the Heart for the Prediction of Pacing Acute Effects in CRT: A Preliminary Clinical Validation”. *Medical Image Analysis*, Vol. 16, No. 1, pp. 201–215, Jan. 2012.
- [Shou 09] G. Shou, L. Xia, M. Jiang, Q. Wei, F. Liu, and S. Crozier. “Solving the ECG Forward Problem by Means of Standard H-and H-hierarchical Adaptive Linear Boundary Element Method: Comparison with Two Refinement Schemes”. *IEEE Transactions on Biomedical Engineering*, Vol. 56, No. 5, pp. 1454–1464, May 2009.
- [Silv 09] D. U. Silverthorn. *Human Physiology: an Integrated Approach*. Pearson/Benjamin Cummings, San Francisco, CA, USA, 5th Ed., Oct. 2009.
- [Ster 95] N. Stergiopoulos, J. Meister, and N. Westerhof. “Evaluation of Methods for Estimation of Total Arterial Compliance”. *American Journal of Physiology – Heart and Circulatory Physiology*, Vol. 268, No. 4, pp. 1540–1548, Apr. 1995.
- [Sutt 98] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, USA, 1st Ed., 1998.

- [Talb 14] H. Talbot. *Interactive Patient-Specific Simulation of Cardiac Electrophysiology*. PhD thesis, Université des Sciences et Technologies de Lille, Lille, France, Dec. 2014.
- [Ten 04] K. Ten Tusscher, D. Noble, P. Noble, and A. Panfilov. “A Model for Human Ventricular Tissue”. *American Journal of Physiology – Heart and Circulatory Physiology*, Vol. 286, No. 4, pp. 1573–1589, Apr. 2004.
- [Tesa 95] G. Tesauro. *TD-Gammon: A Self-Teaching Backgammon Program*, pp. 267–285. Springer Science+Business Media, Boston, MA, USA, 1995.
- [Thal 10] M. S. Thaler. *The Only EKG book You’ll Ever Need*. Lippincott Williams & Wilkins, Philadelphia, PA, USA, May 2010.
- [Tous 13] N. Toussaint, C. T. Stoeck, T. Schaeffter, S. Kozerke, M. Sermesant, and P. G. Batchelor. “In Vivo Human Cardiac Fibre Architecture Estimation using Shape-based Diffusion Tensor Processing”. *Medical Image Analysis*, Vol. 17, No. 8, pp. 1243–1255, Dec. 2013.
- [Tray 11] N. A. Trayanova. “Whole-heart Modeling Applications to Cardiac Electrophysiology and Electromechanics”. *Circulation Research*, Vol. 108, No. 1, pp. 113–128, Jan. 2011.
- [Tu 05] Z. Tu. “Probabilistic Boosting-Tree: Learning Discriminative Models for Classification, Recognition, and Clustering”. In: *IEEE International Conference on Computer Vision*, pp. 1589–1596, IEEE, Beijing, China, Dec. 2005.
- [Urba 01] M. G. Urbanchek, E. B. Picken, L. K. Kalliainen, and W. M. Kuzon. “Specific Force Deficit in Skeletal Muscles of Old Rats is Partially Explained by the Existence of Denervated Muscle Fibers”. *The Journals of Gerontology Series A: Biological Sciences and Medical Sciences*, Vol. 56, No. 5, pp. 191–197, May 2001.
- [Van 11] G. Van Rossum and F. L. Drake. *The Python Language Reference Manual*. Network Theory Ltd., 2011.
- [Van 12] H. Van Hasselt. “Reinforcement Learning in Continuous State and Action Spaces”. In: *Reinforcement Learning: State-of-the-Art*, pp. 207–251, Springer, Berlin, Heidelberg, Germany, Oct. 2012.
- [Verc 09] T. Vercauteren, X. Pennec, A. Perchant, and N. Ayache. “Diffeomorphic Demons: Efficient Non-parametric Image Registration”. *NeuroImage*, Vol. 45, No. 1, pp. 61–72, March 2009.
- [Vett 98] F. J. Vetter and A. D. McCulloch. “Three-dimensional Analysis of Regional Cardiac Function: a Model of Rabbit Ventricular Anatomy”. *Progress in Biophysics and Molecular Biology*, Vol. 69, No. 2, pp. 157–183, March 1998.
- [Vodo 14] Y. Vodovotz and G. An. *Translational Systems Biology: Concepts and Practice for the Future of Biomedical Research*. Academic Press, Oxford, UK, 1st Ed., Oct. 2014.
- [Wall 12] M. Wallman, N. P. Smith, and B. Rodriguez. “A Comparative Study of Graph-based, Eikonal, and Monodomain Simulations for the Estimation of Cardiac Activation Times”. *IEEE Transactions on Biomedical Engineering*, Vol. 59, No. 6, pp. 1739–1748, June 2012.

- [Wall 14] M. Wallman, N. P. Smith, and B. Rodriguez. “Computational Methods to Reduce Uncertainty in the Estimation of Cardiac Conduction Properties from Electroanatomical Recordings”. *Medical Image Analysis*, Vol. 18, No. 1, pp. 228–240, Jan. 2014.
- [Wang 13] Y. Wang, B. Georgescu, T. Chen, W. Wu, P. Wang, X. Lu, R. Ionasec, Y. Zheng, and D. Comaniciu. “Learning-based Detection and Tracking in Medical Imaging: a Probabilistic Approach”. In: *Deformation Models*, pp. 209–235, Springer, Dordrecht, The Netherlands, Oct. 2013.
- [Wang 15] V. Wang, P. Nielsen, and M. Nash. “Image-based Predictive Modeling of Heart Mechanics”. *Annual Review of Biomedical Engineering*, Vol. 17, pp. 351–383, Dec. 2015.
- [Webe 10] F. M. Weber, C. Schilling, G. Seemann, A. Luik, C. Schmitt, C. Lorenz, and O. Dössel. “Wave-Direction and Conduction-Velocity Analysis from Intracardiac Electrograms: a Single-Shot Technique”. *IEEE Transactions on Biomedical Engineering*, Vol. 57, No. 10, pp. 2394–2401, Oct. 2010.
- [West 71] N. Westerhof, G. Elzinga, and P. Sipkema. “An Artificial Arterial System for Pumping Hearts”. *Journal of Applied Physiology*, Vol. 31, No. 5, pp. 776–781, Nov. 1971.
- [Wigg 24] C. J. Wiggers. “Modern Aspects of the Circulation in Health and Disease”. *The Canadian Medical Association Journal*, Vol. 14, No. 2, pp. 185–186, Feb. 1924.
- [Wong 15] K. C. Wong, M. Sermesant, K. Rhode, M. Ginks, C. A. Rinaldi, R. Razavi, H. Delingette, and N. Ayache. “Velocity-based Cardiac Contractility Personalization from Images using Derivative-free Optimization”. *Journal of the Mechanical Behavior of Biomedical Materials*, Vol. 43, pp. 35–52, March 2015.
- [Wu 09] M.-T. Wu, M.-Y. M. Su, Y.-L. Huang, K.-R. Chiou, P. Yang, H.-B. Pan, T. G. Reese, V. J. Wedeen, and W.-Y. I. Tseng. “Sequential Changes of Myocardial Microstructure in Patients Postmyocardial Infarction by Diffusion-Tensor Cardiac MR Correlation with Left Ventricular Structure and Function”. *Circulation: Cardiovascular Imaging*, Vol. 2, No. 1, pp. 32–40, Jan. 2009.
- [Xi 11] J. Xi, P. Lamata, J. Lee, P. Moireau, D. Chapelle, and N. Smith. “Myocardial Transversely Isotropic Material Parameter Estimation from In-silico Measurements based on a Reduced-order Unscented Kalman Filter”. *Journal of the Mechanical Behavior of Biomedical Materials Materials*, Vol. 4, No. 4, pp. 1090–1102, 2011.
- [Xi 13] J. Xi, P. Lamata, S. Niederer, S. Land, W. Shi, X. Zhuang, S. Ourselin, S. G. Duckett, A. K. Shetty, C. A. Rinaldi, D. Rueckert, R. Razavi, and N. P. Smith. “The Estimation of Patient-specific Cardiac Diastolic Functions from Clinical Measurements”. *Medical Image Analysis*, Vol. 17, No. 2, pp. 133–146, Feb. 2013.
- [Yang 08] L. Yang, B. Georgescu, Y. Zheng, P. Meer, and D. Comaniciu. “3D Ultrasound Tracking of the Left Ventricle using One-Step Forward Prediction and Data Fusion of Collaborative Trackers”. In: *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–8, IEEE, Anchorage, AK, USA, Aug. 2008.

- [Zett 13] O. Zettinig, T. Mansi, B. Georgescu, S. Rapaka, A. Kamen, J. Haas, K. S. Frese, F. Sedaghat-Hamedani, E. Kayvanpour, A. Amr, *et al.* “From Medical Images to Fast Computational Models of Heart Electromechanics: an Integrated Framework Towards Clinical Use”. In: *International Conference on Functional Imaging and Modeling of the Heart*, pp. 249–258, Springer, Berlin, Heidelberg, Germany, June 2013.
- [Zett 14] O. Zettinig, T. Mansi, D. Neumann, B. Georgescu, S. Rapaka, P. Seegerer, E. Kayvanpour, F. Sedaghat-Hamedani, A. Amr, J. Haas, H. Steen, H. Katus, B. Meder, N. Navab, A. Kamen, and D. Comaniciu. “Data-Driven Estimation of Cardiac Electrical Diffusivity from 12-Lead ECG Signals”. *Medical Image Analysis*, Vol. 18, No. 8, pp. 1361–1376, Dec. 2014.
- [Zhen 07] Y. Zheng, A. Barbu, B. Georgescu, M. Scheuering, and D. Comaniciu. “Fast Automatic Heart Chamber Segmentation from 3D CT Data using Marginal Space Learning and Steerable Features”. In: *IEEE International Conference on Computer Vision*, pp. 1–8, IEEE, Rio de Janeiro, Brazil, Oct. 2007.
- [Zhen 08] 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, pp. 1668–1681, Nov. 2008.
- [Zhen 14] Y. Zheng and D. Comaniciu. “Marginal Space Learning”. In: *Marginal Space Learning for Medical Image Analysis*, pp. 25–65, Springer Science+Business Media, New York, NY, USA, March 2014.

