
Multilevel Methods and Uncertainty Quantification in Cardiac Electrophysiology

Doctoral Dissertation submitted to the
Faculty of Informatics of the Università della Svizzera Italiana
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy

presented by
Seif Ben Bader

under the supervision of
Prof. Rolf Krause

July 2021

Dissertation Committee

Prof. Michael Multerer Università della Svizzera Italiana, Switzerland
Prof. Stefan Wolf Università della Svizzera Italiana, Switzerland
Prof. Helmut Harbrecht University of Basel, Switzerland
Prof. Luca Pavarino Università di Pavia, Italy

Dissertation accepted on July 2021



Research Advisor

Prof. Rolf Krause

PhD Program Director

Prof. Walter Binder

I certify that except where due acknowledgement has been given, the work presented in this thesis is that of the author alone; the work has not been submitted previously, in whole or in part, to qualify for any other academic award; and the content of the thesis is the result of work which has been carried out since the official commencement date of the approved research program.



Seif Ben Bader
Lugano, July 2021

Lorsqu'un jour le peuple veut
vivre, force est pour le destin de
répondre. Force est pour les
ténèbres de se dissiper, force est
pour les chaînes de se briser ...

When the people will to live,
destiny must surely respond.
Oppression shall then vanish, and
fetters are certain to break ...

Se un giorno il popolo vorrà
vivere, il destino deve
assecondarlo. La notte deve
dissiparsi e le catene devono
spezzarsi ...

Abou el Kacem Chebbi

Abstract

The objective of this thesis is to build an efficient and powerful computational framework for addressing large scale Uncertainty Quantification (UQ) problems arising in cardiac electrophysiology.

Cardiac models for reproducing the electrophysiological activity may suffer from a significant amount of uncertainties. These may be associated to the inherent physical and geometrical parameters of the model, as much as they could be induced by the imprecision of the model itself (e.g. uncertainty in the ion channel model). We here rely exclusively on the monodomain equation that represents a widely used model for computing the activation map of the cardiac cells. This model is highly anisotropic, as it strongly depends on the heart fibers' directions.

In this thesis, we evaluate the effect of spatially correlated perturbations of the heart fibers on the statistics of the resulting activation map. We rely on sampling methods that entail solving a numerous amount of times the computationally expensive monodomain model. The feasibility of such an approach is highly dependent on the efficiency of the computational framework. For this purpose, we introduce a new methodology which relies on a close integration of multilevel Monte-Carlo methods, parallel iterative solvers and a space-time all-at-once approach. This design allows for a fully parallelized framework in space, time, and stochasticity. The following setting is further improved by taking advantage of space-time solutions of past samples in the initialization of the Newton's method, resulting in a tremendous reduction of iterations for convergence.

We numerically assess the described framework on three dimensional geometries. More precisely, we evaluate the convergence properties and compare the performances of classical Monte Carlo and quasi-Monte Carlo methods, as well as multilevel quadrature strategies. In particular, we consider multilevel Monte Carlo and multilevel quasi-Monte Carlo on hierarchies of nested meshes for the time-dependent non linear equation represented by the monodomain model. Finally, and as a synthesis of the following work, we use the reversed construction of the multilevel quadrature in order to adapt the multilevel (quasi-) Monte Carlo to the configuration of non-nested meshes of a realistic heart geometry.

Acknowledgements

My sincere thanks go in the first place to my advisor Prof. Rolf Krause. His supervision, his advice and his dedication made me progress and evolve greatly during the last few years. I am deeply thankful to him for the opportunity he offered me to work in a high-level scientific environment, and for always offering all the conditions of success.

I also wish to express my deep gratitude to my initial supervisor Dr. Alessio Quaglino. My collaboration with him was an extremely pleasant experience, as I would always enjoy his dedication, his brilliant ideas and his unconditional support. I can only feel thankful for all the time he dedicated to me, even after having left for exploring other horizons.

I wish to warmly thank my final supervisor Prof. Michael Multerer. I thank him for insisting on details and quality when all around performance mattered the most. I thank him for all the time he managed to find even when in rush, for all the precious discussions and for proof-reading of this thesis.

I would also like to thank Prof. Helmut Harbrecht. Our collaboration had been intensely beneficial to me. Our meetings around Basel and Lugano, with Marc Schmidlin, were always challenging and extremely helpful. I also deeply thank the other members of my thesis committee Prof. Luca Pavarino and Prof. Stefan Wolf, for having cheerfully accepted to be part of it and dedicate some of their time to its assessment.

I would like further to greatly thank Dr. Patrick Zulian for his great help when it came to technical support and further scientific collaboration. I thank Dr. Pietro Benedusi for the collective work and the great input. I thank all my institute colleagues for the friendly working atmosphere, scientific help and further great friendships. A special thanks to the Decanato team for making my life easier, and particularly to Marisa Clemenz whom kindness and support had no bounds.

Finally, to those always supportive standing on the other side of the Mediterranean sea, to the unique her, to these two friends that years only made us closer, to this brother God sent me from nowhere right in Lugano and to the Gandria family, I send all my love.

Contents

Contents	ix
List of Figures	xiii
List of Tables	xvii
1 Introduction	1
1.1 Brief history of the cardiac research	2
1.2 Modern cardiac research in clinical practice	3
1.3 Computational mathematics in cardiac electrophysiology	4
1.4 Uncertainty Quantification (UQ) motivation	5
1.5 Work contributions and organization	6
2 Heart Electrophysiology & Mathematical Modelling	9
2.1 Heart anatomy and functioning	9
2.2 Mathematical modeling of cardiac electrophysiology	12
2.2.1 Cellular level models	13
2.2.2 Tissue level models	16
2.3 Medical and mathematical tools interaction in clinical practice	18
2.3.1 ECG and the PQRST cycle	18
2.3.2 Forward problem	19
2.3.3 Inverse problem	21
2.4 UQ motivation in cardiac electrophysiology	23
2.4.1 General sources of uncertainty	23
2.4.2 Fiber uncertainty	24
3 Discretization	27
3.1 Heat equation	30
3.1.1 Finite Elements in space	32
3.1.2 Finite Difference in time	38

3.1.3	Space–time approach	40
3.2	The Monodomain equation	42
3.2.1	Weak formulation	42
3.2.2	Newton’s method	43
3.2.3	The linear problem	46
3.2.4	Space–time assembly of the linear problem	47
3.3	Comment on the error estimates	48
4	Uncertainty Quantification	51
4.1	Foundations for UQ	51
4.1.1	Random fields	52
4.1.2	The Karhunen–Loève expansion	54
4.1.3	Discretization of the eigenvalue problem	55
4.1.4	Low rank approximation	57
4.1.5	Truncation error	59
4.1.6	The covariance in brief	62
4.2	Problem setting	70
4.2.1	Stochastic formulation	70
4.2.2	Well-posedness	71
4.2.3	Parametric formulation	73
4.3	Single-level sampling methods	74
4.3.1	Monte Carlo (MC)	74
4.3.2	Quasi–Monte Carlo (QMC)	75
4.4	Multilevel sampling methods	78
4.4.1	Multilevel Monte Carlo (MLMC)	78
4.4.2	Multilevel inverse construction	79
4.4.3	Multilevel Quasi–Monte Carlo (MLQMC)	80
4.5	Convergence properties	81
4.5.1	Discretization error	81
4.5.2	Stochastic error	82
4.5.3	General error	85
4.6	Sampling strategy	88
4.7	Work comparison	91
4.8	Transfer of discrete fields	94
4.8.1	Nested meshes	94
4.8.2	Non–nested meshes	95

5	Uncertainty Quantification for Heart Electrophysiology	97
5.1	Fiber uncertainty	97
5.2	Scar fibrosis uncertainty	101
5.3	Quantities-of-interest	102
5.4	Prior insights on numerical experiments	104
6	Parallelization and Performance Considerations	109
6.1	Preconditioned Generalized Minimal Residual Method (PGMRES)	109
6.2	Parallelization strategy	110
6.2.1	Concurrent tasks and parallelization paradigm	110
6.2.2	Parallelization algorithm	111
6.3	Resources management	112
6.3.1	Discussion over algorithm speed	112
6.3.2	Discussion over simultaneous termination	113
6.3.3	HPC oriented job scheduling strategy	113
6.4	Newton initial guess	115
6.4.1	Multiple time blocks strategy	115
6.4.2	Eikonal initial guess	116
6.4.3	Samples based initial guess	117
6.4.4	Performance comparison	117
7	Numerical Experiments	121
7.1	Setup	121
7.2	Isotropic random diffusion for simple geometries and nested meshes	123
7.2.1	Cube	123
7.2.2	Idealized ventricle	125
7.3	Anisotropic random diffusion for a complex geometry on non-nested meshes	128
7.3.1	Action potential	129
7.3.2	Activation time	133
7.4	Scar uncertainty	136
8	Conclusion	141
A	Short review of normed vector spaces	143
B	Basic notions of measure theory	146
B.1	Measurable spaces	146
B.2	Probability and Bochner spaces	147
B.3	Lebesgue Measure and Integral on \mathbb{R}	148

C Sobolev Spaces	151
Bibliography	155

Figures

1.1	Statistics on death causes worldwide. Published online at Our-WorldInData.org, Source: Institute for Health Metrics and Evaluation.	1
2.1	Portrait of the rib cage demonstrating the heart positioning.	9
2.2	Representation of the veins and arteries of the heart.	10
2.3	Different components of the heart on a vertical cross section.	11
2.6	Representation of the PQRST signal.	19
2.7	A schematization of the forward problem: The simulated model provides solution that is mapped on the torso domain. This is used for reconstructing an ECG signal.	20
2.8	Schematization of the inverse problem: The measured ECG is mapped on the pericardial surface.	23
2.9	UQ-based general pipeline of computational cardiology assistance in clinical practice.	25
2.10	Transmural coordinate computed by means of the Laplace problem (2.22) for a synthetic heart geometry.	26
2.11	Mathematical reconstruction of the fiber vector field surrounding the geometry shown in Figure 2.10	26
3.1	Heat propagation in a cubical domain.	31
3.2	Graphical illustration of Newton’s method through few iterations.	44
4.1	Effect of the correlation length for squared exponential kernel on the local variations. On the left, a random field field obtained for $\sigma = 0.1$. On the right, a random field field obtained for $\sigma = 0.5$	63
4.2	Eigenfunctions (left) and eigenvalues (right) of the squared Gaussian kernel weighted with normal density.	64
4.3	Example of isotropic ($d_r = 1$) random fields on more complex geometries.	65

4.4	Example of anisotropic ($d_r = d$) fiber fields.	65
4.5	Weighted isotropic eigenfunctions (left) and their contours (right) in decreasing order.	66
4.6	Weighted anisotropic eigenfunctions in decreasing order for a diagonal covariance matrix, i.e. $k_{st} = 0$ for $s \neq t$	67
4.7	Unscaled vector field variability arising from KL expansion.	69
4.8	Scaled vector field variability arising from KL expansion with a factor of $s_i = 2$ for each coordinate.	69
5.1	Example of first (mean, left) and centered second (variance, right) moments of the solution to (5.2).	98
5.2	Wafefront propagation of transmembrane potential.	103
5.3	Activation map of a heart. Negative value signifies non activated region.	104
5.4	L^2 -error convergence measured over space-time uniform refinement.	105
5.5	Mean (left) and variance (right) at the final time state for the cube.	105
5.6	Mean (left) and variance (right) at the final time state for the ventricle.	106
5.7	Mean at coarse level and corrections between successive levels for a cube geometry.	107
5.8	Mean at coarse level and corrections between successive levels for an idealized ventricle geometry.	107
5.9	Convergence rate for the MC estimator for a cube and an idealized ventricle geometries.	108
5.10	Convergence rate for the MC and QMC estimators for a cube geometry.	108
6.1	The different batches of samples executing concurrently organized through the range of levels.	114
6.2	MLMC tasks scheduling and job requests step. $T_{l,i}$ denotes the i -th task of level l	114
6.3	Schematization of the multiple time blocks strategy.	116
6.4	Comparison of the different Newton initial guess strategies for a sample on the cube.	118
6.5	Comparison of the different Newton initial guess strategies for a sample on the ventricle.	118
6.6	Comparison of the different Newton initial guess strategies for a sample on a realistic heart geometry.	119

7.1	Controlled convergence graphs for the cube in L^2 (left) and H^1 (right) norms.	124
7.2	Work comparison for MC/MLMC/QMC/MLQMC. Work is computed on the basis of execution time on a single thread for all levels.	126
7.3	Nested mesh hierarchy for the idealized ventricle.	126
7.4	Controlled convergence graphs for the idealized ventricle in L^2 -norm.	127
7.5	Work comparison for MC/MLMC/QMC/MLQMC. Work is computed on the basis of execution time.	127
7.6	Initial state for fibers $\mathbb{E}[\mathbf{V}](\mathbf{x})$	128
7.7	Non-nested mesh hierarchy for the realistic heart.	129
7.8	Locations selected along the wall separating the left and right ventricles.	130
7.9	Action potential behaviour given different mesh level discretizations for the points specified in Figure 7.8.	130
7.10	Convergence in the H^1 -norm of the action potential at the locations specified in Figure 7.8.	131
7.11	Locations selected at the circumference of a horizontal cut of the heart surface.	132
7.12	Convergence of the action potential at the locations specified in Figure 7.11.	133
7.13	Locations selected at equivalent geodesic distance from the stimulus.	133
7.14	Convergence of the activation time at the locations specified in Figure 7.13.	134
7.15	Locations selected at the periphery of the heart surface.	135
7.16	Convergence of the activation time at the locations specified in Figure 7.15.	136
7.17	Mean (left) and variance (right) over time after the inclusion of a scar.	138
7.18	Controlled convergence graphs for the cube after inclusion of a scar in the L^2 -norm.	138
7.19	Work comparison for MC/MLMC/QMC/MLQMC. Work is computed on the basis of execution time.	139

Tables

6.1	Number of iterations required by each method for the considered geometries. The acronym NC stands for Non Converging.	119
7.1	Details about the considered mesh hierarchy for the cube geometry.	124
7.2	Details about the mesh hierarchy for the idealized ventricle geometry.	126
7.3	Details about the considered mesh hierarchy for the realistic heart geometry.	128

Chapter 1

Introduction

Due to its very central role, the heart plays the greatest partition in the human body symphony and the smallest misregulation of its activity might unluckily lead to dramatic consequences. As a matter of fact, cardiovascular diseases carve up the lion's share in terms of mortality causes and lethal cardiac incidents are reported to be responsible of about 17 millions deaths worldwide every year, c.f. Figure 1.1. Faced with this observation, the human ingenuity is left with no other option than to rival of creativity in order to comprehend the inherent complex mechanism of heart functioning, as a first step towards the implementation of medical treatments.

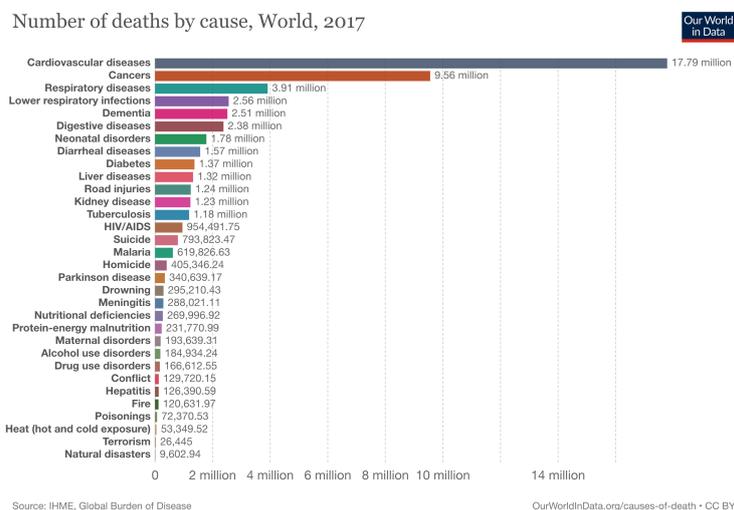


Figure 1.1. Statistics on death causes worldwide. Published online at OurWorldInData.org, Source: Institute for Health Metrics and Evaluation.

1.1 Brief history of the cardiac research

From the early stages of the 19th century up to nowadays, many major discoveries have been achieved contributing largely not only to the understanding of the mysterious heart operating mechanism, but also to the modern clinical practice as we know it today, with its increasingly complex screening instruments made at our disposal. The long pathway always needs to start somewhere and in 1842, the pioneer Dr. Carlo Matteucci (Italian professor of physics at the University of Pisa) was the first to record the heart electrical activity of a frog, see [Mat42]. His outstanding discovery was further brought to the human perspective, as Augustus Waller (British physiologist at St Mary's Medical School) used a capillary electrometer and electrodes to reveal for the first time in the history the electrical activity of a human heart in 1887, see [Wal87]. Greatly inspired by this experiment, Dr. Willem Einthoven (Dutch physiologist) took up the torch and distinguished five significant deviations in the curve of the electrical activity known until today as the PQRS sequence. This represented in a way the first mathematical intervention in electrophysiology, as one of Einthoven's addition consisted in a mathematical correction of the recorded electrical curve. The cardiac research community also owes Einthoven the terminology of electrocardiogram, shortened initially as EKG, to become the famous ECG in reference to its English orthography, see [Ein95]. Back then, the electrocardiograph used by Einthoven was a quite insolite machinery that weighted hundreds of kilos, where two cylinders of electrolyte solution constituted rather rustic electrodes. Combined to a third mouth electrode, this constituted the first three-lead approach in electrocardiography to be used in clinical practice, eventually after further improvement of the maneuverability of the initial electrocardiogram. The next five decades have witnessed a constant evolution of the apparel, in terms of clinical practicality and measuring efficiency. In 1942, Emanuel Goldberger (American cardiologist) adds three additional limb leads. Combined to the six pericordial leads introduced by Dr. Frank Wilson (American cardiologist) and the three initial leads of Einthoven, this constituted the 12-leads electrocardiogram as known today, see [AL12] for more details. Probably the most noticeable discovery that summarizes the effort made on the constant improvement of the measuring system was the invention brought by Normann Jeff Holter (American biophysician at the University of California) in 1949, consisting in a portable backpack device able to record the ECG and transmit the signal to a monitor. A device that

by now is as small as to hold in a pocket, equipped with digital recording.¹ In our modern days, the ECG remains an elementary yet powerful instrument for detecting electrical disorders. Nonetheless, the ECG alone might encounter limitations, being unable to provide the right diagnosis in some situations. The ECG needs therefore to be supplied with additional technologies.

1.2 Modern cardiac research in clinical practice

Historical research on electrocardiography, as shortly narrated above, was of a major importance as it provided scientists with the main interaction tool for a quantitative evaluation of patients hearts' electrical activity. On the way to the most modern and elaborate electrocardiograms of nowadays, and as the experimenting process progressed, the heart functioning mechanism was no longer a mystery and its role in the circulatory system became well understood. This has prepared the ground for new tracks of exploration for medical treatment prospects, since the ability of detecting heart dysfunctioning is a complete separate task from providing curative solutions.

While the responsibility of the healing step accounts in its major part to doctors, cardiologists and biologists, the improvement of screening capabilities in clinical practice definitely facilitate this process. Quantitative and qualitative screening becomes naturally the new adage, as one would like to picture the biggest set of information in the least invasive way.

In that context, cardiac imaging has opened a whole new era. With the arrival of Magnetic Resonance Imaging (MRI) and Cardiac Computed Tomography (CCT), the reach of non-invasive fully complete diagnosis had never become so close. The MRI constituted a revolution as it provided information in tissue characterization, while the CCT refashioned coronary artery disease detection as it allowed visualizing the coronary arteries without an invasive angiogram. For a more detailed resume on cardiac imaging and their different modalities, we refer the reader to [LP16].

Nonetheless, with all the nowadays techniques at hand, some heart diseases remain complicated to comprehend and to heal. Atrial fibrillation for example is still a very challenging phenomenon to combat on a patient-specific basis. In the meantime, the progress made in mathematical modeling, numerical analysis for partial differential equations and the growing computing capabilities opened the way for a new expedition. Mathematicians and computational scientists are

¹See <https://ecglibrary.com/ecghist.html> for a more detailed timeline on history of eletrocardiography

slowly but surely exploring the idea of patient-specific in silico virtual therapy. The whispered hope of a heart digital twin is not that silent anymore.

1.3 Computational mathematics in cardiac electrophysiology

All the dynamics governing the heart functioning are physiological phenomena and consequently mathematically modelable. Whether we refer to the heart movement or its blood pumping function, very well established models exist deriving from the fields of mechanics and hemodynamics. Regarding the electrophysiological aspect, much research around electrical potential propagation inside an excitable cell has been successfully conducted. This was at first initiated by [HH52] followed by [Fit61], who proposed a cellular model for describing a prototype of an excitable cell. Many detailed membrane models followed since then, with increasing precision in terms of number of currents and variables. From the local (heart cell) to the global (heart tissue) scale, [Hen93] summarized the electrical potential propagation inside the cardiac muscle, within the bidomain set of equations that differentiated the behavior of the intra and extracellular potentials. A more exhaustive and detailed description about models at the cellular and tissue level will be provided in Chapter 2.

From that moment on, investigating the heart dynamics through numerically solved mathematical equations got materialized, in what could be referred to as part of computational medicine. This research field became very well established during the two last decades, reaching already the clinical level in some cases. Several models for cardiac electrophysiology, mechanics and hemodynamics had been developed, with even remarkable coupling systems benchmarks. (see [NNN⁺11, Qua15, Tra11]).

Computational medicine can assist cardiology in many aspects, ranging from diagnosis to in silico virtual therapy. Inverse problems in cardiac electrophysiology can even improve data acquisition by inferring models' parameters, while a good data-model integration makes the conceit of digital twin feasible. The contribution of computational medicine to cardiological clinical practice is expected to grow given the continuous and tremendous advances realized in terms of growing computational resources.

1.4 Uncertainty Quantification (UQ) motivation

Despite the very promising aspects of computational medicine, severe hurdles stand in the way. The lack of data that models rely on is definitely a major issue in tailoring patient-specific therapies. We refer to this data deficit with the terminology of uncertainties.

The source of these uncertainties can be very broad and diverse. For instance, some model parameters are hardly quantifiable due to simplifications of the experimental procedure. Indeed, clinical experiments for estimating model input parameters may reveal in some cases to be either very costly, or quite invasive. Limitations of the measurement apparatus can also induce uncertainties, or simply fail to provide the needed information.

Another potential source of uncertainties may arise from the inter and intra-subject variability. Those refer to differences originating from an individual to another (inter-variability), but also among the same individual himself (intra-variability) as indeed humans' bodies represent dynamical systems. Probably the most evident and significant variability in that sense derives from the heart geometry, but one can also evoke uncertainties with respect to initial and boundary conditions.

Among all those possible sources of uncertainty, we are confronted to a major one whenever we consider the heart fibers. Those are comprised in its myocardium, and accomplish a fundamental task when it comes to diffuse the electrical potential all over the cardiac muscle. Unfortunately and even with all the modern technology at hand regarding CCT and MRI imaging, the fibers orientations as well as the conductivities they carry are extremely hard to determine precisely. Moreover, these can vary from one patient to another (inter-variability), leading to huge uncertainties if standard orientations or conductivities are used in patient-specific simulations. It is therefore imperative to quantify the effects of the fibers' uncertainties on the output quantities-of-interest.

Scope of the thesis

Mathematically, the fibers are modeled by means of a tensor diffusion field. In this work, we focus on modeling the fibers' uncertainties (orientation and conductivities) as a spatially-correlated random perturbation applied to the diffusion field. We concentrate exclusively on the monodomain model, that is a simplification of the bidomain model, slightly modified to contain uncertainty represented

as a stochastic variable. We evaluate the effect of these uncertainties on the statistics of the solution, that we approximate by means of stochastic quadrature methods. Those are typically sampling methods requiring many solves of the initial problem for various perturbations. In this context, a single simulation can already take a large amount of time even when considering important computational resources. It follows that, e.g., UQ for such models is currently unfeasible, or extremely inefficient, with plain Monte Carlo methods.

The central point of our approach is to consider the multilevel Monte Carlo method, see [Gil15], which extends the idea of control variates, see [Fis03]. The key idea is to perform most of the simulations on a sequence of low-resolution models. The direct use of the highest-resolution level guarantees convergence, while a significant portion of the computational load is offset to the low-resolution hierarchy. In order to allow for arbitrary coarsening in space and time, we further rely on a space-time discretization of the monodomain equation. By doing so, it also becomes possible to use the already computed samples for the initialization of Newton's method, which solves the space-time equations, thereby significantly reducing the amount of iterations required for convergence.

1.5 Work contributions and organization

Contributions of the thesis

The contributions of the following work stand at the interface of the two fields of cardiac electrophysiology and UQ. They are listed below:

1. We suggest an efficient and fully parallelized computational framework for addressing large scale uncertainty quantification problems arising in cardiac electrophysiology, with high-dimensional quantity of interest. By considering a parallel distribution of the total work among discretization levels, batches of samples, and domain decomposition technique in space and time for a given sample, we obtain a so-called three-layer parallelization in space, time and stochastics.
2. We propose a multiple time-blocks strategy for resolving the problem of non convergence of the Newton's method for long time intervals. Additionally, we take advantage of the computed samples and set them as Newton initial guesses on every single time-block of newer samples, accelerating remarkably the convergence in terms of iterations number. We furthermore consider the global solution in time of the computed samples, to be

set as a global Newton initial guess for future samples. This methodology proved to be very efficient, resulting in a convergence in few iterations, for long time intervals.

3. By applying known stochastic quadrature methods such as (quasi) Monte Carlo and their multilevel construction on the time-dependent non linear monodomain equation, we experimentally evaluate their convergence rates proving in a reverse-engineered way their compliance with the already established convergence properties for a Poisson problem.

Organization of the thesis

This thesis informally comprises three parts: mathematical modeling of cardiac electrophysiology, uncertainty quantification in cardiac electrophysiology and performance improvement considerations with a numerical validation.

More formally, Chapter 2 is dedicated to the description of the heart functioning, with a focus on the underlying electrophysiological aspect. We introduce the common used models for electrical potential simulation in the cardiac muscle, proceeding from the cell level to the tissue level.

In Chapter 3, we introduce the discretization in space and time with respectively finite elements and finite differences. We describe the all-at-once approach, demonstrating how the the large space-time system is assembled. Proceeding from the model problem represented by the heat equation, we describe step-by-step the Newton's method when a non linear term is integrated, expliciting every Newton iteration problem to be solved in the case of the monodomain equation.

In Chapter 4, we proceed to the UQ part by first introducing the foundations, such as random fields generation with Karhunen-Loève expansion, the discretization of its associated eigenvalue problem. We also evaluate the truncation error produced by the discretization and low rank approximation processes. We furthermore define the single and multilevel stochastic quadrature methods. We evaluate their theoretical convergence properties in the case of the heat equation model problem, and establish a theoretical comparison of the work produced by different methods. We also invoke the transfer of discrete fields for a multilevel setting.

In Chapter 5, we treat the uncertainty in the heart fibers, as the core part of this thesis. We additionally consider a second test case for a heart presenting a scar region, also modeled with an uncertainty in terms of width and conductivity. These uncertainties are combined to that of the fibers, in a test case for which numerical experiments will follow in Chapter 7.

In Chapter 6, we invoke all the parallelization concepts and performance improvements considerations used in our developed computational framework. Starting from the three-layer parallelization strategy, to the solver used for solving the space–time system, passing by the local and the global Newton initial guess strategy, we make a full report on every computational aspect of the thesis.

In Chapter 7, we finally provide numerical experiments to support our method's computational efficiency. Mainly, we always go through the same procedure of evaluating the controlled convergence of the invoked stochastic quadrature methods, as well as assessing their general effective work in terms of actual computation time.

Finally, appendices are provided in the very end for the purpose of reader's convenience. We present there basic notions on normed vector spaces (Appendix A), measure theory, Bochner spaces (Appendix B) and Sobolev spaces (Appendix C). This is done in a synthetic way without further details on the proofs of the presented lemmas, propositions and theorems. For a more in depth understanding, the reader is directed towards the following references, from which the appendices are inspired; [[Che08](#), [CK13](#), [Coh13](#)].

Disclaimers. This thesis is partly based on the reviewed paper,

- Ben Bader, S., Benedusi, P., Quaglino, A., Zulian, P., and Krause, R. (2021). Space–time multilevel Monte Carlo methods and their application to cardiac electrophysiology. *Journal of Computational Physics*, 433, 110164.

and the preprint

- Ben Bader, S., Harbrecht, H., Krause, R., Multerer, M., Quaglino, A., and Schmidlin, M. (2021). Space–time multilevel quadrature methods and their application for cardiac electrophysiology. *preprint arXiv:2105.02007*.

which were produced along the thesis.

Chapter 2

Heart Electrophysiology & Mathematical Modelling

2.1 Heart anatomy and functioning

Role and anatomy: The heart is by all means one of the most complex and fascinating organs in the human body. It constitutes a muscular organ, that is located in the middle compartment of the chest, protected by the rib cage (c.f. Figure 2.1¹).

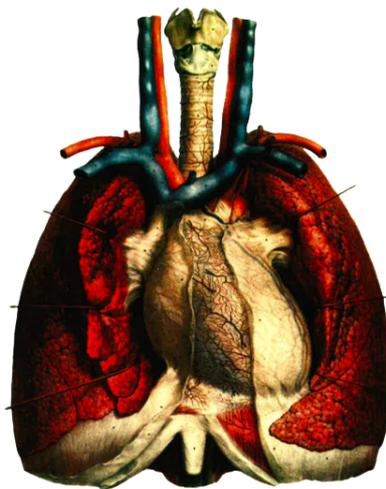


Figure 2.1. Portrait of the rib cage demonstrating the heart positioning.

¹After an artistic representation found in the Est Milan market. The artist remains unfortunately unknown to the thesis' author.

Its combination with the surrounding lungs constitutes a fundamental alliance for the body's well functioning. While the lungs are held responsible for the elimination of all the toxic substances carried by the blood, the heart is the central component of the circulatory system, harmoniously orchestrating the body activity through the vital supply of blood to all of its components. This is achieved by means of a multitude of veins and arteries connecting it to the entire body. These can be seen under different angles in Figure 2.2.²

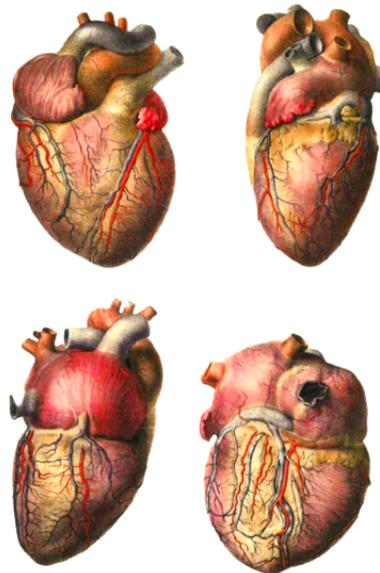


Figure 2.2. Representation of the veins and arteries of the heart.

The heart resembles a thin wall that comprises three distinct layers. The outer layer is called the pericardium and encompasses the heart and its big vessels (c.f. Figure 2.2). The inner layer is called the endocardium. The latter underlies under a much more voluminous middle layer, called the myocardium that represents the actual heart muscle. It is constituted by cardiac cells called the cardiomyocytes.

Heart functioning: The heart is made of four different chambers, two atria in the upper part and two ventricles in the lower part (c.f. Figure 2.3). The inferior vena cava and superior vena cava carry deoxygenated blood to the right atrium. When the latter becomes full, it compresses and the blood flows from the right atria to the right ventricle via the tricuspid valve. The right ventricle in turn contracts

²After an artistic representation found in the Est Milan market. The artist remains unfortunately unknown to the thesis' author.

and the blood is ejected into the pulmonary artery through the pulmonary valve. It is then directed to the lungs, where the blood releases the carbon dioxide and other toxic substances, eliminated by the breathing process. Once oxygenated, the blood is carried by means of the pulmonary veins to the left atria. Its passage to the right ventricle is made possible by the mitral valve. From there, it passes through the aortic valve and travels all over the body to supply its components. This represents a full heart cycle, from arriving deoxygenated blood to leaving oxygenated blood. This process overlaps as when a quantity of blood starts its journey in the right part of the heart, another one is arriving to the end in the left part. This is demonstrated in Figure 2.4. For a full review on the blood circulation system and many technical aspects about intervening components, see [QMV⁺19].

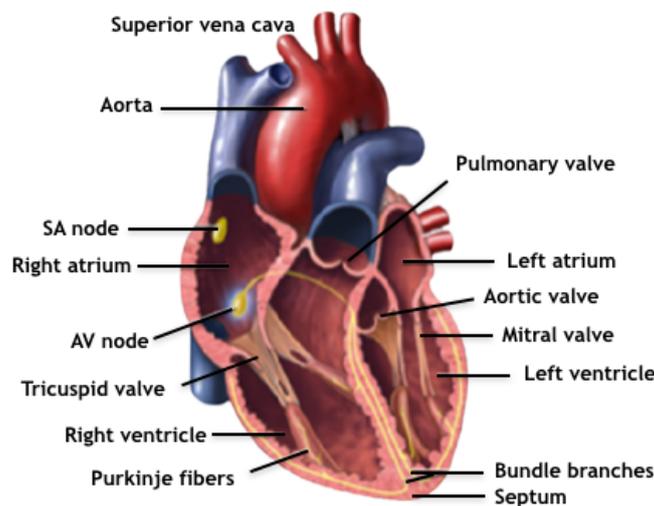


Figure 2.3. Different components of the heart on a vertical cross section.³

Electrical activity role: The heart functioning and the blood flowing from a chamber to another is made possible by a very complex contraction and relaxation cycle occurring in the cardiac cells. This mechanical pattern is induced by an electrical signal that originates at the sinoatrial (SA) node, diffusing along the two atria and causing them to contract. The impulse travels to the atrioventricular node (AV) node, that acts as a gate, delaying further electrical spreading in the ventricles, up until the atria fully released the blood. When this is achieved,

³Found in <https://www.empowher.com/media/reference/heart-murmur> and adapted. Copyright: Nucleus Medical Media, Inc.

the AV fires an electrical signal that travels to the septum by means of the bundle branches. From there, the signal is diffused to the whole myocardium through the Purkinje fibers, see [Kla11]. What follows from the electrical diffusion inside the heart muscle is a succession of opening and closing of the ion channels for every cell, which allows ions such as calcium and potassium to pass back and forth. This biomechanical process creates a transmembrane potential that causes a contraction and relaxation of every single cardiomyocyte. This occurs through the whole myocardium, in a travelling wave fashion, inducing the contraction and relaxation of the whole system. A more complete review can be found in [Kat10].

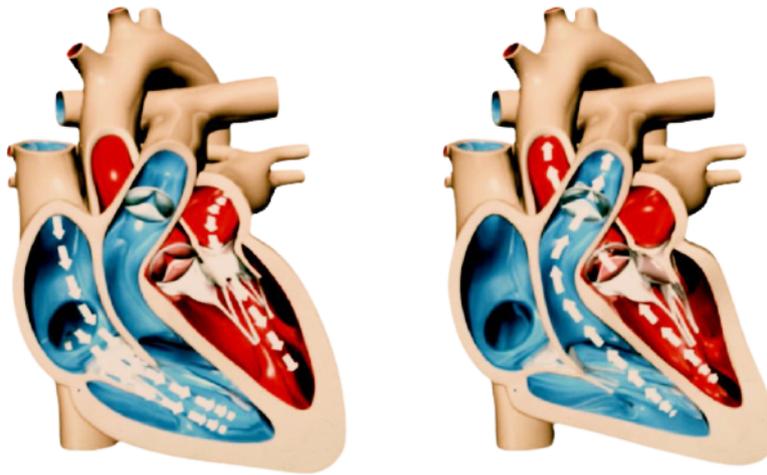


Figure 2.4. Simultaneous blood circulation inside the cardiac muscle. The left figure shows two simultaneous actions; the deoxygenated blood (blue) enters the right part of the heart while a similar process is occurring in the left part with oxygenated blood (red). The right figure shows the progression towards the exiting process.⁴

2.2 Mathematical modeling of cardiac electrophysiology

The heart functions in a great harmony of an electrical potential activity, a mechanical contraction-relaxation cycle, and fluid (blood) exchanges within the sur-

⁴Found in the educational videos provided by <http://www.bodhaguru.com> and adapted.

rounding environment, i.e. the different body organs.

The aim to model the heart behaviour therefore necessitates a close integration of electrical activity, mechanical deformation and fluid dynamics models. These components can however be addressed separately, following a classical divide-and-conquer strategy for solving sophisticated scientific problem. This work falls in this scope, focusing exclusively on some of the electrophysiological aspects.

When considering the phenomenon of cardiac electrophysiology, it is important to recall that local perturbations of the heart cells' action potential, lie at the basis of the electrical travelling wave along the myocardium. It is therefore necessary to differentiate between processes occurring at the microscopic and macroscopic levels, equivalently designated as cellular and tissue levels.

2.2.1 Cellular level models

Cellular models are mathematical models quantifying the membrane potential variation of an excitable cell. What is referred to as the membrane potential is the potential difference between intracellular and extracellular levels, that is

$$u = u_i - u_e,$$

where $u_i = u_i(t)$ and $u_e = u_e(t)$ are respectively intracellular and extracellular potentials.

The potential variation between intracellular and extracellular levels is induced by voltage-gated ionic channels opening and closing, allowing for ions to pass back and forth across the membrane, thus creating an electrical activity. [HH52] modeled the action potential time evolution in the squid giant axon. The particularity of their work consisted in treating each cell component as an electrical element inside a consistent circuit depicting the membrane architecture (in terms of ion channels). The notions of cell membrane capacitance and ion channel conductance derive from this perspective.

Characterized as such, models for simulating the action potential of a cell could be written in the following form, see [KS98],

$$C_m \frac{u(t)}{dt} + I_{\text{ion}}(u) = I_{\text{app}}(t), \quad (2.1)$$

where C_m is the membrane capacitance, I_{app} an external applied stimulus on the cell and $I_{\text{ion}}(u)$ the current arising from the passage of ions, in which the ion channels conductances are to be taken into account.

As an example, the ionic current generated by a cell membrane that contains a unique sodium ion channel can be expressed with the following model:

$$I_{\text{ion}}(u) = g_{\text{Na}}(u - u_{\text{Na}}), \quad (2.2)$$

where g_{Na} is the sodium ion channel conductance and u_{Na} is the sodium reversal potential. Expression (2.2) provides a model of the ionic current generated by a leak channel, characterized by a fixed conductance. Whenever considering a voltage-gated ion channel, the conductance becomes a function of time and voltage, therefore more challenging to model.

In many contemporary works, see [MS03, TTP06b, BOCF08], the variable conductance of voltage-gated ion channels was expressed by means of the so-called gate variables and the concentrations of the ions involved. The gate variables and the ion concentrations are themselves subject to interconnected time evolutionary processes. This implies that the ionic current provided by means of N ion channels can be summed up as, see [Poz19],

$$I_{\text{ion}}(u, \mathbf{w}, \mathbf{c}) = \sum_{k=1}^N I_{\text{ion},k}(u, \mathbf{w}, \mathbf{c}), \quad (2.3)$$

where for $k = 1, \dots, N$, we have that

$$I_{\text{ion},k}(u, \mathbf{w}, \mathbf{c}) = g_k(u, \mathbf{w}, \mathbf{c})(u - u_k), \quad (2.4)$$

for a vector $\mathbf{w} = [w_1, \dots, w_{N_g}]^T$ of N_g gated variables $w_i \in (0, 1)$ described by a set of Ordinary Differential Equations (ODEs), see [Nel04],

$$\frac{d\mathbf{w}}{dt} = f(u, \mathbf{w}, \mathbf{c}), \quad (2.5)$$

and a vector $\mathbf{c} = [c_1, \dots, c_{N_c}]^T$ representing the ionic concentrations inside the cell. Their dynamics is further described by another set of time ODEs, see e.g. [RCG⁺90, NVKN98],

$$\frac{d\mathbf{c}}{dt} = h(u, \mathbf{w}, \mathbf{c}), \quad (2.6)$$

The full cell model reads therefore

$$\begin{cases} C_m \frac{du(t)}{dt} + I_{\text{ion}}(u, \mathbf{w}, \mathbf{c}) = I_{\text{app}}(t), \\ \frac{d\mathbf{w}}{dt} = f(u, \mathbf{w}, \mathbf{c}), \\ \frac{d\mathbf{c}}{dt} = h(u, \mathbf{w}, \mathbf{c}). \end{cases} \quad (2.7)$$

Hence, the ionic current $I_{\text{ion}}(u)$ depends on the characteristics of the considered membrane and the constituent ion-channels. In the context of cardiomyocytes, a large family of ion channel models exist, see e.g. [FPT05, FPS14]. Among those, the Fitz–Hugh Nagumo (FHN) model, see [Fit61], was one of the first. It is a simplified model, still used in some cases for numerical experiments and validation. It has the great advantage of reproducing the action potential, without involving many other ionic species that might represent a computational burden. The FHN model is given by:

$$\begin{cases} C_m \frac{du(t)}{dt} + I_{\text{ion}}(u, w) = I_{\text{app}}(t), \\ \frac{\partial w}{\partial t} = \gamma u^* - \beta w, \end{cases} \quad (2.8)$$

where

$$I_{\text{ion}}(u, w) = \alpha(u - u_{\text{rest}})(u - u_{\text{th}})(u - u_{\text{peak}}) + \eta w(u - u_{\text{rest}}). \quad (2.9)$$

Here, α , η , γ and β are all positive parameters affecting the action potential general shape. The values u_{rest} , u_{th} and u_{peak} are characteristic potential values of the electrical activation process. They represent the resting, the threshold and the peak potential values characterizing respectively the initial, the activation and the maximal states the membrane goes through. Finally, we have

$$u^* = \frac{u - u_{\text{rest}}}{u_{\text{peak}} - u_{\text{rest}}}.$$

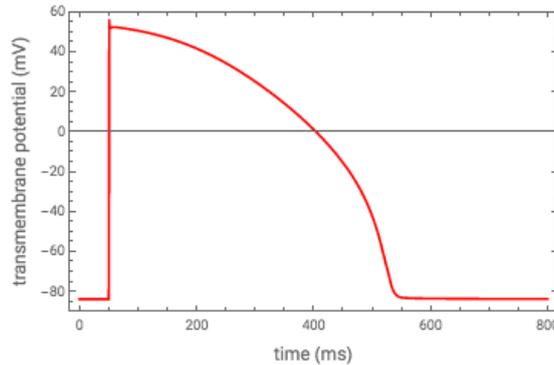


Figure 2.5. Action potential of a cardiomyocyte located at the mid-endocardium.⁵

⁵Copyright: Graph generated through WOLFRAM Demonstrations Project platform, licensed under the Creative Commons Attribution - NonCommercial - ShareAlike 3.0 Unported License.

Remark. *In presence of a stimulus important enough for activating the cell membrane (the cell potential reaches the threshold), the latter goes over the process of an instant depolarization characterized by a sudden raise of electrical potential up to a peak value. This is afterwards followed by a repolarization where the electrical potential slowly recovers its initial resting value. This process is called the action potential and is depicted in Figure 2.5.*

2.2.2 Tissue level models

Tissue models are concerned with the macro-scale electrical activity inside the myocardium as a whole. They consist on integrating the cell models, that describe each cell individually, into a continuum model represented by average properties of the considered group of cells constituting the tissue. They therefore incorporate the space variable additionally to the time variable and are consequently defined as coupled PDE systems.

The bidomain model: A very-well established tissue model consists of the so-called bidomain model, see [Tun78, Hen93, KS98]. It simulates the electrical potential activity inside the cardiac muscle (myocardium), differentiating the two interpenetrating domains consisting of the intracellular and extracellular domains. Each domain is characterized by an anisotropic conductivity tensor describing the different propagating directions and their intensities. These are typically reproducing the fibers and their orientations, see [LSC⁺95].

The unknowns for such a model are therefore naturally two, namely the intracellular $v_i = v_i(\mathbf{x}, t)$ and the extracellular $v_e = v_e(\mathbf{x}, t)$ potentials. However, it is generally defined with respect to the extracellular potential and the transmembrane potential by setting the intracellular potential to, see [PBC05],

$$u_i = u + u_e. \quad (2.10)$$

Based in the conservation of total current, see [PBC05], the bidomain model set of equations can be expressed likewise $\forall(\mathbf{x}, t) \in D \times (0, T]$:

$$\begin{cases} \chi \left(C_m \frac{\partial u}{\partial t} + I_{\text{ion}}(u, \mathbf{w}, \mathbf{c}) \right) - \nabla \cdot (G_i \nabla u) - \nabla \cdot (G_i \nabla u_e) = \chi I_{\text{app}}, \\ -\nabla \cdot ((G_i + G_e) \nabla u_e) - \nabla \cdot (G_i \nabla u) = 0, \\ \frac{\partial \mathbf{w}}{\partial t} = f(u, \mathbf{w}, \mathbf{c}), \\ \frac{\partial \mathbf{c}}{\partial t} = h(u, \mathbf{w}, \mathbf{c}), \end{cases} \quad (2.11)$$

where $G_i = G_i(\mathbf{x})$ and $G_e = G_e(\mathbf{x})$ are respectively intracellular and extracellular conductivity tensors, χ the surface-to-volume ratio, D the heart tissue domain

and $[0, T]$ a given time interval. Notice that we abbreviated with $u = u(\mathbf{x}, t)$, $u_e = u_e(\mathbf{x}, t)$, $\mathbf{w} = \mathbf{w}(\mathbf{x}, t)$, $\mathbf{c} = \mathbf{c}(\mathbf{x}, t)$ and $I_{\text{app}} = I_{\text{app}}(\mathbf{x}, t)$ for the sake of readability. The lasting parameters are exactly as defined in (2.8), with respect to cellular models. The system of equation further encodes homogeneous Neumann boundary conditions given by, see [Tun78]:

$$\begin{cases} G_i \nabla(u + u_e) \cdot \mathbf{n} = 0, \\ (G_i \nabla u + (G_i + G_e) \nabla u_e) \cdot \mathbf{n} = 0, \end{cases} \quad (2.12)$$

for every $(\mathbf{x}, t) \in \partial D \times (0, T]$. Here, $\mathbf{n} = \mathbf{n}(\mathbf{x})$ is the regular surface outer-normal usually considered in boundary conditions. The set of equation is further provided adequate initial conditions:

$$\begin{cases} u(\mathbf{x}, 0) = u_0(\mathbf{x}), & \mathbf{x} \in D, \\ u_e(\mathbf{x}, 0) = u_{e,0}(\mathbf{x}), & \mathbf{x} \in D, \\ \mathbf{w}(\mathbf{x}, 0) = \mathbf{w}_0(\mathbf{x}), & \mathbf{x} \in D, \\ \mathbf{c}(\mathbf{x}, 0) = \mathbf{c}_0(\mathbf{x}), & \mathbf{x} \in D. \end{cases} \quad (2.13)$$

The bidomain model represents a well-established model, yet complex and computationally expensive. A *relative* simplification can be provided by considering instead the monodomain model.

The monodomain model: Assuming that the following expression

$$G_e(\mathbf{x}) = \lambda G_i(\mathbf{x}), \quad \forall \mathbf{x} \in D, \quad \lambda \in \mathbb{R},$$

relates the two conductivity tensors (i.e. a proportional anisotropy ratio), the second equation of (2.11) can be rewritten as, see [PBC05],

$$\nabla \cdot (G_i \nabla u_e) = -\frac{1}{1 + \lambda} \nabla \cdot (G_i \nabla u).$$

This induces that:

$$\begin{aligned} -\nabla \cdot (G_i \nabla u) - \nabla \cdot (G_i \nabla u_e) &= -\nabla \cdot (G_i \nabla u) + \frac{1}{1 + \lambda} \nabla \cdot (G_i \nabla u) \\ &= -\frac{\lambda}{1 + \lambda} \nabla \cdot (G_i \nabla u) \end{aligned}$$

Therefore by setting $G_m = \frac{\lambda}{1 + \lambda} G_i$, we recover the monodomain model given by the following set of equations:

$$\left\{ \begin{array}{l} \chi \left(C_m \frac{\partial u}{\partial t} + I_{\text{ion}}(u, \mathbf{w}, \mathbf{c}) \right) - \nabla \cdot (G_m \nabla u) = \chi I_{\text{app}}, \\ \frac{\partial \mathbf{w}}{\partial t} = f(u, \mathbf{w}, \mathbf{c}), \quad \frac{\partial \mathbf{c}}{\partial t} = h(u, \mathbf{w}, \mathbf{c}), \\ G_m \nabla u \cdot \mathbf{n} = 0, \\ u(\mathbf{x}, 0) = u_0(\mathbf{x}), \\ \mathbf{w}(\mathbf{x}, 0) = \mathbf{w}_0(\mathbf{x}), \quad \mathbf{c}(\mathbf{x}, 0) = \mathbf{c}_0(\mathbf{x}), \end{array} \right. \quad (2.14)$$

which holds for every $(\mathbf{x}, t) \in D \times (0, T]$.

Remark. *Anisotropy scaling assumption between intracellular and extracellular fields (i.e. $G_e(\mathbf{x}) = \lambda G_i(\mathbf{x})$ for some $\lambda \in \mathbb{R}$) is practically never achieved nor verifiable, yet it offers a considerable reduction in computational cost and memory with respect to the bidomain. Many studies demonstrated that the monodomain, depending on the final quantity of interest, remains reliable even considering that it is based on the non-necessarily holding assumption, see [PDR⁺06a].*

2.3 Medical and mathematical tools interaction in clinical practice

2.3.1 ECG and the PQRST cycle

In Chapter 1, we have largely talked about the ECG as the first clinically acquired measure, yet still one of the most important. It represents by far the least invasive procedure for recording the heart electrical activity, if we compare with other methods such as cardiac catheterization, for example, see [DB97].

The typical electrical signal recorded by the ECG is represented by the PQRST cycle. The latter characterizes electrically a heart beat. A PQRST recorded through a healthy heart demonstrates the following distinctive phases:

- The P-wave: it corresponds to the atria depolarization as soon as they are full of blood (SA activation).
- The P-Q segment: it corresponds to when the signal is travelling from the SA to the AV node.
- The QRS-complex: marks the AV node firing and the ventricles depolarization. It comprises itself three distinctive phases:
 - (i) The Q-wave: depolarization of the inter ventricular septum

- (ii) The R–wave: main mass ventricles depolarization
- (iii) The S–wave: last phase of ventricular depolarization, also accompanied with atrial repolarization
- The S–T segment: the plateau in the action potential of the myocardium when ventricles contract to pump blood
- T-wave: characterizes ventricular repolarization after ventricles relaxation, i.e. the diastole

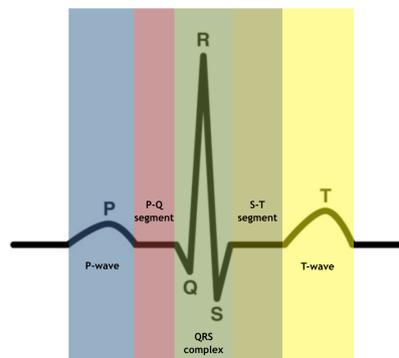


Figure 2.6. Representation of the PQRST signal.

Heart anomalies can, in some fortunate cases, already be detected from analyzing the ECG and every phase of its PQRST cycle. However, sometimes this might be invisible to the naked eye, even for experimented cardiologists. Resorting to computational techniques can in those cases be of a great use. So far however, we have only been concerned about simulating the electrical activity at the heart surface, which cannot be compared to the signal measured clinically on a torso. The forward problem can be solved for achieving this purpose, i.e. reconstructing ECG signals from tissue models simulations.

2.3.2 Forward problem

The monodomain/bidomain models simulate the electrical activity inside the cardiac muscle. This electrical activity further propagates inside the torso up to the body surface (c.f. Figure 2.7).

In the absence of external current applied on the body, the latter can be considered just as a static regular conductor. The conservation of currents therefore

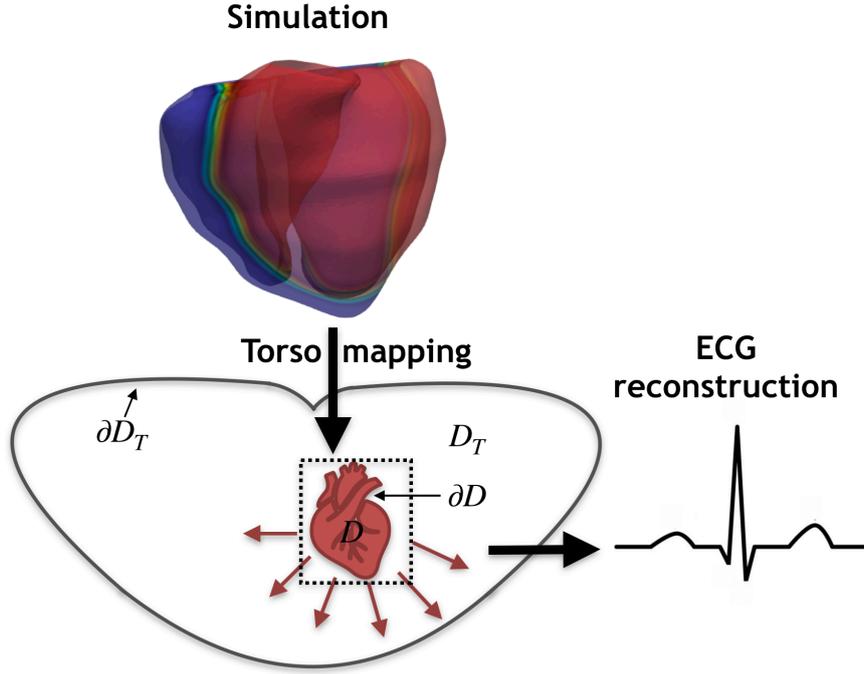


Figure 2.7. A schematization of the forward problem: The simulated model provides solution that is mapped on the torso domain. This is used for reconstructing an ECG signal.

applies and a natural boundary condition is to consider the simulated electrical activity on the heart surface. Accounting for the following statements, we recover the formulation of the forward problem, expressed by means of a simple Laplacian diffusion problem, with non-homogeneous Robin boundary conditions:

$$\begin{cases} \nabla \cdot (G_T \nabla u_T) = 0, & \forall (\mathbf{x}, t) \in D_T \times (0, T] \\ G_T \nabla u_T \cdot \mathbf{n}_T = 0, & \forall (\mathbf{x}, t) \in \partial D_T \times (0, T] \\ u_T = u_e, & \forall (\mathbf{x}, t) \in \partial D \times (0, T] \\ (G_e \nabla u_e + G_T \nabla u_T) \cdot \mathbf{n} = 0 & \forall (\mathbf{x}, t) \in \partial D \times (0, T], \\ u_T(\mathbf{x}, 0) = u_{T,0}(\mathbf{x}) & \forall \mathbf{x} \in D_T, \end{cases} \quad (2.15)$$

where D_T is the torso domain, $G_T = G_T(\mathbf{x})$ is the conductivity tensor inside the torso and $n_T = n_T(\mathbf{x})$ is the outer normal to the torso surface. The quantity of interest is $u_T = u_T(\mathbf{x}, t)$ is the electrical potential inside the torso, that can eventually be calculated at its surface. It therefore provides the potential at the body surface, reconstructing a signal similar to what is measured by the ECG,

see [BCF⁺10, FZ10]. The system of equations (2.15) can further be coupled with (2.11) in a compact form, see [LGT03].

2.3.3 Inverse problem

Combination of simulations at the heart tissue level and forward problem can serve for ECG reconstruction and might therefore be helpful for predictive behaviour of experimental treatment. However, for tailoring patient specific virtual therapies in a reliable manner, one step back has to be taken as initially model parameters need to be set correctly, or at least approximated well enough.

Model parameters can be inferred from clinical data by means of inverse problems. Originally, those were designed for achieving the exact reciprocal operation with respect to the forward problem, that is, mapping the torso potential measurements on the epicardial surface. Nowadays also referred to as the ECG Inverse problem (ECGI), this is applied in clinical practice as a non-invasive epicardial potential reconstruction method, used e.g. for preventing arrhythmias, see [RGJ⁺04, CZC⁺17]. Its formulation is very closely related to that of the forward problem, and can be expressed as, see [FGTV85],

$$\begin{cases} \nabla \cdot (G_T \nabla u_T) = 0, & \forall (\mathbf{x}, t) \in D_T \times (0, T] \\ G_T \nabla u_T \cdot \mathbf{n}_T = 0, & \forall (\mathbf{x}, t) \in \partial D_T \times (0, T] \\ u_T = u_{\text{meas}}, & \forall (\mathbf{x}, t) \in \partial \Sigma_{\text{meas}} \times (0, T] \\ u_T(\mathbf{x}, 0) = u_{T,0}(\mathbf{x}) & \forall \mathbf{x} \in D_T, \end{cases} \quad (2.16)$$

where u_{meas} and $\Sigma_{\text{meas}} \subset \partial D_T$ are respectively the measured potential and the measurement surface. We also recall that D is the heart domain, ∂D the heart surface, whereas D_T is the torso domain and ∂D_T its outer surface. The epicardial distribution can then be recovered by restricting the solution u_T on the heart surface, i.e. $u_T|_{\partial D}$.

The inverse problem (2.16) is well-known to be an ill-posed problem in the sense of Hadamard as its solution does not depend continuously on the data, see [FGTV85]. Consequently, a small perturbation, e.g. of the boundary condition such as that of the measured potential on the thorax surface, might lead to a great variation with respect to the solution of interest. Given the noise affecting the measurements, this scenario is very likely to occur, and therefore the inverse problem is generally formulated in an alternative manner.

Closely related to the inverse problem, the forward problem can be solved by computing a transfer matrix A , interpreted as an observation operator such that

$$Au_e|_{\partial D} = u_T|_{\partial D_T}, \quad (2.17)$$

where u_e is provided by solving the bidomain model for example. Similarly, the inverse problem can then also be reformulated as: provided u_{meas} , find $u_T|_{\partial D}$ such that

$$Au_T|_{\partial D} = u_{\text{meas}}. \quad (2.18)$$

By introducing a cost function $J(u) = \int_{\Sigma} |Au - u_{\text{meas}}|^2 d\mathbf{x}$, the solution can be recovered by solving the minimization problem consisting of finding u^* such that

$$u^* = \inf_u J(u). \quad (2.19)$$

The instability of this problem, straightforward consequence of the ill-posed nature of (2.16), can be overcome by means of regularization methods, such as Tikhonov or L^1 -regularization. For a review of the main regularization methods and their comparison, see [KBMZ18].

The formulation (2.19) further inspired model parameters estimation, when the cost functional is instead set to

$$J(v(\sigma)) = \int_{\Sigma} |v(\sigma) - u_{\text{meas}}|^2 d\mathbf{x}, \quad (2.20)$$

where $v(\sigma)$ is the simulated thorax potential with respect to the parameter to estimate σ . This formulation is known after the terminology of optimal control problems, where σ represents the control, i.e. the parameter to infer, and solving

$$\sigma^* = \inf_{\sigma} J(v(\sigma)), \quad (2.21)$$

is equivalent to answer the question of which value of σ results in the observed u_{meas} . Such optimal control problems can be used for instance in recovering the extracellular potential from solving the monodomain equation, see [NK11], termination of re-entry waves, see [NKP13], inverse identification of arterial wall material parameters, see [LMQR13], or detecting infarcted areas, see [CZC⁺17, BB15].

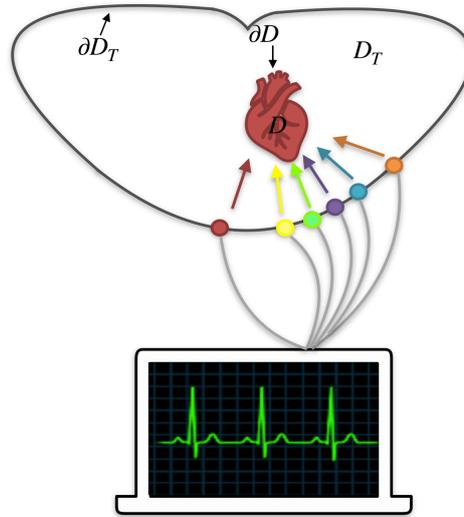


Figure 2.8. Schematization of the inverse problem: The measured ECG is mapped on the pericardial surface.

Remark. *It is also possible to consider a group of parameters to estimate, represented then by a vector σ . This is naturally a more demanding optimization problem.*

2.4 UQ motivation in cardiac electrophysiology

2.4.1 General sources of uncertainty

Forward and inverse problems represent the bridges connecting heart tissue models and clinical data. They allow the interaction between mathematical modeling and clinical practice. They are therefore essential since clinical acquired data represents evidently a crucial ingredient for tailored computational cardiology.

The measurement data is however generally affected by noise. This can occur for many reasons, among which one can cite the possible limited accuracy of the measuring device. These uncertainties, if not accounted for, might lead to great variations and possibly heavily impact on the accuracy of the obtained solution.

The integration of UQ in cardiac electrophysiology can therefore be motivated by the aforementioned observation. It translates to the integration of a

stochasticity for the considered uncertain parameters of the used model, leading a PDE formulation which queries about statistical moments of the solution, for given random distributions of the parameters of interest.

The estimation of the first moment of the solution, i.e. the expectation, can therefore provide a more general idea of the system behaviour, accounting for several selected values of the uncertain parameters. The information on the second moment, i.e. the variance, provides an insight of the solution variability with respect to a given parameter. This approach is related to sensitivity analysis, and allows to evaluate the random parameters of main impact on the solution model. This information can in turn be used for reducing the uncertainties dimensionality in an optimal control optimization for example, as considering a high number of control variables for those leads to more complex problems. Resorting to sensitivity analysis in these cases might constitute a great benefit.

The usage of UQ in electrophysiology is very broad. Whether employed for myocardial infarction detection, see [Pag17], or for quantification of the error induced by the imaging imprecision with respect to the torso MRI, see [GKMP20], UQ has been clearly adopted and identified to be essential to integrate inside a framework pipeline for in-silico virtual therapy. A full review on many other UQ usages and sources of uncertainties can be found in [Pag17].

2.4.2 Fiber uncertainty

Many sources of uncertainty can also find their root in the limitation of the mathematical model with respect to the phenomenon of interest, or the absence of measurements instruments for a particular aspect of it. As an example, this can be the case for heart fibers. As even with modern imaging techniques the fibers pattern can still not be precisely identified, one typically resorts to their mathematical construction, provided a given heart geometry.

The local fiber orientations are generally described by means of an orthonormal set of vectors, constituted by the fiber direction, the sheet direction and the sheet normal direction, see [LSC⁺95, LHS97, BPP⁺18]. These can be computed by means of the transmural coordinate $\rho = \rho(\mathbf{x})$ providing the distance from the endocardium to the epicardium, calculated for every spatial point $\mathbf{x} \in D$, through solving the Laplace problem, [BBPT12],

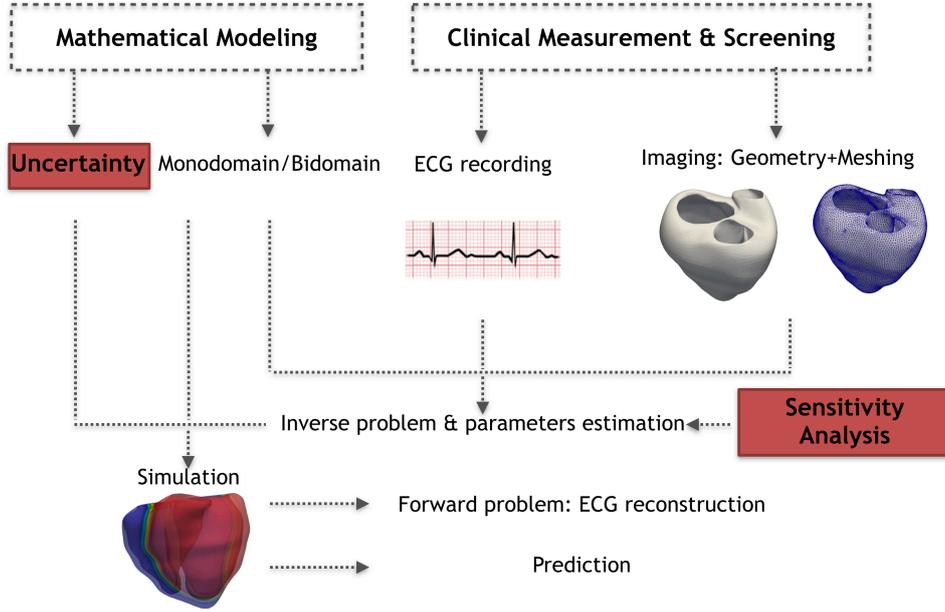


Figure 2.9. UQ-based general pipeline of computational cardiology assistance in clinical practice.

$$\begin{cases} \Delta \rho(\mathbf{x}) = 0, & \forall \mathbf{x} \in D, \\ \rho(\mathbf{x}) = -1, & \forall \mathbf{x} \in \partial LV, \\ \rho(\mathbf{x}) = 0, & \forall \mathbf{x} \in D, \\ \rho(\mathbf{x}) = 1, & \forall \mathbf{x} \in \partial RV, \end{cases} \quad (2.22)$$

where LV and RV denote respectively the left and right ventricles.

Given that the fibers demonstrate a helical orientation from endocardium to epicardium along the wall thickness, they are reconstructed by means of a function of their rotation with respect to the transmural coordinate, see [PDR⁺06b, Poz19]. An example of mathematical fiber construction is shown in Figure 2.11.

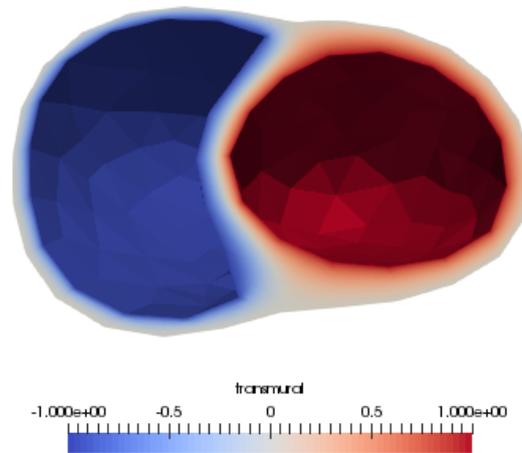


Figure 2.10. Transmural coordinate computed by means of the Laplace problem (2.22) for a synthetic heart geometry.

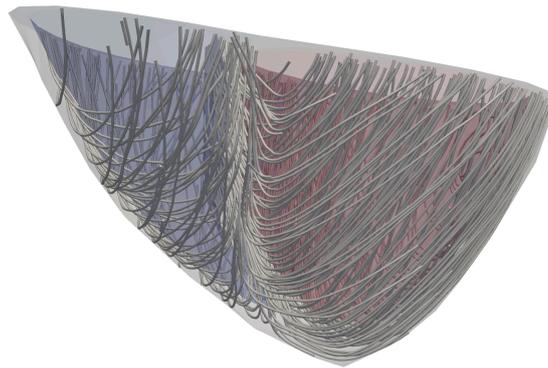


Figure 2.11. Mathematical reconstruction of the fiber vector field surrounding the geometry shown in Figure 2.10

The mathematical reconstruction of fibers may fail to reproduce the exact fiber pattern for the heart of a given patient. Moreover, fibers naturally exhibit an inter and intra patients variability. An alternative approach consists in considering fiber uncertainties and perform UQ.

Chapter 3

Discretization

The illusion of closed form solutions The vast majority of PDEs arising in natural or physiological phenomena may at best be very effortful to solve analytically, if not at all. This even generalizes to the set of PDEs deriving from general purpose mathematical modeling, as the less academical example of complex financial models. In general, given the very large definition of a PDE, as any partial derivatives combination of a quantity of interest in a closed form, it might be quite intuitive that there should not be any magic wand and that analytical solution for a PDE is more a matter of case by case.

Some combinations of smart techniques such as the separation of variables, the principle of superposition and the Fourier decomposition could be used in the past to find the solution of PDEs such as the heat equation under particular forms of initial conditions, see [bdF22]. The change of variables is also a famous technique for resolving some wave equations settings, see [d'A47]. Yet, the generalization of such methods to the wide set of PDEs with a larger combination of initial conditions remains undoubtedly unsatisfying.

The existence of an analytical solution can even depend on the parameters defining the same PDE problem. Let us consider the famous example of the Poisson problem in 1D, given by

$$-u''(x) = f, \quad x \in (0, 1),$$

$$u(0) = u(1) = 0,$$

for a given function f . Whenever the function f is given by a constant $C > 0$, the analytical solution $u(x) = \frac{C}{2}x(1-x)$ can very easily be found by integrating twice the given PDE and applying the boundary conditions. At the contrary, when

f is given by a more general formulation, see e.g. [LB13], analytical techniques such as the ones described above do not come to the rescue.

Towards numerical methods The fatality that follows from the aforementioned findings is that alternative approaches are no unnecessary luxury but rather an imperative, and this serves as a good motivation into directing the research focus on numerical techniques for solving PDEs. The researchers community had since then concentrated almost exclusively on understanding PDEs' behaviour under the different paradigm of numerics, diverting the solution from the continuous analytical scope to that of the discrete approximated one. Several numerical methods such as Finite Difference (FD), Finite Volumes (FV) or Finite Elements (FE) have been developed providing different advantages.

FE methodology The philosophy of the FE method is to approach the PDE solutions with approximation spaces (usually piecewise continuous functions), turning the initial problem into an elegant variational formulation. The latter is defined through the paradigm of seeing the original problem under the prism of virtual applied loads, see e.g. [ZT00, Bat06]. It is characterized by multiplying the PDE with a so-called test function, followed by an integration in the sense of Lebesgue. This has the advantage of relaxing the regularity of the sought solution, as the original smoothness for which a problem is defined might undoubtedly constitute an additional difficulty for finding a solution. Discretization of the domain into small elements, hence the name of the method, constitutes then a key tool of this approach allowing for characterizing the space domain with a finite set of basis functions. These will serve to formalize the solution under the form of a linear combination. After further mathematical manipulations, the discretized solution is retrieved by solving a linear system.

The FE method can be and has been employed in several applications throughout many decades with a particular success, see e.g. [KC71, RMRM02, MG15]. The theory developed around this method allows in addition to a priori estimate the solution error approximation, with even possible techniques for enhancing the method's precision, such as a posteriori adaptive refinement, see [Löh87, ZZ88]. Literature furthermore continuously address the performance concern, either by designing efficient linear system solvers (iterative solvers of low complexity), or by taking advantage of the growing computations resources to retrieve the final solution through distributing the total set of computations to multiple computing units (domain decomposition), see e.g. [TW06, PT12].

Time-stepping vs parallel-in-time methods Whenever invoking non-stationary

problems, the statements above relatively to the FE method still hold, yet it remains to consider the discretization of the time variable. The space and time entities have mainly been seen for decades as two separate ones. It is quite natural indeed to intuitively differentiate what can be seen (space) from what can be sensed (time). It therefore appeared of a good common sense to combine FE in space with FD in time in order to solve time-dependent problems. As the FD method is well established for structured grids, it is rather straightforward to use in the case of a time interval, with no further physical abstraction of an additional dimensionality. This way of proceeding has been summarized under the terminology of time-stepping methods and is widely used in the research community.

Time-stepping methods are divided in two categories, explicit and implicit. The latter are usually preferred due to their stability, however, they require solving one linear system at the time for every single time step (for a linear problem). This consideration unavoidably leads to slowing down the solving process, particularly when a single linear system ends up being particularly expensive (in terms of computations) to solve. Of course, every linear system can be solved in parallel using many computational resources. However, this parallelization process saturates once a limit of the number of involved computing units with respect to the dimension of the problem, is reached. Even when efficient load balanced domain decomposition strategies are adopted, communication becomes the natural bottleneck of the overall performance.

This has motivated a growing interest of so-called parallel-in-time methods. Those methods, thanks to brilliant mathematical formalizations, developed concurrency also in the time dimension, allowing therefore for parallelization in time, hence the name. Around 60 years ago, [Nie64] considered a predictor for different time intervals. Many other works have emerged since then, see [LMT01], and more recently by [GKZ18] where it was merged into a multi-grid solver setting. Methods involving geometrical space-time discretization, where time was considered as a physical dimension, had also been developed. Such a paradigm allowed for space-time adaptivity, time moving domains (such as a heart for example) among other advantageous features. A non-exhaustive of parallel-in-time methods can be found in [Gan15, Ben20].

The all-at-once method In this work, we employ a monolithic approach in space and time, where a large spacetime system is assembled and solved in parallel. Such a methodology has been introduced by [McD16] and helps achieving time parallelism in a quasi immediate way.

We illustrate this method in the following chapter by initially introducing it for the heat equation, a historical parabolic problem that differs from the monodomain equation by the absence of the nonlinear reaction term derived from the ion channel model (c.f. Equation (2.14)). The monodomain equation is afterwards being treated with Newton's method which we detail also in this chapter.

3.1 Heat equation

When it comes to the study of time-dependent PDEs, the heat equation had been widely used in the literature as the main benchmark. The heat equation models the time evolution of an initial heat distribution in a given domain, provided a certain thermal conductivity characterizing the domain, possibly under the presence of an additional heat source provided to the domain.

Let us consider a domain $D \in \mathbb{R}^d$ and a time interval given by $[0, T]$. We define the heat equation as follows

$$\begin{aligned} \frac{\partial u(\mathbf{x}, t)}{\partial t} - \nabla \cdot (k(\mathbf{x})\nabla u(\mathbf{x}, t)) &= f(\mathbf{x}, t), \quad \forall (\mathbf{x}, t) \in D \times (0, T] \\ u(\mathbf{x}, t) &= 0, \quad \forall (\mathbf{x}, t) \in \partial D \times (0, T] \\ u(\mathbf{x}, 0) &= g(\mathbf{x}), \quad \forall \mathbf{x} \in D, \end{aligned} \quad (3.1)$$

where (\mathbf{x}, t) represents the space and time variables, $u = u(\mathbf{x}, t)$ is the distribution of heat in the domain D along the time interval $[0, T]$, $k \in \mathbb{R}^{d \times d}$ is the anisotropic thermal conductivity, f a heat source supplied at point \mathbf{x} for a time t and $g(\mathbf{x})$ an initial distribution of heat in the domain. Here, the operator ∇ is the usual gradient operator given by

$$\nabla = \left(\frac{\partial}{\partial x_1}, \dots, \frac{\partial}{\partial x_d} \right).$$

Notice that we here defined the problem with homogeneous Dirichlet boundary conditions, but this should not be considered as a restriction and the problem may also be defined for homogeneous/inhomogeneous Neumann or Robin boundary conditions. In Figure 3.1, we propose a visualization for the evolution of the heat in a cube in the particular setting of having an initial heat distribution in the upper corner, and a heat source applied (after a certain amount of time) on the opposite lower corner.

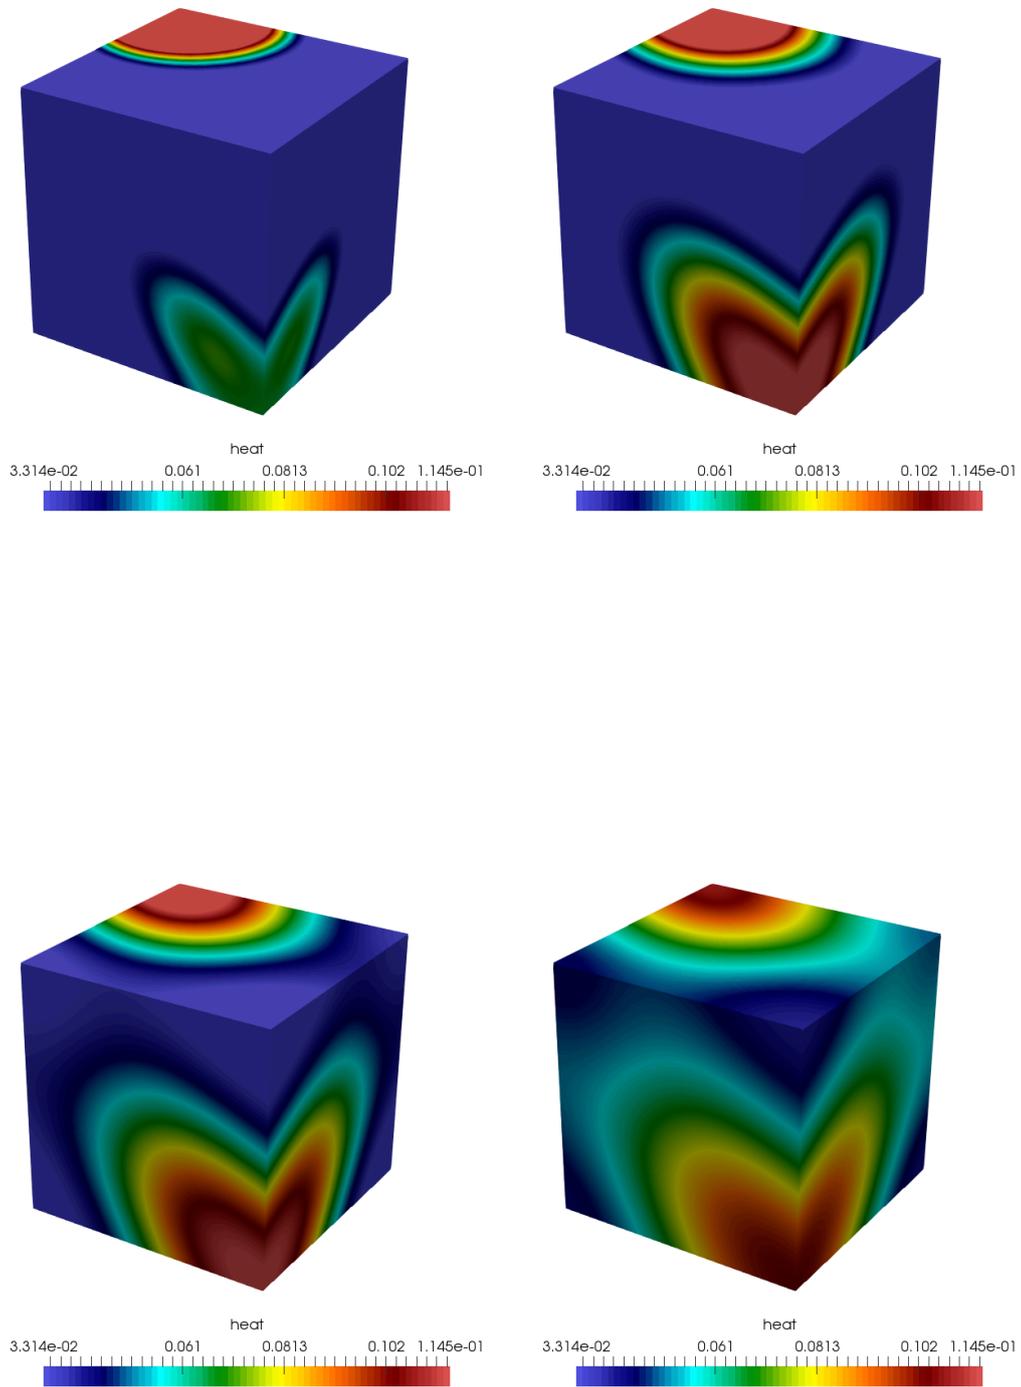


Figure 3.1. Heat propagation in a cubical domain.

3.1.1 Finite Elements in space

Weak Formulation: The FE method is derived by first rewriting the initial PDE problem under its weak formulation, also called the variational formulation. This is achieved by means of a multiplication with a test function, followed by an integration over the considered domain. Let us consider a simplified time-independent version of the heat equation (c.f. Equation (3.1)), i.e. the Poisson equation given by

$$-\nabla \cdot (k(\mathbf{x})\nabla u(\mathbf{x})) = f(\mathbf{x}), \quad \forall \mathbf{x} \in D, \quad (3.2)$$

for $u \in \mathcal{V}$ (where \mathcal{V} is a functional space to be determined) subject to the same set of homogeneous Dirichlet boundary conditions as in (3.1). This motivates the additional assumption that every function $v \in \mathcal{V}$ vanishes on the boundary, that is, $v|_{\partial D} = 0, \forall v \in \mathcal{V}$. We furthermore assume isotropic diffusion, i.e. $k(\mathbf{x}) \in \mathbb{R}$, and that there exists constants k_{\min}, k_{\max} such that

$$0 < k_{\min} \leq k(\mathbf{x}) \leq k_{\max} < \infty \quad (3.3)$$

almost everywhere. The variational formulation goes by

$$-\int_D \nabla \cdot (k(\mathbf{x})\nabla u(\mathbf{x}))v(\mathbf{x}) d\mathbf{x} = \int_D f(\mathbf{x})v(\mathbf{x}) d\mathbf{x}, \quad (3.4)$$

for all $v \in \mathcal{V}$, that is the same functional space where the solution u lies. By using Green's first identity (see Appendix C), the Equation (3.4) can further be written as follows

$$\int_D (k(\mathbf{x})\nabla u(\mathbf{x})) \cdot \nabla v(\mathbf{x}) d\mathbf{x} - \int_{\partial D} (k(\mathbf{x})\nabla u(\mathbf{x}) \cdot \mathbf{n}(\mathbf{x}))v(\mathbf{x}) ds = \int_D f(\mathbf{x})v(\mathbf{x}) d\mathbf{x}. \quad (3.5)$$

When applying the boundary conditions, the variational formulation (3.5) simplifies to

$$\int_D (k(\mathbf{x})\nabla u(\mathbf{x})) \cdot \nabla v(\mathbf{x}) d\mathbf{x} = \int_D f(\mathbf{x})v(\mathbf{x}) d\mathbf{x}. \quad (3.6)$$

for all $v \in \mathcal{V}$. In order for the weak formulation (3.6) to make sense, we must ensure that the solution and the test function are well behaved enough, so to guarantee the existence of the underlying integrals. In the following case, assuming that the function f is sufficiently smooth, it is enough to consider

$\mathcal{V} = H_0^1(D) = \{v \in H^1(D) | v|_{\partial D} = 0\}$ (see Appendix C). Indeed, the Cauchy–Schwarz inequality can be used in this case to demonstrate the boundedness of the two integrals of (3.6):

$$\begin{aligned} \int_D (k(\mathbf{x})\nabla u(\mathbf{x})) \cdot \nabla v(\mathbf{x}) \, d\mathbf{x} &\leq k_{\max} \left| \int_D \nabla u(\mathbf{x}) \cdot \nabla v(\mathbf{x}) \, d\mathbf{x} \right| \\ &\leq k_{\max} \|\nabla u\|_{L^2(D)} \|\nabla v\|_{L^2(D)} < \infty, \end{aligned} \quad (3.7)$$

and

$$\int_D f(\mathbf{x})v(\mathbf{x}) \, d\mathbf{x} \leq \left| \int_D f(\mathbf{x})v(\mathbf{x}) \, d\mathbf{x} \right| \leq \|f\|_{L^2(D)} \|v\|_{L^2(D)} < \infty. \quad (3.8)$$

In the above discussion about the functional space to consider resides the benefit of the variational formulation. Indeed, the PDE strong form requirement of a solution in $C^2(D)$ is softened to the space $H^1(D)$, relaxing therefore the condition on the existence of a solution.

Furthermore, by observing that the left hand side of (3.6) is linear in u and v , whereas the right hand side is linear in v , the weak formulation can be rewritten in the following form:

Find $u \in \mathcal{V} = H_0^1(D)$ such that

$$a(u, v) = l(v), \quad \forall v \in \mathcal{V} = H_0^1(D), \quad (3.9)$$

where a is a bilinear form on $\mathcal{V} \times \mathcal{V}$ defined by

$$\begin{aligned} a : \mathcal{V} \times \mathcal{V} &\rightarrow \mathbb{R} \\ (u, v) &\mapsto a(u, v) := \int_D (k(\mathbf{x})\nabla u(\mathbf{x})) \cdot \nabla v(\mathbf{x}) \, d\mathbf{x}, \end{aligned}$$

and l a linear form on \mathcal{V} defined by

$$\begin{aligned} l : \mathcal{V} &\rightarrow \mathbb{R} \\ v &\mapsto l(v) := \int_D f(\mathbf{x})v(\mathbf{x}) \, d\mathbf{x}. \end{aligned}$$

Solution existence and uniqueness: Rewriting (3.6) in the form of (3.9) is not purely a matter of elegance as much as it enables a whole set of theory on linear operators and their Riesz representation, see e.g. [Kre78]. This leads to a primordial theorem about the existence and the uniqueness of a solution to the PDE weak formulation.

Theorem 3.1 (Babuska–Lax–Milgram). *Let \mathcal{H} be a Hilbert space, endowed with its scalar product (\cdot, \cdot) and its associated norm $\|\cdot\|$. Let $a(\cdot, \cdot)$ be a bilinear form on \mathcal{H} that is:*

- *continuous on $\mathcal{H} \times \mathcal{H}$, i.e. $\exists C_a > 0$ such that $|a(u, v)| \leq C_a \|u\| \|v\|$.*
- *coercive, i.e. $\exists \alpha > 0$ such that $a(u, u) \geq \alpha \|u\|^2$.*

Let furthermore $l(\cdot)$ be a linear form on \mathcal{H} that is

- *continuous on \mathcal{H} , i.e. $\exists C_l > 0$ such that $|l(v)| \leq C_l \|v\|$.*

Under the following assumptions, there exists a unique $u \in \mathcal{H}$ such that

$$a(u, v) = l(v), \quad \forall v \in \mathcal{H}.$$

Proof. See [Bab71]. □

The assumptions of the Lax–Milgram theorem hold for our variational formulation. Considering the Hilbert space $\mathcal{H} = H^1(D)$ endowed with its scalar product $(\cdot, \cdot)_{H^1(D)}$ and its induced norm $\|\cdot\|_{H^1(D)}$ (see Appendix C), the continuity of $a(\cdot, \cdot)$ is straightforward from (3.7), for $C_a = k_{\max}$:

$$\begin{aligned} |a(u, v)| &= \left| \int_D (k(\mathbf{x}) \nabla u(\mathbf{x})) \cdot \nabla v(\mathbf{x}) \, d\mathbf{x} \right| \leq k_{\max} \|\nabla u\|_{L^2(D)} \|\nabla v\|_{L^2(D)} \quad (3.10) \\ &\leq k_{\max} \|u\|_{H^1(D)} \|v\|_{H^1(D)}. \end{aligned}$$

The coercivity of $a(\cdot, \cdot)$ follows from the following inequality, for $\alpha = \frac{k_{\min}}{C_D^2 + 1}$:

$$\begin{aligned} a(u, u) &\geq \int_D k(\mathbf{x}) (\nabla u(\mathbf{x}))^2 \, d\mathbf{x} \geq k_{\min} \int_D \nabla u(\mathbf{x})^2 \, d\mathbf{x} = k_{\min} \|\nabla u\|_{L^2(D)}^2 \quad (3.11) \\ &\geq \frac{k_{\min}}{C_D^2 + 1} \|u\|_{H^1(D)}^2, \end{aligned}$$

where $C_D > 0$ is the constant from the Poincaré inequality (see Appendix C). The continuity of $l(\cdot)$ follows from (3.8), for $C_l = \|f\|_{L^2(D)}$, assuming that $f \in L^2(D)$:

$$l(v) = \int_D f(\mathbf{x}) v(\mathbf{x}) \, d\mathbf{x} \leq \|f\|_{L^2(D)} \|v\|_{L^2(D)} \leq \|f\|_{L^2(D)} \|v\|_{H^1(D)}. \quad (3.12)$$

Consequently, the weak formulation does have a unique solution $u \in \mathcal{H} = H^1(D)$. Moreover,

$$C_l \|u\|_{H^1(D)} \geq l(u) = a(u, u) \geq \alpha \|u\|_{H^1(D)}^2, \quad (3.13)$$

from which we deduce that

$$\|u\|_{H^1(D)} \leq \frac{C_l}{\alpha} = \frac{\|f\|_{L^2(D)} (C_D^2 + 1)}{k_{\min}}. \quad (3.14)$$

Discretization: The FE methodology relies furthermore on a discretization of the spatial domain D into smaller elements, aiming at locally approximating the solution by piecewise continuous functions.

Let \mathcal{T} represent a tetrahedralization of $D \subset \mathbb{R}^d$ ($d = 2, 3$), or equivalently a mesh, characterized by a set of d -simplices elements $\{K_i\}_{i=1}^{N_e}$ (triangles for $d = 2$, tetrahedrons for $d = 3$) such that $D = \bigcup_{i=1}^{N_e} K_i$ and $K_i \cap K_j$ for $i \neq j$ is either empty, a corner, an edge ($d = 2$), or a face ($d = 3$).

Furthermore, every element $K \in \mathcal{T}$ is characterized by a length h_K , that is defined to be the diameter, or equivalently the longest edge, of K . The mesh step size is then given by $h = \max_{K \in \mathcal{T}} h_K$. We introduce the space of continuous piecewise linear polynomials with

$$V_h = \{v \in C^0(D) \mid v_K \in \mathbb{P}_1(K), \forall K \in \mathcal{T}\}. \quad (3.15)$$

In the definition of V_h above, $C^0(D)$ denotes the space of continuous functions on D , whereas $\mathbb{P}_1(K)$ is the set of linear polynomials on K .

The space V_h is uniquely determined by the mesh nodal values $\{\mathbf{x}_i\}_{i=1}^n$ (that is the spatial points defining the elements K of \mathcal{T}), and conversely every set of nodal values uniquely determines a piecewise continuous linear function in V_h , see [LB13]. We therefore define $\{\phi_i\}_{i=1}^n \subset V_h$ to be a basis for the space V_h defined around the degrees of freedom (Dofs) $\{\mathbf{x}_i\}_{i=1}^n$, i.e.

$$\phi_i(\mathbf{x}_j) = \delta_{ij}, \quad i, j = 1, \dots, n.$$

We now can define the Galerkin approximation of the variational formulation (3.9)

Find $u_h \in V_h \subset \mathcal{V}$ such that

$$a(u_h, v_h) = l(v_h), \quad \forall v_h \in V_h \subset \mathcal{V}. \quad (3.16)$$

Since $V_h \subset \mathcal{V}$, all the Lax–Milgram assumptions still hold since they were established for the more constraining case of functions $v \in \mathcal{V}$. We therefore can

affirm that the Galerkin approximation problem still have a solution and that it is unique.

FE solution: Let us recall the Galerkin approximation in (3.16). Since $\{\phi_j\}_{j=1}^n$ represents a basis for V_h , every element in V_h can be written as a linear combination of $\phi_j, j = 1, \dots, n$. In particular,

$$u_h(\mathbf{x}) = \sum_{j=1}^n u_j \phi_j(\mathbf{x}), \quad (3.17)$$

where $u_j \in \mathbb{R}$ for $j = 1, \dots, n$. By inserting (3.17) in (3.16), the problem translates to:

Find $u_j \in \mathbb{R}$ for $j = 1, \dots, n$ such that

$$a\left(\sum_{j=1}^n u_j \phi_j(\mathbf{x}), v_h\right) = l(v_h), \quad \forall v_h \in V_h \subset \mathcal{V}. \quad (3.18)$$

Furthermore, since V_h is completely represented by its basis, it is reasonable (actually equivalent) to replace test functions $v_h \in V_h$ with $\phi_i(\mathbf{x})$ for $i = 1, \dots, n$. The formulation of the Galerkin problem becomes

Find $u_j \in \mathbb{R}$ for $j = 1, \dots, n$ such that

$$a\left(\sum_{j=1}^n u_j \phi_j(\mathbf{x}), \phi_i(\mathbf{x})\right) = l(\phi_i(\mathbf{x})), \quad \forall i = 1, \dots, n. \quad (3.19)$$

Taking advantage of the bilinearity and symmetry of $a(\cdot, \cdot)$, we obtain:

$$\sum_{j=1}^n a(\phi_i(\mathbf{x}), \phi_j(\mathbf{x})) u_j = l(\phi_i(\mathbf{x})), \quad \forall i = 1, \dots, n. \quad (3.20)$$

Therefore, if we set $A \in \mathbb{R}^{n \times n}$ and $\mathbf{f} \in \mathbb{R}^n$ such that

$$(A)_{ij} = a(\phi_i(\mathbf{x}), \phi_j(\mathbf{x})) = \int_D k(\mathbf{x}) \nabla \phi_i(\mathbf{x}) \cdot \nabla \phi_j(\mathbf{x}) d\mathbf{x}, \quad (3.21)$$

and

$$(\mathbf{f})_i = l(\phi_i(\mathbf{x})) = \int_D f(\mathbf{x}) \phi_i(\mathbf{x}) d\mathbf{x}, \quad (3.22)$$

it holds that

$$\mathbf{A}\mathbf{u} = \mathbf{f}, \quad (3.23)$$

where $\mathbf{u} = [u_1, \dots, u_n]^T$. The solution to the Galerkin approximation can be recovered by solving the above linear system in (3.23) and reconstruct $u(\mathbf{x}) \in V_h$

as in (3.17).

Approximation property: We recall the original weak formulation for the solution $u \in V$

$$a(u, v) = l(v), \quad \forall v \in \mathcal{V}. \quad (3.24)$$

The Galerkin approximation on the other hand yields

$$a(u_h, v_h) = l(v_h), \quad \forall v_h \in V_h. \quad (3.25)$$

Since $V_h \subset \mathcal{V}$, we can apply (3.24) for $v_h \in V_h$ and subtract (3.24) from (3.25). By using the linearity of $a(\cdot, \cdot)$, we obtain the Galerkin orthogonality:

$$a(u - u_h, v_h) = 0, \quad \forall v_h \in V_h. \quad (3.26)$$

From this equality, the following holds:

$$\begin{aligned} a(u - u_h, u - v_h) &= a(u - u_h, u - u_h + \underbrace{u_h - v_h}_{\in V_h}) \\ &= a(u - u_h, u - u_h) + \underbrace{a(u - u_h, u_h - v_h)}_{=0} \\ &= a(u - u_h, u - u_h). \end{aligned} \quad (3.27)$$

By using the coercivity and the continuity of $a(\cdot, \cdot)$, we have that

$$\begin{aligned} \alpha \|u - u_h\|_{H^1(D)}^2 &\leq a(u - u_h, u - u_h) = a(u - u_h, u - v_h) \\ &\leq C_a \|u - u_h\|_{H^1(D)} \|u - v_h\|_{H^1(D)}, \end{aligned} \quad (3.28)$$

where α and C_a are respectively the coercivity and the continuity constants. This leads to having:

$$\|u - u_h\|_{H^1(D)} \leq \frac{C_a}{\alpha} \|u - v_h\|_{H^1(D)} \quad \forall v_h \in V_h. \quad (3.29)$$

This equality is known as Céa's Lemma. Since it holds $\forall v_h \in V_h$, we can rewrite it as

$$\|u - u_h\|_{H^1(D)} \leq \frac{C_a}{\alpha} \inf_{v_h \in V_h} \|u - v_h\|_{H^1(D)}. \quad (3.30)$$

which we can interpret as u_h being the quasi best approximation of u in V_h . This indicates why the whole theory developed above works, as long as V_h approximates V , i.e. $V_h \xrightarrow{h \rightarrow 0} V$.

3.1.2 Finite Difference in time

Let us now consider the original model problem, that is the heat equation. This equation only differs from the previous Poisson problem with a dependency in time provided in the form of a time derivative.

The derivation of the Galerkin approximation emanating of the weak formulation follows therefore the exact methodology and can be phrased as:

$\forall 0 < t \leq T$, find $u_h = u_h(\mathbf{x}, t)$ such that

$$\int_D \frac{\partial u_h(\mathbf{x}, t)}{\partial t} v_h(\mathbf{x}) d\mathbf{x} + a(u_h, v_h) = l(v_h), \quad \forall v_h = v_h(\mathbf{x}) \in V_h. \quad (3.31)$$

In the above, we underline the fact that the test function depends exclusively on the space variable, whereas the solution u depends also on the time variable. The Galerkin approximation in that context needs to be interpreted as several approximations for many fixed time frames.

We can introduce the space V_h exactly as defined in (3.15) with its inherent set of basis functions $\{\phi_j(\mathbf{x})\}_{j=1}^n$. To (3.31), we seek a solution of the form

$$u_h(\mathbf{x}, t) = \sum_{j=1}^n u_j(t) \phi_j(\mathbf{x}), \quad (3.32)$$

in which $\{u_j(t)\}_{j=1}^n$ for $0 < t \leq T$ are to be determined. The ansatz (3.32) can further be injected in the Galerkin approximation and v_h substituted with ϕ_i for $i = 1, \dots, n$:

$$\int_D \frac{\partial}{\partial t} \left(\sum_{j=1}^n u_j(t) \phi_j(\mathbf{x}) \right) \phi_i(\mathbf{x}) d\mathbf{x} + a \left(\sum_{j=1}^n u_j(t) \phi_j, \phi_i \right) = l(\phi_i). \quad (3.33)$$

Since the unknowns $u_j(t)$ are independent from the variable \mathbf{x} , they can be taken out from the space integral (and consequently also considered as scalars for $a(\cdot, \cdot)$). By exploiting the bilinearity of $a(\cdot, \cdot)$ and the linearity of the integral operator, we obtain

$$\sum_{j=1}^n \left(\int_D \phi_i(\mathbf{x}) \phi_j(\mathbf{x}) d\mathbf{x} \right) \frac{\partial u_j(t)}{\partial t} + \sum_{j=1}^n a(\phi_i(\mathbf{x}), \phi_j(\mathbf{x})) u_j(t) = l(\phi_i(\mathbf{x})). \quad (3.34)$$

When setting $M \in \mathbb{R}^{n \times n}$ defined by

$$(M)_{ij} = \int_D \phi_i(\mathbf{x}) \phi_j(\mathbf{x}) d\mathbf{x}, \quad (3.35)$$

$A \in \mathbb{R}^{n \times n}$ and $\mathbf{f} \in \mathbb{R}^n$ respectively as in (3.21) and (3.22), we obtain the following ODE system:

$$M\dot{\mathbf{u}}(t) + A\mathbf{u}(t) = \mathbf{f}(t) \quad (3.36)$$

in which we rewrote $\frac{\partial \mathbf{u}(t)}{\partial t} = \frac{\partial}{\partial t} [u_1(t), \dots, u_n(t)]^T$ as $\dot{\mathbf{u}}(t)$.

FD time discretization: In order to solve the ODE system (3.36), we need to proceed to a discretization of the time interval $[0, T]$. For this purpose, let us consider $m-1$ uniform subdivisions of $[0, T]$ providing equally sized intervals of length $\Delta t = T/(m-1)$. We therefore obtain m discrete time modes characterized by $t_k = (k-1)/\Delta t$ for $k = 1, \dots, m$.

Integrating (3.36) over $[t_k, t_{k+1}]$, $\forall k \in 1, \dots, m-1$ yields the following:

$$M \int_{t_k}^{t_{k+1}} \dot{\mathbf{u}}(\tau) d\tau + A \int_{t_k}^{t_{k+1}} \mathbf{u}(\tau) d\tau = \int_{t_k}^{t_{k+1}} \mathbf{f}(\tau) d\tau, \quad (3.37)$$

since the matrices $M, A \in \mathbb{R}^{n \times n}$ are time independent. Naturally, this can also be rewritten as:

$$M(\mathbf{u}(t_{k+1}) - \mathbf{u}(t_k)) + A \int_{t_k}^{t_{k+1}} \mathbf{u}(\tau) d\tau = \int_{t_k}^{t_{k+1}} \mathbf{f}(\tau) d\tau. \quad (3.38)$$

Let us furthermore denote the approximation of $\mathbf{u}(t_k)$ to be given by \mathbf{u}_k , while similar notation \mathbf{f}_k indicates the exact value of $\mathbf{f}(t_k)$. One can use a quadrature rule to approximate the remaining integrals of (3.38), as for example the trapezoidal rule:

$$M(\mathbf{u}_{k+1} - \mathbf{u}_k) + \Delta t A \left(\frac{\mathbf{u}_{k+1} + \mathbf{u}_k}{2} \right) = \Delta t \left(\frac{\mathbf{f}_{k+1} + \mathbf{f}_k}{2} \right). \quad (3.39)$$

A more general weighting for the quadrature can even be considered for recovering a larger class of quadrature formulas, by setting $\theta \in [0, 1]$ such that:

$$M(\mathbf{u}_{k+1} - \mathbf{u}_k) + \Delta t A(\theta \mathbf{u}_{k+1} + (1-\theta)\mathbf{u}_k) = \Delta t(\theta \mathbf{f}_{k+1} + (1-\theta)\mathbf{f}_k). \quad (3.40)$$

Rearranging in the following form,

$$M \frac{\mathbf{u}_{k+1} - \mathbf{u}_k}{\Delta t} + A(\theta \mathbf{u}_{k+1} + (1-\theta)\mathbf{u}_k) = (\theta \mathbf{f}_{k+1} + (1-\theta)\mathbf{f}_k). \quad (3.41)$$

yields the so-called θ -method. Particular values of θ define different FD schemes methods, each of which presenting distinct advantages.

- $\theta = 0$: is called the Forward (or explicit) Euler method given by

$$M\mathbf{u}_{k+1} = (M - \Delta t A)\mathbf{u}_k + \mathbf{f}_k. \quad (3.42)$$

This method is known to be fast, but comes with a stability condition on the time discretization step Δt .

- $\theta = 1$: is called the Backward (or implicit) Euler method given by

$$(M + \Delta t A)\mathbf{u}_{k+1} = M\mathbf{u}_k + \mathbf{f}_{k+1}. \quad (3.43)$$

The method is characterized by good stability properties.

- $\theta = 1/2$: is called the Crank-Nicolson method given by

$$\left(M + \frac{\Delta t}{2}A\right)\mathbf{u}_{k+1} = \left(M - \frac{\Delta t}{2}A\right)\mathbf{u}_k + \frac{\Delta t}{2}(\mathbf{f}_{k+1} + \mathbf{f}_k). \quad (3.44)$$

This method has the advantage of being second order accurate in time.

Let us recall that

$$\mathbf{u}_k \approx \mathbf{u}(t_k) = [u_1(t_k), \dots, u_n(t_k)]^T, \quad (3.45)$$

where n is the number of space discretization nodes. Starting from $k = 1$ (corresponding to $t_1 = 0$) and $u(0)$ provided by the PDE initial condition, all of these methods require solving recursively $m - 1$ linear systems. In the case where the number of time steps is important, this undoubtedly might represent a hindrance to the overall performance of these methods, especially when every time iteration itself requires a high number of space degrees of freedom.

The purpose of the space–time approach (or all–at–once method) developed in the next subsection is to collect all the linear systems of interest in a larger global one, so to exploit parallelism also in the time dimension.

3.1.3 Space–time approach

The heat equation (3.1) can be numerically solved, as demonstrated in the previous sections, by means of FE in space and FD in time, using a sequential time-stepping method such as the θ –method. The parallel scalability of this approach can however be limited by the time integration process. In fact, when the parallelization in space saturates, sequential time integration becomes the natural bottleneck for the parallel scalability of the solution process. This can be overcome by parallel-in-time methods, among which we here focus on the all–at–once

approach, a monolithic approach where we assemble a large space-time system that is solved in parallel.

The derivation of this approach is step-by-step similar to the classical one, for which the weak formulation of the PDE, recall

$$\int_D \frac{\partial u(\mathbf{x}, t)}{\partial t} v(\mathbf{x}) d\mathbf{x} + \int_D (k(\mathbf{x}) \nabla u(\mathbf{x}, t)) \nabla v(\mathbf{x}) d\mathbf{x} = \int_D f(\mathbf{x}, t) v(\mathbf{x}) d\mathbf{x}.$$

leads to the semi-discretized problem in space:

$$M \frac{\partial \mathbf{u}(t)}{\partial t} + A \mathbf{u}(t) = \mathbf{f}(t). \quad (3.46)$$

Here, $M \in \mathbb{R}^{n \times n}$ and $A \in \mathbb{R}^{n \times n}$ are the standard mass and stiffness matrices obtained using the usual n linear nodal basis functions $\{\phi_i\}_{i=1}^n \subset \mathbb{P}_1$, i.e.

$$M := \left[\int_D \phi_i(\mathbf{x}) \phi_j(\mathbf{x}) d\mathbf{x} \right]_{i,j=1}^n, \quad A := \left[\int_D (G(\mathbf{x}) \nabla \phi_i(\mathbf{x})) \nabla \phi_j(\mathbf{x}) d\mathbf{x} \right]_{i,j=1}^n, \quad (3.47)$$

$$\mathbf{f}(t) := \left[\int_D f(\mathbf{x}, t) \phi_i(\mathbf{x}) d\mathbf{x} \right]_{i=1}^n, \quad (3.48)$$

arising from the approximation:

$$u(\mathbf{x}, t) \approx \sum_{i=1}^n u_i(t) \phi_i(\mathbf{x}), \quad \text{with } \mathbf{u}(t) = [u_1(t), u_2(t), \dots, u_n(t)]^T.$$

Let us now consider a uniform partition of the time interval $[0, T]$ in m nodes, such that $\Delta t = T/(m-1)$ and $t_k = (k-1)\Delta t$, with $k = 1, \dots, m$. We apply the θ -method for the time discretization of (3.46) and obtain, for $k = 1, \dots, m-1$

$$(M + \theta \Delta t A) \mathbf{u}_{k+1} + (-M + (1 - \theta) \Delta t A) \mathbf{u}_k = \mathbf{f}_k \quad \text{and} \quad \mathbf{u}_k \approx \mathbf{u}(t_k), \quad (3.49)$$

$$\text{with } \mathbf{f}_k := \Delta t (\theta \mathbf{f}(t_{k+1}) + (1 - \theta) \mathbf{f}(t_k)). \quad (3.50)$$

If we define $D_{t,h} := M + \theta \Delta t A$ and $B_{t,h} := -M + (1 - \theta) \Delta t A$, the system of equations (3.49) can be summarized in the compact form:

$$\begin{bmatrix} D_{t,h} & & & & \\ B_{t,h} & D_{t,h} & & & \\ & \ddots & \ddots & \ddots & \\ & & B_{t,h} & D_{t,h} & \\ & & & & \ddots & \ddots \end{bmatrix} \begin{pmatrix} \mathbf{u}_1 \\ \mathbf{u}_2 \\ \vdots \\ \mathbf{u}_m \end{pmatrix} = \begin{pmatrix} \mathbf{f}_1 \\ \mathbf{f}_2 \\ \vdots \\ \mathbf{f}_m \end{pmatrix} \iff \mathbf{C} \mathbf{u} = \mathbf{f}, \quad (3.51)$$

where $C \in \mathbb{R}^{nm \times nm}$ is a large space-time system that can be distributed and solved in parallel and

$$\mathbf{u} := [\mathbf{u}_1, \mathbf{u}_1, \dots, \mathbf{u}_m]^T \quad \text{and} \quad \mathbf{f} := [\mathbf{f}_1, \mathbf{f}_1, \dots, \mathbf{f}_m]^T.$$

3.2 The Monodomain equation

In the previous section, we have been through the methodology of solving the heat equation. The latter model problem is undoubtedly a good starting point for the solution of the monodomain equation, as both problems only differ by the reaction term modeling the cells interaction. Recall the monodomain equation given by:

$$\begin{aligned} \frac{\partial u(\mathbf{x}, t)}{\partial t} - \nabla \cdot (G(\mathbf{x}) \nabla u(\mathbf{x}, t)) + I_{\text{ion}}(u(\mathbf{x}, t)) &= I_{\text{app}}(\mathbf{x}, t), \quad \forall (\mathbf{x}, t) \in D \times (0, T] \\ G(\mathbf{x}) \nabla u(\mathbf{x}, t) \cdot \mathbf{n} &= 0, \quad \forall (\mathbf{x}, t) \in \partial D \times (0, T] \\ u(\mathbf{x}, 0) &= 0, \quad \forall \mathbf{x} \in D. \end{aligned} \quad (3.52)$$

The difficulty arises however from the fact that this additional reaction term I_{ion} is non linear, preventing from deriving with ease the final linear system to be solved. In this section, we explore the weak formulation of the monodomain equation, show how to make use of Newton's method for overcoming the non linearity of the problem and express the final system to be assembled.

3.2.1 Weak formulation

We consider the Galerkin approximation of the weak formulation to the monodomain equation, proceeding in a similar way to what have been done previously: $\forall 0 < t \leq T$, find $u_h = u_h(\mathbf{x}, t)$ such that

$$\int_D \frac{\partial u_h(\mathbf{x}, t)}{\partial t} v_h(\mathbf{x}) d\mathbf{x} + a(u_h, v_h) = l(v_h), \quad \forall v_h = v_h(\mathbf{x}) \in V_h. \quad (3.53)$$

where

$$a(u_h, v_h) = \int_D (G(\mathbf{x}) \nabla u_h(\mathbf{x}, t)) \cdot \nabla v_h(\mathbf{x}) d\mathbf{x} + \int_D I_{\text{ion}}(u_h(\mathbf{x}, t)) v_h(\mathbf{x}) d\mathbf{x}, \quad (3.54)$$

and

$$l(v_h, t) = \int_D I_{\text{app}}(\mathbf{x}, t) v_h(\mathbf{x}) d\mathbf{x}. \quad (3.55)$$

This provides us with the semi-discretized formulation of the monodomain equation in space. Since $a(\cdot, \cdot)$ is not anymore a bilinear form due to the presence of I_{ion} , injecting the solution ansatz

$$u_h(\mathbf{x}, t) = \sum_{j=1}^n u_j(t) \phi_j(\mathbf{x}), \quad (3.56)$$

in (3.54), does not produce the desired linear system to solve. Instead, we will proceed slightly differently by considering a time-discretization scheme, say Backward Euler for simplicity, to obtain the following fully-discretized problem:

for any $k \geq 0$, given $u_{h,k} \in V_h$, find $u_{h,k+1} \in V_h$ such that

$$\begin{cases} \frac{1}{\Delta t} \int_D u_{h,k+1}(\mathbf{x}) v_h(\mathbf{x}) d\mathbf{x} + a(u_{h,k+1}, v_h) = l(v_h, t_{k+1}) + \frac{1}{\Delta t} \int_D u_{h,k}(\mathbf{x}) v_h(\mathbf{x}) d\mathbf{x}, \\ u_{h,0} = u_0, \end{cases} \quad (3.57)$$

where $u_{h,k}(\mathbf{x}) = u_h(\mathbf{x}, t_k)$. We furthermore introduce the residual, defined as the difference between the left and the right hand side of the equality defined in (3.57), that is

$$\begin{aligned} R(u_{h,k+1})(v_h) &= \frac{1}{\Delta t} \int_D u_{h,k+1}(\mathbf{x}) v_h(\mathbf{x}) d\mathbf{x} + a(u_{h,k+1}, v_h) \\ &\quad - l(v_h, t_{k+1}) - \frac{1}{\Delta t} \int_D u_{h,k}(\mathbf{x}) v_h(\mathbf{x}) d\mathbf{x}. \end{aligned} \quad (3.58)$$

We are now left with the new problem formulation of finding $u_{h,k+1}$ that cancels the residual.

3.2.2 Newton's method

The idea behind Newton's method is to produce successively better approximations, in order to find the root of a given function.

Graphical interpretation: Starting from an initial guess close enough to the root of a function, Newton's method produces successive iterates based on the tangent to the function calculated at the considered iteration. The tangent intersection with the x -axis defines the next iteration point, that is meant to be a better approximation of the root.

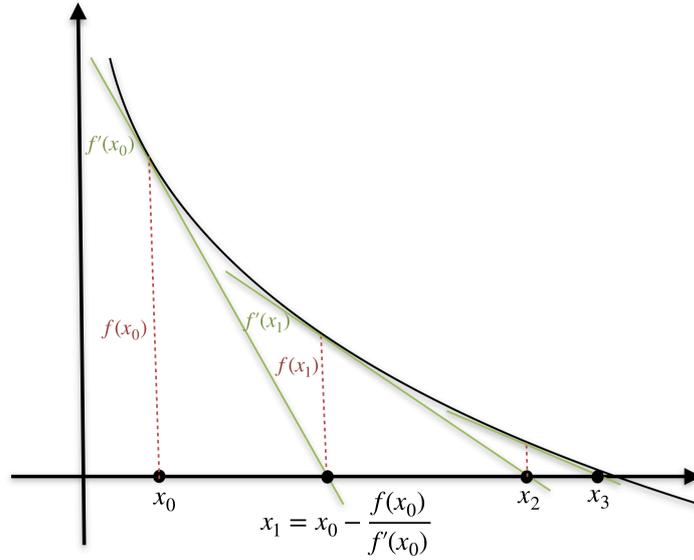


Figure 3.2. Graphical illustration of Newton's method through few iterations.

Mathematical formulation: Mathematically, the Newton's method acts as a linearization of a given function $f : (a, b) \rightarrow \mathbb{R}$, using its Taylor expansion for a given point $x_0 \in (a, b)$ and a neighborhood $h > 0$

$$f(x_0 + h) \approx f(x_0) + hf'(x_0), \quad (3.59)$$

provided that the function is differentiable in (a, b) . If the desired incremented value $x_1 := x_0 + h$ is a root of f , then we get $f(x_1) = f(x_0 + h) = 0$, from which one can deduce the value of the incrementation h , given by

$$0 \approx f(x_0) + hf'(x_0) \implies h \approx -\frac{f(x_0)}{f'(x_0)} \quad (3.60)$$

Newton's method for functional spaces: The situation for finding the root of the residual (3.58) is slightly different since we are dealing with functional spaces and the derivative is not to be understood in the usual sense. Let us recall that if $F : X \rightarrow Y$ is an application of two normed linear spaces X and Y , we call the Gâteaux (or weak) differential of F at a point $\mathbf{x}_0 \in X$ in the direction of \mathbf{h} , the limit

$$DF(\mathbf{x}_0, \mathbf{h}) = \lim_{\lambda \rightarrow +\infty} \frac{F(\mathbf{x}_0 + \lambda\mathbf{h}) - F(\mathbf{x}_0)}{\lambda} \quad (3.61)$$

for $\lambda \in \mathbb{R}$. If the weak differential is linear, it can further be expressed as:

$$DF(\mathbf{x}_0, \mathbf{h}) = F'_G(\mathbf{x}_0)\mathbf{h}, \quad (3.62)$$

in which $F'_G(\mathbf{x}_0)$ is called the *Fréchet derivative* of F in \mathbf{x}_0 in the direction of \mathbf{h} . From this, we deduce Newton's method for finding the root of (3.58), that is updating $u_{h,k+1}^{(i)} \in V_h$ with $\delta u_{h,k+1}^{(i)} \in W_h$ such that

$$DR(u_{h,k+1}^{(i)}, \delta u_{h,k+1}^{(i)})(v_h) = -R(u_{h,k+1}^{(i)}), \quad \forall v_h \in V_h, \quad (3.63)$$

in which similarly to the notation introduced previously, $DR(u_{h,k+1}^{(i)}, \delta u_{h,k+1}^{(i)})(v_h)$ is the differential of $R(u_{h,k+1}^{(i)})(v_h)$ along $\delta u_{h,k+1}^{(i)}$. The latter differential is obtained by computing the limit:

$$DR(u_{h,k+1}^{(i)}, \delta u_{h,k+1}^{(i)})(v_h) = \lim_{\lambda \rightarrow +\infty} \frac{R(u_{h,k+1}^{(i)} + \lambda \delta u_{h,k+1}^{(i)})(v_h) - R(u_{h,k+1}^{(i)})(v_h)}{\lambda}. \quad (3.64)$$

The motivation of setting $\delta u_{h,k+1}^{(i)}$ in W_h , that might represent a different space from V_h , can be explained by eventual boundary conditions requirements that limit the set of admissible updates, as for examples when non-homogeneous Dirichlet boundary conditions are required.

We will later on show how the condition (3.63) serves as a linearization of the weak formulation (3.57), from which the update $\delta u_{h,k+1}^{(i)}$ can be recovered. We propose in the meantime, based on what have been developed above, the following algorithm for solving (3.57)

1. Initialize with $u_{h,k+1}^{(0)} = u_{h,k}$
2. Set $i = 0$ and residual = $\epsilon + 1$
3. While ($i < \max_{\text{iter}}$) and (residual $> \epsilon$)
 - (i) Find $\delta u_{h,k+1}^{(i)}$ by solving (3.63)
 - (ii) Update with $u_{h,k+1}^{(i+1)} = u_{h,k+1}^{(i)} + \delta u_{h,k+1}^{(i)}$
 - (iii) Set $i = i + 1$
 - (iv) Compute residual = $\left\| \delta u_{h,k+1}^{(i)} \right\| / \left\| u_{h,k+1}^{(i+1)} \right\|$
4. Set $u_{h,k+1} = u_{h,k+1}^{(i)}$

3.2.3 The linear problem

Let us now attempt to solve (3.63) in order to recover the suitable update $\delta u_{h,k+1}^{(i)}$. Using the definition of the Gâteaux derivative, we can compute:

$$\begin{aligned} DR(u_{h,k+1}^{(i)}, \delta u_{h,k+1}^{(i)})(v_h) &= \frac{1}{\Delta t} \int_D \delta u_{h,k+1}^{(i)}(\mathbf{x}) v_h(\mathbf{x}) d\mathbf{x} + \int_D G(\mathbf{x}) \nabla \delta u_{h,k+1}^{(i)}(\mathbf{x}) \cdot \nabla v_h(\mathbf{x}) d\mathbf{x} \\ &+ \int_D I'_{\text{ion}}(u_{h,k+1}^{(i)}(\mathbf{x})) (\delta u_{h,k+1}^{(i)}(\mathbf{x})) v_h(\mathbf{x}) d\mathbf{x}, \end{aligned} \quad (3.65)$$

where

$$\begin{aligned} I'_{\text{ion}}(u_{h,k+1}^{(i)})(\delta u_{h,k+1}^{(i)}) &= ((u_{h,k+1}^{(i)} - u_{\text{peak}})(u_{h,k+1}^{(i)} - u_{\text{th}}) \\ &+ (u_{h,k+1}^{(i)} - u_{\text{peak}})(u_{h,k+1}^{(i)} - u_{\text{rest}}) \\ &+ (u_{h,k+1}^{(i)} - u_{\text{th}})(u_{h,k+1}^{(i)} - u_{\text{rest}})) \delta u_{h,k+1}^{(i)}. \end{aligned} \quad (3.66)$$

Using the derivative expression (3.65) in (3.63), we recover the following formulation to be solved at every Newton iteration i for $\delta u_{h,k+1}^{(i)}$, given $u_{h,k+1}^{(i)}$ and $u_{h,k}$:

$$\begin{aligned} &\frac{1}{\Delta t} \int_D \delta u_{h,k+1}^{(i)}(\mathbf{x}) v_h(\mathbf{x}) d\mathbf{x} + \int_D G(\mathbf{x}) \nabla \delta u_{h,k+1}^{(i)}(\mathbf{x}) \cdot \nabla v_h(\mathbf{x}) d\mathbf{x} + \\ &\int_D I'_{\text{ion}}(u_{h,k+1}^{(i)}(\mathbf{x})) (\delta u_{h,k+1}^{(i)}(\mathbf{x})) v_h(\mathbf{x}) d\mathbf{x} = \\ &-\frac{1}{\Delta t} \int_D u_{h,k+1}^{(i)}(\mathbf{x}) v_h(\mathbf{x}) d\mathbf{x} - a(u_{h,k+1}^{(i)}, v_h) + l(v_h, t_{k+1}) + \frac{1}{\Delta t} \int_D u_{h,k}(\mathbf{x}) v_h(\mathbf{x}) d\mathbf{x}. \end{aligned} \quad (3.67)$$

Rewriting (3.67) by inserting a linear combination of $\delta u_{h,k+1}^{(i)}$ in terms of V_h basis elements, i.e.

$$\delta u_{h,k+1}^{(i)}(\mathbf{x}) = \sum_{j=1}^n \delta u_j \phi_j(\mathbf{x}), \quad (3.68)$$

yields the linear system to be solved. In this context, we have that the functions $u_{h,k+1}^{(i)}$ and $u_{h,k}$ are known by means of their nodal values which we denote with:

$$\mathbf{u}_{k+1}^{(i)} = [(u_{k+1}^{(i)})_1, \dots, (u_{k+1}^{(i)})_n]^T, \quad \mathbf{u}_k = [(u_k)_1, \dots, (u_k)_n]^T. \quad (3.69)$$

Based on that, the formulation (3.67) yields the linear system:

$$\left(\frac{1}{\Delta t} M + A \right) \delta \mathbf{u}^{(i)} + M I'_{\text{ion}}(\mathbf{u}_{k+1}^{(i)}) = \mathbf{b}, \quad (3.70)$$

where

$$\mathbf{b} := -\left(\frac{1}{\Delta t}M + A\right)\mathbf{u}_{k+1}^{(i)} + MI_{\text{ion}}(\mathbf{u}_{k+1}^{(i)}) + \mathbf{f} + \frac{1}{\Delta t}M\mathbf{u}_k. \quad (3.71)$$

Here, $I_{\text{ion}}(\mathbf{u}_{k+1}^{(i)}) \in \mathbb{R}^n$ and $I'_{\text{ion}}(\mathbf{u}_{k+1}^{(i)}) \in \mathbb{R}^n$ are given by

$$I_{\text{ion}}(\mathbf{u}_{k+1}^{(i)}) = [I_{\text{ion}}((\mathbf{u}_{k+1}^{(i)})_1), \dots, I_{\text{ion}}((\mathbf{u}_{k+1}^{(i)})_n)]^T, \quad (3.72)$$

and

$$I'_{\text{ion}}(\mathbf{u}_{k+1}^{(i)}) = [I'_{\text{ion}}((\mathbf{u}_{k+1}^{(i)})_1), \dots, I'_{\text{ion}}((\mathbf{u}_{k+1}^{(i)})_n)]^T, \quad (3.73)$$

whereas

$$\mathbf{f} := \left[\int_D I_{\text{app}}(\mathbf{x}, t_{k+1}) \phi_i(\mathbf{x}) \right]_{i=1}^n. \quad (3.74)$$

3.2.4 Space–time assembly of the linear problem

In the previous subsection, we have shown how to derive the linear system to be solved at every Newton iteration, for every time step. Nonetheless, we aim at assembling the system in an all–at–once fashion given the benefit from profiting also from parallelization in time. The space–time discretization of the monodomain equation (3.52) is an extension of the space–time assembly procedure described in Subsection 3.1.3 for the heat equation. In particular, the linear system (3.51) is modified to contain the discretization of the non-linear reaction term I_{ion} :

$$\mathbf{R}(\mathbf{u}) = C\mathbf{u} + \mathbf{r}(\mathbf{u}) - \mathbf{f} = \mathbf{0}, \quad (3.75)$$

where C and \mathbf{f} are respectively defined as in (3.51) and (3.50). The reaction term $\mathbf{r}(\mathbf{u}) \in \mathbb{R}^{nm}$ is given by

$$\mathbf{r}(\mathbf{u}) := (\Delta t I_m \otimes M) \mathbf{I}_{\text{ion}}(\mathbf{u}) \quad \text{with} \quad \mathbf{I}_{\text{ion}}(\mathbf{u}) := [I_{\text{ion}}(u_1), \dots, I_{\text{ion}}(u_{nm})]^T,$$

where n and m are respectively the space and time degrees of freedom. The non-linear equation (3.75) is further solved using the Newton method with a strategy of setting the space–time initial guess described in Chapter 6, Section 6.4. The residual derivative, or Jacobian, $\mathbf{DR}(\mathbf{u}) \in \mathbb{R}^{nm \times nm}$ of the non-linear residual $\mathbf{R}(\mathbf{u})$ (3.75) is given by

$$\mathbf{DR}(\mathbf{u}) = I_m \otimes (M + \Delta t A) + \Delta t I_m \otimes M \cdot \mathbf{J}_{\text{ion}}(\mathbf{u}).$$

with $\mathbf{J}_{\text{ion}}(\mathbf{u}) \in \mathbb{R}^{nm \times nm}$ the diagonal matrix defined by

$$\mathbf{J}_{\text{ion}}(\mathbf{u}) := \text{diag}([I'_{\text{ion}}(u_1), I'_{\text{ion}}(u_2), \dots, I'_{\text{ion}}(u_{nm})]).$$

$$\mathbf{J}_{\text{ion}}(\mathbf{u}) := \begin{bmatrix} I'_{\text{ion}}(u_1) & & & \\ & I'_{\text{ion}}(u_2) & & \\ & & \ddots & \\ & & & I'_{\text{ion}}(u_{nm}) \end{bmatrix}.$$

Newton's method requires solving the following space–time system at every iteration

$$\mathbf{DR}(\delta \mathbf{u}^{(i)}) = -\mathbf{R}(\mathbf{u}^{(i)}), \quad (3.76)$$

in which both the current iteration $\mathbf{u}^{(i)}$ and the update $\delta \mathbf{u}^{(i)}$ are provided as space–time vectors, i.e.

$$\mathbf{u}^{(i)} = [\mathbf{u}_1^{(i)}, \dots, \mathbf{u}_m^{(i)}]^T, \quad \delta \mathbf{u}^{(i)} = [\delta \mathbf{u}_1^{(i)}, \dots, \delta \mathbf{u}_m^{(i)}]^T. \quad (3.77)$$

The update for the Newton scheme is given by

$$\mathbf{u}^{(i+1)} = \mathbf{u}^{(i)} + \delta \mathbf{u}^{(i)}. \quad (3.78)$$

3.3 Comment on the error estimates

The combination of FE in space and FD in time, more specifically with a Crank-Nicolson scheme, naturally introduces a discretization error. It is important to be able to quantify this error, and have an a priori knowledge on how good a used method approximates the final solution.

Regarding the heat equation, many works in the literature have addressed this matter. In [QSS10] for instance, based on [QV08], it is shown that the error for the solution approximation of the heat equation can be bounded as follows:

$$\|u(\cdot, t_k) - u_{h,k}(\cdot)\|_{L^2(D)} \leq C(u_0, f, u)(\Delta t^2 + h^2), \quad (3.79)$$

where C is a function of its argument, and the solution approximated with linear FE and Crank-Nicolson. The error can also be approximated in the H^1 -norm yielding the following estimate

$$\|u(\cdot, t_k) - u_{h,k}(\cdot)\|_{H^1(D)} \leq C(u_0, f, u)(\Delta t^2 + h). \quad (3.80)$$

It is possible to rely on techniques such as the ones used in [AM89] to rewrite the following estimates in the Bochner norms, i.e.

$$\|u - u_h\|_{L^2([0,T];L^2(D))} \leq C(u_0, f, u, D, T)(\Delta t^2 + h^2), \quad (3.81)$$

$$\|u - u_h\|_{L^2([0,T];H^1(D))} \leq C(u_0, f, u, D, T)(\Delta t^2 + h). \quad (3.82)$$

Based on [RV19], these error estimates can be generalized to the case of the monodomain equation.

Chapter 4

Uncertainty Quantification

As motivated in Chapters 1 and 2, we are interested in equations that are subject to uncertain or unreliable data in the aim of evaluating the information deriving from the process of simulations. This concept has been summarized under the terminology of *uncertainty quantification* (UQ) and is concerned about extracting robust and qualitative information despite the presence of this variability. In the next section, we set the theoretical ground lying behind the characterization of the randomness to consider, and introduce the reader to the notation we rely on for the rest of the chapter. This part is largely based on the work introduced in [Pet14]. Many of the theoretical aspects presented in the aforementioned work, are reintroduced here for a matter of completion of the present document.

4.1 Foundations for UQ

Randomness in our problem of interest is defined through (random) fields defined all over the considered domain. For mathematical purposes, these fields need to be explicitly derived in terms of adequate functional spaces for our problem to be well-defined, mainly in terms of the smoothness that should characterize them.

In this section, we go through the construction of these so-called random fields, showing how they can be explicitly derived by means of the Karhunen-Loève expansion, that is a decomposition of the random fields in terms of orthonormal bases both for the stochastic and the spatial variables.

4.1.1 Random fields

Solving PDE problems with FE systematically relies on adequate functional spaces for all the interfering components. It is therefore of a crucial matter to address the question of what should be a suitable functional space for our random fields. Let $(\Omega, \mathcal{F}, \mathbb{P})$ denote a complete and separable probability space, where Ω is the sample space (or equivalently the universe of all possible outcomes), \mathcal{F} the σ -algebra and \mathbb{P} the probability measure.

Additionally, let us introduce $(\mathcal{B}, \|\cdot\|_{\mathcal{B}})$ to be a Banach space over the space \mathbb{R} . Following the conceptual idea of having a random event $\omega \in \Omega$ realizing a particular random field $V(\cdot, \omega) \in \mathcal{B}$, the natural environment for the random fields becomes therefore defined by means of the Lebesgue–Bochner spaces.

Definition 4.1 (Lebesgue–Bochner spaces). *The Lebesgue–Bochner spaces denoted by $L^p(\Omega, \mathcal{F}, \mathbb{P}; \mathcal{B})$ (here written as $L^p(\Omega; \mathcal{B})$ for simplicity), for $1 \leq p \leq \infty$, are defined to be the equivalence classes of strongly measurable maps $v : \Omega \rightarrow \mathcal{B}$ such that $\|v\|_{L^p(\Omega; \mathcal{B})} < \infty$, where the norm $\|\cdot\|_{L^p(\Omega; \mathcal{B})}$ is defined as follows*

$$\|v\|_{L^p(\Omega; \mathcal{B})} := \begin{cases} \left(\int_{\Omega} \|v(\cdot, \omega)\|_{\mathcal{B}}^p d\mathbb{P}(\omega) \right)^{1/p} & \text{if } 1 \leq p < \infty, \\ \text{ess sup}_{\omega \in \Omega} \|v(\cdot, \omega)\|_{\mathcal{B}} & \text{if } p = \infty. \end{cases}$$

An important property for such a Bochner space is that whenever $p = 2$ and \mathcal{B} is a separable Hilbert space, it is isomorphic to the tensor product space $L^2(\Omega) \otimes B$ ([BBN14]). This property is crucial in two aspects. Being a tensor product of two Hilbert spaces, the Lebesgue–Bochner space is also a Hilbert space. Additionally, as we will later detail, this property allows to decouple the Lebesgue–Bochner space and express every one of its elements as a tensorized decomposition of orthonormal basis spanning both the probabilistic space Ω and the Banach space \mathcal{B} .

Let us now be more explicit about the space \mathcal{B} . We consider a sufficiently smooth domain $D \in \mathbb{R}^d$ and set \mathcal{B} to be the Hilbert space $\mathcal{B} = [L^2(D)]^{d_r}$. Here, d_r is the dimension deciding for whether we are treating scalar or vectorial random fields and is typically set to either $d_r = 1$ or $d_r = d$.

Let $\mathbf{V} \in L^2(\Omega; [L^2(D)]^{d_r}) \cong L^2(\Omega) \otimes [L^2(D)]^{d_r}$ be a representation of a random field defined as

$$\mathbf{V}(\mathbf{x}, \omega) = [V_1(\mathbf{x}, \omega), \dots, V_{d_r}(\mathbf{x}, \omega)]^T. \quad (4.1)$$

By means of the Bochner integral, we additionally introduce the mean

$$\mathbb{E}[\mathbf{V}](\mathbf{x}) = \int_{\Omega} \mathbf{V}(\mathbf{x}, \omega) d\mathbb{P}(\omega) \in [L^2(D)]^{d_r}, \quad (4.2)$$

as well as the centered random field provided with

$$\mathbf{V}_0 = \mathbf{V} - \mathbb{E}[\mathbf{V}] \in L^2(\Omega; [L^2(D)]^{d_r}) \cong L^2(\Omega) \otimes [L^2(D)]^{d_r}. \quad (4.3)$$

By a result emanating from the theory on tensor products of Hilbert spaces and their representation, see [Pet14], \mathbf{V}_0 exhibits an orthogonal decomposition, that is

$$\mathbf{V}_0 = \sum_{k \in \mathcal{A}} \sigma_k X_k \otimes \phi_k, \quad (4.4)$$

where $\{X_k\}_{k \in \mathcal{A}} \subset L^2(\Omega)$ and $\{\phi_k\}_{k \in \mathcal{A}} \subset [L^2(D)]^{d_r}$ are respectively orthonormal in $L^2(\Omega)$ and $[L^2(D)]^{d_r}$. On the other hand, $\{\sigma_k\}_{k \in \mathcal{A}} \subset \mathbb{R}$. The centered random field is therefore completely determined by means of $\{\sigma_k\}_{k \in \mathcal{A}}$, $\{X_k\}_{k \in \mathcal{A}}$ and $\{\phi_k\}_{k \in \mathcal{A}}$.

Let us now define the covariance of \mathbf{V} to be given by

$$\begin{aligned} \text{Cov}[\mathbf{V}](\mathbf{x}, \mathbf{y}) &= \mathbb{E}[(\mathbf{V}(\mathbf{x}) - \mathbb{E}[\mathbf{V}(\mathbf{x})])(\mathbf{V}(\mathbf{y}) - \mathbb{E}[\mathbf{V}(\mathbf{y})])] \\ &= \mathbb{E}[\mathbf{V}_0(\mathbf{x})\mathbf{V}_0^T(\mathbf{y})] \\ &= \int_{\Omega} \mathbf{V}_0(\mathbf{x}, \omega)\mathbf{V}_0^T(\mathbf{y}, \omega) d\mathbb{P}(\omega) \in [L^2(D \times D)]^{d_r \times d_r} \end{aligned} \quad (4.5)$$

The boundedness of the covariance, i.e. $\text{Cov}[\mathbf{V}] \in [L^2(D \times D)]^{d_r \times d_r}$, derives from the element-wise boundedness of $(\text{Cov}[\mathbf{V}])_{ij}(\mathbf{x}, \mathbf{y}) = \int_{\Omega} V_{0,i}(\mathbf{x}, \omega)V_{0,j}(\mathbf{y}, \omega) d\mathbb{P}(\omega)$ for $1 \leq i, j \leq d_r$, that is a consequence of the Cauchy–Schwarz inequality.

Following the definition of the covariance, we introduce the Hilbert–Schmidt operator \mathcal{C} associated to the covariance, defined by:

$$\begin{aligned} \mathcal{C} : [L^2(D)]^{d_r} &\rightarrow [L^2(D)]^{d_r} \\ (\mathcal{C}\mathbf{u})(\mathbf{x}) &\mapsto (\mathcal{C}\mathbf{u})(\mathbf{x}) = \int_D \text{Cov}[\mathbf{V}](\mathbf{x}, \mathbf{y})\mathbf{u}(\mathbf{y}) d\mathbf{y}. \end{aligned} \quad (4.6)$$

\mathcal{C} is a linear operator, and we can therefore evoke its eigenvalues and eigenfunctions. If additionally $\text{Cov}[\mathbf{V}]$ is symmetric non-negative, then by the spectral theorem for compact operators, the covariance has a representation given by

$$\text{Cov}[\mathbf{V}] = \sum_{k \in \mathcal{A}} \lambda_k \phi_k \otimes \phi_k, \quad (4.7)$$

where $\{\phi_k\}_{k \in \mathcal{A}}$ and non-negative $\{\lambda_k\}_{k \in \mathcal{A}}$ are respectively eigenvalues and eigenfunctions of \mathcal{C} . Furthermore, the $\{\phi_k\}_{k \in \mathcal{A}}$ correspond to those used in \mathbf{V}_0 decomposition (c.f. (4.4)) in which we also have that $\lambda_k = \sigma_k^2$. The latter affirmation can be elegantly proven by interpreting (with analogy to Hilbert–Schmidt

operators on tensor product of Hilbert spaces) the covariance as the element corresponding to the trace–class operator defined around the representative \mathbf{V}_0 and its adjoint, c.f. [Pet14].

4.1.2 The Karhunen–Loève expansion

In theory, random fields are now very well characterized by what has been developed in the previous subsection. Indeed, from the orthogonal decomposition of \mathbf{V}_0 , the random fields are well characterized by the orthonormal basis $\{X_k\}_{k \in \mathcal{A}}$ and $\{\phi_k\}_{k \in \mathcal{A}}$, from which we could retrieve the random fields:

$$\mathbf{V}(\mathbf{x}, \omega) = \mathbb{E}[\mathbf{V}](\mathbf{x}) + \sum_{k \in \mathcal{A}} \sigma_k X_k(\omega) \phi_k(\mathbf{x}). \quad (4.8)$$

This expression is known as the Karhunen–Loève (KL) expansion. Assuming that the covariance operator defined in (4.5) is known, the KL expansion for the random fields can be retrieved by solving the eigenvalue problem related to the covariance operator. In practice however, the information on the covariance might be unknown. Therefore, the random fields in our case are actually modeled by reverse engineering the machinery described above, via an empirical correlation kernel prescribing the covariance matrix.

More concretely, this is achieved by setting $(\text{Cov}[\mathbf{V}])_{ij}(\mathbf{x}, \mathbf{y}) = k_{ij}(\|\mathbf{x} - \mathbf{y}\|_2)$, $1 \leq i, j \leq d_r$ for functions k_{ij} , where $\|\mathbf{x} - \mathbf{y}\|_2$ is the euclidian distance between \mathbf{x} and \mathbf{y} .

For the random variables $\{X_k\}_{k \in \mathcal{A}}$, we make the reasonable assumptions to have them independent and identically distributed, centered, i.e. $\mathbb{E}[X_k] = 0$, with bounded values in $[-1, 1]$ ($-1 \leq X_k(\omega) \leq 1$, $\forall k \in \mathcal{A}$ and $\forall \omega \in \Omega$). The latter assumption is compensated by introducing a scaling factor $\mathbf{s} = [s_1, \dots, s_{d_r}]^T$, with $s_i \in \mathbb{R}$ for $1 \leq i \leq d_r$. We introduce the matrix $S \in \mathbb{R}^{d_r \times d_r}$ with

$$S = \text{diag}(\mathbf{s}) \begin{bmatrix} s_1 & & \\ & \ddots & \\ & & s_{d_r} \end{bmatrix} \quad (4.9)$$

and we rewrite the KL expansion as:

$$\mathbf{V}(\mathbf{x}, \omega) = \mathbb{E}[\mathbf{V}](\mathbf{x}) + S \sum_{k \in \mathcal{A}} \sigma_k X_k(\omega) \phi_k(\mathbf{x}). \quad (4.10)$$

In the scalar case ($d_r = 1$), this will serve to preserve the uniform ellipticity condition (c.f. Subsection 4.2.2), whereas in the vectorial case ($d_r = d$) this can be used to impose a reasonable variability with respect to $\mathbb{E}[\mathbf{V}](\mathbf{x})$.

4.1.3 Discretization of the eigenvalue problem

In the context of the methodology introduced above, generating the random fields relies on solving the eigenvalue problem. So far, we have only been evoking the covariance as lying in a given infinite dimensional functional space, i.e. $\text{Cov}[\mathbf{V}] \in [L^2(D)]^{d_r \times d_r}$. However, in order to have the problem numerically feasible, a discretization must come into play.

The eigenvalue problem resides in finding $\{(\lambda_k, \boldsymbol{\phi}_k)\}_{k=1}^\infty$ such that

$$\begin{cases} \mathcal{C} \boldsymbol{\phi}_k = \lambda_k \boldsymbol{\phi}_k \\ (\boldsymbol{\phi}_k, \boldsymbol{\phi}_{k'})_{[L^2(D)]^{d_r}} = \delta_{kk'}, \end{cases} \quad (4.11)$$

where \mathcal{C} is the covariance operator defined in (4.6). We express the weak formulation of this problem by considering test functions $\mathbf{v} \in [L^2(D)]^{d_r}$, such that

$$\begin{cases} (\mathcal{C} \boldsymbol{\phi}_k, \mathbf{v})_{[L^2(D)]^{d_r}} = \lambda_k (\boldsymbol{\phi}_k, \mathbf{v})_{[L^2(D)]^{d_r}}, \quad \forall \mathbf{v} \in [L^2(D)]^{d_r} \\ (\boldsymbol{\phi}_k, \boldsymbol{\phi}_{k'})_{[L^2(D)]^{d_r}} = \delta_{kk'}. \end{cases} \quad (4.12)$$

Let us now consider a finite dimensional subspace $U_h \subset [L^2(D)]^{d_r}$ of dimension N_h , with its spanning basis $\{\boldsymbol{\xi}_k\}_{k=1}^{N_h}$. The Galerkin approximation deriving from (4.12) is now to find $\{(\lambda_{k,h}, \boldsymbol{\phi}_{k,h})\}_{k=1}^{N_h}$ such that

$$\begin{cases} (\mathcal{C} \boldsymbol{\phi}_{k,h}, \mathbf{v}_h)_{[L^2(D)]^{d_r}} = \lambda_{k,h} (\boldsymbol{\phi}_{k,h}, \mathbf{v}_h)_{[L^2(D)]^{d_r}}, \quad \forall \mathbf{v}_h \in U_h \\ (\boldsymbol{\phi}_{k,h}, \boldsymbol{\phi}_{k',h})_{[L^2(D)]^{d_r}} = \delta_{kk'}. \end{cases} \quad (4.13)$$

Since $\{\boldsymbol{\xi}_k\}_{k=1}^{N_h}$ represent a basis for U_h , every $\boldsymbol{\phi}_{k,h}$ can be written as the linear combination $\boldsymbol{\phi}_{k,h}(\mathbf{x}) = \sum_{j=1}^{N_h} \chi_{k,j} \boldsymbol{\xi}_j(\mathbf{x})$. By inserting this in (4.13), and substituting the test function \mathbf{v}_h with $\boldsymbol{\xi}_i$ for $i = 1, \dots, N_h$, we recover:

$$\begin{cases} (\mathcal{C} [\sum_{j=1}^{N_h} \chi_{k,j} \boldsymbol{\xi}_j], \boldsymbol{\xi}_i)_{[L^2(D)]^{d_r}} = \lambda_{k,h} (\sum_{j=1}^{N_h} \chi_{k,j} \boldsymbol{\xi}_j, \boldsymbol{\xi}_i)_{[L^2(D)]^{d_r}} \\ (\sum_{j=1}^{N_h} \chi_{k,j} \boldsymbol{\xi}_j, \sum_{j'=1}^{N_h} \chi_{k',j'} \boldsymbol{\xi}_{j'})_{[L^2(D)]^{d_r}} = \delta_{kk'}. \end{cases} \quad (4.14)$$

Using the linearity and the symmetry of the scalar product in $[L^2(D)]^{d_r}$, we end up with

$$\begin{cases} \sum_{j=1}^{N_h} (\boldsymbol{\xi}_i, \mathcal{C} \boldsymbol{\xi}_j)_{[L^2(D)]^{d_r}} \chi_{k,j} = \lambda_{k,h} \sum_{j=1}^{N_h} (\boldsymbol{\xi}_i, \boldsymbol{\xi}_j)_{[L^2(D)]^{d_r}} \chi_{k,j} \\ \sum_{j=1}^{N_h} \sum_{j'=1}^{N_h} \chi_{k,j} (\boldsymbol{\xi}_j, \boldsymbol{\xi}_{j'})_{[L^2(D)]^{d_r}} \chi_{k',j'} = \delta_{kk'}. \end{cases} \quad (4.15)$$

Finally, if we set $C \in \mathbb{R}^{N_h \times N_h}$ and $M \in \mathbb{R}^{N_h \times N_h}$ to be the matrices defined by

$$(C)_{ij} = (\xi_i, \mathcal{C} \xi_j)_{[L^2(D)]^{d_r}} \quad (4.16)$$

and

$$(M)_{ij} = (\xi_i, \xi_j)_{[L^2(D)]^{d_r}}, \quad (4.17)$$

we end up solving the following system:

for $k = 1, \dots, N_h$, find $\boldsymbol{\chi}_k = [\chi_{k,1}, \dots, \chi_{k,N_h}]^T$ and $\lambda_{k,h}$ such that:

$$\begin{cases} C \boldsymbol{\chi}_k = \lambda_{k,h} M \boldsymbol{\chi}_k \\ \boldsymbol{\chi}_k^T M \boldsymbol{\chi}_{k'} = \delta_{kk'} \end{cases} \quad (4.18)$$

The obtained system represents a discretization of the initial eigenvalue problem (4.11). Yet, the matrix C in (4.16) is defined with respect to the covariance operator \mathcal{C} that is the Hilbert–Schmidt operator for the covariance $\text{Cov}[\mathbf{V}]$, recall

$$(\mathcal{C} \mathbf{u})(\mathbf{x}) = \int_D \text{Cov}[\mathbf{V}](\mathbf{x}, \mathbf{y}) \mathbf{u}(\mathbf{y}) d\mathbf{y} \quad (4.19)$$

We therefore still need a suitable discretization for expressing $\text{Cov}[\mathbf{V}](\mathbf{x}, \mathbf{y}) \in [L^2(D \times D)]^{d_r \times d_r}$. Given the tensor product structure of $\text{Cov}[\mathbf{V}]$ (c.f. (4.7)) and with respect to the canonical map

$$\begin{aligned} [L^2(D)]^{d_r} \otimes [L^2(D)]^{d_r} &\rightarrow [L^2(D \times D)]^{d_r \times d_r} \\ \xi_k \otimes \xi_l &\mapsto \xi_k(\mathbf{x}) \xi_l^T(\mathbf{y}), \end{aligned}$$

we can approximate the covariance by means of the following expression:

$$\text{Cov}[\mathbf{V}](\mathbf{x}, \mathbf{y}) \approx \sum_{k,l=1}^{N_h} c_{kl} \xi_k(\mathbf{x}) \xi_l^T(\mathbf{y}). \quad (4.20)$$

We also define the matrix $\tilde{C} \in \mathbb{R}^{N_h \times N_h}$ with

$$(\tilde{C})_{kl} = c_{kl}. \quad (4.21)$$

This results in having:

$$\begin{aligned} \mathcal{C} \xi_j(\mathbf{x}) &\approx \int_D \left(\sum_{k,l=1}^{N_h} c_{kl} \xi_k(\mathbf{x}) \xi_l^T(\mathbf{y}) \right) \xi_j(\mathbf{y}) d\mathbf{y} \\ &= \sum_{k,l=1}^{N_h} c_{kl} \xi_k(\mathbf{x}) \int_D \xi_l^T(\mathbf{y}) \xi_j(\mathbf{y}) \\ &= \sum_{k,l=1}^{N_h} c_{kl} \xi_k(\mathbf{x}) M_{lj}, \end{aligned}$$

and therefore,

$$\begin{aligned} (\xi_i, \mathcal{C} \xi_j(\mathbf{x}))_{[L^2(D)]^{d_r}} &\approx \int_D \xi_i^T(\mathbf{x}) \left(\sum_{k,l=1}^{N_h} c_{kl} \xi_k(\mathbf{x}) M_{lj} \right) d\mathbf{x} \\ &= \sum_{k,l=1}^{N_h} c_{kl} M_{lj} \int_D \xi_i^T(\mathbf{x}) \xi_k(\mathbf{x}) d\mathbf{x} \\ &= \sum_{k,l=1}^{N_h} M_{ik} c_{kl} M_{lj}. \end{aligned}$$

This leads to the approximation

$$M\tilde{C}M \approx C. \quad (4.22)$$

By replacing C with the above approximation in (4.18), we obtain the eigenvalue problem defined by

$$\begin{cases} M\tilde{C}M \boldsymbol{\chi}_k = \lambda_{k,h} M \boldsymbol{\chi}_k \\ \boldsymbol{\chi}_k^T M \boldsymbol{\chi}_{k'} = \delta_{kk'} \end{cases} \quad (4.23)$$

Approximating the initial problem (4.11) with the generalized eigenvalue problem (4.23) introduces an error that is the result of the approximation of the eigenvalues and the eigenfunctions of the covariance operator, combined to an additional error deriving from the spatial approximation of the covariance matrix. In [Pet14], it is shown that provided an extra spatial regularity of the random field, i.e. $\mathbf{V} \in L^2(\Omega) \otimes [H^q(D)]^{d_r}$ with $q > 0$, then the total error brought by the process of discretization is bounded by $h^{2\min\{2,q\}}$.

4.1.4 Low rank approximation

The process of approximation of the infinite dimensional space $[L^2(D)]^{d_r}$ with the finite dimensional space U_h leads to solve the eigenvalue problem (4.23)

of size N_h . In general, one should expect N_h to be given by a high number, as to capture the spatial resolution of the solution. This evidently leads not only to performance issues due to the size of the problem, but also eventually to memory limits due to the density of the matrix \tilde{C} (therefore $M\tilde{C}M$).

One might therefore think of approximating the eigenvalue problem (4.23) with a smaller one of size $M \ll N_h$. This suggests that the random fields would be retrieved by means of a truncated KL expansion

$$\mathbf{V}(\mathbf{x}, \omega) = \mathbb{E}[\mathbf{V}](\mathbf{x}) + \sum_{k=1}^M \sigma_k X_k(\omega) \boldsymbol{\phi}_k(\mathbf{x}). \quad (4.24)$$

Undoubtedly, for (4.24) to be a good approximation of the original infinite sum KL expansion, a significant proportion of the information deriving from the eigenvalue problem (λ_k and $\boldsymbol{\phi}_k$) needs to be carried by the first M terms. It is actually sufficient whenever the eigenvalues exhibit a certain decay, to not only achieve the desired objective, but furthermore rigorously control the truncation error. We leave the following argument to be addressed later.

Since \tilde{C} is a symmetric positive semi-definite matrix, it can exactly be expressed in terms of its Cholesky decomposition, i.e. $\tilde{C} = LL^T$ for $L \in \mathbb{R}^{N_h \times N_h}$. A variant of this would instead be to approximate \tilde{C} with a low-rank Cholesky decomposition, that is $\tilde{C} \approx LL^T$ for $L \in \mathbb{R}^{N_h \times M}$ with $M \ll N_h$. The new eigenvalue problem to consider becomes therefore:

for $k = 1, \dots, N_h$, find $\mathbf{v}_k = [v_{k,1}, \dots, v_{k,N_h}]^T$ and $\lambda_{k,h}$ such that:

$$\begin{cases} MLL^T M \mathbf{v}_k = \lambda_{k,h} M \mathbf{v}_k \\ \mathbf{v}_k^T M \mathbf{v}_{k'} = \delta_{kk'} \end{cases} \quad (4.25)$$

Let us consider the subsidiary eigenvalue problem given by:

for $k = 1, \dots, M$, find the M eigenfunctions $\tilde{\mathbf{v}}_k = [\tilde{v}_{k,1}, \dots, \tilde{v}_{k,M}]^T$ and the M eigenvalues $\tilde{\lambda}_{k,h}$ such that:

$$\begin{cases} L^T M L \tilde{\mathbf{v}}_k = \tilde{\lambda}_{k,h} \tilde{\mathbf{v}}_k \\ \tilde{\mathbf{v}}_k^T M \tilde{\mathbf{v}}_{k'} = \delta_{kk'} \end{cases} \quad (4.26)$$

If $\tilde{\mathbf{v}}_k$ is an eigenvector of (4.26), then $\mathbf{v}_k := L \tilde{\mathbf{v}}_k$ is an eigenvector of (4.25) for the same eigenvalue $\lambda_{k,h} := \tilde{\lambda}_{k,h}$ since

$$MLL^T M (L \tilde{\mathbf{v}}_k) = ML (L^T M L) \tilde{\mathbf{v}}_k = \tilde{\lambda}_{k,h} M (L \tilde{\mathbf{v}}_k). \quad (4.27)$$

Furthermore,

$$(L \tilde{\mathbf{v}}_k)^T M L \tilde{\mathbf{v}}_{k'} = \tilde{\mathbf{v}}_k^T L^T M L \tilde{\mathbf{v}}_{k'} = \tilde{\lambda}_{k',h} \tilde{\mathbf{v}}_k^T \tilde{\mathbf{v}}_{k'} = \tilde{\lambda}_{k',h} \delta_{kk'}. \quad (4.28)$$

Therefore, setting $\mathbf{v}_k = L\tilde{\mathbf{v}}_k$ leads to already scaled eigenfunctions by $\sqrt{\tilde{\lambda}_{k,h}}$.

Note about the low rank approximation: The low rank approximation can be computed by means of Pivoted Cholesky Decomposition (PCD) as described in [HPS15]. This algorithm presents the main advantage of not requiring to fill the matrix \tilde{C} with all of its entries, but instead computing one row at the time when this is needed.

4.1.5 Truncation error

The process of low rank approximation of the discretized covariance matrix introduces an additional truncation error to that provided at the end of Subsection 4.1.3. The behaviour of this error is closely related to the eigenvalues decay rate of the covariance operator.

Eigenvalues decay rate: We start by characterizing the random field \mathbf{V} with an additional smoothness, i.e. $\mathbf{V} \in L^2(\Omega) \otimes [H^q(D)]^{d_r}$ for $q \geq 1$. This implies that $\text{Cov}[\mathbf{V}] \in [H^q(D \times D)]^{d_r \times d_r}$. Furthermore, $\text{Cov}[\mathbf{V}]$ can still be expressed as in (4.7), i.e. $\text{Cov}[\mathbf{V}] = \sum_{k \in \mathcal{A}} \lambda_k \boldsymbol{\phi}_k \otimes \boldsymbol{\phi}_k$, and we recall that:

$$\|\mathbf{V}_0\|_{L^2(\Omega) \otimes [L^2(D)]^{d_r}} = (\text{Tr} \mathcal{C})^{1/2} = \sqrt{\sum_{k \in \mathcal{A}} \sigma_k^2} = \sqrt{\sum_{k \in \mathcal{A}} \lambda_k}. \quad (4.29)$$

In [Pet14] is developed a generalization of the results in [GH14] and [ST06] to the case of vector fields, showing that the eigenvalues of the covariance operator decay as:

$$\lambda_k \lesssim \left(\frac{k}{d_r}\right)^{-2q/d} = \left(\frac{d_r}{k}\right)^{2q/d}, \text{ as } k \rightarrow \infty. \quad (4.30)$$

This suggests an asymptotical decay of the eigenvalues, accelerated by the smoothness characterizing the random fields.

Truncation error for the original (non-discretized) eigenvalue problem:

The effect of the truncation error on the original problem is characterized by the following error:

$$\begin{aligned} \left\| \mathbf{V}_0 - \sum_{k=1}^M \sigma_k X_k \otimes \boldsymbol{\phi}_k \right\|_{L^2(\Omega) \otimes [L^2(D)]^{d_r}} &= \left\| \sum_{k=M+1}^{\infty} \sigma_k X_k \otimes \boldsymbol{\phi}_k \right\|_{L^2(\Omega) \otimes [L^2(D)]^{d_r}} \\ &= \left(\sum_{k=M+1}^{\infty} \sigma_k^2 \right)^{1/2} \\ &= \left(\sum_{k=M+1}^{\infty} \lambda_k \right)^{1/2}. \end{aligned}$$

Recall (4.30),

$$\begin{aligned} \sum_{k=M+1}^{\infty} \lambda_k &\lesssim \sum_{k=M+1}^{\infty} \left(\frac{d_r}{k}\right)^{\alpha(q)}, \text{ for } \alpha(q) = \frac{2q}{d} > 0 \\ &\leq \int_M^{\infty} \left(\frac{d_r}{x}\right)^{\alpha(q)} dx, \text{ since } f(x) = \left(\frac{d_r}{x}\right)^{\alpha(q)} \text{ is decreasing on } [M, \infty] \\ &= d_r^{\alpha(q)} \frac{x^{-\alpha(q)+1}}{-\alpha(q)+1} \Big|_M^{\infty} = d_r^{\alpha(q)} \frac{M^{1-\alpha(q)}}{\alpha(q)-1}, \text{ for } \alpha(q) \neq 1. \end{aligned}$$

Therefore, if we would require a controlled accuracy for an arbitrary $\epsilon > 0$, i.e.

$$\left\| \mathbf{v}_0 - \sum_{k=1}^M \sigma_k X_k \otimes \phi_k \right\|_{L^2(\Omega) \otimes [L^2(D)]^{d_r}} = \left(\sum_{k=M+1}^{\infty} \lambda_k \right)^{1/2} \leq \left(d_r^{\alpha(q)} \frac{M^{1-\alpha(q)}}{\alpha(q)-1} \right)^{1/2} \approx \epsilon, \quad (4.31)$$

we would have:

$$d_r^{\alpha(q)} \frac{M^{1-\alpha(q)}}{\alpha(q)-1} \approx \epsilon^2 \implies M^{1-\alpha(q)} \approx \frac{(\alpha(q)-1)\epsilon^2}{d_r^{\alpha(q)}}. \quad (4.32)$$

This is possible only when $\alpha(q) > 1$, for which we would have

$$M \approx \left(\frac{(\alpha(q)-1)\epsilon^2}{d_r^{\alpha(q)}} \right)^{1/(1-\alpha(q))}. \quad (4.33)$$

Truncation error for the discretized eigenvalue problem:

As detailed in Subsection 4.1.3, the eigenvalues and eigenfunctions $\{(\lambda_k, \phi_k)\}_{k=1}^M$ are obtained after discretization, leading to approximations $\{(\lambda_{k,h}, \phi_{k,h})\}_{k=1}^M$ of the original quantities of interest. We therefore would like to investigate the decay rate of the approximated eigenvalues, as well as ensure the existence of a rank M for which the truncation error becomes arbitrarily small.

We aim to provide a bound for

$$I = \left\| \sum_{k=1}^M \sigma_k X_k \otimes \phi_k - \sum_{k=1}^M \sigma_{k,h} X_k \otimes \phi_{k,h} \right\|_{L^2(\Omega) \otimes [L^2(D)]^{d_r}}. \quad (4.34)$$

For this purpose, we use the following identity:

$$\sigma_k X_k \otimes \phi_k - \sigma_{k,h} X_k \otimes \phi_{k,h} = \sigma_k X_k \otimes (\phi_k - \phi_{k,h}) + (\sigma_k - \sigma_{k,h}) X_k \otimes \phi_{k,h}, \quad (4.35)$$

in such way that we get the following inequality:

$$I \leq \left\| \sum_{k=1}^M \sigma_k X_k \otimes (\phi_k - \phi_{k,h}) \right\| + \left\| \sum_{k=1}^M (\sigma_k - \sigma_{k,h}) X_k \otimes \phi_{k,h} \right\|, \quad (4.36)$$

where the norm stands for that in $L^2(\Omega) \otimes [L^2(D)]^{d_r}$. In [GH14], the following inequalities are shown to hold:

$$\|\phi_k - \phi_{k,h}\|_{[L^2(D)]^{d_r}} \lesssim \frac{h^q}{\sqrt{\lambda_k}}, \quad (4.37)$$

and

$$0 \leq \lambda_k - \lambda_{k,h} \lesssim \lambda_k \|\phi_k - \phi_{k,h}\|_{[L^2(D)]^{d_r}}^2 \lesssim h^{2q}. \quad (4.38)$$

We therefore get for the first part of the inequality (4.36):

$$\begin{aligned} \left\| \sum_{k=1}^M \sigma_k X_k \otimes (\phi_k - \phi_{k,h}) \right\| &\lesssim \sqrt{\sum_{k=1}^M \lambda_k \underbrace{\|X_k\|_{L^2(\Omega)}^2}_{=1} \|\phi_k - \phi_{k,h}\|^2} \\ &= \sqrt{\sum_{k=1}^M h^{2q}} = \sqrt{M} h^q. \end{aligned} \quad (4.39)$$

With respect to the second part of the inequality (4.36), we have:

$$\begin{aligned} \left\| \sum_{k=1}^M (\sigma_k - \sigma_{k,h}) X_k \otimes \phi_{k,h} \right\| &\lesssim \sqrt{\sum_{k=1}^M (\sigma_k - \sigma_{k,h})^2} \\ &\leq \sqrt{\sum_{k=1}^M (\lambda_k - \lambda_{k,h})} \lesssim \sqrt{\sum_{k=1}^M h^{2q}} = \sqrt{M} h^q \end{aligned} \quad (4.40)$$

Consequently, we get the condition under which I becomes arbitrarily small:

$$I \leq 2\sqrt{M} h^q \approx \epsilon. \quad (4.41)$$

Mixed condition for total truncation error control:

The total truncation error can be interpreted as the superposition of the truncation error on the non-discretized eigenvalue problem, and the error emanating from the approximation of the eigenvalues and eigenfunctions. This is motivated by the following inequality:

$$\begin{aligned} \left\| \mathbf{V}_0 - \sum_{k=1}^M \sigma_{k,h} X_k \otimes \phi_{k,h} \right\| &= \left\| \mathbf{V}_0 - \sum_{k=1}^M \sigma_k X_k \otimes \phi_k + \sum_{k=1}^M \sigma_k X_k \otimes \phi_k - \sum_{k=1}^M \sigma_{k,h} X_k \otimes \phi_{k,h} \right\| \\ &\leq \left\| \mathbf{V}_0 - \sum_{k=1}^M \sigma_k X_k \otimes \phi_k \right\| + \left\| \sum_{k=1}^M \sigma_k X_k \otimes \phi_k - \sum_{k=1}^M \sigma_{k,h} X_k \otimes \phi_{k,h} \right\| \end{aligned}$$

Consequently, for the eigenvalue problem approximated solution to be arbitrarily smaller than a given tolerance $\epsilon' = 2\epsilon$, the equality (4.33) provides the condition on the truncation number. The condition on the mesh size is provided by (4.41). Insuring both condition guarantees a good overall truncated discretized approximation of the solution to the eigenvalue problem.

4.1.6 The covariance in brief

We now formulate a synthesis of all the primarily presented theoretical material on random fields, meant as a concrete methodology for practical usage. It has been argued in Subsection 4.1.2 that the random fields are expressed by means of the KL expansion, in the form of a superposition of an empirical mean field $\mathbb{E}[\mathbf{V}](\mathbf{x})$, and a perturbation provided in terms of the sum given by $\sum_{k \in \mathcal{A}} \sqrt{\lambda_k} X_k(\omega) \phi_k(\mathbf{x})$, where λ_k and ϕ_k are respectively eigenvalues and eigenfunctions of the covariance operator. The latter is also provided empirically, by means of correlation kernel functions $k_{ij}(r)$, for $1 \leq i, j \leq d$, where $r = r(\mathbf{x}, \mathbf{x}')$ is a distance measure between two given spatial points \mathbf{x} and \mathbf{x}' , of the considered domain D .

On the other hand, the efficiency of the low-rank approximation and the control of the truncation error is strongly dependent on the eigenvalues decay, that in turns is determined by the smoothness of the covariance as suggested by (4.30) and (4.41). Consequently, a special class of functions constitutes a preferential choice for correlation kernels, namely the Matérn class of functions. These can be expressed by means of the following formulation, see [Ras03],

$$k_{p+1/2}(r) = \exp\left(-\frac{\sqrt{2(p+1/2)r}}{\sigma}\right) \frac{p!}{(2p)!} \sum_{k=0}^p \frac{(p+k)!}{k!(p-k)!} \left(\frac{\sqrt{8(p+1/2)r}}{\sigma}\right)^{p-k} \quad (4.42)$$

for $p \in \mathbb{N}$ and $\sigma \in (\mathbb{R}^+)^*$. The reason for which this class of functions is particularly attractive for numerical examples is that their eigenvalues decay is known in advance, as shown in [GKN⁺15]

$$\lambda_k \leq C k^{-(1+2(p+1)/d)}. \quad (4.43)$$

In particular, for $p = 0$ and $p = \infty$ we have

$$k_{1/2}(r) = \exp\left(-\frac{r}{\sigma}\right), \quad (4.44)$$

and

$$k_{\infty}(r) = \exp\left(-\frac{r^2}{2\sigma^2}\right). \quad (4.45)$$

Written as such, the case $p = \infty$ defines a Gaussian bell in which the parameter σ represents the standard deviation. In the context of covariance kernels, this defines a correlation length that prescribes the magnitude of local variations. Using the Matérn class functions therefore prescribes spatial correlation on the random fields, defined by the parameter σ .

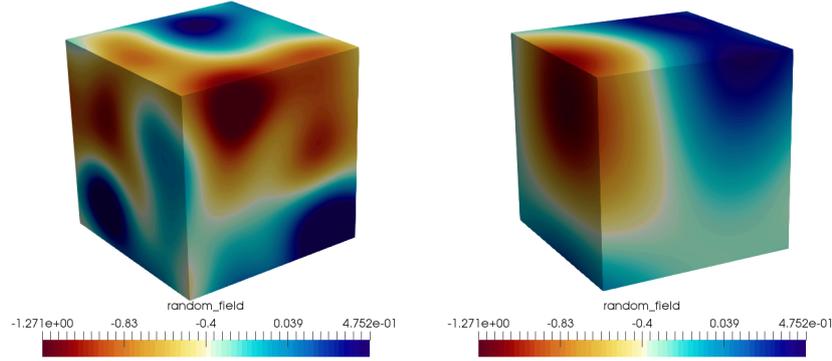


Figure 4.1. Effect of the correlation length for squared exponential kernel on the local variations. On the left, a random field field obtained for $\sigma = 0.1$. On the right, a random field field obtained for $\sigma = 0.5$.

1D Analytical example: For $D = \mathbb{R}$, an interesting analytical example can be found in [RW06]. This is a different situation from that described along this work, as it shows the eigenfunctions and eigenvalues of

$$\int_D k(x, x') \phi(x) d\mu(x) = \lambda \phi(x'),$$

for when there is a density function $p(x)$ such that $d\mu(x) = p(x)dx$. Nonetheless, it can be of interest to notice the behaviour of the eigenfunctions, and the decay of eigenvalues through this example. Provided a kernel given by $k(x, x') = \exp\left(-\frac{(x-x')^2}{2\sigma^2}\right)$ and a Gaussian density function given by $p(x) = \mathcal{N}(0, l^2)$, the eigenvalues and eigenfunctions are given by:

$$\lambda_k = \sqrt{\frac{2a}{A}} B^k, \quad k = 0, 1, 2, \dots \quad (4.46)$$

$$\phi_k(x) = \exp(-(c-a)x^2) H_k(\sqrt{2c}x), \quad k = 0, 1, 2, \dots \quad (4.47)$$

where

$$a = (4l^2)^{-1}, \quad b = (2\sigma^2)^{-1}, \quad (4.48)$$

$$c = \sqrt{a^2 + 2ab}, \quad A = a + b + c, \quad B = b/A, \quad (4.49)$$

and H_k represents the Hermite polynomial of order k . We report in Figure 4.2 the plot some of the eigenvalues and the eigenfunctions provided by (4.46) and (4.48) for the values of $\sigma = 0.4$ and $l = 0.5$.

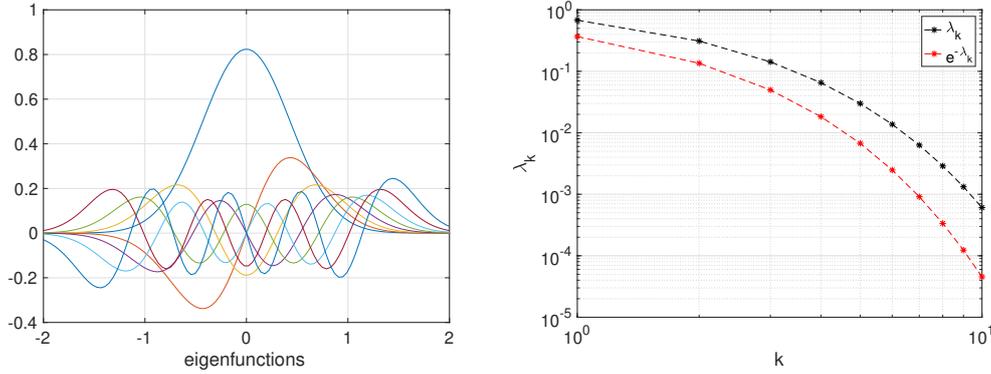


Figure 4.2. Eigenfunctions (left) and eigenvalues (right) of the squared Gaussian kernel weighted with normal density.

The noticeable fact about the Figure 4.2 is the shape of the eigenfunctions, as they exhibit properties of (anti)symmetry and periodicity. This is very close to what can be observed for Fourier modes for example. One can also see how the eigenvalues follow an exponential decay, in accordance to (4.43) for $p = \infty$.

Remark. *It is important to recall that the latter example is not directly concerned with the theory developed in this chapter, and that it does not apply to bounded domains. However, the eigenfunctions and eigenvalues behaviour is very similar in general, demonstrating respectively symmetry and decay properties. Figures 4.5 and 4.6 show an example of weighted eigenfunctions (by means of the eigenvalues square roots) in the cases $d_r = 1$ and $d_r = 3$ respectively.*

Methodology-from continuous correlation kernels to discrete random fields:

Considering a sufficiently fine mesh discretization resulting in a set of nodal vertices $\{\mathbf{x}_i\}_{i=1}^{N_h}$, the associated covariance matrix $C \in \mathbb{R}^{d_r N_h \times d_r N_h}$ is given by:

$$(C)_{mn} = k_{st}(\mathbf{x}_i, \mathbf{x}_j), \quad (4.50)$$

for $1 \leq m, n \leq d_r N_h$, where $s = \lfloor m/N_h \rfloor$, $t = \lfloor n/N_h \rfloor$ and $i \equiv m \pmod{N_h}$, $j \equiv n \pmod{N_h}$. Here, k_{st} are correlation kernels in the form of (4.44) or (4.45), for $1 \leq s, t \leq d_r$.

The PCD algorithm can be used as a low-rank approximation technique for matrix C as detailed in Subsection 4.1.4. In this procedure, the covariance matrix need not to be computed fully, providing a matrix $L \in \mathbb{R}^{d_r N_h \times M}$ such that

$$\|C - LL^T\|_{\text{tr}} \leq \epsilon, \quad (4.51)$$

where $\epsilon > 0$ is a given tolerance and $\|\cdot\|_{\text{tr}}$ denotes the trace norm. As suggested by (4.26), solving the M -sized new eigenvalue problem results in getting M eigenfunctions $\tilde{\mathbf{v}}_k \in \mathbb{R}^M$. The nodal value of one particular realization of random fields can finally recovered by means of the following expression:

$$\mathbf{V}(\mathbf{x}_i, \mathbf{y}) = \mathbb{E}[\mathbf{V}](\mathbf{x}_i) + \sum_{k=1}^M [(L\tilde{\mathbf{v}}_k)_i, (L\tilde{\mathbf{v}}_k)_{i+N_h} \cdots, (L\tilde{\mathbf{v}}_k)_{i+(d_r-1)N_h}]^T \mathbf{y}_k, \quad (4.52)$$

where $\mathbf{y} = (y_1, \dots, y_M)$ is an M -tuple of random numbers sampled from an a priori chosen distribution.

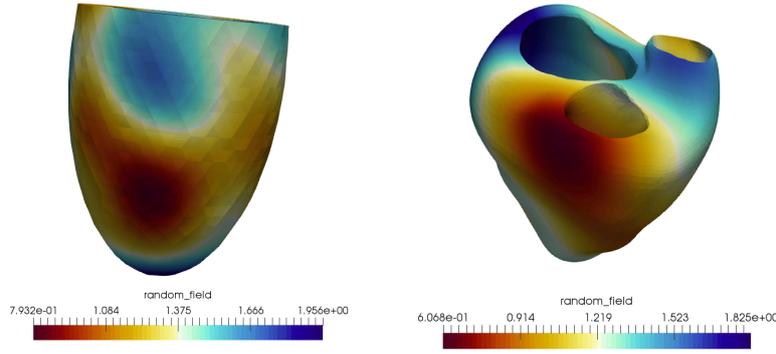


Figure 4.3. Example of isotropic ($d_r = 1$) random fields on more complex geometries.

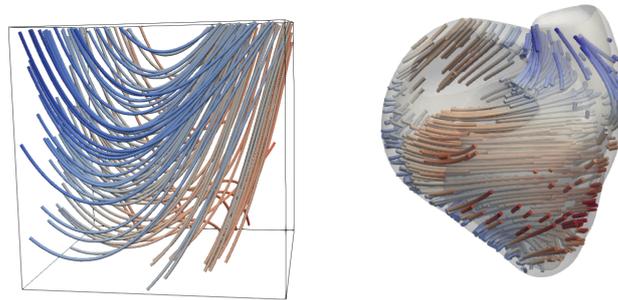


Figure 4.4. Example of anisotropic ($d_r = d$) fiber fields.

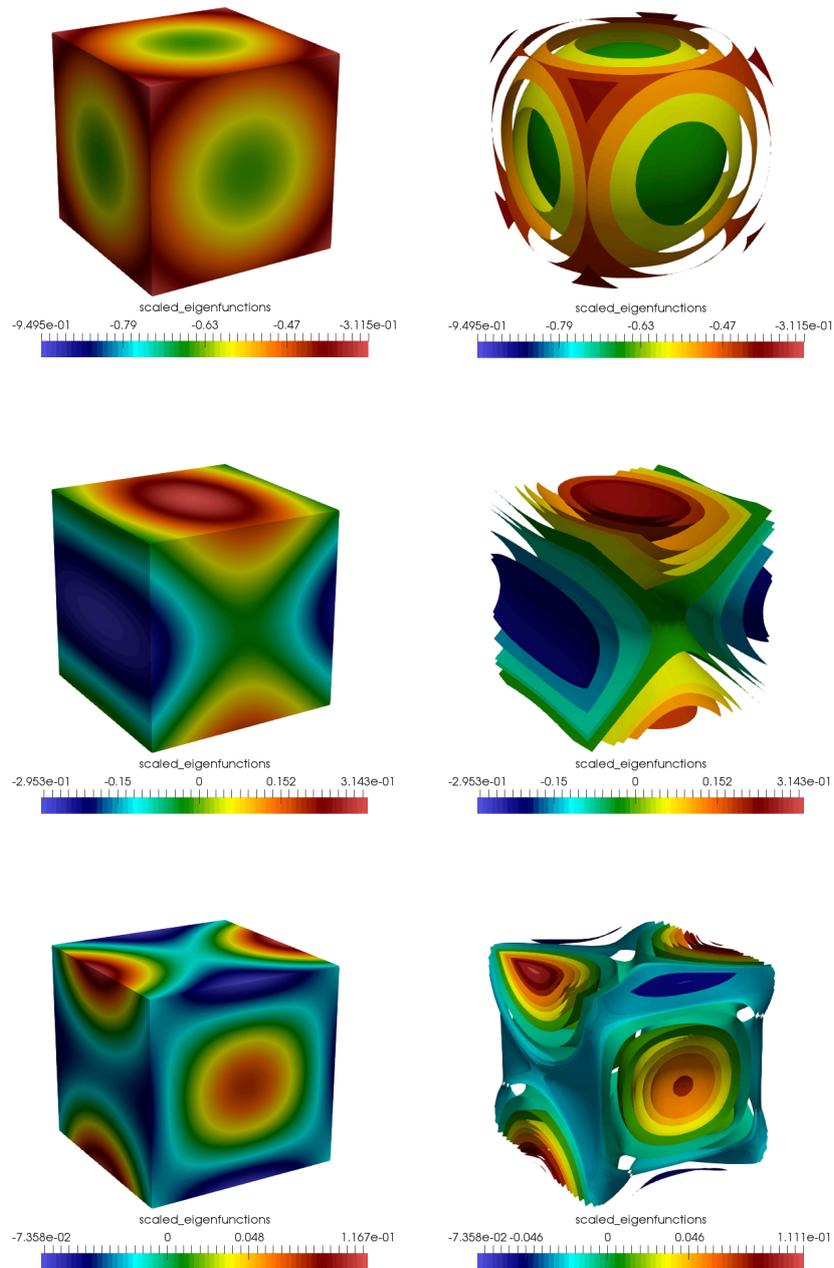


Figure 4.5. Weighted isotropic eigenfunctions (left) and their contours (right) in decreasing order.

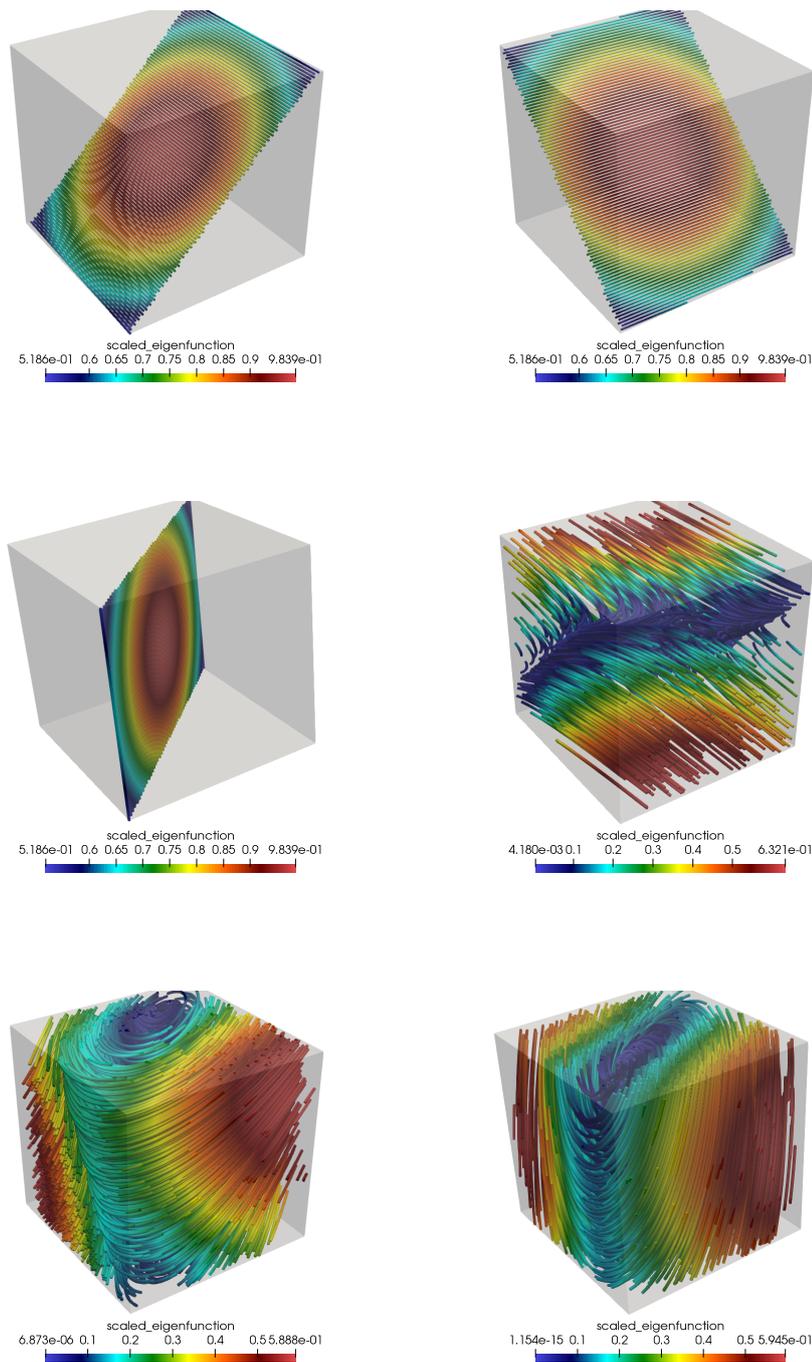


Figure 4.6. Weighted anisotropic eigenfunctions in decreasing order for a diagonal covariance matrix, i.e. $k_{st} = 0$ for $s \neq t$.

Scaling factor for tuning components variability: Let us recall the scaled KL

version given in (4.10), under its truncated form:

$$\mathbf{V}(\mathbf{x}, \omega) = \mathbb{E}[\mathbf{V}](\mathbf{x}) + S \sum_{k=1}^M \sigma_k X_k(\omega) \boldsymbol{\phi}_k(\mathbf{x}), \quad (4.53)$$

where $S \in \mathbb{R}^{d_r \times d_r}$ is the diagonal matrix defined in (4.9). Let $\mathbf{x}_j \in D$ be a given spacial point of the domain D . The i -th component, $1 \leq i \leq d_r$, of the KL expansion in \mathbf{x}_j is given by:

$$[\mathbf{V}(\mathbf{x}_j, \omega)]_i = [\mathbb{E}[\mathbf{V}](\mathbf{x}_j)]_i + s_i \sum_{k=1}^M \sigma_k X_k(\omega) [\boldsymbol{\phi}_k(\mathbf{x}_j)]_i. \quad (4.54)$$

The variance of the KL expansion applied on \mathbf{x}_j on its i -th component is given by:

$$\text{Var}([\mathbf{V}(\mathbf{x}_j, \omega)]_i) = s_i^2 \sum_{k=1}^M \sigma_k^2 \text{Var}(X_k(\omega)) [\boldsymbol{\phi}_k(\mathbf{x}_j)]_i^2, \quad (4.55)$$

since $\text{Var}([\mathbb{E}[\mathbf{V}](\mathbf{x}_j)]_i) = 0$. The variance of the random variable X_k for $k = 1, \dots, M$ shall be known in advance, e.g. $\text{Var}(X_k) = 1/3$ for $X_k \sim \mathcal{U}(-1, 1)$. In which case we would obtain the following:

$$\text{Var}([\mathbf{V}(\mathbf{x}_j, \omega)]_i) = \frac{1}{3} s_i^2 \sum_{k=1}^M \sigma_k^2 [\boldsymbol{\phi}_k(\mathbf{x}_j)]_i^2. \quad (4.56)$$

The standard deviation for the KL expansion applied on \mathbf{x}_j on its i -th component is therefore given by:

$$\text{std}([\mathbf{V}(\mathbf{x}_j, \omega)]_i) = \sqrt{\text{Var}([\mathbf{V}(\mathbf{x}_j, \omega)]_i)} = s_i \sqrt{\frac{1}{3} \sum_{k=1}^M \sigma_k^2 [\boldsymbol{\phi}_k(\mathbf{x}_j)]_i^2}. \quad (4.57)$$

This expression shows that it is possible to tune the coefficient s_i to obtain a desired amplitude for the variability around $[\mathbb{E}[\mathbf{V}](\mathbf{x}_j)]_i$. Obviously, this variability differs from a spatial points to another. It is however possible to choose s_i to guarantee a certain maximal amplitude variability throughout the domain, as reported in Figures 4.7 and 4.8. On a given geometry, Figure 4.7 reports the vector field variability on each coordinate separately, for a given initial state vector field and a corresponding perturbation arising from a non-scaled ($s_i = 1$) KL expansion. Figure 4.7 reports the same perturbation on the same configuration, with a scaling factor $s_i = 2$ on each coordinate.

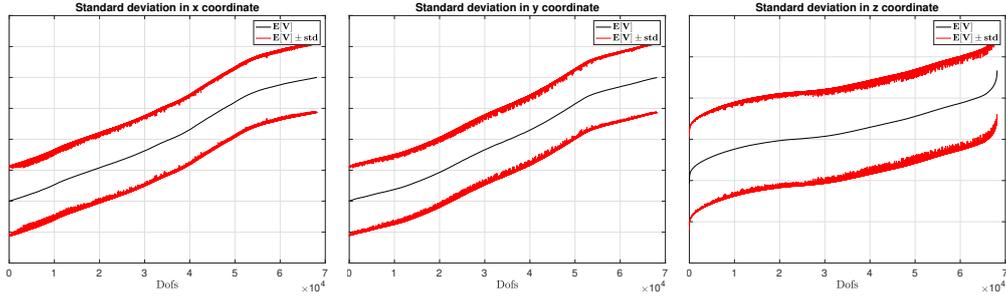


Figure 4.7. Unscaled vector field variability arising from KL expansion.

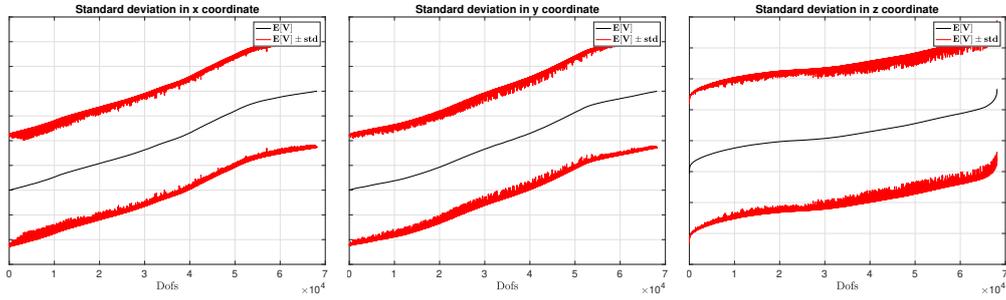


Figure 4.8. Scaled vector field variability arising from KL expansion with a factor of $s_i = 2$ for each coordinate.

Further a posteriori truncation: After a truncated KL expansion is computed by means of Low-rank approximation of the covariance matrix, it is possible to perform an a posteriori truncation. This can be done by selecting a truncation number $M_{\text{trunc}} < M$ such that:

$$M_{\text{trunc}} = \underset{m}{\operatorname{argmin}} \left\{ m : \sum_{k=1}^m \sigma_k / \sum_{k=1}^M \sigma_k > 1 - \epsilon_{\text{trunc}} \right\}, \quad (4.58)$$

typically for $\epsilon_{\text{trunc}} \rightarrow 0$. The chosen tolerance corresponds likewise to the portion of information to carry from the previous KL, in terms of eigenvalues. The a posteriori may allow to reduce significantly the initial M , while carrying most of the information. A qualitative measurement of this statement will be provided in Chapter 7 in a concrete example involving the numerical results for this work.

4.2 Problem setting

4.2.1 Stochastic formulation

Let us recall $(\Omega, \Sigma, \mathbb{P})$ to be a complete and separable probability space and a bounded Lipschitz domain given by $D \subset \mathbb{R}^d$, where $d = 1, 2, 3$. We additionally consider a time domain $[0, T]$ and define $\mathcal{V} = L^2([0, T]; H^1(D))$ to be the Bochner space equipped with the norm in $H^q(D)$ for $q = 0, 1$, given by:

$$\|v\|_{\mathcal{V}} := \left(\int_{[0, T]} \|v(\cdot, t)\|_{H^q(D)}^2 dt \right)^{1/2}. \quad (4.59)$$

We have already argued (c.f. discussion in Subsection 4.1.1) that since $H^1(D)$ is a Hilbert space, $\mathcal{V} = L^2([0, T]; H^1(D))$ is also a Hilbert space. With a similar reasoning, we arrive to the natural conclusion that $L^2(\Omega; \mathcal{V})$ is also characterized as a Hilbert space. We introduce our model problem, that is formulated as a stochastic heat equation:

find $u \in L^2(\Omega; \mathcal{V})$ such that for almost every $\omega \in \Omega$,

$$\begin{aligned} \partial_t u(\mathbf{X}, \omega) - \nabla \cdot (G(\mathbf{x}, \omega) \nabla u(\mathbf{X}, \omega)) &= f(\mathbf{X}), \quad \forall \mathbf{X} = (\mathbf{x}, t) \in D \times (0, T] \\ u(\mathbf{X}) &= 0, \quad \forall \mathbf{X} = (\mathbf{x}, t) \in \partial D \times (0, T] \\ u(\mathbf{x}, 0) &= g(\mathbf{x}), \quad \forall \mathbf{x} \in D, \end{aligned} \quad (4.60)$$

where $G \in L^\infty(\Omega; [L^\infty(D)]^{d_r \times d_r})$ is the random diffusion field for $d_r = 1$ (isotropic case) or $d_r = d$ (anisotropic case), i.e. satisfying the uniform ellipticity condition as a guarantee of the well-posedness of the PDE problem.

Remark. If the forcing function f was depending on $u(\mathbf{X})$ and further written as $f(\mathbf{X}) = I_{\text{app}}(\mathbf{X}) - I_{\text{ion}}(u(\mathbf{X}, \omega))$ (with respect to the notation introduced by (3.52)), we would recover a stochastic formulation of the monodomain equation. The latter undoubtedly represents our main equation of interest with respect to all the numerical experiments. Nonetheless, the stochastic heat model equation will serve as a referential point with respect to the theoretical error analysis.

More specifically, the formulation (4.60) suggests a heat equation subject to the uncertainty of the diffusion field, to be provided by a function of the truncated KL expansion as detailed in Section 4.1, recall

$$\mathbf{V}(\mathbf{x}, \omega) = \mathbb{E}[\mathbf{V}](\mathbf{x}) + \sum_{k=1}^M \sqrt{\lambda_k} X_k(\omega) \boldsymbol{\phi}_k(\mathbf{x}). \quad (4.61)$$

We demonstrate how this is done in a way that ensures the well-posedness of the problem.

4.2.2 Well-posedness

We separate for convenience the two cases: isotropic and anisotropic.

Isotropic: When $d_r = 1$, the uniform ellipticity condition translates to the existence of two constants $0 < G_{\min} < G_{\max} < \infty$ such that:

$$G_{\min} \leq \operatorname{ess\,inf}_{\mathbf{x} \in D} G(\mathbf{x}, \omega) \leq \operatorname{ess\,sup}_{\mathbf{x} \in D} G(\mathbf{x}, \omega) \leq G_{\max} \quad \mathbb{P}\text{-almost surely.} \quad (4.62)$$

In this case, we simply set $G(\mathbf{x}, \omega) = V(\mathbf{x}, \omega)$. Requiring that

$$G_{\min} \leq \operatorname{ess\,inf}_{\mathbf{x} \in D} V(\mathbf{x}, \omega) \leq \operatorname{ess\,sup}_{\mathbf{x} \in D} V(\mathbf{x}, \omega) \leq G_{\max} \quad \mathbb{P}\text{-almost surely} \quad (4.63)$$

is therefore a sufficient condition for the well-posedness of the problem.

Remark. *In accordance to the discussion in Subsection 4.1.2, the condition (4.63) can always be ensured by considering a scaling factor s ahead of the perturbation provided by the KL expansion, as in (4.10). The scaling factor can be selected such that the highest variation is dumped to remain bounded by $\mathbb{E}[V]$, i.e.*

$$s \left| \sum_{i=1}^M \sqrt{\lambda_i} \phi_k(\mathbf{x}) \right| \leq \mathbb{E}[V](\mathbf{x}), \quad \forall \mathbf{x} \in D.$$

Anisotropic: When $d_r = d$, the uniform ellipticity condition transforms to:

$$G_{\min} \leq \operatorname{ess\,inf}_{\mathbf{x} \in D} \lambda_{\min}(G(\mathbf{x}, \omega)) \leq \operatorname{ess\,sup}_{\mathbf{x} \in D} \lambda_{\max}(G(\mathbf{x}, \omega)) \leq G_{\max} \quad \mathbb{P}\text{-almost surely,} \quad (4.64)$$

with $0 < G_{\min} < G_{\max} < \infty$, and where λ_{\min} and λ_{\max} refer respectively to the smallest and largest eigenvalues. In this case, vector-valued random fields $\mathbf{V}(\mathbf{x}, \omega)$ can be used for designing tensor random fields by means of a tensor product construction given by

$$G(\mathbf{x}, \omega) := aI_d + (\|\mathbf{V}(\mathbf{x}, \omega)\|_2 - a) \frac{\mathbf{V}(\mathbf{x}, \omega)\mathbf{V}^T(\mathbf{x}, \omega)}{\mathbf{V}^T(\mathbf{x}, \omega)\mathbf{V}(\mathbf{x}, \omega)}, \quad (4.65)$$

as suggested in [HPS17b], for a given $a > 0$. By doing so, we are modeling a diffusion of strength $\|\mathbf{V}(\mathbf{x}, \omega)\|_2$ in the direction of $\mathbf{V}(\mathbf{x}, \omega)$, and perpendicularly with strength a .

Remark. In the context of electrophysiology, this suggests that $\mathbf{V}(\mathbf{x}, \omega)$ represent the fiber field orientations comprised in the myocardium, diffusing electrical potential with conductivity $\|\mathbf{V}(\mathbf{x}, \omega)\|_2$ in the fibers' directions, and a transversal conductivity of strength a .

The following lemma as demonstrated in [HPS17b], ensures the well-posedness of the tensor model provided by (4.65).

Lemma 4.1 ([HPS17b]). *Given the existence of constants $b_{\min} \leq 1$ and $b_{\max} \geq 1$ such that $b_{\min} \leq a \leq b_{\max}$ and*

$$b_{\min} \leq \operatorname{ess\,inf}_{\mathbf{x} \in D} \|\mathbf{V}(\mathbf{x}, \omega)\|_2 \leq \operatorname{ess\,sup}_{\mathbf{x} \in D} \|\mathbf{V}(\mathbf{x}, \omega)\|_2 \leq b_{\max} \quad \mathbb{P}\text{-almost surely,} \quad (4.66)$$

the conductivity tensor (4.65) is symmetric and satisfies a uniform ellipticity condition (4.64) for $G_{\min} = b_{\min}$ and $G_{\max} = b_{\max}$

Proof. For almost every $\omega \in \Omega$ and almost every $\mathbf{x} \in D$ we have that $G(\mathbf{x}, \omega)$ is well-formed, because of

$$\mathbf{V}^T(\mathbf{x}, \omega)\mathbf{V}(\mathbf{x}, \omega) = \|\mathbf{V}(\mathbf{x}, \omega)\|_2^2 \geq b_{\min}^2 > 0,$$

and clearly symmetric. Furthermore, we can choose $\mathbf{u}_2, \dots, \mathbf{u}_d \in \mathbb{R}^d$ that are perpendicular to $\mathbf{V}(\mathbf{x}, \omega)$ and linearly independent. Thus, it holds, for $i = 2, \dots, d$, that

$$G(\mathbf{x}, \omega)\mathbf{u}_i = a\mathbf{u}_i \quad \text{and} \quad G(\mathbf{x}, \omega)\mathbf{V}(\mathbf{x}, \omega) = \|\mathbf{V}(\mathbf{x}, \omega)\|_2\mathbf{V}(\mathbf{x}, \omega).$$

Consequently, we obtain for almost every $\omega \in \Omega$ and almost every $\mathbf{x} \in D$ that

$$\begin{aligned} \lambda_{\min}(G(\mathbf{x}, \omega)) &= \min\{a, \|\mathbf{V}(\mathbf{x}, \omega)\|_2\} \geq b_{\min}, \\ \lambda_{\max}(G(\mathbf{x}, \omega)) &= \max\{a, \|\mathbf{V}(\mathbf{x}, \omega)\|_2\} \leq b_{\max}. \end{aligned}$$

□

Solution existence and uniqueness: Due to (4.63) (respectively (4.64)), the problem defined in (4.60) is elliptic and admits a unique solution for almost every $\omega \in \Omega$. This claim can be proven by using the Banach-Nečas-Babuška, a generalization of the Babuška-Lax-Milgram (Theorem 3.1), c.f. [CMO18]. The proof can be found e.g. in [LM12], or [Bai69]. Furthermore, the solution satisfies

$$\|u(\cdot, \omega)\|_{\gamma} \leq C(C_a, \alpha, u_0) \|f\|_{V^*}, \quad (4.67)$$

where $C > 0$ is a constant depending on C_a and α respectively the boundedness and coercivity constants, and the initial condition u_0 , c.f. [Bai69, Joh19]

4.2.3 Parametric formulation

Regarding (4.60), we are interested in computing the statistics of the stochastic solution u , with respect to a functional \mathcal{F} , i.e.

$$\mathbb{E}[\mathcal{F}(u)](\mathbf{X}) = \int_{\Omega} \mathcal{F}(u(\mathbf{X}, \omega)) d\mathbb{P}(\omega). \quad (4.68)$$

As described in the previous subsection, the characterization of G is provided by the truncated KL expansion

$$\mathbf{V}(\mathbf{x}; \omega) = \mathbb{E}[\mathbf{V}](\mathbf{x}) + \sum_{k=1}^M \sqrt{\lambda_k} \phi_k(\mathbf{x}) \psi_k(\omega), \quad (4.69)$$

where λ_k and $\phi_k(\mathbf{x})$ are respectively eigenvalues and eigenvectors of the covariance induced by the spatial correlation kernel. The $\{\psi_k(\omega)\}_{k=1}^M$ are stochastically uncorrelated random variables of the joint density function $\rho(\mathbf{y}) = \prod_{i=1}^M \rho_k(y_k)$. We use here a uniform distribution, i.e. $\psi_i \sim \mathcal{U}(-1, 1)$. We are therefore able to identify the stochastic space Ω with its image $[-1, 1]^M$ using the mapping

$$\psi : \Omega \rightarrow [-1, 1]^M, \quad (4.70)$$

$$\omega \mapsto (\psi_1(\omega), \dots, \psi_M(\omega)). \quad (4.71)$$

Likewise by inserting a realization $\mathbf{y} = (y_1, \dots, y_M) = (\psi_1(\omega), \dots, \psi_M(\omega))$, we could reformulate the truncated KL expansion in a deterministic form as

$$\mathbf{V}(\mathbf{x}, \mathbf{y}) = \mathbb{E}[\mathbf{V}](\mathbf{x}) + \sum_{k=1}^M \sqrt{\lambda_k} \phi_k(\mathbf{x}) y_k. \quad (4.72)$$

Consequently, we can also rewrite the stochastic equation (4.60) in a parametric form: find $u \in L^2([-1, 1]^M; L^2([0, T]; H^1(D)))$

$$\partial_t u(\mathbf{X}, \mathbf{y}) - \operatorname{div}(G(\mathbf{X}, \mathbf{y}) \nabla u(\mathbf{X}, \mathbf{y})) = f(\mathbf{X}) \text{ in } D \times [0, T], \quad (4.73)$$

for all $\mathbf{y} \in [-1, 1]^M$. The expectation of the considered functional of the solution for the stochastic heat equation (4.73) can be rewritten as a high dimensional integral

$$\mathbb{E}[\mathcal{F}(u)](\mathbf{X}) = \int_{\Omega} \mathcal{F}(u(\mathbf{X}, \omega)) d\mathbb{P}(\omega) \approx \int_{[-1, 1]^M} \mathcal{F}(u(\mathbf{X}, \mathbf{y})) \rho(\mathbf{y}) d\mathbf{y}. \quad (4.74)$$

For the sake of simplicity, we consider from now on the functional corresponding to the PDE solution, that is given by the identity function $\mathcal{F}(u) = u$.

4.3 Single-level sampling methods

In order to numerically approximate the integral (4.74), one should rely on a quadrature formula of the type

$$\mathcal{Q}u(\cdot) = \sum_{i=1}^N w_i u(\cdot, \xi^{(i)}) \rho(\xi^{(i)}). \quad (4.75)$$

Approximating a stochastic integral likewise can further be interpreted as a sampling method, where the samples represent the quadrature points.

4.3.1 Monte Carlo (MC)

A well known sampling method for this type of problem is provided by the Monte Carlo (MC) estimator. It is characterized by a fairly easy implementation and is straightforward to use. The approximation provided by the MC estimator is furthermore independent of the considered stochastic dimension and the method requires only weak regularity assumptions. Nonetheless, the disadvantage of such a method is characterized by its low convergence rate, requiring therefore a large number of samples, i.e. simulations. In some applications, this might represent a heavy burden, especially when a single simulation is already computationally demanding. Combination of multilevel techniques and deterministic methods, as suggested further in this work, may represent an alternative by drastically lowering the computational cost of the MC method.

The MC estimator's quadrature formula is given by the following expression

$$\mathcal{Q}_{\text{MC}}^N u(\cdot) = N^{-1} \sum_{i=1}^N u^i(\cdot, \xi^{(i)}), \quad (4.76)$$

i.e. the mean over solutions u^i , for $i = 1, \dots, N$, computed for N independent samples $\xi^{(i)} \in [-1, 1]^M$ corresponding to realizations of independent and identically distributed random vectors.

Remark. Realizations of random vectors, with uniform distributions e.g., require the generation of an absolute random number between two provided bounds. The computer, however, is unable to generate such numbers completely randomly and relies on parameter dependent algorithms. We refer to these numbers as pseudo-random numbers.

The definition of the MC estimator (4.76) would be completely auto-sufficient if a solution u^i to the sample $\xi^{(i)}$ could be computed exactly. However in the context the monodomain equation, as in many other applications, an exact and analytical solution can not be provided and we rely on finite elements to numerically approximate it. Let us consider an $l \geq 0$ and a corresponding tetrahedralization \mathcal{T}_l of the spatial domain D of mesh size $h_l \sim 2^{-l}$. We define the continuous, piecewise linear finite element spaces

$$\mathcal{S}_l = \{v_l \in C^0(D) : v_{l|T} \in \mathbb{P}_1(T), \forall T \in \mathcal{T}_l\}. \quad (4.77)$$

We now introduce the notation

$$\mathcal{Q}_{\text{MC}}^{N_l} u_l(\cdot) = N_l^{-1} \sum_{i=1}^{N_l} u_l^i(\cdot, \xi^{(i)}), \quad (4.78)$$

where u_l^i is the finite element approximation of the solution u^i for the sample $\xi^{(i)}$, lying in the space \mathcal{S}_l (N_l replaces N in (4.76) for designating the number of samples for the discretization level l). Therefore, the solution obtained through the MC method is subject not only to the stochastic error depending on the number of samples used as quadrature points, but also on the discretization error inherent to solving the PDE within the context of finite elements. The latter concern is discussed in Section 4.5.

4.3.2 Quasi-Monte Carlo (QMC)

The Quasi-Monte Carlo (QMC) estimator is a deterministic quadrature method that might represent an alternative to MC, under the condition of an extra regularity exhibited by the integrand. The lack of such regularity may induce a higher upper bound in the method's convergence. Furthermore, the method may suffer from the curse of dimensionality depending on the considered deterministic sequencing, causing an exponential growth of the number of required evaluations. The method is however still widely used in many applications, leading in fortunate but not so rare cases to a faster convergence towards the desired quantity.

The inherent idea motivating the QMC method is to rely on a deterministic sequence of quadrature points that exhibits a low discrepancy, defined to be as such if the proportion of points from the considered sequence falling into a containing set is close to its measure. In other terms, the QMC typically relies on a set of points rendering a better coverage of a given (stochastic) space. In order

to fully identify the features of this method, we shall proceed more formally by introducing the following definitions.

Definition 4.2 (Discrepancy). *The discrepancy of a set $S^N = \{\xi^{(1)}, \dots, \xi^{(N)}\} \subset [0, 1]^M$ is defined to be given by $D_N(S^N) = \sup_{\mathcal{B} \in \mathcal{J}} \left| \frac{1}{N} \sum_{i=1}^N \mathbb{1}_{\mathcal{B}}(\xi^{(i)}) - \lambda(\mathcal{B}) \right|$ where $\mathbb{1}_{\mathcal{B}}$ is the indicator function of the set \mathcal{B} , that is*

$$\mathbb{1}_{\mathcal{B}}(\xi) = \begin{cases} 1 & \text{if } \xi \in \mathcal{B}, \\ 0, & \text{otherwise,} \end{cases}$$

$\lambda(\mathcal{B})$ the Lebesgue measure of the set \mathcal{B} , and \mathcal{J} the set of M -dimensional boxes of the form

$$\prod_{i=1}^M [a_i, b_i) = \{\mathbf{x} \in \mathbb{R}^M : a_i \leq x_i \leq b_i\},$$

for $0 \leq a_i < b_i \leq 1$, where $1 \leq i \leq M$.

Remark. *The number $\frac{1}{N} \sum_{i=1}^N \mathbb{1}_{\mathcal{B}}(\xi^{(i)})$ represents the proportion of S^N set points that are contained in \mathcal{B} .*

Definition 4.3 (star-Discrepancy). *The star-Discrepancy is defined similarly for $a_i = 0$ and $b_i < 1$ for $1 \leq i \leq M$, that is*

$$D_N^*(S^N) = \sup_{\mathbf{z} \in [0, 1]^M} \left| \frac{1}{N} \sum_{i=1}^N \mathbb{1}_{[0, \mathbf{z})}(\xi^{(i)}) - \lambda([0, \mathbf{z})) \right|.$$

We below define the van der Corput sequence, an example of equidistributed and dense sequence over the interval $[0, 1]$ that exhibits a low discrepancy.

Definition 4.4 (van der Corput sequence). *Let r be a positive integer representing a base, equivalently called a radix. The representation of a positive integer $n \geq 1$ in the radix of r is given by:*

$$n = \sum_{k=0}^{K-1} d_k(n) r^k,$$

where $0 \leq d_k(n) < r$ is the k -th digit in the representation of n in radix r . The n -th number of the van der Corput sequence with respect to the radix r is obtained by the following transformation:

$$g_r(n) = \sum_{k=0}^{K-1} d_k(n) r^{-k-1}.$$

Defined as such, the van der Corput sequence suggests to reverse the representation of the sequence of natural numbers with respect to a given radix. This means that if a number n can be written as $n = \cdots d_3 d_2 d_1$ in the r -base, its transformation with respect to the van der Corput sequence is obtained by $g_r(n) = 0.d_1 d_2 d_3 \cdots$. We illustrate this with the following example of the first five numbers expressed in the 2-base:

$$\begin{aligned} 1 &= 1 \star 2^0 = 1_2, \\ 2 &= 1 \star 2^1 + 0 \star 2^0 = 10_2, \\ 3 &= 1 \star 2^1 + 1 \star 2^0 = 11_2, \\ 4 &= 1 \star 2^2 + 0 \star 2^1 + 0 \star 2^0 = 100_2, \\ 5 &= 1 \star 2^2 + 0 \star 2^1 + 1 \star 2^0 = 101_2, \end{aligned}$$

for which the five corresponding numbers of the van der Corput sequence are provided by:

$$\begin{aligned} 0.1_2 &= 1 \star 2^{-1} = 1/2, \\ 0.01_2 &= 0 \star 2^{-1} + 1 \star 2^{-2} = 1/4, \\ 0.11_2 &= 1 \star 2^{-1} + 1 \star 2^{-2} = 3/4, \\ 0.001_2 &= 0 \star 2^{-1} + 0 \star 2^{-2} + 1 \star 2^{-3} = 1/8, \\ 0.101_2 &= 1 \star 2^{-1} + 0 \star 2^{-2} + 1 \star 2^{-3} = 5/8. \end{aligned}$$

The van der Corput sequence is a low-discrepancy sequence over the unit interval, that can be generalized to M dimensions, for any M . Provided the first M prime numbers, this gives rise to the Halton sequence defined as follows.

Definition 4.5 (Halton sequence). *The M -dimensional Halton sequence is given by*

$$\xi_{HP}^{(i)} = [g_{p_1}(i), \cdots, g_{p_M}(i)]^T, \text{ for } i \in \mathbb{N},$$

where p_1, \cdots, p_M are M first prime numbers, and $g_{p_k}(i)$ denotes the i -th element of the van der Corput sequence with respect to the base p_k .

In most applications, the integration interval is not necessarily given by $[0, 1]^M$. In the case of an integration interval given $[a, b]^M$, the Halton sequence can easily be mapped by means of an affine transformation of the form

$$\begin{aligned} h : [0, 1] &\rightarrow [a, b] \\ x &\mapsto h(x) := cx + d, \end{aligned}$$

such that $h(0) = a$ and $h(1) = b$, for which the low discrepancy with respect to $[0, 1]^M$ is preserved in $[a, b]^M$. We consider for the current work an integration interval defined by $[-1, 1]^M$, which yields a transformation $h(x) = 2x - 1$.

Remark. For the more complex case where the integration domain is given by \mathbb{R}^M , the mapping can be performed by means of the cumulative normal distribution and its inverse, see [HPS17a].

We therefore have everything at hand for defining formally the QMC estimator in the case of our application,

$$\mathcal{Q}_{\text{QMC}}u(\cdot) = N^{-1} \sum_{i=1}^N u(\cdot, 2\xi_{\text{HP}}^{(i)} - 1). \quad (4.79)$$

We also introduce the notation related to the use of the QMC estimator in the context of approximating the solution u on a discretization level l

$$\mathcal{Q}_{\text{QMC}}^{N_l}u_l(\cdot) = N_l^{-1} \sum_{i=1}^{N_l} u_l(\cdot, 2\xi_{\text{HP}}^{(i)} - 1). \quad (4.80)$$

4.4 Multilevel sampling methods

Multilevel methods were developed to overcome the high computational cost of MC, in view of its slow convergence (to be discussed in Section 4.5). They were introduced as a variance reduction technique for the approximation of integrals in [Hei01], generalized to stochastic ODEs in [Gil08]. Multilevel techniques have become since then widely adopted and many works brought different contributions, see e.g. [BSZ11, Gil13a, HPS12b, MS12]. The main idea is to rely on a multilevel splitting that creates low-resolution models. Provided a balancing of samples priorily set on the basis of the method's convergence properties, the major portion of samples can be offset to the low-resolution models, while the overall accuracy still preserved. The sampling strategy, which we discuss in Section 4.6, is therefore of paramount importance in the multilevel setup.

4.4.1 Multilevel Monte Carlo (MLMC)

The Multilevel Monte Carlo (MLMC) estimator relies on a hierarchy of meshes priorily defined for a given spatial domain. Let us therefore start by considering a nested sequence of shape regular tetrahedralizations $\{\mathcal{T}_l\}_{l \geq 0}$ of the spatial

domain D , where each \mathcal{T}_l is of mesh size $h_l \sim 2^{-l}$. For all levels $l \geq 0$, we reconsider the continuous, piecewise linear finite element spaces \mathcal{S}_l , defined in (4.77). To each of the finite element spaces \mathcal{S}_l , we furthermore associate the subscript corresponding time-discretization step from $\{\Delta t_l\}_{l \geq 0}$ where $\Delta t_l \sim h_l$. The MLMC estimator can be defined as follows:

$$\mathcal{Q}_{\text{MLMC}} u_L(\cdot) = \sum_{l=0}^L \mathcal{Q}_{\text{MC}}^{N_l}(\Delta u_l(\cdot)) = \sum_{l=0}^L \mathcal{Q}_{\text{MC}}^{N_l}(u_l(\cdot) - u_{l-1}(\cdot)), \quad (4.81)$$

where $\{\mathcal{Q}_{\text{MC}}^{N_l}\}_{l=0}^L$, involving N_l samples, is a sequence of MC quadratures of precision ϵ_l with respect to its stochastic error, therefore defined by the sampling strategy. These quadratures are expressed as priorly defined in the case of MC, applied to the difference of FE spaces. We can write them in the following way:

$$\mathcal{Q}_{\text{MC}}^{N_l}(u_l(\cdot) - u_{l-1}(\cdot)) = N_l^{-1} \sum_{i=1}^{N_l} (u_l(\cdot, \xi^{(i)}) - u_{l-1}(\cdot, \xi^{(i)})). \quad (4.82)$$

where $u_l(\cdot, \xi^{(i)})$ is the solution obtained at level l for the sample $\xi^{(i)}$, with the corresponding mesh and time discretization steps. We further define $u_{-1} \equiv 0$.

Remark. Defined as in (4.81), the MLMC estimator is given by a telescopic sum that algebraically is equivalent to a MC estimator applied for the fine level, assuming that identical samples are solved at every level.

Remark. By looking closer at (4.82), one can deduce that an essential component of the MLMC resides in the intergrid transfer of the data. The samples represent random fields to be interpreted at different mesh levels. The quantities-of-interest, here the samples' solutions for simplicity, are to be transferred on the fine level. This is discussed in Section 4.8.

4.4.2 Multilevel inverse construction

The MLMC as introduced by the expression (4.81) can be interpreted in an abstract way as the sum of quadrature formulas applied to the difference of FE functions. In this context, the FE spaces nestedness requirement is the consequence of the necessity of having the difference of FE functions well-defined. In [HPS12b], based on the work developed in [GH13b], it is shown that the multilevel estimator resembles a sparse grid tensor representation of product spaces, in which the complementary spaces are provided by the differences of FE spaces. This observation has motivated the so-called inverse construction of the MLMC

estimator, by inverting the roles of the quadrature and the FE spaces. The MLMC estimator can therefore also be written in the form of differences in quadrature, applied to FE functions:

$$\mathcal{Q}_{\text{MLMC}} u_L(\cdot) = \sum_{l=0}^L \Delta \mathcal{Q}_{\text{MC}}^{N_l}(u_l(\cdot)) = \sum_{l=0}^L (\mathcal{Q}_{\text{MC}}^{N_l} - \mathcal{Q}_{\text{MC}}^{N_{l+1}})(u_l(\cdot)), \quad (4.83)$$

where we set $\mathcal{Q}_{\text{MC}}^{N_{l+1}} \equiv 0$. The benefit from this representation is that it allows to give up the FE spaces nestedness requirement. This represents a non negligible advantage, as in the case of complex 3D geometries, it might be impossible to construct a nested sequence of meshes. However, such a representation does not allow for recovering the full solution as such a process introduces an additional error. The expression in (4.83) is abusive in that sense and should rather be defined for a pointwise functional

$$\mathcal{Q}_{\text{MLMC}} u_L(\cdot) = \sum_{l=0}^L \Delta \mathcal{Q}_{\text{MC}}^{N_l} \mathcal{F}(u_l(\cdot)) = \sum_{l=0}^L (\mathcal{Q}_{\text{MC}}^{N_l} - \mathcal{Q}_{\text{MC}}^{N_{l+1}}) \mathcal{F}(u_l(\cdot)), \quad (4.84)$$

where \mathcal{F} is the functional.

The difference in quadrature is well-defined beyond any nestedness condition. It turns out however to be an advantage to consider nested quadrature as this reduces the overall work of the MLMC estimator, see [GHM15]. Provided that the MLMC involves solving more samples on coarse levels (implying therefore $N_l > N_{l+1}$, $\forall l = 0, \dots, L-1$), the use of nested quadrature results in solving N_l samples on every level l , summed up with different weights:

$$(\mathcal{Q}_{\text{MC}}^{N_l} - \mathcal{Q}_{\text{MC}}^{N_{l+1}})(u_l(\cdot)) = \frac{N_{l+1} - N_l}{N_l N_{l+1}} \sum_{i=0}^{N_{l+1}} u_l(\cdot, \xi^{(i)}) + \frac{1}{N_0} \sum_{i=N_{l+1}+1}^{N_l} u_l(\cdot, \xi^{(i)}). \quad (4.85)$$

In this process, furthermore, all samples $\xi^{(i)}$ are drawn only once at all the involved MC quadratures.

4.4.3 Multilevel Quasi-Monte Carlo (MLQMC)

The Multilevel Quasi-Monte Carlo (MLQMC) estimator is evidently the combination of the MLMC estimator and QMC, i.e. samples selected with a deterministic approach as in the case of the Halton sequence for example. It can be formally represented with the following formula

$$\mathcal{Q}_{\text{MLQMC}} u_L(\cdot) = \sum_{l=0}^L \mathcal{Q}_{\text{QMC}}^{N_l}(u_l(\cdot) - u_{l-1}(\cdot)) = \sum_{l=0}^L (\mathcal{Q}_{\text{QMC}}^{N_l} - \mathcal{Q}_{\text{QMC}}^{N_{l+1}})(u_l(\cdot)), \quad (4.86)$$

where the second equality is provided by the inverse multilevel construction. Here, $\{\mathcal{Q}_{\text{QMC}}^{N_l}\}_{l=0}^L$ is a sequence of QMC quadratures with precision ϵ_l . This precision is induced by the number of samples used on every level, as in the case of the MLMC.

4.5 Convergence properties

In the following, we discuss in detail the convergence behaviour for each of the introduced sampling methods. As argued previously, the use of stochastic quadrature methods in the context of FE approximation entails stochastic and discretization contributions to the general (or total) error.

The study of the convergence behaviour is crucial for the understanding of these methods, providing moreover a clear idea on the sampling strategy to be adopted (discussed in Section 4.6). The latter in turn prepares the ground for the comparison of the performance of these methods. This will be addressed based on the work evaluation (c.f. Section 4.7).

4.5.1 Discretization error

Let us first recall the approximation errors deriving from the FE discretization process, already introduced in Section 3.3, combined to the solution boundedness provided by (4.67):

$$\|u - u_l\|_{L^2([0,T];L^2(D))} \leq C(h_l^2 + \Delta t_l^2) \|f\|_{L^2([0,T];L^2(D))}, \quad (4.87)$$

and

$$\|u - u_l\|_{L^2([0,T];H^1(D))} \leq C(h_l + \Delta t_l^2) \|f\|_{L^2([0,T];H^1(D))}, \quad (4.88)$$

for the case of linear FE as relied upon all along this work. In order to make further error analysis not burdensome, we suggest an expression of these error estimates only in terms of the space discretization step h_l , by considering $h_l \sim \Delta t_l$. This does not affect the error order as we perform equal refinement

(or coarsening) in space and time between different levels. Hence, we recover simplified error estimates given by

$$\|u - u_l\|_{L^2([0,T];L^2(D))} \leq Ch_l^2 \|f\|_{V^*}, \quad (4.89)$$

and

$$\|u - u_l\|_{L^2([0,T];H^1(D))} \leq Ch_l \|f\|_{V^*}. \quad (4.90)$$

We furthermore write the aforementioned error estimates in a single compact form, by setting

$$\|u - u_l\|_{\mathcal{V}_q} \leq Ch_l^{2-q} \|f\|_{V^*}, \quad (4.91)$$

for $q = 0, 1$ and $\mathcal{V}_q = L^2([0, T]; H^q(D))$.

4.5.2 Stochastic error

The stochastic error emanating from the MC process of estimating the expectation of a random variable u converges as suggested by the following proposition.

Proposition 4.1 ([BSZ11]). *Recall the MC method for approximating the expectation $\mathbb{E}[u] \in \mathcal{V}_q$ ($u \in L^2([-1, 1]^M; \mathcal{V}_q)$) given by*

$$\mathcal{Q}_{MC}^N u(\cdot) = N^{-1} \sum_{i=1}^N u^i(\cdot). \quad (4.92)$$

For any $N \in \mathbb{N}$, it holds that

$$\left\| \mathbb{E}[u] - \mathcal{Q}_{MC}^N u \right\|_{L^2([-1,1]^M; \mathcal{V}_q)} \leq N^{-1/2} \|u\|_{L^2([-1,1]^M; \mathcal{V}_q)}, \quad (4.93)$$

Proof. For a proof of this statement, see [BSZ11]. □

We have presented the MC quadrature method as a relatively easy method to employ, with nonetheless a low convergence rate. The bound given in Proposition 4.1 gives a quantitative measure of this statement, since it suggests that the error converges as $\mathcal{O}(N^{-1/2})$.

The QMC method instead, is designed to improve the convergence rate of the MC method. Through a deterministic choice of quadrature points presenting a lower discrepancy, a faster convergence of the stochastic error can be recovered. The success of such a procedure is not automatic, but depends instead on some

intrinsic features of the problem, mainly in terms of the smoothness of the considered integrand. The study of the convergence for the QMC requires therefore more analytical effort than the MC. For this purpose, we define two important notions in the treatment of the QMC convergence error: the variation in the sense of Vitali, and the variation in the sense of Hardy and Krause.

Definition 4.6 (Variation in the sense of Vitali). *Let $f : \prod_{k=1}^M [a_k, b_k] \rightarrow \mathbb{R}$. We define:*

$$\Delta_{h_k}(f, \mathbf{x}) = f(x_1, \dots, x_k + h_k, \dots, x_M) - f(x_1, \dots, x_k, \dots, x_M),$$

for $h_k > 0$. We define recursively:

$$\Delta_{h_1, \dots, h_k}(f, \mathbf{x}) = \Delta_{h_k}(\Delta_{h_1, \dots, h_{k-1}}, \mathbf{x}).$$

We also create in every direction the partition π_k , defined by N_{k+1} points t_k^i such that $a_k = t_k^1 < \dots < t_k^{N_{k+1}} = b_k$, to which we associate the spacings $h_k^i = t_k^{i+1} - t_k^i$. The variation in the sense of Vitali of f is given by:

$$V^M(f) = \sup_{(\pi_1, \dots, \pi_M)} \sum_{k=1}^M \sum_{i_k=1}^{N_k} \left| \Delta_{h_1^{i_1}, \dots, h_k^{i_k}}(f, (x_1^{i_1}, \dots, x_k^{i_k})) \right|.$$

Definition 4.7 (Variation in the sense of Hardy and Krause). *The variation in the sense of Hardy and Krause is given by:*

$$V_{HK}(f) = \sum V^M(f),$$

where the sum is provided on every sub-interval of the M -dimensional box provided by $\prod_{k=1}^M [a_k, b_k] \subset \mathbb{R}^M$, of dimension inferior or equal to M . Additionally, a function is said of bounded variation in the sense of Hardy and Krause if $V_{HK}(f) < \infty$.

For functions $f : [0, 1]^M \mapsto \mathbb{R}$ that are bounded in the sense of Hardy and Krause, the approximation error for the integral of f is provided by the Koksma-Hlawka inequality,

$$|\mathcal{Q}_{\text{QMC}}^N f - I(f)| \leq D_N^*(S^N) V_{HK}(f), \quad (4.94)$$

where $I(f) = \int_{[0,1]^M} f(\mathbf{x}) d\mathbf{x}$ and S^N the samples set. This inequality was established in [Hla61], generalized to any dimension in [Zar68].

In our case however, we are interested in the integration of a function u belonging to the Bochner space $L^2([-1, 1]^M; \mathcal{Y}_q)$. Let us assume an extra smoothness of the function u such as $u \in W_{\text{mix}}^{1,1}([-1, 1]^M; \mathcal{Y}_q)$, for which the norm defined as

$$\|u\|_{W_{\text{mix}}^{1,1}([-1, 1]^M; \mathcal{Y}_q)} := \sum_{\|\alpha\|_{\infty} \leq 1} \int_{[0, 1]^{|\alpha|}} \left\| \partial_{\mathbf{y}}^{\alpha} u_{|\alpha|}(\cdot) \right\|_{\mathcal{Y}_q} \quad (4.95)$$

is bounded, where $\alpha = (\alpha_1, \dots, \alpha_M)$ is the usual multi-index used in the partial derivatives notations with $\|\alpha\|_{\infty} = \max_{k=1, \dots, M} |\alpha_k|$, and $u_{|\alpha|}$ the projection of u on the $|\alpha|$ -dimensional face defined by the set $\{\mathbf{y} \in \mathbb{R}^M \text{ such that } y_j = 1 \text{ for } \alpha_j = 0\}$. We may obtain an analogous bound to that given in (4.94), see [HPS17a],

$$\left\| \mathcal{Q}_{\text{QMC}}^N u - \mathbb{E}[u] \right\|_{\mathcal{Y}_q} \leq \|u\|_{W_{\text{mix}}^{1,1}([-1, 1]^M; \mathcal{Y}_q)} D_N^*(S^N). \quad (4.96)$$

The upper bound provided by (4.96) is very important for understanding the convergence behaviour of the QMC estimator. It is provided as a product of two separate terms.

The first is independent of the considered set of integration points and is given in terms of the integrand's mixed first order derivatives. It is therefore crucial to guarantee that the latter exist and are bounded, in what is referred to as the mixed regularity. This is not an obvious task and is a matter of the solution specificity and the distribution of the random variables. It has been proven e.g. to hold in the case of the Poisson equation, for log-normally distributed random diffusion coefficients, see [HPS17a]. It also holds for the same problem when uniformly elliptic (an)isotropic random diffusion coefficients are used, see [HPS17b]. This represents the closest configuration to our problem of interest.

Remark. *A similar result has not been proven yet for the monodomain equation, and is more a matter of work in progress. In this work, we assume that such as a mixed regularity holds and suggest to show by experiments that the theoretical convergence rates can be attained.*

The second bounding term in (4.96) is given by the discrepancy of the considered sampling points, as a measure of its quality. It is therefore a term depending exclusively on the considered stochastic domain (in terms of all the intervening components of the inducing KL expansion) and the dispersion of sample points therein. In the case of the Halton sequence, it has been shown that it is estimated to be of order $\mathcal{O}(N^{-1}(\log N)^M)$, see e.g. [Nie92]. This bound is very

interesting to multiple extents, as it suggest an almost linear convergence for a small stochastic dimension M , but also an eventual exponential growth in this dimension. This is referred to as the curse of dimensionality. It can however be demonstrated that provided a certain decay of the eigenvalues, the convergence is almost linear whatever the considered stochastic dimension.

We define $\gamma_k := \|\sigma_k \phi_k\|_{(L^\infty(D))^{d_r}}$, with respect to the KL notation used all along this work. The proposition 4.2 states the conditions under which this is achieved.

Proposition 4.2. *Assuming that the mixed regularity of the integrand holds, and that there exists $\epsilon > 0$ such that $\gamma_k \leq ck^{-3-\epsilon}$ for some positive constant $c > 0$, the QMC error converges as*

$$\left\| \mathbb{E}[u] - \mathcal{Q}_{\text{QMC}}^N u \right\|_{L^2([-1,1]^M; \gamma_q)} \leq C_\delta N^{-1+\delta} \|u\|_{W_{\text{mix}}^{1,1}([-1,1]^M; \gamma_q)}, \quad (4.97)$$

for any $\delta > 0$, where C_δ is a constant depending on it.

Proof. See [HPS16], as a consequence of the result proven in [Wan02]. \square

Notice that given the deterministic selection of quadrature points, different realizations in the $[-1, 1]^M$ space provide exactly the same set of samples. As a consequence, we have that $\left\| \mathbb{E}[u] - \mathcal{Q}_{\text{QMC}}^N u \right\|_{L^2([-1,1]^M; \gamma)} = \left\| \mathbb{E}[u] - \mathcal{Q}_{\text{QMC}}^N u \right\|_{\mathcal{V}}$ and we could also restrict ourselves to the norm in \mathcal{V} . For a matter of homogeneity between the error estimates given for all the quadrature methods, we will however keep the use of the plain norm.

4.5.3 General error

The convergence bounds established by Propositions 4.1 and 4.2 combined to the discretization error estimate (4.91) can be used in order to evaluate the general (or total) error behaviour for all the quadrature methods of interest.

Theorem 4.1 (MC approximation). *Given a discretization level l and $q = 0, 1$, the following error bound holds*

$$\left\| \mathbb{E}[u] - \mathcal{Q}_{\text{MC}, l}^{N_l} u \right\|_{L^2([-1,1]^M; \gamma_q)} \leq C(N_l^{-1/2} + h_l^{2-q}) \|f\|_{V^*}, \quad (4.98)$$

for approximating the expectation of the random variable $u \in L^2([-1, 1]^M; \gamma_q)$ with the MC method.

Proof. For a matter of spacing, we denote $[-1, 1]^M$ with \square . Then,

$$\begin{aligned} \left\| \mathbb{E}[u] - \mathcal{Q}_{MC}^{N_l} u_l \right\|_{L^2(\square; \gamma_q)} &= \left\| \mathbb{E}[u] - \mathbb{E}[u_l] + \mathbb{E}[u_l] - \mathcal{Q}_{MC}^{N_l} u_l \right\|_{L^2(\square; \gamma_q)} \\ &\leq \left\| \mathbb{E}[u] - \mathbb{E}[u_l] \right\|_{L^2(\square; \gamma_q)} + \left\| \mathbb{E}[u_l] - \mathcal{Q}_{MC}^{N_l} u_l \right\|_{L^2(\square; \gamma_q)} \\ &= \left\| \mathbb{E}[u - u_l] \right\|_{\gamma_q} + \left\| \mathbb{E}[u_l] - \mathcal{Q}_{MC}^{N_l} u_l \right\|_{L^2(\square; \gamma_q)}. \end{aligned}$$

The first inequality is induced by the norm triangular inequality, whereas in the second equality we used the linearity of the expectation operator and the fact that $\mathbb{E}[u]$ and $\mathbb{E}[u_l]$ do not depend on the random variable.

The first term of the inequality can be bounded by using the FE discretization error estimate given in (4.91). Proposition 4.1 can be used in turn to bound the second term of the inequality, fulfilling the bound proposed by the theorem. \square

Theorem 4.2 (QMC approximation). *Given a discretization level l and $q = 0, 1$, the following error bound holds*

$$\left\| \mathbb{E}[u] - \mathcal{Q}_{QMC, l}^{N_l} u \right\|_{L^2([-1, 1]^M; \gamma_q)} \leq C(N_l^{-1+\delta} + h_l^{2-q}) \|f\|_{V^*}, \quad (4.99)$$

for approximating the expectation of the random variable $u \in W_{mix}^{1,1}([-1, 1]^M; \gamma_q)$ with the QMC method.

Proof. The sketch of the proof is identical to that of Theorem 4.1, where Proposition 4.2 with the bound provided by (4.97) is used instead. \square

Theorems 4.1 and 4.2 indicate that the general error for the single level methods is additive with stochastic and discretization error components. The situation for the multilevel methods is slightly different. As the inner structure of these alternative estimators is given by the stochastic quadrature applied to the difference of discretized solutions, each level contribution to the error becomes stochastic and discretization multiplicative. Theorems 4.3 and 4.4 illustrate this fact.

Theorem 4.3 (MLMC approximation). *Recall the MLMC estimator given by*

$$\mathcal{Q}_{MLMC} u_L(\cdot) = \sum_{l=0}^L \mathcal{Q}_{MC}^{N_l} (u_l(\cdot) - u_{l-1}(\cdot)), \quad (4.100)$$

where $\{\mathcal{Q}_{MC}^{N_l}\}_{l=0}^L$ is a sequence of MC quadratures with accuracy $\epsilon_l = N_l^{-1/2}$ for every level $l = 0, \dots, L$. The MLMC approximation for the expectation of the random variable $u \in L^2([-1, 1]^M; \gamma_q)$ yields the following error bound

$$\left\| \mathbb{E}[u] - \mathcal{Q}_{MLMC,L} u \right\|_{L^2([-1,1]^M; \gamma_q)} \leq D(h_L^{2-q} + \sum_{l=0}^L N_l^{-1/2} h_l^{2-q}) \|f\|_{V^*} \quad (4.101)$$

for $q = 0, 1$ and a constant $D > 0$.

Proof. In the following demonstration, we again use $\square = [-1, 1]^M$. We find

$$\begin{aligned} \left\| \mathbb{E}[u] - \mathcal{Q}_{MLMC,L} u \right\|_{L^2(\square; \gamma_q)} &\leq \left\| \mathbb{E}[u] - \mathbb{E}[u_L] + \mathbb{E}[u_L] - \mathcal{Q}_{MLMC,L} u \right\|_{L^2(\square; \gamma_q)} \\ &= \left\| \mathbb{E}[u] - \mathbb{E}[u_L] + \mathbb{E}[u_L] - \sum_{l=0}^L \mathcal{Q}_{MC}^{N_l}(u_l - u_{l-1}) \right\|_{L^2(\square; \gamma_q)} \\ &\leq \left\| \mathbb{E}[u] - \mathbb{E}[u_L] \right\| + \left\| \mathbb{E}[u_L] - \sum_{l=0}^L \mathcal{Q}_{MC}^{N_l}(u_l - u_{l-1}) \right\|_{L^2(\square; \gamma_q)} \end{aligned}$$

The first part of the inequality can be easily bounded by making use of the FE approximation estimate (4.91),

$$\left\| \mathbb{E}[u] - \mathbb{E}[u_L] \right\|_{L^2(\square; \gamma_q)} = \left\| \mathbb{E}[u] - \mathbb{E}[u_L] \right\|_{\gamma_q} \leq Ch_L^{2-q} \|f\|_{V^*}. \quad (4.102)$$

By observing that the expectation of the discretized solution for level L can be rewritten in the form of the following telescopic sum,

$$\mathbb{E}[u_L] = \sum_{l=0}^L \mathbb{E}[u_l] - \mathbb{E}[u_{l-1}] = \sum_{l=0}^L \mathbb{E}[u_l - u_{l-1}], \quad (4.103)$$

the second term of the inequality can be bounded as follows:

$$\begin{aligned} \left\| \sum_{l=0}^L \mathbb{E}[u_l - u_{l-1}] - \sum_{l=0}^L \mathcal{Q}_{MC}^{N_l}(u_l - u_{l-1}) \right\| &\leq \sum_{l=0}^L \left\| (\mathbb{E} - \mathcal{Q}_{MC}^{N_l})(u_l - u_{l-1}) \right\|_{L^2(\square; \gamma_q)} \\ &\leq \sum_{l=0}^L N_l^{-1/2} \|u_l - u_{l-1}\|_{L^2(\square; \gamma_q)}. \quad (4.104) \end{aligned}$$

On the other hand, solution differences are bounded as follows:

$$\begin{aligned}
\|u_l - u_{l-1}\|_{L^2(\square; \mathcal{V}_q)} &\leq \|u_l - u\|_{L^2(\square; \mathcal{V}_q)} + \|u - u_{l-1}\|_{L^2(\square; \mathcal{V}_q)} \\
&\leq C(h_l^{2-q} + h_{l-1}^{2-q}) \|f\|_{V^*} \\
&= C(h_l^{2-q} + (2h_l)^{2-q}) \|f\|_{V^*} \\
&= (1 + 2^{2-q}) C h_l^{2-q} \|f\|_{V^*}.
\end{aligned}$$

Therefore the inequality (4.104) becomes:

$$\left\| \sum_{l=0}^L \mathbb{E}[u_l - u_{l-1}] - \sum_{l=0}^L \mathcal{Q}_{MC}^{N_l}(u_l - u_{l-1}) \right\| \leq (1 + 2^{2-q}) C \left(\sum_{l=0}^L N_l^{-1/2} h_l^{2-q} \right) \|f\|_{V^*} \quad (4.105)$$

Combining (4.105) with (4.102) yields the bound (4.101) for $D = (1 + 2^{2-q})C$. \square

Theorem 4.4 (MLQMC approximation). *Recall the MLQMC estimator given by*

$$\mathcal{Q}_{MLQMC} u_L(\cdot) = \sum_{l=0}^L \mathcal{Q}_{QMC}^{N_l}(u_l(\cdot) - u_{l-1}(\cdot)), \quad (4.106)$$

where $\{\mathcal{Q}_{QMC}^{N_l}\}_{l=0}^L$ is a sequence of MC quadratures with accuracy $\epsilon_l = N_l^{-1+\delta}$ for every level $l = 0, \dots, L$. The MLQMC approximation of the expectation of the random variable $u \in W_{mix}^{1,1}([-1, 1]^M; \mathcal{V}_q)$ yields the following error bound

$$\left\| \mathbb{E}[u] - \mathcal{Q}_{MLQMC, L} u \right\|_{L^2([-1, 1]^M; \mathcal{V}_q)} \leq D(h_L^{2-q} + \sum_{l=0}^L N_l^{-1+\delta} h_l^{2-q}) \|f\|_{V^*}. \quad (4.107)$$

for $q = 0, 1$ and a constant $D > 0$.

Proof. The sketch of the proof is identical to that of Theorem 4.1, where Proposition 4.2 with the bound provided by (4.97) is used instead. \square

4.6 Sampling strategy

The sampling strategy is concerned with the amount of computed samples for a given experiment. As the number of samples to be considered influences the overall computational cost, it seems reasonable not to overcompute unnecessary additional samples which do not increase the precision of the global solution.

It is important to notice that the overall accuracy is always at best bounded by the discretization error deriving from the fine considered level. The number of computed samples in turn, can only impact on the stochastic error. The sampling strategy is therefore based on the convergence rates priorly established for all the considered quadrature methods, with the aim of balancing the stochastic error to be of the same order as the discretization error.

Recall the error bound given in Theorem 4.1 for the MC quadrature method

$$\left\| \mathbb{E}[u] - \mathcal{Q}_{MC,l}^{N_l} u \right\|_{L^2([-1,1]^M; \gamma_q)} \leq C(N_l^{-1/2} + h_l^{2-q}) \|f\|_{V^*}. \quad (4.108)$$

Assuming that $h_l \sim 2^{-l}$, equal stochastic and discretization contributions to the error can be imposed by setting

$$N_l^{-1/2} \sim \mathcal{O}(h_l^{2-q}) \implies N_l \sim \mathcal{O}(h_l^{2q-4}) = \mathcal{O}(2^{(4-2q)l}), \quad (4.109)$$

for $q = 0, 1$. Regarding the QMC method, the error bound provided by Theorem 4.2 is given by

$$\left\| \mathbb{E}[u] - \mathcal{Q}_{QMC,l}^{N_l} u \right\|_{L^2([-1,1]^M; \gamma_q)} \leq C(N_l^{-1+\delta} + h_l^{2-q}) \|f\|_{V^*}. \quad (4.110)$$

Neglecting the $\delta > 0$ (arbitrarily small), we can deduce from the above the ideal number of samples to be computed for the QMC method

$$N_l^{-1} \sim \mathcal{O}(h_l^{2-q}) \implies N_l \sim \mathcal{O}(h_l^{q-2}) = \mathcal{O}(2^{(2-q)l}), \quad (4.111)$$

for $q = 0, 1$. Expressions (4.109) and (4.111) motivate Propositions 4.3 and 4.4.

Proposition 4.3 (MC convergence rate). *The H^q -error ($q = 0, 1$) produced by the MC estimator converges asymptotically with a rate of $2^{(q-2)l}$ provided a number of samples $N_{MC,l} = \mathcal{O}(2^{(4-2q)l})$.*

Proposition 4.4 (QMC convergence rate). *The H^q -error ($q = 0, 1$) produced by the QMC estimator converges asymptotically with a rate of $2^{(q-2)l}$ provided a number of samples $N_{QMC,l} = \mathcal{O}(2^{(2-q)l})$.*

The main observation about Propositions 4.3 and 4.4 is that the two methods converge asymptotically with the same convergence rate, that of the discretization error. MC requires however solving $N_{MC,l}/N_{QMC,l} = 2^2$ factor more samples to produce the same convergence rate.

Regarding the multilevel strategies, the situation is slightly more complex as it requires an optimal sample distribution on each level for a given experiment. Recall the error bound provided by Theorem 4.3 for the MLMC quadrature method

$$\left\| \mathbb{E}[u] - \mathcal{Q}_{\text{MLMC},L} u \right\|_{L^2([-1,1]^M; \mathcal{Y}_q)} \leq D(h_L^{2-q} + \sum_{l=0}^L N_l^{-1/2} h_l^{2-q}) \|f\|_{V^*}. \quad (4.112)$$

When looking closely at the sum, one can observe that the error contribution from each level is stochastic–discretization multiplicative, given by $N_l^{-1/2} h_l^{2-q}$, $\forall l = 0, \dots, L$. A reasonable thought in this case is to impose an equal error contribution from each level by imposing $h_l^{2-q} N_l^{-1/2} = h_{l'}^{2-q} N_{l'}^{-1/2}$, for distinct levels l and l' . Applying this equality for two consecutive levels and taking advantage of the uniform coarsening ($h_l = 2h_{l+1}$), we recover the following:

$$h_l^{2-q} N_l^{-1/2} = h_{l+1}^{2-q} N_{l+1}^{-1/2} \implies h_l^{2-q} N_{l+1}^{1/2} = h_{l+1}^{2-q} N_l^{1/2} \implies 2^{4-2q} N_{l+1} = N_l \quad (4.113)$$

The sampling ratio between consecutive levels is therefore given by

$$N_l / N_{l+1} = 2^{4-2q}, \quad \forall l = 0, \dots, L-1. \quad (4.114)$$

Furthermore, from (4.112), we also notice that the error is at best bounded by h_L^{2-q} . The error contribution from the finest level L is given by $N_L^{-1/2} h_L^{2-q}$, from which we deduce that $N_L \sim \mathcal{O}(1)$. Based on this and (4.114), we deduce that the sampling strategy should be such as

$$N_l \sim \mathcal{O}(2^{(4-2q)(L-l)}), \quad \forall l = 0, \dots, L. \quad (4.115)$$

With a similar reasoning regarding the MLQMC, where levels error contributions become $N_l^{-1/2} h_l^{2-q}$, we obtain a sampling ratio between successive levels given by

$$N_l / N_{l+1} = 2^{2-q}, \quad \forall l = 0, \dots, L-1, \quad (4.116)$$

and a sampling strategy defined by

$$N_l \sim \mathcal{O}(2^{(2-q)(L-l)}), \quad \forall l = 0, \dots, L. \quad (4.117)$$

All these results are summarized in Propositions 4.5 and 4.6.

Proposition 4.5 (MLMC convergence rate). *The H^q -error ($q = 0, 1$) produced by the MLMC estimator converges asymptotically with a rate of $2^{(q-2)L}$ provided a sampling strategy given by*

$$N_{\text{MLMC},l} \sim \mathcal{O}(2^{(4-2q)(L-l)}), \quad \forall l = 0, \dots, L. \quad (4.118)$$

Proposition 4.6 (MLQMC convergence rate). *The H^q -error ($q = 0, 1$) produced by the MLQMC estimator converges asymptotically with a rate of $2^{(q-2)L}$ provided a sampling strategy given by*

$$N_{MLQMC,l} \sim \mathcal{O}(2^{(2-q)(L-l)}), \quad \forall l = 0, \dots, L. \quad (4.119)$$

4.7 Work comparison

It may seem counter-intuitive from Propositions 4.3, 4.4, 4.5 and 4.6 that all the quadrature methods presented in this work converge asymptotically with a similar rate. However, they indeed do so as the error is always bounded by the considered discretization. This, nonetheless, should not be interpreted in the sense that all these methods perform just as equal. The different outcome regarding these methods arises at the level of their performance, and the computational cost they imply.

In order to quantitatively evaluate this difference, we rely here on the notion of work, defined to be a theoretical measure of performance, materialized by a balance of the computational cost of computing a single sample for a given discretization level (in terms of floating points operations), and the number of samples to be computed for reaching a desired accuracy.

For the sake of simplicity, we will separate the cases for the L^2 - and H^1 -error.

Theoretical work for convergence in the L^2 - error:

Let us recall that in order to achieve an error of order 2^{-2l} in the L^2 -norm using linear FE, the cost of solving a sample is given by

$$C_{FE,l} = 2^{\gamma dl}, \quad (4.120)$$

where γ is the complexity of the solver used and d is the dimension of the PDE problem considered (here $d = 4$). Given the convergence rates for MC and QMC, the number of samples to be executed by these methods on a level l to get to the same order of error (2^{-2l}) is provided by Propositions 4.3 and 4.4, respectively given by

$$N_{MC,l} = 2^{4l} \quad (4.121)$$

and

$$N_{QMC,l} = 2^{2l}. \quad (4.122)$$

We can recover the total amount of work required by both methods given a fine discretization level L :

$$W_{MC,L} = C_{FE,L} N_{MC,L} = 2^{\gamma dL} 2^{4L} = 2^{(\gamma d+4)L}, \quad (4.123)$$

and

$$W_{\text{QMC},L} = C_{\text{FE},L} N_{\text{QMC},L} = 2^{\gamma d L} 2^{2L} = 2^{(\gamma d + 2)L}. \quad (4.124)$$

The total work for MLMC with $L + 1$ discretization levels can also be written in terms of the sum of works at every involved level:

$$W_{\text{MLMC},L} = \sum_{l=0}^L C_{\text{FE},l} N_{\text{MLMC},l} = \sum_{l=0}^L 2^{\gamma d l} 2^{4(L-l)} = 2^{4L} \sum_{l=0}^L 2^{(\gamma d - 4)l}. \quad (4.125)$$

In that case, if $\gamma d \neq 4$, we have a geometric sum and we can write

$$W_{\text{MLMC},L} = 2^{4L} \frac{1 - 2^{(\gamma d - 4)(L+1)}}{1 - 2^{\gamma d - 4}} = \frac{2^{4L} - 2^{(\gamma d L + \gamma d - 4)}}{1 - 2^{\gamma d - 4}} = \frac{2^{4L} - A 2^{\gamma d L}}{1 - 2^{\gamma d - 4}}, \quad (4.126)$$

where we set $A = 2^{\gamma d - 4}$. The case where $\gamma d = 4$ is rather trivial, and we can resume the work expression for MLMC in the following form:

$$W_{\text{MLMC},L} = \begin{cases} (L + 1)2^{4L} & \text{if } \gamma d = 4, \\ \frac{A 2^{\gamma d L} - 2^{4L}}{2^{\gamma d - 4} - 1} & \text{if } \gamma d \neq 4. \end{cases} \quad (4.127)$$

Similarly, the total work for MLQMC can be recovered by computing:

$$W_{\text{MLQMC},L} = \sum_{l=0}^L C_{\text{FE},l} N_{\text{MLQMC},l} = \sum_{l=0}^L 2^{\gamma d l} 2^{2(L-l)} = 2^{2L} \sum_{l=0}^L 2^{(\gamma d - 2)l}. \quad (4.128)$$

This yields the following expression:

$$W_{\text{MLQMC},L} = \begin{cases} (L + 1)2^{2L} & \text{if } \gamma d = 2, \\ \frac{B 2^{\gamma d L} - 2^{2L}}{2^{\gamma d - 2} - 1} & \text{if } \gamma d \neq 2, \end{cases} \quad (4.129)$$

for $B = 2^{\gamma d - 2}$.

Remark. The asymptotical work behaviour for MLMC is bounded by $\mathcal{O}(2^{4L})$ when $\gamma d \leq 4$, and by $\mathcal{O}(2^{\gamma d L})$ if $\gamma d > 4$. For MLQMC, the asymptotical work is bounded by $\mathcal{O}(2^{2L})$ when $\gamma d \leq 2$, and by $\mathcal{O}(2^{\gamma d L})$ if $\gamma d > 2$. Regarding the MC and the QMC, the asymptotical work is augmented by few orders of magnitude to reach $\mathcal{O}(2^{(\gamma d + 4)L})$ and $\mathcal{O}(2^{(\gamma d + 2)L})$, respectively.

Theoretical work for convergence in the H^1 -error:

The study of the methods' work for the error reduction in the H^1 -norm is similar

to what has been developed above. The cost of solving samples stays unchanged, instead the number of samples for reaching a 2^{-l} error rate is given by:

$$N_{MC,l} = 2^{2l}, \quad N_{QMC,l} = 2^l, \quad (4.130)$$

and

$$N_{MLMC,l} = 2^{2(L-l)}, \quad N_{MLQMC,l} = 2^{2(L-l)}. \quad (4.131)$$

This provides us with the following expressions:

$$W_{MC,L} = 2^{(\gamma d + 2)L}, \quad W_{QMC,L} = 2^{(\gamma d + 1)L}, \quad (4.132)$$

$$W_{MLMC,L} = \begin{cases} (L+1)2^{2L} & \text{if } \gamma d = 2, \\ \frac{C2^{\gamma d L} - 2^{2L}}{2^{\gamma d - 2} - 1} & \text{if } \gamma d \neq 2, \end{cases} \quad (4.133)$$

and

$$W_{MLQMC,L} = \begin{cases} (L+1)2^L & \text{if } \gamma d = 1, \\ \frac{D2^{\gamma d L} - 2^L}{2^{\gamma d - 1} - 1} & \text{if } \gamma d \neq 1. \end{cases} \quad (4.134)$$

Here, $C = 2^{\gamma d - 2}$ and $D = 2^{\gamma d - 1}$.

General formulation of the work for convergence in H^q - error ($q=0,1$):

For the sake of simplicity, we have treated separately the cases regarding $q = 0$ and $q = 1$. By looking closely at the obtained formulations, one can actually notice that these can be concatenated in the following way:

$$W_{MC,L}^q = 2^{(\gamma d + 4 - 2q)L}, \quad W_{QMC,L}^q = 2^{(\gamma d + 2 - q)L}, \quad (4.135)$$

$$W_{MLMC,L}^q = \begin{cases} (L+1)2^{(4-2q)L} & \text{if } \gamma d = 4 - 2q, \\ \frac{C(q)2^{\gamma d L} - 2^{(4-2q)L}}{2^{\gamma d - (4-2q)} - 1} & \text{if } \gamma d \neq 4 - 2q, \end{cases} \quad (4.136)$$

and

$$W_{MLQMC,L}^q = \begin{cases} (L+1)2^{(2-q)L} & \text{if } \gamma d = 2 - q, \\ \frac{D(q)2^{\gamma d L} - 2^{(2-q)L}}{2^{\gamma d - (2q-1)} - 1} & \text{if } \gamma d \neq 2 - q, \end{cases} \quad (4.137)$$

with $C(q) = 2^{\gamma d - (4-2q)}$ and $D(q) = 2^{\gamma d - (2-q)}$.

4.8 Transfer of discrete fields

In the context of multilevel stochastic quadrature methods, information needs to be transferred from a level to another; random fields are computed on the fine level and need to be interpreted on coarser levels, and reciprocally quantities-of-interest (c.f. Section 5.3, Chapter 5) computed on coarser levels are to be transferred to the original fine level. We here detail how to proceed, and distinguish between two cases with respect to the nestedness or not of the used mesh hierarchy.

4.8.1 Nested meshes

Random fields from fine to coarse: The transfer of random fields from the fine space to the coarse space is done by L^2 -projection as they are proven to be optimal, stable and in general superior to interpolation, see [HB06]. The L^2 -projection we employ is constructed by exploiting the properties of the dual basis, as in [Woh00, PWGW12]. This approach allows for computing transfer matrices on an *offline* basis.

For $l = 0, \dots, L - 1$, let $G_l \in \mathbb{R}^{n_l}$ be the coefficients of the discrete random field on level l , where n_l is the number of spatial nodes on l . Provided the spatial discrete representation $P_{l,h} \in \mathbb{R}^{n_l \times n}$ of the L^2 -projection (or transfer matrix) from the fine level to a coarser level l computed by the method indicated above, the transferred discrete random field G_l is obtained by the matrix-vector multiplication,

$$G_l = P_l^l G_L,$$

where $G_L \in \mathbb{R}^n$ is the discrete random field computed for level L .

Quantities-of-interest from coarse to fine: We start by considering the overall space-time transmembrane potential as a quantity-of-interest. This represents a different situation from the case cited above as it requires a transfer in the space and time dimensions. By exploiting the tensor-product structure of (3.51) we can simplify the implementation and benefit from better computational performance. In this case, the tensor-product structure of the space-time grid allows us to construct space-time operators in a convenient way which requires the assembly of the spatial transfer operator to be performed only once.

For $l = 0, \dots, L - 1$, let $u_l \in \mathbb{R}^{n_l m_l}$ be the coefficients of the discrete space-time solution on level l where m_l represents the number of time steps. The space-time transfer of the overall solution quantity-of-interest is performed from coarse to fine level, by means of a tensor product of the space and time interpolation

matrices. Provided the discrete representations of the spatial and temporal interpolation matrices $I_{l,h}^L \in \mathbb{R}^{n \times n_l}$ and $I_{l,t}^L \in \mathbb{R}^{m \times m_l}$ from coarse level l to fine level L , the transfer matrix is *offline*-obtained by means of the tensor product

$$I_l^L = I_{l,t}^L \otimes I_{l,h}^L \in \mathbb{R}^{nm \times n_l m_l}.$$

The transferred fine field is then obtained by matrix–vector multiplication,

$$u_L = I_l^L u_l,$$

where $u_L \in \mathbb{R}^{nm}$ is the discrete space–time solution obtained at the fine level L . When considering exclusively time–dependent quantities-of-interest, such as the action potential on a given location of the domain, the transfer is obtained by using the time interpolation matrix $I_{l,t}^L \in \mathbb{R}^{m \times m_l}$. That is, given the discrete representation of the time variant quantity-of-interest $u_l \in \mathbb{R}^{m_l}$, the transfer on the fine level is obtained by means of the matrix–vector multiplication

$$u_L = I_{l,t}^L u_l \in \mathbb{R}^m.$$

Regarding scalar quantities-of-interest, such as the activation time for a given location of the domain, no transfer is needed.

4.8.2 Non–nested meshes

Random fields from fine to coarse: According to [Sie15], the computation of the stiffness matrix with respect to the random diffusion field was shown to be consistent with a piecewise linear finite element discretization when the midpoint rule with respect to the actual grid is applied. Therefore, for the transfer of the random fields from the finest level to the coarser levels, we perform an element–wise transfer based on the midpoint rule; that is, for every element of the coarse level, we assign a constant diffusion value corresponding to the fine element containing its center. Therefore, the assembly of the matrix on the coarser levels can be performed with linear complexity relative to the particular level of discretization.

Quantities-of-interest from coarse to fine: We only consider in this case time–dependent and scalar quantities-of-interest. We proceed as indicated in the case of nested meshes.

Chapter 5

Uncertainty Quantification for Heart Electrophysiology

5.1 Fiber uncertainty

The mathematical and numerical instruments developed in decades of research in the electrophysiology field allow in principle for virtual therapy planning. The monodomain equation is by now an established model in cardiac electrophysiology, describing with high accuracy the electrical activity in the myocardium. Nonetheless, patient-specific simulations are still not widely employed as a routine tool in the treatment of patients.

A particular reason for this can be found in the data which is acquired in clinical practice. For instance, the fiber structure still represents a great challenge to be determined from an imaging point of a view. One should therefore account for possible uncertainties of the diffusion field in the monodomain equation, which mathematically models the fiber conductivities and orientations.

Modeling the uncertainty requires to account for an additional stochastic variable in the formulation of the monodomain equation. Let us denote this stochastic variable with $\omega \in \Omega$ where Ω is, as before, the set of all possible outcomes, which in this particular case represents the set of all diffusion fields modelled by the uncertainty. We are now interested in estimating the statistics of $u(\mathbf{x}, t, \omega) = u(\mathbf{X}, \omega)$, i.e. $\mathbb{E}[u(\mathbf{x}, t)] = \mathbb{E}[u(\mathbf{X})]$, as a solution to the monodomain stochastic PDE, which reads: for almost every $\omega \in \Omega$:

$$\frac{\partial u(\mathbf{X}, \omega)}{\partial t} - \nabla \cdot (G(\mathbf{x}, \omega) \nabla u(\mathbf{X}, \omega)) + I_{\text{ion}}(u(\mathbf{X}, \omega)) = I_{\text{app}}(\mathbf{X}), \quad (5.1)$$

for $\mathbf{X} = (\mathbf{x}, t) \in D \times (0, T]$. The stochastic model (5.1) can furthermore be expressed with the parametric formulation:

$$\frac{\partial u(\mathbf{X}, \mathbf{y})}{\partial t} - \nabla \cdot (G(\mathbf{x}, \mathbf{y}) \nabla u(\mathbf{X}, \mathbf{y})) + I_{\text{ion}}(u(\mathbf{X}, \mathbf{y})) = I_{\text{app}}(\mathbf{X}), \quad (5.2)$$

for all $\mathbf{y} \in [-1, 1]^M$, after identifying the stochastic space Ω with its image $[-1, 1]^M$ using the mapping

$$\begin{aligned} \psi : \Omega &\rightarrow [-1, 1]^M, \\ \omega &\mapsto (\psi_1(\omega), \dots, \psi_M(\omega)), \end{aligned}$$

as done in Subsection 4.2.3. One can now approximate statistics of the stochastic solution, as for example the expectation (mean)

$$\mathbb{E}[u](\mathbf{X}) = \int_{[-1, 1]^M} u(\mathbf{X}, \mathbf{y}) \rho(\mathbf{y}) d\mathbf{y}, \quad (5.3)$$

or the variance

$$\mathbb{V}[u](\mathbf{X}) = \int_{[-1, 1]^M} (u(\mathbf{X}, \mathbf{y}) - \mathbb{E}[u](\mathbf{X}))^2 \rho(\mathbf{y}) d\mathbf{y}, \quad (5.4)$$

where $\rho(\mathbf{y}) = \prod_{i=1}^M \rho_k(y_k)$ is the joint density function for M uniformly distributed random variables. The approximation of the aforementioned quantities can be done by means of single and multilevel Monte-Carlo techniques. These methods imply solving the monodomain equation for different independent and identically distributed samples of random diffusion fields.

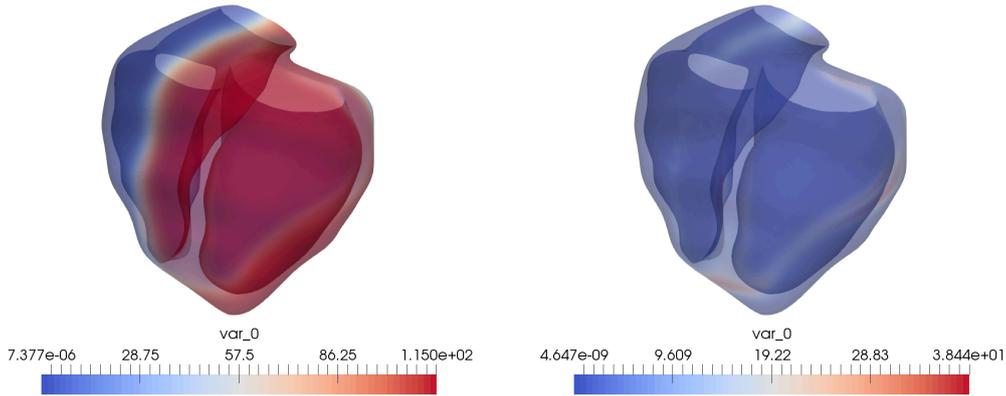


Figure 5.1. Example of first (mean, left) and centered second (variance, right) moments of the solution to (5.2).

In the case where we consider an isotropic diffusion, a random diffusion field is obtained by considering a realization of the KL expansion

$$V(\mathbf{x}, \mathbf{y}) = \mathbb{E}[V](\mathbf{x}) + \sum_{k=1}^M \sqrt{\lambda_k} \phi_k(\mathbf{x}) y_k. \quad (5.5)$$

and setting $G(\mathbf{x}, \mathbf{y}) = V(\mathbf{x}, \mathbf{y}) \in \mathbb{R}$, for all $\mathbf{x} \in D$ and $\mathbf{y} \in [-1, 1]^M$. One can equivalently set

$$G(\mathbf{x}, \mathbf{y}) = \begin{bmatrix} V(\mathbf{x}, \mathbf{y}) & & \\ & \ddots & \\ & & V(\mathbf{x}, \mathbf{y}) \end{bmatrix} \in \mathbb{R}^{d \times d}.$$

This allows for a straightforward physical interpretation of the current scenario, that is an equal diffusion prescribed in each direction $\mathbf{e}_i \in \mathbb{R}^d$, $i = 1, \dots, d$, forming the standard canonical basis. We therefore refer to this setting as the diffusion uncertainty. We can instead consider the more realistic configuration consisting of anisotropic diffusion, by non restricting to diagonal tensors. Given a symmetric tensor diffusion $G(\mathbf{x}) \in \mathbb{R}^{d \times d}$ for $\mathbf{x} \in D$, it is possible to find a corresponding diagonalization

$$G(\mathbf{x}) = S(\mathbf{x})\Lambda(\mathbf{x})S^{-1}(\mathbf{x}), \quad (5.6)$$

for a matrix $S(\mathbf{x})$ consisting of d orthogonal eigenvectors, and a diagonal matrix $\Lambda(\mathbf{x})$ consisting of eigenvalues. Furthermore, whenever $G(\mathbf{x})$ is initially positive definite, the eigenvalues are also positive. In the case where $d = 3$, these positive eigenvalues define three conductivities σ_f , σ_s , and σ_n in the fiber, sheet normal, and fiber normal directions (respectively $\mathbf{e}_f(\mathbf{x})$, $\mathbf{e}_s(\mathbf{x})$ and $\mathbf{e}_n(\mathbf{x})$) provided by the matrix $S(\mathbf{x})$. Hence, a similar scenario has to be accounted for when considering fiber orientation and diffusion uncertainties.

This can be done by means of the vector valued random fields,

$$\mathbf{V}(\mathbf{x}, \mathbf{y}) = \mathbb{E}[\mathbf{V}](\mathbf{x}) + \sum_{k=1}^M \sqrt{\lambda_k} \phi_k(\mathbf{x}) y_k, \quad (5.7)$$

which are used to define the random diffusion tensors, see [HPS17b],

$$G(\mathbf{x}, \mathbf{y}) := aI_d + (\|\mathbf{V}(\mathbf{x}, \mathbf{y})\|_2 - a) \frac{\mathbf{V}(\mathbf{x}, \mathbf{y})\mathbf{V}^T(\mathbf{x}, \mathbf{y})}{\mathbf{V}^T(\mathbf{x}, \mathbf{y})\mathbf{V}(\mathbf{x}, \mathbf{y})}, \quad (5.8)$$

for $a > 0$. It is easily verifiable that such a tensor diffusion matrix is symmetric. It is, moreover, positive definite since we have for all $\mathbf{x} \in D$, $\mathbf{y} \in [-1, 1]^M$ and $\mathbf{z} \in \mathbb{R}^d$:

$$\begin{aligned} \mathbf{z}^T G(\mathbf{x}, \mathbf{y}) \mathbf{z} &= a \mathbf{z}^T \mathbf{z} + (\|\mathbf{V}(\mathbf{x}, \mathbf{y})\|_2 - a) \frac{\mathbf{z}^T \mathbf{V}(\mathbf{x}, \mathbf{y}) \mathbf{V}^T(\mathbf{x}, \mathbf{y}) \mathbf{z}}{\mathbf{V}^T(\mathbf{x}, \mathbf{y}) \mathbf{V}(\mathbf{x}, \mathbf{y})} \\ &= a \left(\mathbf{z}^T \mathbf{z} - \frac{\mathbf{z}^T \mathbf{V}(\mathbf{x}, \mathbf{y}) \mathbf{V}^T(\mathbf{x}, \mathbf{y}) \mathbf{z}}{\mathbf{V}^T(\mathbf{x}, \mathbf{y}) \mathbf{V}(\mathbf{x}, \mathbf{y})} \right) + \underbrace{\|\mathbf{V}(\mathbf{x}, \mathbf{y})\|_2 \frac{\mathbf{z}^T \mathbf{V}(\mathbf{x}, \mathbf{y}) \mathbf{V}^T(\mathbf{x}, \mathbf{y}) \mathbf{z}}{\mathbf{V}^T(\mathbf{x}, \mathbf{y}) \mathbf{V}(\mathbf{x}, \mathbf{y})}}_{\geq 0}. \end{aligned}$$

The first term of the equality can in term be rewritten as:

$$a \left(\mathbf{z}^T \mathbf{z} - \frac{\mathbf{z}^T \mathbf{V}(\mathbf{x}, \mathbf{y}) \mathbf{V}^T(\mathbf{x}, \mathbf{y}) \mathbf{z}}{\mathbf{V}^T(\mathbf{x}, \mathbf{y}) \mathbf{V}(\mathbf{x}, \mathbf{y})} \right) = a \left(\|\mathbf{z}\|_2^2 - \frac{\langle \mathbf{z}, \mathbf{V}(\mathbf{x}, \mathbf{y}) \rangle^2}{\|\mathbf{V}(\mathbf{x}, \mathbf{y})\|_2^2} \right), \quad (5.9)$$

from which we deduce the nonnegativity by the Cauchy–Schwarz inequality. We therefore have:

$$\mathbf{z}^T G(\mathbf{x}, \mathbf{y}) \mathbf{z} \geq 0. \quad (5.10)$$

Furthermore, by considering for every $\mathbf{x} \in D$ an arbitrary vector $\tilde{\mathbf{z}}(\mathbf{x})$ that decomposes in the following form

$$\tilde{\mathbf{z}}(\mathbf{x}) = \alpha \mathbf{V}(\mathbf{x}, \mathbf{y}) + \beta \mathbf{V}_\perp(\mathbf{x}, \mathbf{y}), \quad (5.11)$$

where $\mathbf{V}_\perp(\mathbf{x}, \mathbf{y})$ is perpendicular to $\mathbf{V}(\mathbf{x}, \mathbf{y})$. When multiplying this vector with the diffusion tensor defined in (5.8), we obtain:

$$G(\mathbf{x}, \mathbf{y}) \tilde{\mathbf{z}}(\mathbf{x}) = \alpha G(\mathbf{x}, \mathbf{y}) \mathbf{V}(\mathbf{x}, \mathbf{y}) + \beta G(\mathbf{x}, \mathbf{y}) \mathbf{V}_\perp(\mathbf{x}, \mathbf{y}). \quad (5.12)$$

When using the following identities,

$$\left(\mathbf{V}(\mathbf{x}, \mathbf{y}) \mathbf{V}^T(\mathbf{x}, \mathbf{y}) \right) \mathbf{V}(\mathbf{x}, \mathbf{y}) = \mathbf{V}(\mathbf{x}, \mathbf{y}) \left(\mathbf{V}^T(\mathbf{x}, \mathbf{y}) \mathbf{V}(\mathbf{x}, \mathbf{y}) \right) = \|\mathbf{V}(\mathbf{x}, \mathbf{y})\|_2^2 \mathbf{V}(\mathbf{x}, \mathbf{y}), \quad (5.13)$$

and

$$\left(\mathbf{V}(\mathbf{x}, \mathbf{y}) \mathbf{V}^T(\mathbf{x}, \mathbf{y}) \right) \mathbf{V}_\perp(\mathbf{x}, \mathbf{y}) = \mathbf{V}(\mathbf{x}, \mathbf{y}) \left(\mathbf{V}^T(\mathbf{x}, \mathbf{y}) \mathbf{V}_\perp(\mathbf{x}, \mathbf{y}) \right) = 0, \quad (5.14)$$

we can deduce that

$$G(\mathbf{x}, \mathbf{y}) \mathbf{V}(\mathbf{x}, \mathbf{y}) = \|\mathbf{V}(\mathbf{x}, \mathbf{y})\|_2 \mathbf{V}(\mathbf{x}, \mathbf{y}), \quad (5.15)$$

and

$$G(\mathbf{x}, \mathbf{y}) \mathbf{V}_\perp(\mathbf{x}, \mathbf{y}) = a \mathbf{V}_\perp(\mathbf{x}, \mathbf{y}). \quad (5.16)$$

This implies the following:

$$G(\mathbf{x}, \mathbf{y})\tilde{\mathbf{z}}(\mathbf{x}) = \|\mathbf{V}(\mathbf{x}, \mathbf{y})\|_2 (\alpha\mathbf{V}(\mathbf{x}, \mathbf{y})) + a (\beta\mathbf{V}_\perp(\mathbf{x}, \mathbf{y})), \quad (5.17)$$

from which we deduce that the action of the tensor $G(\mathbf{x}, \mathbf{y})$ amplifies the vector $\tilde{\mathbf{z}}(\mathbf{x})N$ by $\|\mathbf{V}(\mathbf{x}, \mathbf{y})\|_2$ in the direction of $\mathbf{V}(\mathbf{x}, \mathbf{y})$ and by the factor a in the direction of $\mathbf{V}_\perp(\mathbf{x}, \mathbf{y})$. Therefore, when defined in the diffusion operator of the PDE, the tensor $G(\mathbf{x}, \mathbf{y})$ prescribes a conductivity of strength $\|\mathbf{V}(\mathbf{x}, \mathbf{y})\|_2$ in the direction of $\mathbf{V}(\mathbf{x}, \mathbf{y})$, and a conductivity of strength a perpendicularly. For this reason, this model allows for fibers orientation and diffusion uncertainties. The uncertainty with respect to orientation is however only considered for the fibers direction, i.e. $\mathbf{e}_f(\mathbf{x})$. The same holds with respect to diffusion, where the unique considered conductivity uncertainty is given by σ_f . It is however possible to consider the parameter a as an additional stochastic variable, including therefore an uncertainty on $\sigma_n = \sigma_s$, which is imposed by the model.

5.2 Scar fibrosis uncertainty

Scar identification represents a significantly meaningful input for clinical practice. Typically, inverse problems are used to identify parameters modeling a scar (position and width mainly) where the control is updated in order to approximate a target solution (obtained through an ECG for example). However, these inverse problems are usually non convex, and a regularization term is needed. The obtained characterization of the scar is therefore influenced by the regularization term and this can serve as a basis point for considering UQ in this context. Many sources of uncertainties can be considered when modeling a scar. The position, the width, the conductivity inside the scar are all possible choices for an uncertainty. For example, one can keep considering fiber uncertainty over the domain as in (5.1), and combine it to an uncertainty related to the width of the scar and its conductivity. The following setting assumes the position of the scar is known, which can be motivated by prior investigation considering an inverse problem.

Mathematically, the scar region can be modeled by means of a Gaussian function, see e.g. [Pag17, BB15], given by

$$s(\mathbf{x}; \mu) = \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}_0\|^2}{\mu^2}\right),$$

where \mathbf{x}_0 is the center of the scar, $\mathbf{x} \in D \subset \mathbb{R}^d$ a point of the considered domain and $\mu \in \mathbb{R}$ is the width of the scar. The function $s(\mathbf{x}; \mu)$ has therefore the shape of a bell centered in \mathbf{x}_0 and the healthy part of the tissue can be characterized by

$$h(\mathbf{x}; \mu) = 1 - s(\mathbf{x}; \mu) = 1 - \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}_0\|^2}{\mu^2}\right).$$

Let us furthermore denote by $\sigma \in \mathbb{R}$ the conductivity inside the scar. The conductivity of the cardiac tissue $G(\mathbf{x}; \mathbf{y}, \mu, \sigma)$ is therefore subject to the three uncertainties $\mathbf{y} \in \mathbb{R}^M$ (M is the stochastic dimension obtained by means of the KL expansion), $\mu \in \mathbb{R}^+$ and $\sigma \in \mathbb{R}^+$. It can be written under the following form:

$$G(\mathbf{x}; \mathbf{y}, \mu, \sigma) = G_h(\mathbf{x}; \mathbf{y})h(\mathbf{x}; \mu) + \sigma(1 - h(\mathbf{x}; \mu)),$$

where $G_h(\mathbf{x}; \mathbf{y})$ is the conductivity for the healthy portion of the cardiac tissue and is obtained through the KL expansion. The model consists now of the following stochastic PDE: given $(\mathbf{y}, \mu, \sigma) \in \mathbb{R}^{M+2}$, find $u(\cdot, \mathbf{y}, \mu, \sigma) = u(\mathbf{x}, t; \mathbf{y}, \mu, \sigma)$ such that

$$\frac{\partial u(\cdot, \mathbf{y}, \mu, \sigma)}{\partial t} - \nabla \cdot (G(\cdot, \mathbf{y}, \mu, \sigma) \nabla u(\cdot, \mathbf{y}, \mu, \sigma)) + I_{\text{ion}}(u(\cdot, \mathbf{y}, \mu, \sigma)) = I_{\text{app}}(\cdot),$$

in $D \times (0, T]$, where the ion channel term is now written as

$$I_{\text{ion}}(u(\cdot, \mathbf{y}, \mu, \sigma)) = \alpha h(\mathbf{x}; \mu)(u - u_{\text{rest}})(u - u_{\text{th}})(u - u_{\text{peak}}),$$

since the ion exchanges in ill regions are most likely to be impacted. The diffusion field randomness $\mathbf{y} = (y_1, \dots, y_M) \in \mathbb{R}^M$ can still be sampled as previously done, i.e. every y_i is a realization of a uniformly distributed random variable $\psi_i \sim \mathcal{U}(-1, 1)$, whereas μ and σ can be sampled independently from each other following two other uniform distributions $\mathcal{U}(\mu_{\min}, \mu_{\max})$ and $\mathcal{U}(\sigma_{\min}, \sigma_{\max})$ respectively. Here, (μ_{\min}, μ_{\max}) and $(\sigma_{\min}, \sigma_{\max})$ are the considered bounds for the scar width and conductivity.

The aim is now to approximate the expectancy of the solution $u(\mathbf{x}, t)$

$$\mathbb{E}[u](\mathbf{x}, t) = \int_{[-1, 1]^M \times [\mu_{\min}, \mu_{\max}] \times [\sigma_{\min}, \sigma_{\max}]} u(\mathbf{x}, t; \mathbf{y}, \mu, \sigma) \rho_{\mathbf{y}}(\mathbf{y}) \rho_{\mu}(\mu) \rho_{\sigma}(\sigma) d\mathbf{y} d\mu d\sigma,$$

using the usual stochastic quadrature methods.

5.3 Quantities-of-interest

One might be interested in approximating different quantities-of-interest, depending the purpose of the experiment. This can be done by reconsidering the

high-dimensional stochastic integral under the formulation

$$\mathbb{E}[\mathcal{F}(u)](\mathbf{X}) = \int_{[-1,1]^M} \mathcal{F}(u(\mathbf{X}, \mathbf{y})) \rho(\mathbf{y}) d\mathbf{y}, \quad (5.18)$$

\mathcal{F} denotes a functional that computes a particular quantity-of-interest. We mainly consider here three different quantities of interest.

Transmembrane potential The transmembrane potential over the totality of the heart geometry, and its evolution in time is the quantity obtained by solving the monodomain equation. This is demonstrated as an electrical potential wave travelling through the heart (c.f. Figure 5.2). The functional \mathcal{F} in this case is simply the identity function, i.e. $\mathcal{F}(u(\mathbf{X}, \mathbf{y})) = u(\mathbf{X}, \mathbf{y})$. In general and considering a fine discretization in space and time, the full information on the transmembrane potential represents a high dimensional output that might become a burden at the memory level in a context of a UQ study. This is the reason we might consider in some cases other quantities of interest.

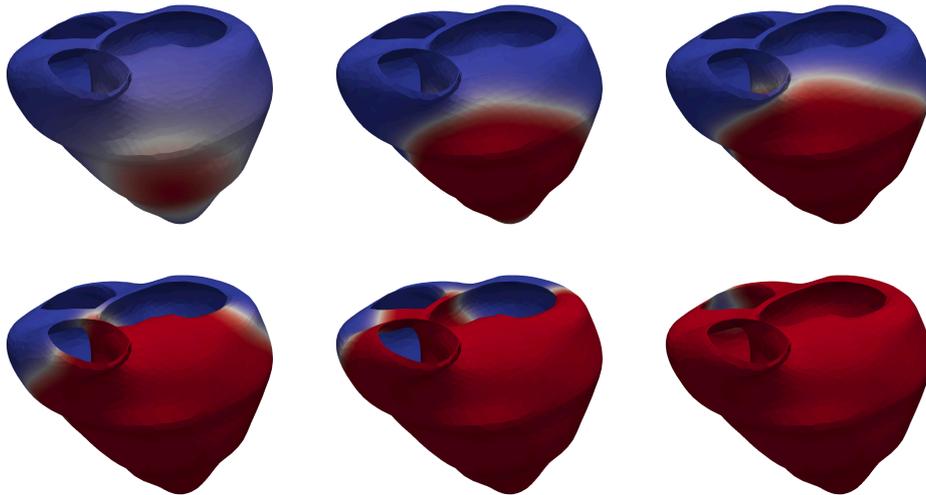


Figure 5.2. Wafefront propagation of transmembrane potential.

Activation map The transmembrane potential can be used in order to extract the activation map of the heart. This reduces the size of the output to that of the space dimension uniquely, indicating the times at which cells are activated. If the activation time in a given location \mathbf{x}_0 is denoted by $a(\mathbf{x}_0)$, then it is obtained by computing $a(\mathbf{x}_0, \mathbf{y}) = \min_t \{t \in [0, T] \text{ such that } u(\mathbf{x}_0, t, \mathbf{y}) \geq u_{\text{th}}\}$. The functional \mathcal{F} can therefore be written in this case as $\mathcal{F}(u(\mathbf{X}, \mathbf{y})) = a(\mathbf{x}, \mathbf{y})$.

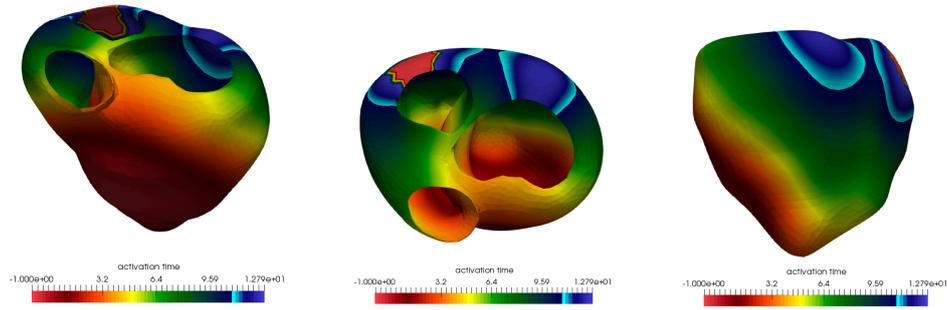


Figure 5.3. Activation map of a heart. Negative value signifies non activated region.

Action potential in a given location Another quantity of interest that can be considered is the evaluation of the transmembrane time evolution at a given location. This is typically what is clinically monitored when electrodes are placed on the torso of a patient. For cardiologists, this information is of a high degree of relevance as several rhythm troubles or dysfunctions can be revealed exclusively from that. Mathematically, for every point in space \mathbf{x}_0 , a functional $\mathcal{F}_{\mathbf{x}_0} = u(\mathbf{X}, \mathbf{y})|_{\mathbf{x}=\mathbf{x}_0}$.

Regarding the approximation of the integral (5.18) for a given functional \mathcal{F} , we rely on stochastic quadrature methods that require solving the mono-domain equation in quadrature points represented by different realization of $\mathbf{y} \in [-1, 1]^M$. The following procedure is labeled as sampling method and requires a finite element solve for every sample.

5.4 Prior insights on numerical experiments

It may be considered good practice to verify certain properties of the problem at hand before further deep investigations on the behaviour of the stochastic quadrature methods for quantities-of-interest approximation.

In view of samples computation at different levels of space-time discretizations, one may verify that the usual error estimates hold when considering the mono-domain equation. In Figure 5.4, we plot the L^2 -error evaluation of $L = 5$ levels of nested cube meshes in space and time, with respect to a reference solution computed at a fine level. The plot shows that the error estimates arising from this comparison yield the usual convergence rate of $\mathcal{O}(h^2) = \mathcal{O}(\Delta t^2)$.

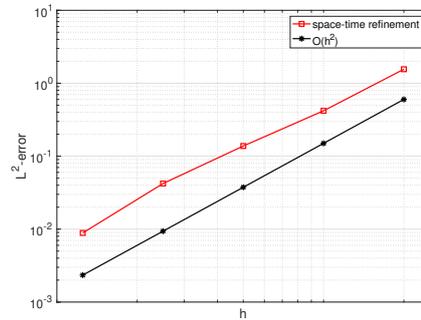


Figure 5.4. L^2 -error convergence measured over space-time uniform refinement.

In Figures 5.5 and 5.6, we conduct an MC mean and variance estimation for a cube and an idealized ventricle. We notice that in the absence of repolarization in the considered monodomain model, the variance is localized around the electrical potential traveling wave. Therefore, the unique disparity caused by the diffusion uncertainty is with respect to the speed of the travelling wave.

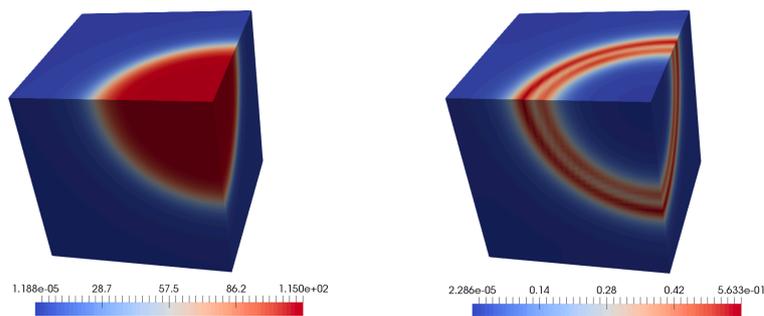


Figure 5.5. Mean (left) and variance (right) at the final time state for the cube.

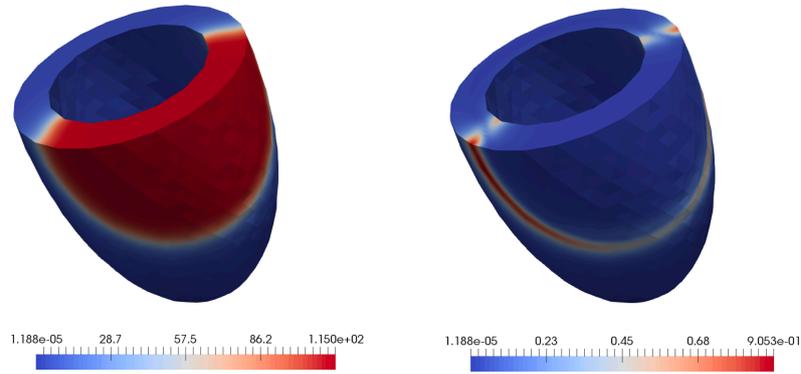


Figure 5.6. Mean (left) and variance (right) at the final time state for the ventricle.

We next provide in Figures 5.7 and 5.8 a visualization of the MLMC estimator (mean at coarse level and corrections between successive levels) for nested meshes hierarchies of a cube and an idealized ventricle.

It is interesting to notice the decreasing amplitude of corrections as levels are refined. One may also observe that these are localized around the travelling wave, but are more concerned by the discretization correction as one can clearly see from demonstrated shapes, i.e. size-decreasing cubes surrounding the travelling wave.

We also test the stochastic convergence rate for the single-level methods. In Figure 5.9, we report the MC convergence of the two aforementioned geometries on a given space-time discretization level using a reference solution computed with 10000 samples at that same level. The convergence rate seems to be consisting to the theoretical one, i.e. $\mathcal{O}(N^{-0.5})$. In Figure 5.10, we compare the convergence behaviour of the MC and QMC methods with respect to the same setting. We may notice a faster convergence rate for the QMC method provided by $\mathcal{O}(N^{-1})$.

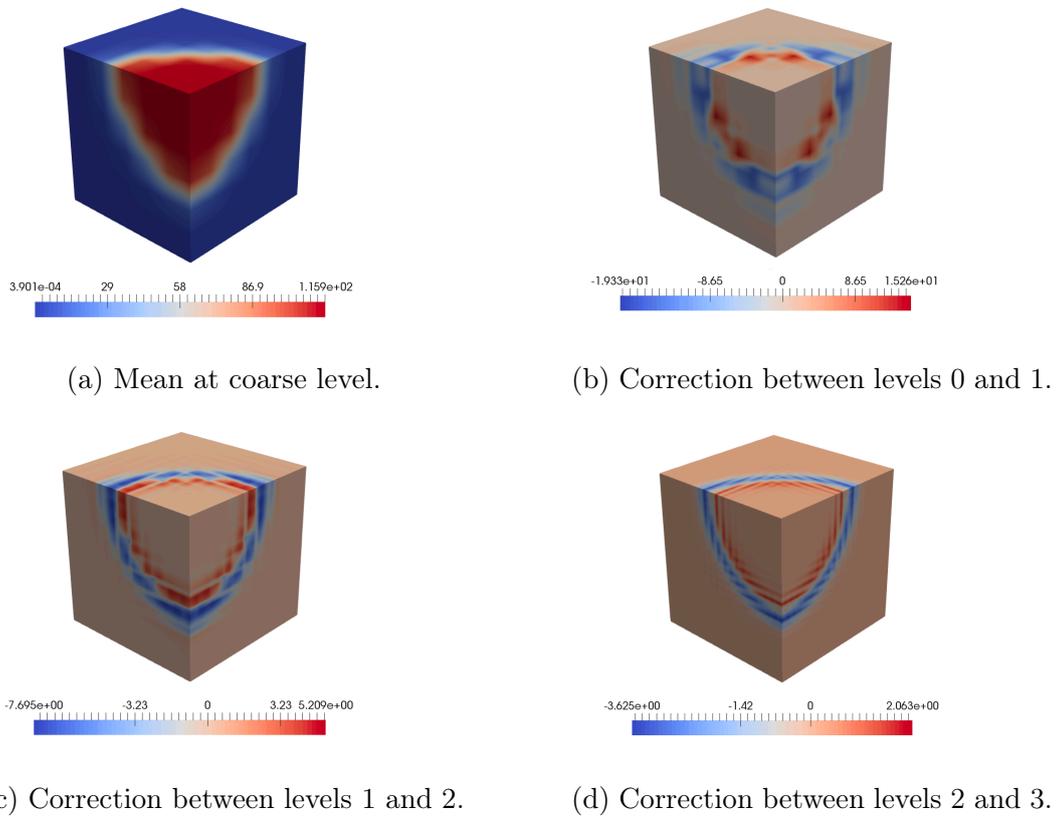


Figure 5.7. Mean at coarse level and corrections between successive levels for a cube geometry.

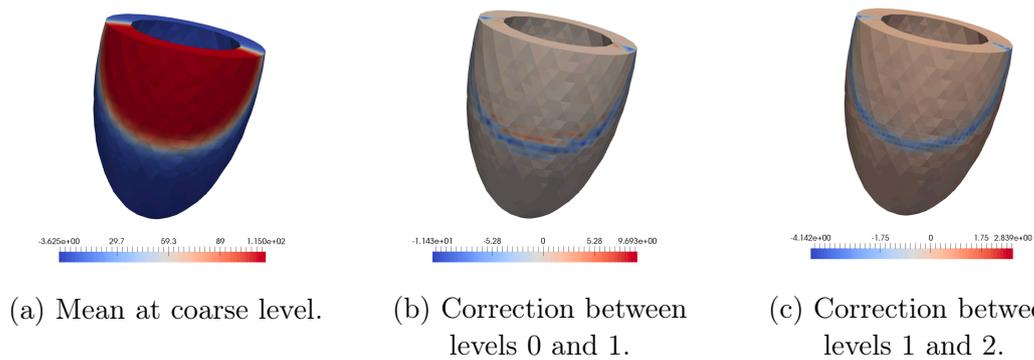


Figure 5.8. Mean at coarse level and corrections between successive levels for an idealized ventricle geometry.

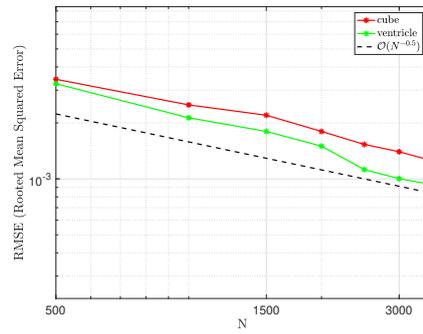


Figure 5.9. Convergence rate for the MC estimator for a cube and an idealized ventricle geometries.

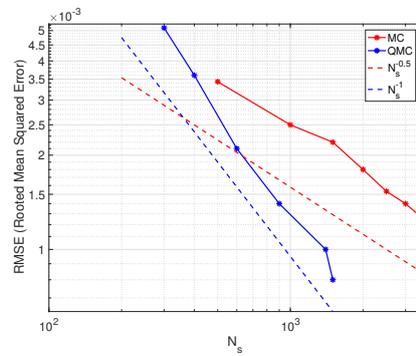


Figure 5.10. Convergence rate for the MC and QMC estimators for a cube geometry.

Chapter 6

Parallelization and Performance Considerations

For UQ problems, applying stochastic quadrature methods is particularly expensive in terms of needed computations. Even for a demonstrated efficiency of multilevel and deterministic techniques, a UQ study may remain unfeasible with respect to the total execution time, provided a non-efficient solve strategy for samples computations (e.g. sequential execution). Moreover, the generated data may require an extremely high amount of memory that is not affordable on a single machine. For this reason, one usually resorts to High Performance Computing (HPC) facilities, in order to fully profit from parallelization and adequate memory management. Yet, it remains an imperative to design a clever strategy that fully takes advantage of parallelizing concurrent tasks, as much as reducing their own execution time.

In what follows, we describe the computational framework designed in order to efficiently tackle the aforementioned obstacles. This framework follows from the work produced in [BBBQ⁺21]. It is characterized by a space-time and stochastics parallel execution, consisting of a so-called three-layer parallelization strategy for multilevel methods, combined to a space-time solver with a special care given to the Newton initial guess.

6.1 Preconditioned Generalized Minimal Residual Method (PGMRES)

The non-linear problem (3.75) is solved with Newton's method. For each of the arising linear problems in the form (3.51), we employ a space-time parallel

PGMRES, with the preconditioner given by an ILU(0) factorization within the PETSc [BAA⁺16] framework. If multiple processors are used, the corresponding block-Jacobi preconditioner is employed, and the PGMRES is applied. As default by PETSc, the GMRES solver is restarted after 30 iterations. It is important to recall that the first iterate plays an important role for the convergence of Newton's method. The strategy for setting the Newton initial guess is discussed in Section 6.4.

6.2 Parallelization strategy

6.2.1 Concurrent tasks and parallelization paradigm

Multilevel techniques were designed in the first place to reduce the high computational complexity demonstrated by a standard MC approach, carried on a fine domain discretization. However, in order to fully exploit their capabilities, an adequate parallel environment should be designed in a way that fully takes advantage of all different layers under which concurrent work can be executed.

Multilevel methods entail computing different numbers of samples on distinct discretization levels, independently from each other, both in terms of the considered level and sample. From this observation, one may draw two main conclusions. On the one hand, the work load across levels is completely independent and can be executed concurrently. On the other hand, different tasks (samples computation) among distinct levels do not require the same amount of parallel resources in order to execute in a reasonable time.

This second observation is important to many extents. First of all, it is rather preferable not to compute samples with overly many resources than they would actually require. But also, this would allow for a better use of the saved computational resources as they can be deployed on coarse levels which have a more important number of samples instead.

To illustrate this, let us suppose we are considering an MLMC with $L + 1$ levels and L^2 as an error metric. We recall that in this case, the number of samples to compute on each level is provided by (c.f. Section 4.6)

$$N_l \sim \mathcal{O}(2^{4(L-l)}).$$

Let us furthermore consider a situation where we would dispose of a certain amount of HPC nodes with $P = 2^{4L}$ processors for every single node (a very

unlikely scenario for a single node, especially when L is big). One can totally compensate for the higher number of samples on coarse levels by using $P_l = 2^{4l}$ processors for a single sample on level l , while taking advantage of the remaining processors to concurrently compute $P/P_l = 2^{4(L-l)}$ samples, which corresponds exactly to the number of samples required for each level l . If additionally, the solver complexity was provided by $\gamma = 1$ and $d = 4$ such that the cost for solving a sample at level l is given by

$$C_{\text{FE},l} = 2^{\gamma d l} = 2^{4l},$$

we would have all nodes for different levels terminating at the same time (assuming perfect scalability and neglecting communication latencies). The combination of these assumptions might be very unlikely, yet the example explains the philosophy to be adopted for the parallelization strategy.

6.2.2 Parallelization algorithm

The ideas developed above can be conceptualized in a so-called three layer parallelization strategy, similar to what has been developed in [DGR⁺17], where many other strategies are considered and compared. The strategy consists of initially distributing the required amount of work for every level on the available resources, as a first parallelization layer. On every one of these levels, multiple samples need to be computed. Those are completely independent and we can consider distributing computing resources across batches of samples as a second parallelization layer. The third layer, in the context of finite elements, would naturally be a parallelization across the tempo-spatial grid, using the space-time parallel solver described in Section 6.1.

More formally, let us consider a MLMC setting of $L + 1$ levels with corresponding N_l samples to compute. The following algorithm defines the steps of the above described strategy.

Algorithm 1 Three-layer parallelization strategy for multilevel methods

- 1: Decide on t_l batches of samples of size N_l/t_l on each level l
 - 2: Decide on the number of processors P_l to use for a single sample on each level l
 - 3: Proceed in parallel for $l = 0, \dots, L$ and $k = 1, \dots, t_l$ by:
 recursively solve the N_l/t_l samples of batches T_{l,t_k} using P_l processors
-

6.3 Resources management

Clearly, the overall performance of the described strategy is very much dependent on the parameters t_l and P_l , provided the resources at hand, i.e. the number of available nodes and the number of processors on each one of those. By performance, we intend the speed of execution and simultaneous termination for all batches. We propose here formal discussions about the order of priorities in setting the values of these parameters, as well as suggest a strategy for simultaneous termination in the aim of avoiding wasting computing time resources.

6.3.1 Discussion over algorithm speed

In the presence of a fixed amount of resources (computing nodes denoted by K), parallelization over level and sample layers should always be preferred to parallelization over the space–time grid. Indeed, it is not guaranteed that the latter would perfectly scale when it comes to solving two completely separate samples. That means that in general, one should privilege the highest possible number of sample batches t_l for a corresponding number of processors P_l , while having

$$\sum_{l=0}^L P_l t_l \leq KP.$$

This observation could lead to the conclusion that the optimal choice of P_l is given by $P_l = 1$, for $l = 0, \dots, L$. This is not a reasonable choice as fine samples require consequent memory allocation, and are more likely to execute in an unbearable time on a single processor. Usually, a full node with all processors is required for computing a single sample on the finest level.

On the other hand, ideally a single processor would be able to compute efficiently samples on coarse levels. Provided the following facts experimentally verified in many settings, we came to the conclusion that the most reasonable choice is given by

$$P_l = 2^l,$$

for a processor ratio of $\alpha_{\text{Proc}} = P_{l+1}/P_l = 2$. This can also be motivated by the space–time coarsening between different levels, as the parallel space–time solver we are using is more likely to be effective when the number of time steps is a multiple the number of used processors.

6.3.2 Discussion over simultaneous termination

In the following, we suggest a strategy that aims at having all tasks terminating at the same time. This obviously can never be guaranteed given the nature of the problem, and all the uncertainties in this given setting. Let us denote by α , α_{Cost} , and α_{Proc} respectively the ratios of samples, computational cost and processors between levels. We intend to understand the relationship between these factors, in order to assign subsequent number of nodes, denoted by K_l , in order to achieve simultaneous termination.

On fine levels, the cost acts as the only disadvantage with respect to the total time execution. This is being compensated on the other hand by the higher samples number on coarse levels, and the number of processors choice (see e.g. discussion in Subsection 6.3.1). This leads to consider the evaluation of the following parameter β , defined by:

$$\beta = (\alpha_{\text{Cost}}/\alpha_{\text{Proc}})/\alpha.$$

Indeed, we have the following scenarios with respect to β 's value:

- $\beta > 1$: the cost on fine samples is still dominating and we shall set e.g.

$$K_l = \beta^l \text{ for } l = 0, \dots, L.$$

- $\beta = 1$: the balance of the processors and number of samples exactly compensate for the cost on fine level. We can set an equal number of nodes, e.g.

$$K_l = 1 \text{ for } l = 0, \dots, L.$$

- $\beta < 1$: The cost on coarse levels is dominating and we shall consider using more nodes for those, by setting e.g.

$$K_l = \beta^{l-L} \text{ for } l = 0, \dots, L.$$

6.3.3 HPC oriented job scheduling strategy

We furthermore suggest an adaptation of the parallelization strategy to standard allocation procedures on current High Performance Computing (HPC) systems, in which a request for a small amount of resources is more likely to be granted in a shorter amount of time. We issue consequently different job calls involving different amount of resources for every parallel task consisting of a batches of

samples to solve (c.f. Figure 6.2). Such a procedure allows for being granted resources as they become available, and a shorter absolute time to solution. In this situation, the discussion about simultaneous termination of tasks becomes obsolete, but the real advantage of this scheduling approach is to avoid wasting resources when one do not carefully set the parameters t_l, P_l and K_l .

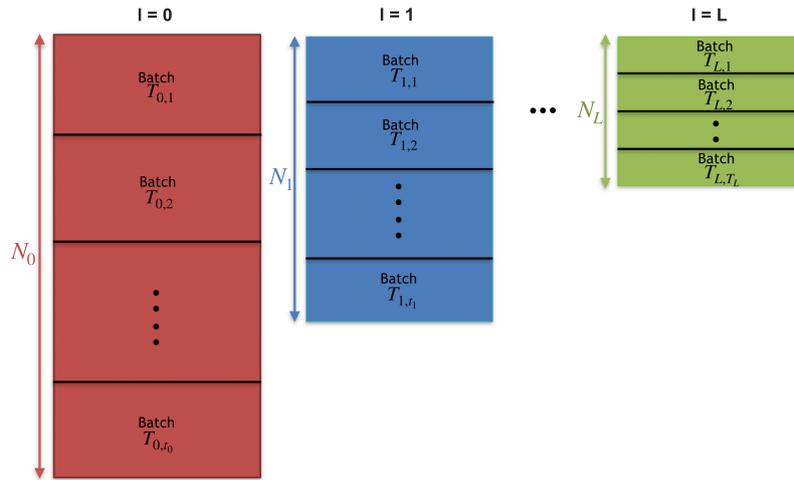


Figure 6.1. The different batches of samples executing concurrently organized through the range of levels.

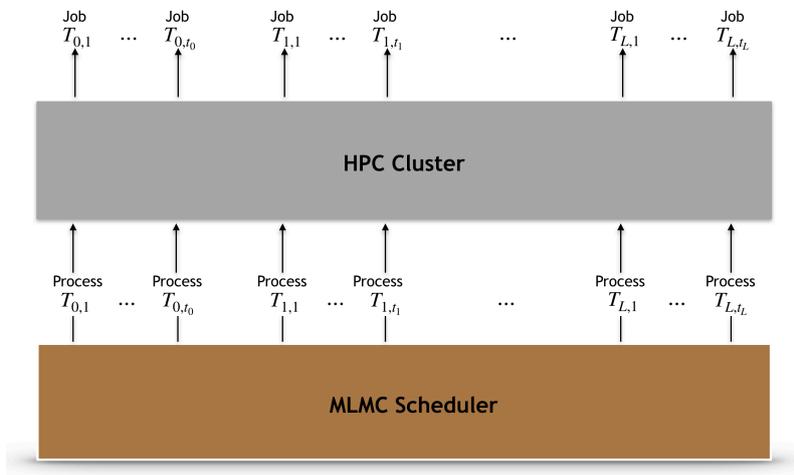


Figure 6.2. MLMC tasks scheduling and job requests step. $T_{l,i}$ denotes the i -th task of level l .

6.4 Newton initial guess

Combining the space–time all–at–once approach with the Newton’s strategy can result in convergence and performance issues. The Newton’s method indeed is highly dependent on the choice of the initial guess, as it is known to converge quadratically only in a vicinity of the solution, eventually diverging when initially too far from it. Considering a big time interval enlarges the gap between the final solution and the provided initial guess, with not only possible consequences in terms of performance, but also in terms of convergence to the final solution.

It is therefore an imperative to consider an efficient Newton initial guess strategy. In general, this can be achieved either by means of an a priori knowledge of the solution behaviour, or by solving lower order models of the considered phenomenon. In the context of UQ, the spectrum of possible workarounds expands naturally, as one can profit from the availability of computed solutions. Other strategies relying on surface response models can also be exploited. We here propose and compare many different strategies for improving the nonlinear solver performance, mainly in terms of iterations to convergence.

6.4.1 Multiple time blocks strategy

The all–at–once methodology described in Subsection 3.1.3 fails to converge when a high number of time steps is considered. A possible immediate fix is to consider a multiple time blocks strategy that consists in decomposing the original time interval into smaller chunks, for which a local number of time steps is used. The equation is then solved sequentially on each time block, provided that each block receives the final state solution from the precedent one, to serve as an initial condition. This initial condition can be further generalized to all the time steps of the current time block, and serve a constant initial guess.

More formally, let us consider a time interval $[0, T]$, uniformly partitioned in m time steps. For the sake of simplicity, let us consider an integer M and assume that there exists $T' \in \mathbb{R}$ and $m' \in \mathbb{N}$ such that $T = MT'$ and $m = Mm'$. We denote the solution on the k –th time block with

$$\mathbf{u}_{[k]} = [\mathbf{u}_{[k],1}, \dots, \mathbf{u}_{[k],m'}]^T \in \mathbb{R}^{nm'}, \quad (6.1)$$

where n is the number of spatial Dofs. Conforming to the original given problem, we consider the initial condition of setting:

$$\mathbf{u}_{[k+1],0} = \mathbf{u}_{[k],m'}, \quad \forall k = 1, \dots, M-1. \quad (6.2)$$

The Newton initial guess for every new time block is further set to be the final time step solution of the previous block, generalized to all the time steps, that is

$$\mathbf{u}_{[k+1]}^{(0)} = [\mathbf{u}_{[k],m'}, \dots, \mathbf{u}_{[k],m'}]^T \in \mathbb{R}^{nm'}. \quad (6.3)$$

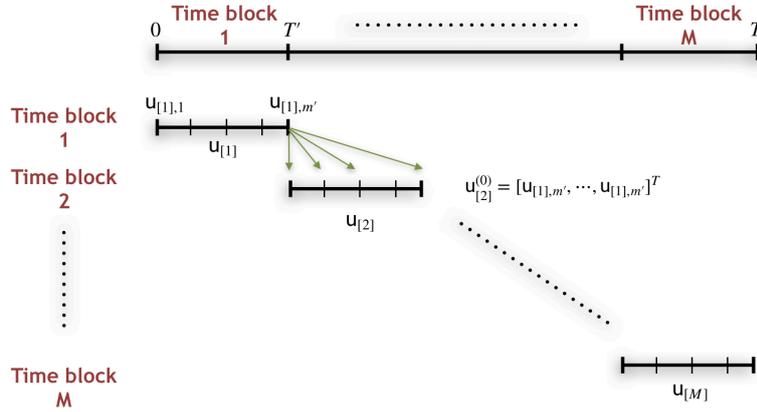


Figure 6.3. Schematization of the multiple time blocks strategy.

Remark. Resorting to a multiple time blocks strategy naturally reduces the parallelization scalability in the time dimension. It is therefore preferable to use the least number of time blocks for a given time interval, in a way that guarantees convergence.

6.4.2 Eikonal initial guess

In the context of cardiac electrophysiology, the Eikonal equation is a low order model that is often used to investigate the heart activation map. It represents a reasonable alternative to the monodomain or the bidomain models, that comes with a significantly lower computational cost. We here rely on a the Eikonal-Diffusion model, see [PTH02a], given by:

$$\frac{1}{\tau} \nabla \cdot (G(\mathbf{x}) \nabla a(\mathbf{x})) + \frac{c_0}{\tau} \sqrt{\nabla a(\mathbf{x}) \cdot G(\mathbf{x}) \nabla a(\mathbf{x})} = 1, \quad (6.4)$$

where $a(\mathbf{x})$ is the activation time, $G(\mathbf{x})$ the conductivity tensor, c_0 and τ respectively local velocity and time parameters. An initial guess state for the monodomain equation can be recovered by means of the following expression:

$$u(\mathbf{x}, t) = \begin{cases} u_{\text{rest}} & \text{if } a(\mathbf{x}) > t, \\ u_{\text{peak}} & \text{otherwise.} \end{cases}$$

Although potentially representing a good compromise for setting initial guess, this method did not show in general the desired effect. The choice on τ and c_0 which should depend on a priori knowledge of the travelling wave speed, is not at all trivial. We have noticed however a relative acceleration of the convergence, as shown in [Ben20].

6.4.3 Samples based initial guess

Starting from first samples computed with a combination of multiple time block strategy and initial guess provided by the eikonal equation, one can take advantage of the computed solutions as they become available. Indeed, it might be reasonable to think that these represent good approximations even for a different random field realization.

(i) Locally: A possibility is to keep using multiple time blocks and replace the Newton initial guess with the solution provided by an available reference sample, while still applying the initial condition 6.3 for every time block. This resorts to solve recursively for every time block, after setting the Newton initial guess:

$$\mathbf{u}_{[k]}^{(0)} = \left[\mathbf{u}_{[k-1],m'}, \mathbf{u}_{[k],2}^{\text{ref}}, \dots, \mathbf{u}_{[k],m'}^{\text{ref}} \right]^T \in \mathbb{R}^{nm'},$$

for $k = 1, \dots, M$ and $\mathbf{u}_{[k]}^{\text{ref}}$ the reference computed sample, with corresponding solution on the k -th time block. We additionally set $\mathbf{u}_{[0],m'}$ to the discretized initial condition of the problem u_0 .

(ii) Globally: The Newton initial guess is provided by the computed reference sample all over the time domain, and the problem solved without considering multiple time blocks. That is, we simply set the initial guess:

$$\mathbf{u}^{(0)} = \mathbf{u}^{\text{ref}} = \left[\mathbf{u}_{[1]}^{\text{ref}}, \dots, \mathbf{u}_{[M]}^{\text{ref}} \right]^T \in \mathbb{R}^{nm}.$$

6.4.4 Performance comparison

In the following, we compare the performances obtained with the different strategies described above, by solving the monodomain equation on six time blocks. We test them on our three main reference geometries: a cube, an idealized ventricle and a realistic heart geometry. The acronyms DM, MTB, LNIG and GNIG

designate respectively the Direct Method, the Multiple Time Block, the Local Newton Initial Guess and the Global Newton Initial Guess.

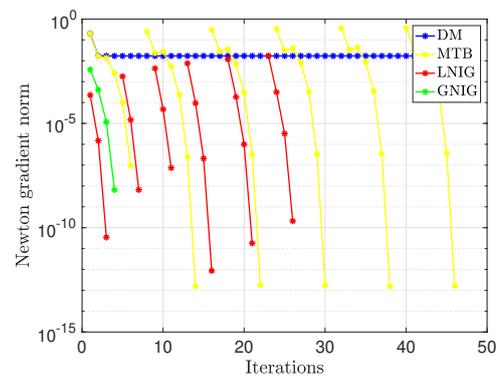


Figure 6.4. Comparison of the different Newton initial guess strategies for a sample on the cube.

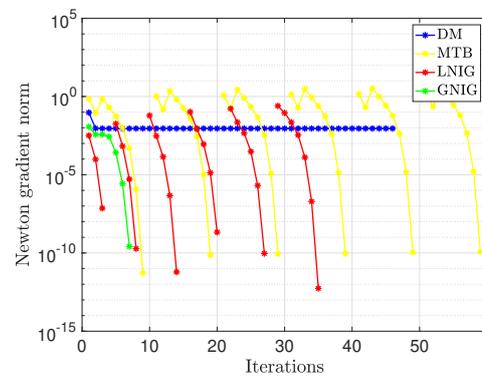


Figure 6.5. Comparison of the different Newton initial guess strategies for a sample on the ventricle.

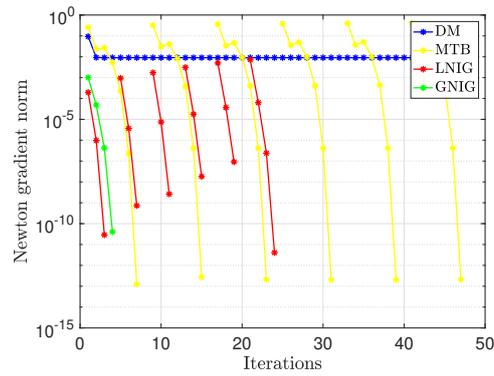


Figure 6.6. Comparison of the different Newton initial guess strategies for a sample on a realistic heart geometry.

We resume the results arising from the Figures 6.4, 6.5 and 6.6 in the following table.

	DM	MTB	LNIG	GNIG
Cube	NC	41	21	4
Ventricle	NC	54	30	7
Heart	NC	42	19	4

Table 6.1. Number of iterations required by each method for the considered geometries. The acronym NC stands for Non Converging.

Chapter 7

Numerical Experiments

7.1 Setup

The numerical experiments have been conducted for three test-case geometries: a cube, an idealized ventricle and a heart geometry (atria excluded) that was acquired through real patient CCT data. We will refer to the latter geometry with the term “realistic heart”. The simulations have been realized using SLOTH, see [QBBK17], a UQ Python library developed at the Institute of Computational Science (ICS) in Lugano. For this work, we extended it to the monodomain equation (and in general to all types of 3D+1 dimensional PDEs) by employing Utopia, see [ZKN⁺16], for the finite element formulation.

Parameters for the monodomain equation. Regarding the model (3.52) and (2.3), we will always rely on the following parameters:

- The values for the ion channel model $I_{\text{ion}}(u)$ in (2.3) are set as $\alpha = 1.4 \cdot 10^{-3} \text{ mV}^{-2} \text{ ms}^{-1}$, $u_{\text{rest}} = 0 \text{ mV}$, $u_{\text{th}} = 28 \text{ mV}$, and $u_{\text{peak}} = 115 \text{ mV}$.
- We choose

$$I_{\text{app}}(\mathbf{x}, t) = \left(u_{\text{rest}} + u_{\text{peak}} \exp\left(-\frac{(\mathbf{x} - \mathbf{x}_0)^2}{\sigma^2}\right) \right) \chi_{[0, t_1]}(t),$$

where $t_1 = \Delta t = 0.005 \text{ ms}$ is the function we rely on for the applied stimulus. Parameters σ and \mathbf{x}_0 represent respectively the power and the location of the stimulus. They are geometry dependent.

Parameters for KL expansion For the numerical experiments, we accompany

the perturbation on the KL expansion with a scaling factor, i.e. $\forall \mathbf{x} \in D$,

$$\mathbf{V}(\mathbf{x}, \omega) = \mathbb{E}[\mathbf{V}](\mathbf{x}) + S \sum_{k=1}^M \sqrt{\lambda_k} \psi_k(\mathbf{x}) Y_k(\omega).$$

where $S \in \mathbb{R}^{d_r \times d_r}$ is given by:

$$S = \begin{bmatrix} s_1 & & \\ & \ddots & \\ & & s_{d_r} \end{bmatrix}$$

where $1 \leq d_r \leq d$ is the dimension of the considered random vector field. The parameters used in the KL expansion differ from an experiment to another and are listed below:

- **Cube:** We consider isotropic diffusion, i.e. $d_r = 1$, and a scaling factor $s_1 = 0.3$. The covariance matrix is induced by the correlation kernel $\text{Cov}_{1,1}[V](\mathbf{x}, \mathbf{y}) = e^{-\frac{d(\mathbf{x}, \mathbf{y})^2}{\sigma_{\text{KL}}}}$, for $\sigma_{\text{KL}} = 0.25$. The low rank Cholesky approximation of the covariance matrix yielded a stochastic dimension $M = 66$. The mean diffusion is set to $\mathbb{E}[V](\mathbf{x}) = 3.325 \cdot 10^{-3} \text{mm}^2 \text{ms}^{-1}$, $\forall \mathbf{x} \in D$.
- **Idealized ventricle:** We consider isotropic diffusion and a scaling factor $s_1 = 0.3$. The correlation kernel is given by $\text{Cov}_{1,1}[V](\mathbf{x}, \mathbf{y}) = e^{-\frac{d(\mathbf{x}, \mathbf{y})^2}{\sigma_{\text{KL}}}}$, for $\sigma_{\text{KL}} = 0.5$. The stochastic dimension is given by $M = 87$. The mean diffusion is set to $\mathbb{E}[V](\mathbf{x}) = 3.325 \cdot 10^{-3} \text{mm}^2 \text{ms}^{-1}$, $\forall \mathbf{x} \in D$.
- **Heart geometry:** We consider anisotropic diffusion, i.e. $d_r = d$ with a block-diagonal covariance matrix given by $\text{Cov}_{i,j}[V](\mathbf{x}, \mathbf{y}) = \delta_{ij} e^{-\frac{d(\mathbf{x}, \mathbf{y})^2}{\sigma_{\text{KL}}}}$, for $1 \leq i, j \leq d_r$ where δ is the Kronecker delta. We set $\sigma_{\text{KL}} = 0.16$. The scaling factor is set to $s_i = 0.3$ for $i = 1, \dots, d_r$. The recovered stochastic dimension for the described setting is given by $M = 135$. An additional a posteriori truncation (see Subsection 4.1.6 and in particular Equation (4.58)) is performed with $\epsilon_{\text{trunc}} = 0.05$, leading to $M_{\text{trunc}} = 48$. Furthermore, the perpendicular diffusion in (4.65) is set to $a = 1.625 \cdot 10^{-3} \text{mm}^2 \text{ms}^{-1}$.

We furthermore use a KL truncation error of $\epsilon = 10^{-2}$.

Reference solution and error metrics. In all the convergence and work comparison graphs that follow, the referenced H^q norm ($L^2 = H^0$ and H^1 for respectively $q = 0$ and $q = 1$) of the mean square error is given by

$$e_l = \mathbb{E} \left[\|u_l(\cdot) - \mathbb{E}[u_{\text{ref}}(\cdot)]\|_{L^2((0,T);H^q(D))}^2 \right] \quad (7.1)$$

in case of a space–time quantity of interest. In case of a quantity of interest which is associated with a specific point \mathbf{x}_0 in space, we consider instead the error

$$e_l = \mathbb{E}[\|u_l(\mathbf{x}_0, \cdot) - \mathbb{E}[u_{\text{ref}}(\mathbf{x}_0, \cdot)]\|_{H^q(0,T)}^2], \quad (7.2)$$

The reference solution u_{ref} is computed by using $N = 10'000$ samples drawn from the Halton sequence. Regarding the Monte Carlo quadrature, the expectation for both, the single-level and multilevel runs, are approximated by averaging over 10 simulations at each level of precision for the nested case study, and 5 simulations for the non–nested example.

Quadrature methods. In our experiments, we will consider the Monte Carlo (MC) and quasi-Monte Carlo (QMC) quadrature method and their multilevel pendants MLMC and MLQMC. Let us recall that the H^q –error ($q = 0, 1$) is of order $2^{(q-2)l}$ when using linear finite elements of mesh size $h_l = 2^{-l}$. Therefore, and as suggested by Propositions 4.3 and 4.4, the number of samples to be executed by these methods on a level l to get the same order of error $2^{(q-2)l}$ is respectively given by

$$N_{\text{MC},l} = \mathcal{O}(2^{(4-2q)l}) \quad \text{and} \quad N_{\text{QMC},l} = \mathcal{O}(2^{(2-q)l}). \quad (7.3)$$

Regarding the multilevel strategies, the number of samples to be executed on each level $l = 0, \dots, L$ to get the same order of error $2^{(q-2)l}$ is given by

$$N_{\text{MLMC},l} \sim \mathcal{O}(2^{(4-2q)(L-l)}) \quad \text{and} \quad N_{\text{MLQMC},l} \sim \mathcal{O}(2^{(2-q)(L-l)}), \quad (7.4)$$

as suggested by 4.5 and 4.6

7.2 Isotropic random diffusion for simple geometries and nested meshes

In these first experiments, we consider isotropic diffusion for the sake of simplicity. The mean diffusion field $G_0(\mathbf{x})$ is set to the scalar value $G_0(\mathbf{x}) = 3.325 \cdot 10^{-3} \text{mm}^2 \text{ms}^{-1}$. The experiments are conducted on the cube and the idealized ventricle geometry.

7.2.1 Cube

We use a hierarchy of $L + 1 = 6$ nested levels, for which every coarse level $l = 0, 1, \dots, L - 1$ is obtained from the prior fine level $l + 1 = 1, 2, \dots, L$ by uniformly

coarsening in space and time, starting from the finest level $l = L$. The number of space-time degrees of freedom (Dof's), the space and time discretization steps of all the different levels are reported in Table 7.1.

l	0	1	2	3	4	5
Dof's	16	256	4'096	65'536	1'048'576	16'777'216
h	0.5	0.25	0.125	0.0625	0.03125	0.015625
Δt	0.16	0.08	0.04	0.02	0.01	0.005

Table 7.1. Details about the considered mesh hierarchy for the cube geometry.

(i) Controlled convergence of the general error: We intend to estimate and verify the convergence rate for the quadrature methods under consideration. The number of samples on each level is determined by the sampling strategy for controlling the error (c.f. Section 4.6) by using the sample numbers (7.3) and (7.4). We report in figure 7.1 the controlled convergence in L^2 and H^1 norms.

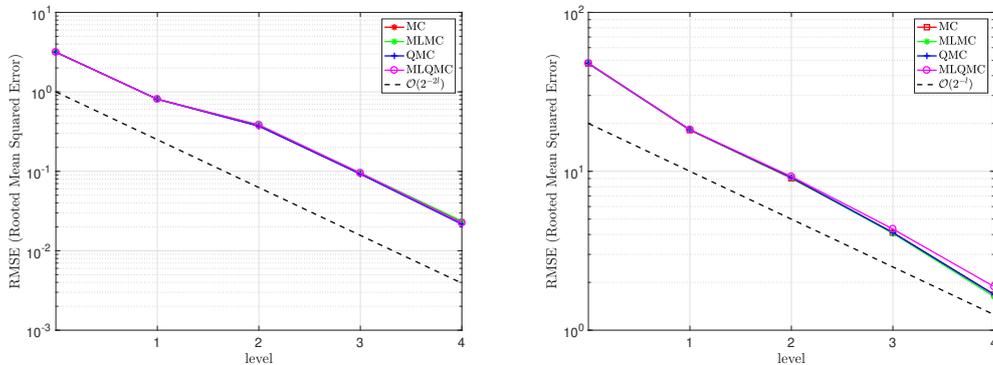


Figure 7.1. Controlled convergence graphs for the cube in L^2 (left) and H^1 (right) norms.

The plots show that we recover the expected convergence rate of the general error for all the quadrature methods tested. This should not be interpreted as all of these methods are equivalent, but that these instead get to the same precision with a different balancing of samples on every level. This gets demonstrated in a better way once converting these plots into work comparison plots.

(ii) Asymptotical work behaviour: Equivalently to the controlled convergence concept with the previously introduced sampling strategy, we would like now to study the work in the context of the controlled error. Let us recall that in order to achieve an error of order $2^{(q-2)l}$ in the H^q norm ($q = 0, 1$) using linear FE, the work produced by all the considered quadrature methods is given by:

$$W_{MC,L}^q = 2^{(\gamma d + 4 - 2q)L}, \quad W_{QMC,L}^q = 2^{(\gamma d + 2 - q)L}, \quad (7.5)$$

$$W_{MLMC,L}^q = \begin{cases} (L+1)2^{(4-2q)L} & \text{if } \gamma d = 4 - 2q, \\ \frac{C(q)2^{\gamma d L} - 2^{(4-2q)L}}{2^{\gamma d - (4-2q)} - 1} & \text{if } \gamma d \neq 4 - 2q, \end{cases} \quad (7.6)$$

and

$$W_{MLQMC,L}^q = \begin{cases} (L+1)2^{(2-q)L} & \text{if } \gamma d = 2 - q, \\ \frac{D(q)2^{\gamma d L} - 2^{(2-q)L}}{2^{\gamma d - (2q-1)} - 1} & \text{if } \gamma d \neq 2 - q, \end{cases} \quad (7.7)$$

with γ representing the solver complexity, $C(q) = 2^{\gamma d - (4-2q)}$ and $D(q) = 2^{\gamma d - (2-q)}$.

The complexity parameter γ is therefore of a major importance in the asymptotical work behaviour of the quadrature methods at hand. In our case, given the Newton initial guess strategy, the solver preconditioning and the difference in parallel resources used from a level to another do not allow to give this parameter a concise value over all levels (c.f. Chapter 6). We suggest however to evaluate the work in terms of total execution time, in which the cost for solving a sample at a level l is given by the time to solution (averaged over a hundred samples). The plot of work comparison between the different methods is reported in Figure 7.2, where the error is measured in the L^2 -norm.

7.2.2 Idealized ventricle

In this second test case, we rely on a mesh hierarchy of $L+1 = 3$ levels. The main reason for the limitation of levels number for this geometry is essentially due to the nestedness condition. Indeed, for this geometry as opposed to the simple cube one, we proceed in an inverted way, i.e. refining an initial given mesh. This procedure becomes very quickly demanding at the memory level. In general, this limitation can be very often encountered when dealing with nested meshes for a realistic geometry. This is the main motivation for relying on non nested meshes for the last test-case (c.f. Section 7.3).

The number of space-time degrees of freedom (Dof's), the space and time discretization steps of the three different levels are reported in Table 7.2. The meshes are demonstrated in Figure 7.3.

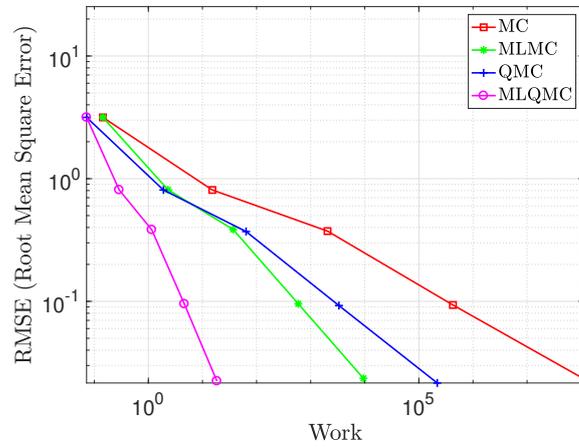


Figure 7.2. Work comparison for MC/MLMC/QMC/MLQMC. Work is computed on the basis of execution time on a single thread for all levels.

l	0	1	2
Dof's	154'546	2'120'420	31'184'747
h	0.1	0.05	0.025
Δt	0.02	0.01	0.005

Table 7.2. Details about the mesh hierarchy for the idealized ventricle geometry.

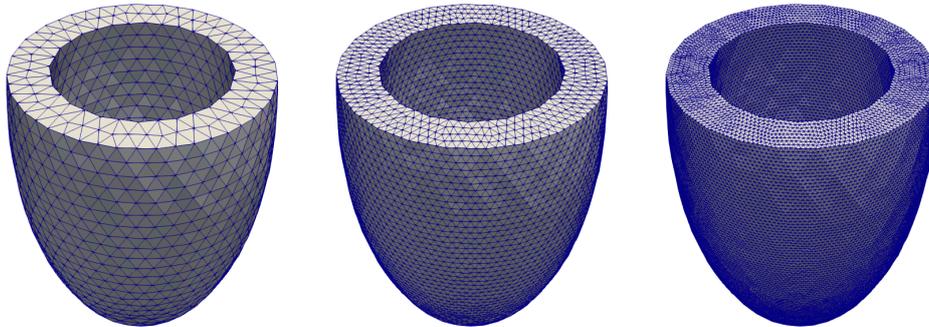


Figure 7.3. Nested mesh hierarchy for the idealized ventricle.

(i) Controlled convergence of the general error:

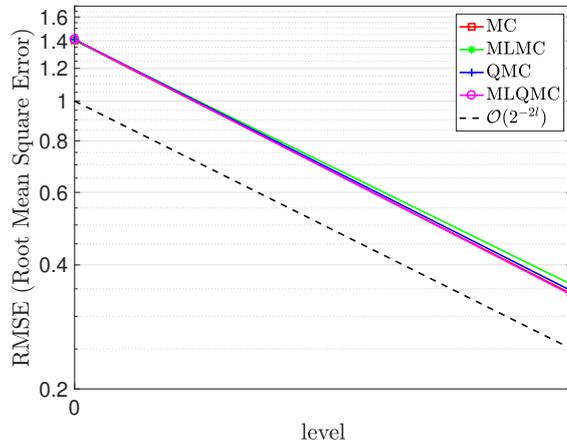


Figure 7.4. Controlled convergence graphs for the idealized ventricle in L^2 -norm.

(ii) Asymptotical work behaviour:

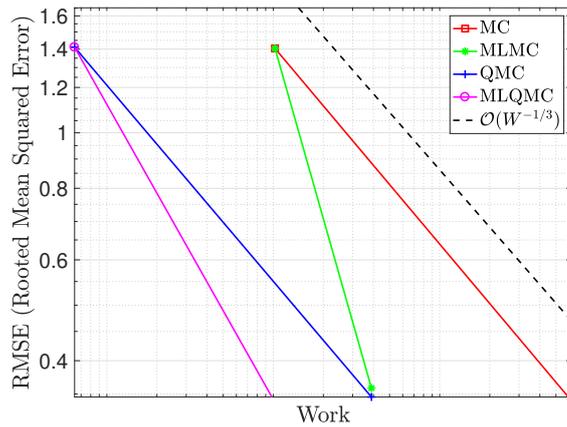


Figure 7.5. Work comparison for MC/MLMC/QMC/MLQMC. Work is computed on the basis of execution time.

7.3 Anisotropic random diffusion for a complex geometry on non-nested meshes

The last test case concerns a realistic heart geometry with data acquired from clinical practice. As this is meant to be the synthesis of this work, we also account for anisotropic diffusion defined in (4.65). The background conductivity coefficient a in (4.65) is set to the value $a = 1.625 \cdot 10^{-3} \text{mm}^2 \text{ms}^{-1}$. The associated expected fiber field $\mathbb{E}[\mathbf{V}](\mathbf{x})$ is shown in Figure 7.6. It is obtained from a mathematical reconstruction using transmural coordinates [PDR⁺06a]. The transmural coordinates are derived by initially solving a diffusion problem with adapted boundary conditions at the contour of the left and right ventricles [BBPT12].

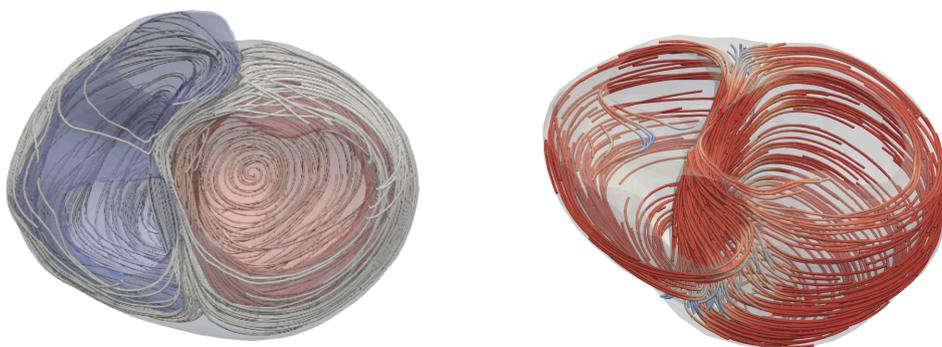


Figure 7.6. Initial state for fibers $\mathbb{E}[\mathbf{V}](\mathbf{x})$.

We also relax the nestedness condition by considering non-nested mesh levels. As we have previously argued, the nestedness condition becomes very quickly a burden in considering a high number of levels. We rely on 6 mesh levels in this example. They are shown in Figure 7.7. The details on the space time Dof's and discretization steps are reported in Table 7.3.

l	0	1	2	3	4	5
Dof's	18'480	113'312	583'104	1'740'800	8'777'728	34'894'848
h	0.16	0.08	0.04	0.03	0.02	0.01
Δt	0.16	0.08	0.04	0.02	0.01	0.005

Table 7.3. Details about the considered mesh hierarchy for the realistic heart geometry.

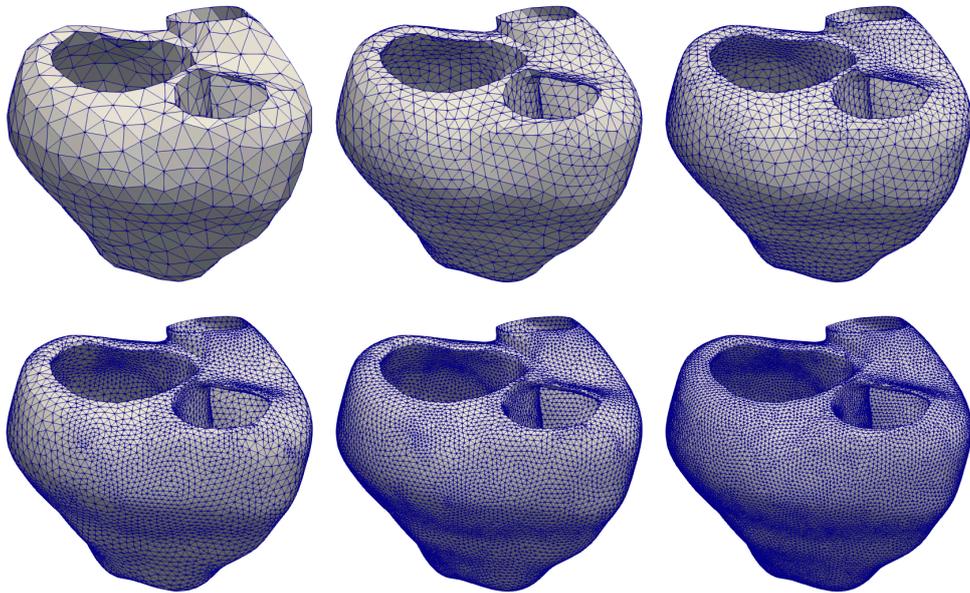


Figure 7.7. Non-nested mesh hierarchy for the realistic heart.

As we do not have nested finite element spaces, we rely here on the multilevel estimator (4.83). Moreover, we evaluate the convergence for the two quantities of interest, namely the action potential and the activation times for given location in the domain.

7.3.1 Action potential

We evaluate the evolution of the action potential in several locations of the heart domain. The first example considers a set of points that are placed along the wall separating the left and right ventricles. These points are shown in Figure 7.8.

As one can see, these points have been selected so they trace the behaviour of locations at different distance from the stimulus center, starting from very close (the very bottom point) to relatively far (the very top point). We report in Figure 7.9 the action potential obtained by the MC quadrature method for different discretization levels. The order goes from top left for the very bottom point defined in Figure 7.8, to below right for the upper point.

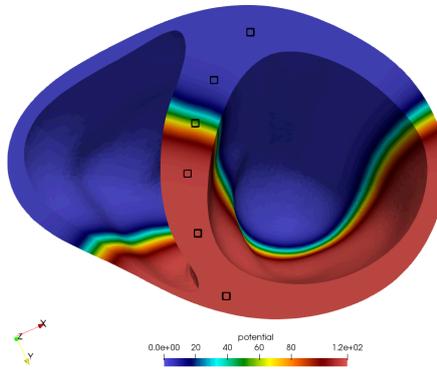


Figure 7.8. Locations selected along the wall separating the left and right ventricles.

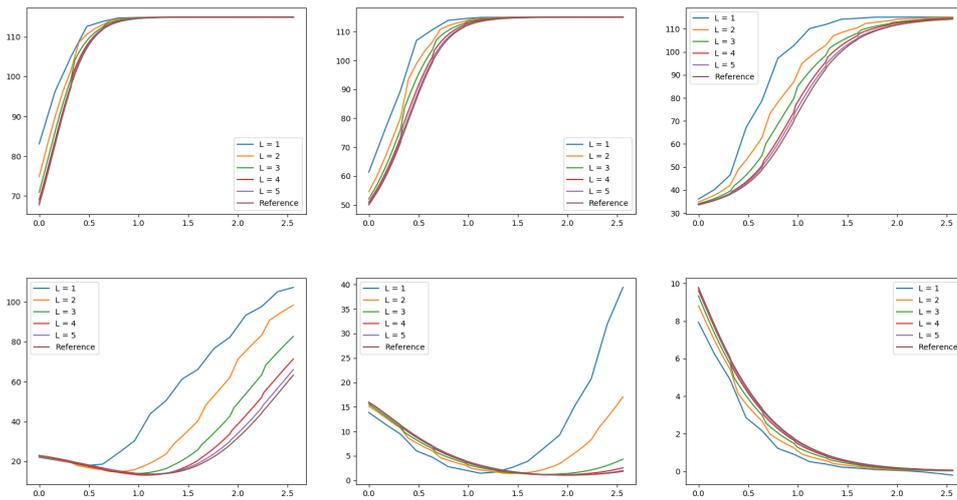


Figure 7.9. Action potential behaviour given different mesh level discretizations for the points specified in Figure 7.8.

In Figure 7.10, we report the convergence graphs of the (pointwise) error (7.2) for the action potential at the locations introduced in Figure 7.8 and $q = 1$. Notice that the graphs report the root mean square errors.

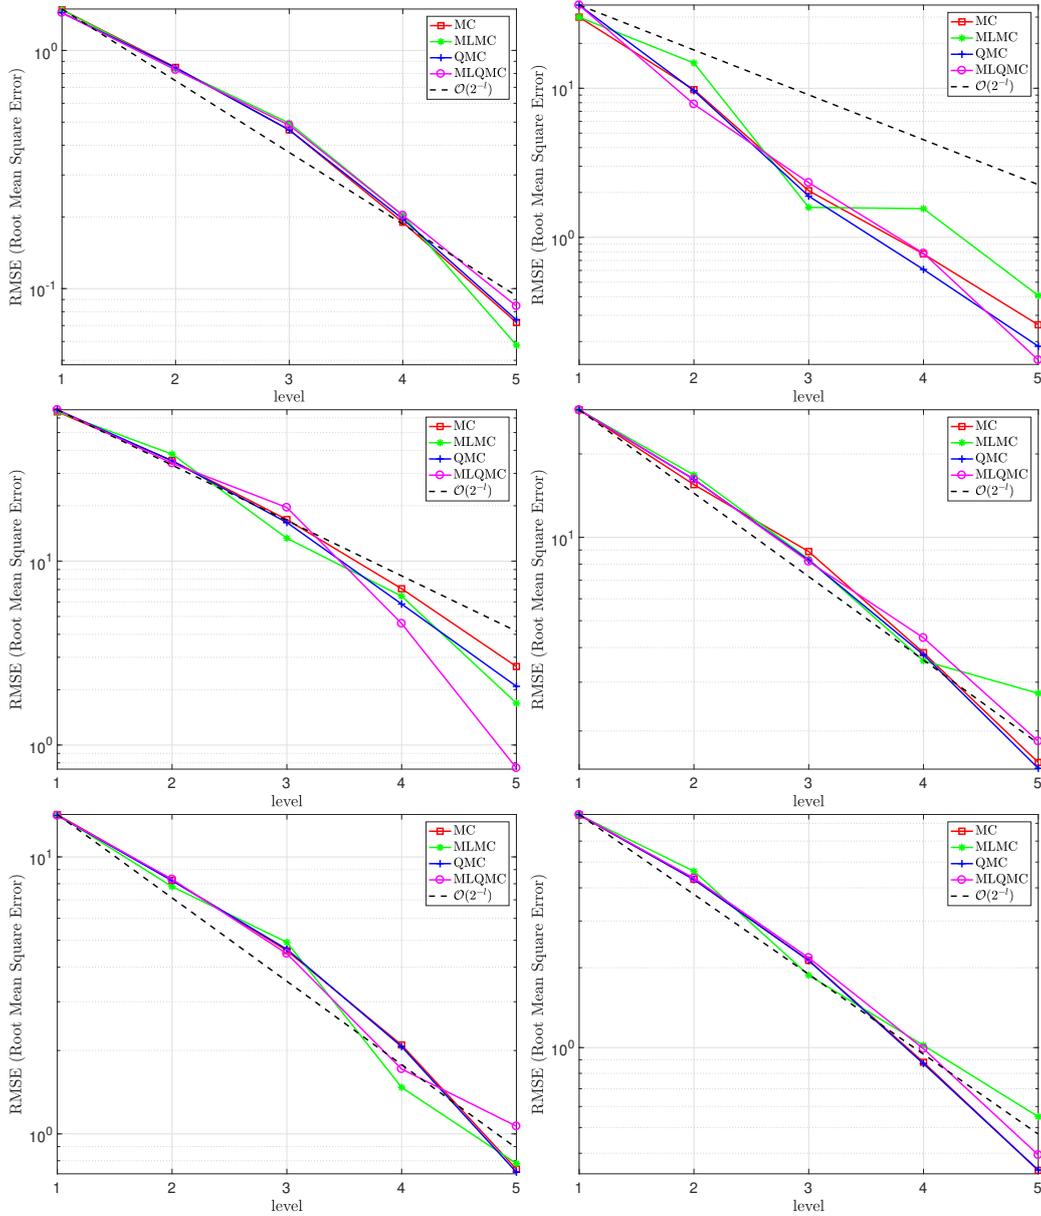


Figure 7.10. Convergence in the H^1 -norm of the action potential at the locations specified in Figure 7.8.

The second test is concerned with points located at the circumference of a horizontal cut of the heart surface. These are shown in Figure 7.11.

Since the behaviour of the action potential on these points follows a similar pattern than that of the previously shown ones (c.f. Figure 7.9), we only show the convergence graphs of the action potential for these points.

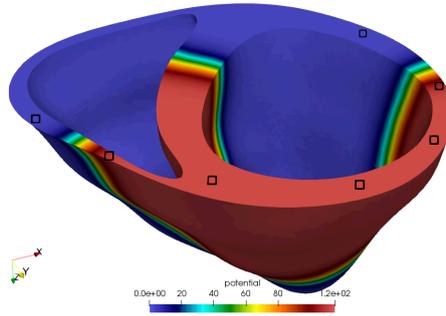
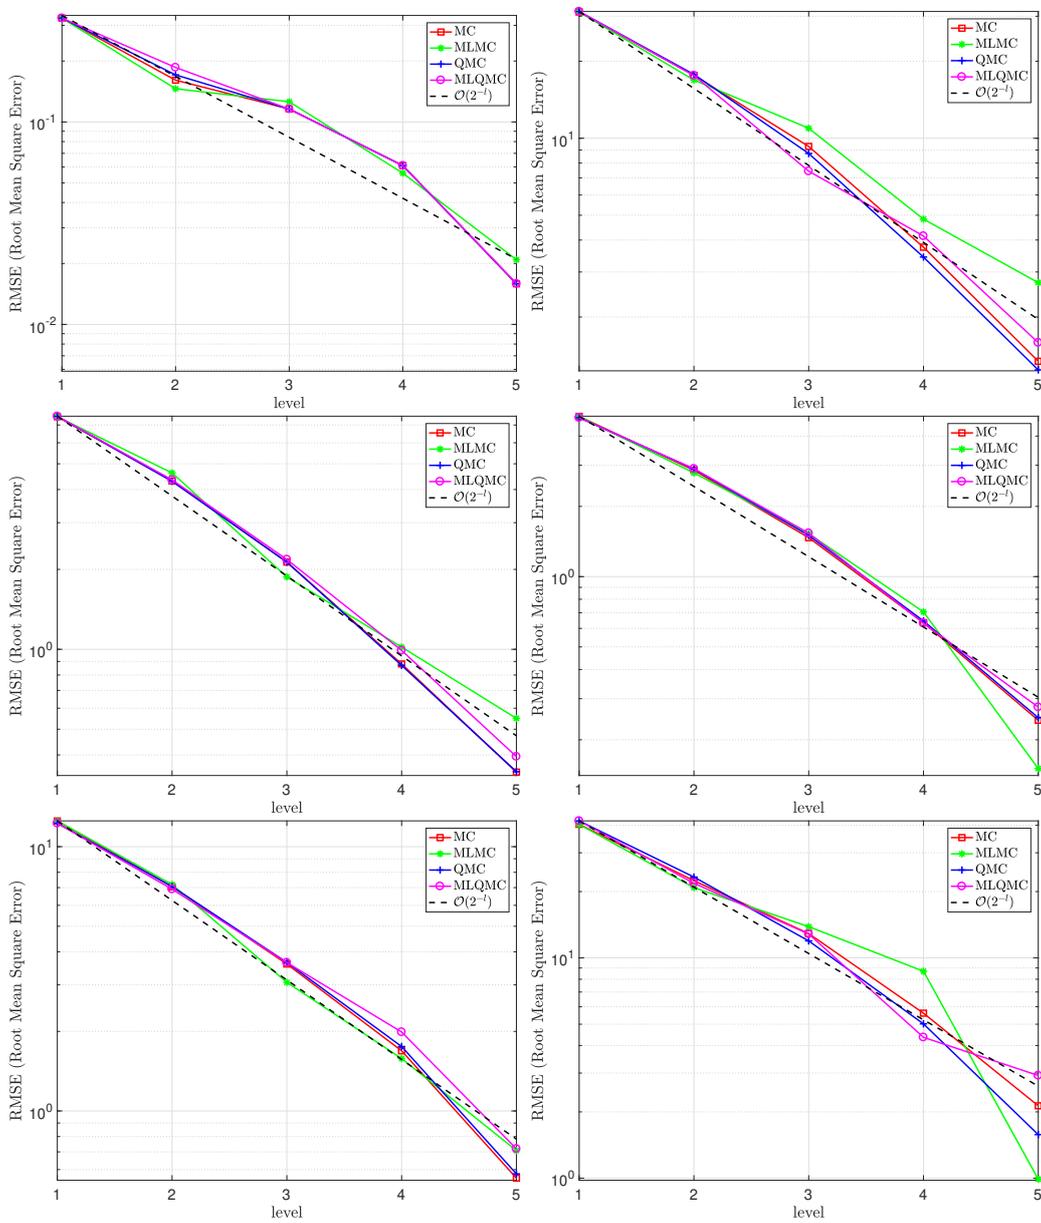


Figure 7.11. Locations selected at the circumference of a horizontal cut of the heart surface.



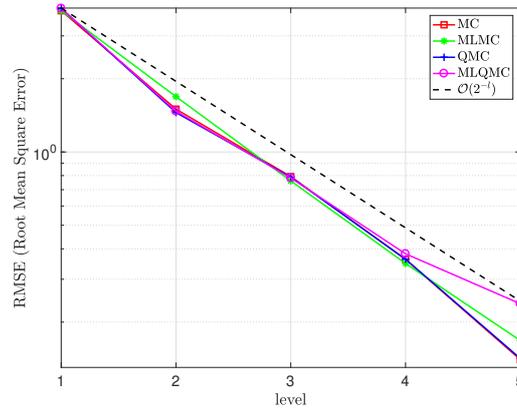


Figure 7.12. Convergence of the action potential at the locations specified in Figure 7.11.

7.3.2 Activation time

We start by selecting points at equivalent geodesic distance from the stimulus location. These points are shown in Figure 7.13. The geodesic distance is calculated by solving an eikonal problem with a zero initial condition on the originating point [PQP19], i.e. the stimulus in our case.

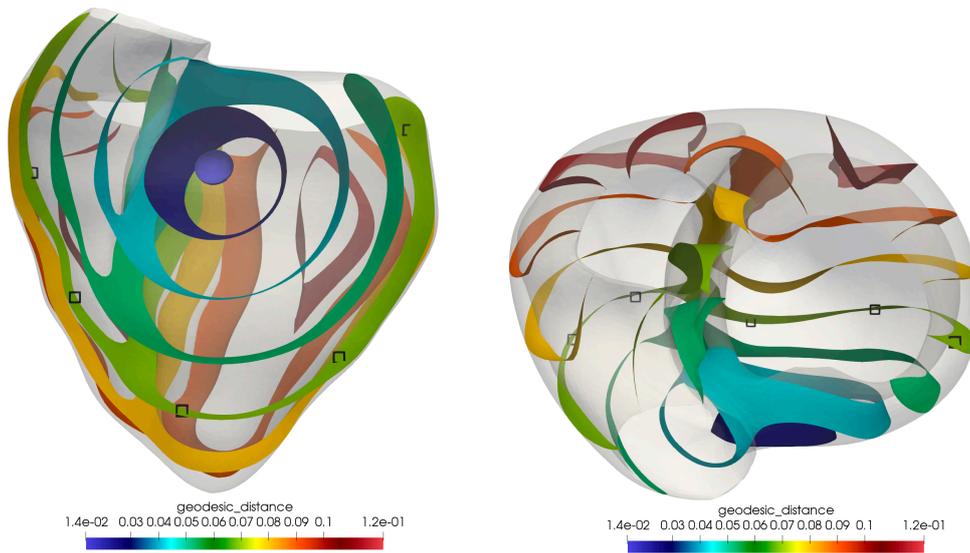


Figure 7.13. Locations selected at equivalent geodesic distance from the stimulus.

The graphs on the convergence of the activation times for these locations are reported in Figure 7.14.

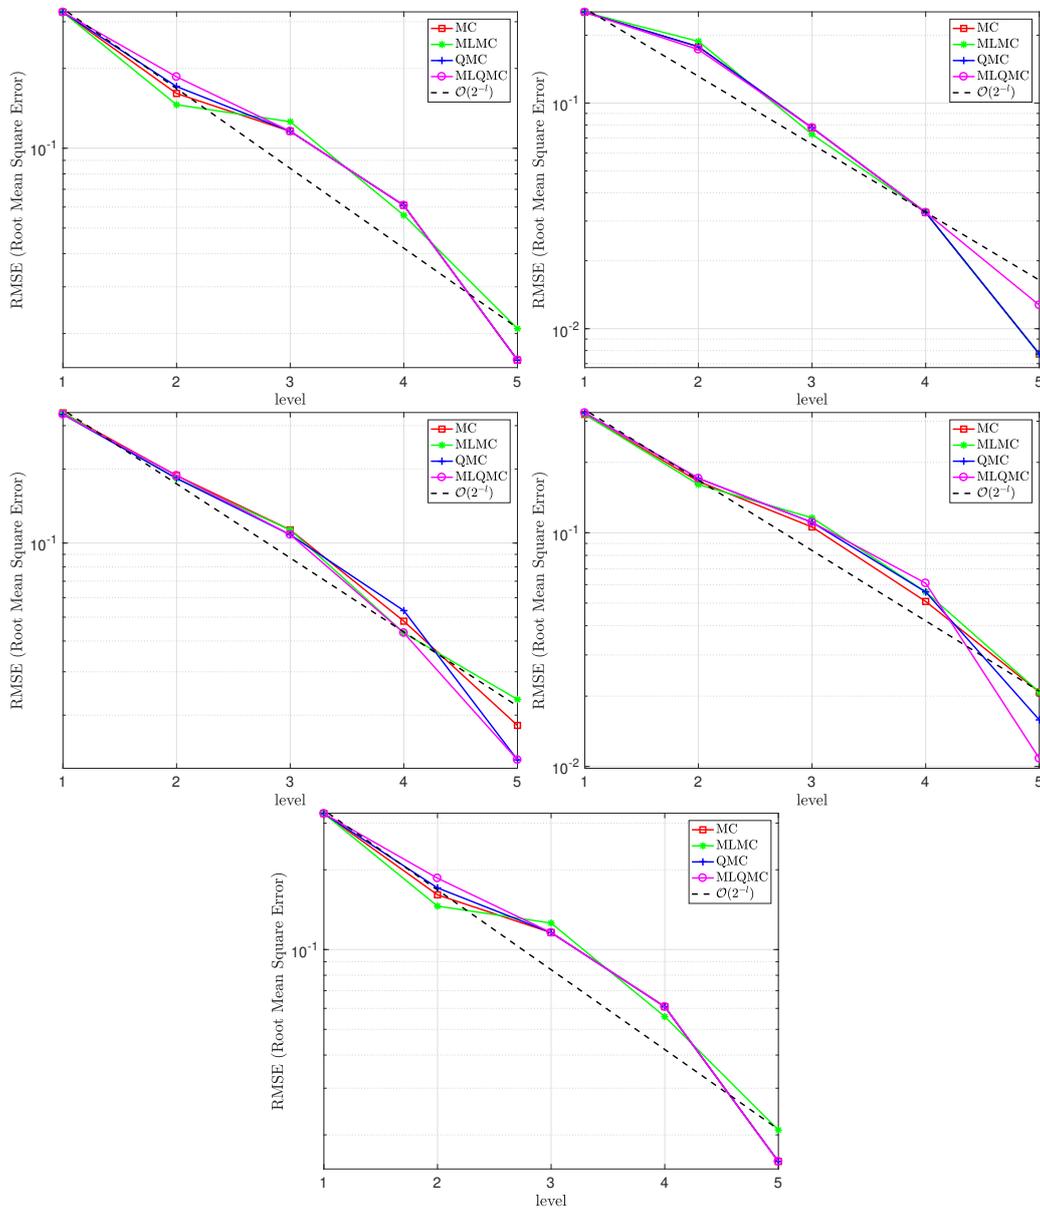


Figure 7.14. Convergence of the activation time at the locations specified in Figure 7.13.

We now select locations at the circumference of the left ventricle. These are shown in Figure 7.15. This electrical signal, when propagated to the chest, is

exactly what is perceived clinically (on a electrocardiogram monitor). Mathematically, it is possible to map the surface potential to the chest by solving an additional diffusion problem, as described in Subsection 2.3.2 (see further [FZ10, BCF⁺10]). Notice that these are not placed exactly at the surface of heart. As we have several discretization levels, we need to ensure that these points are well-defined on each one of them.

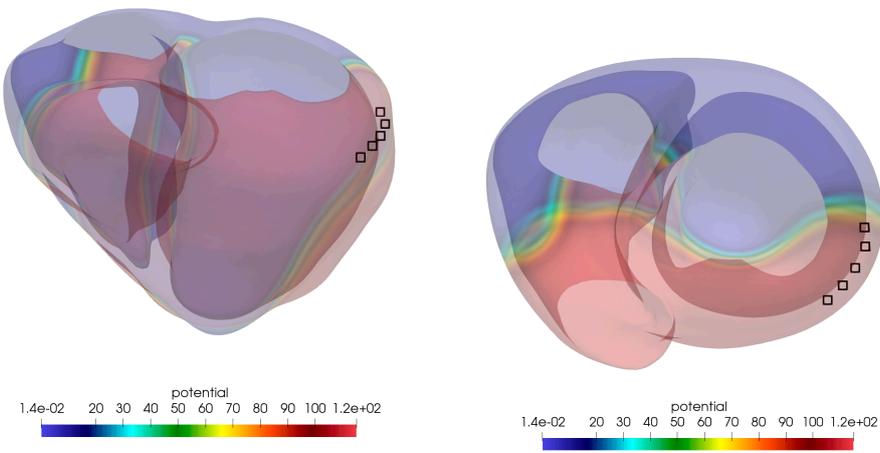
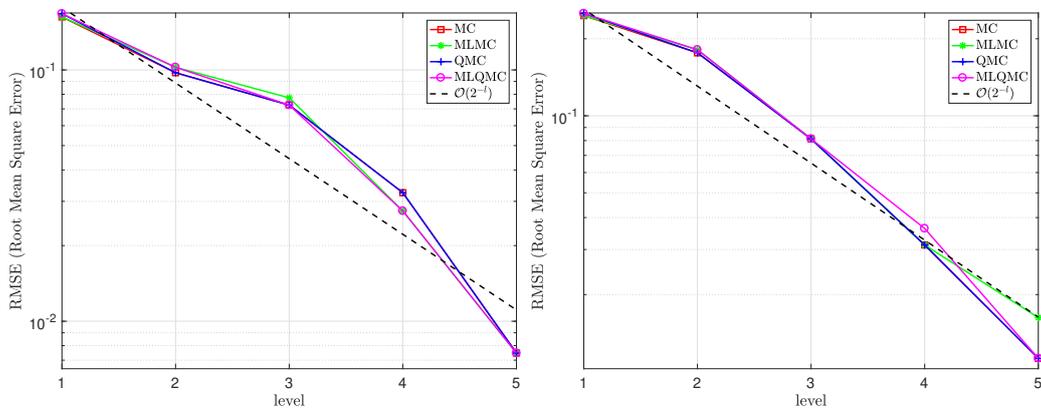


Figure 7.15. Locations selected at the periphery of the heart surface.



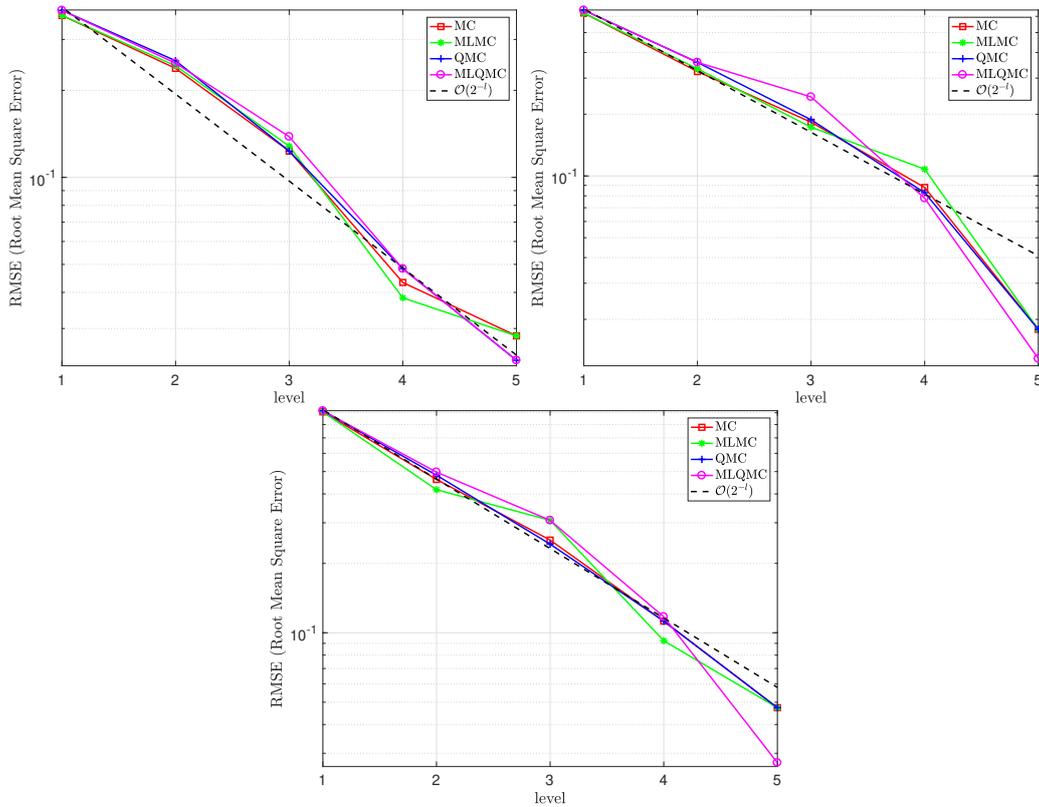


Figure 7.16. Convergence of the activation time at the locations specified in Figure 7.15.

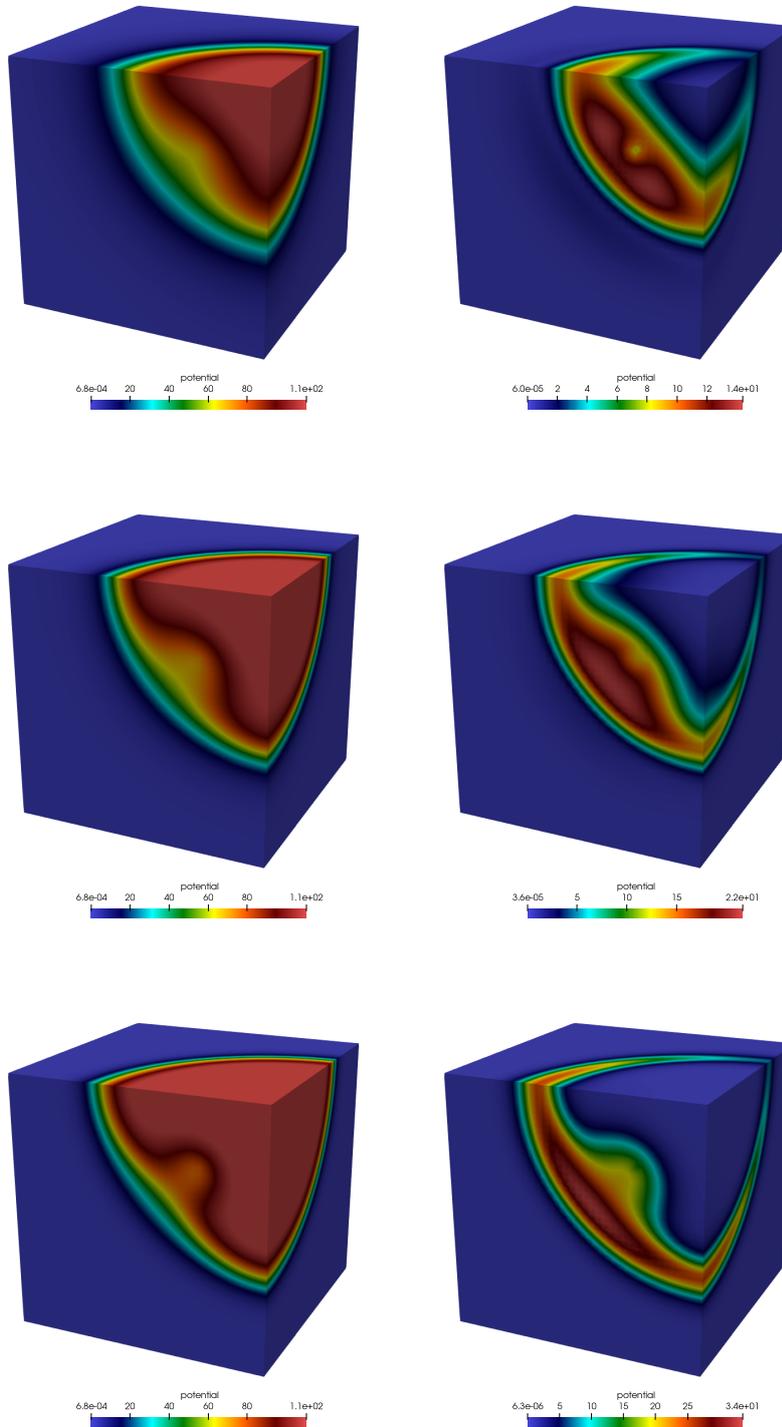
7.4 Scar uncertainty

As a final numerical experiment, we consider a setting where a scar fibrosis is included in the computational domain and define additional uncertainties consisting of the scar width and the conductivity, c.f. Section 5.2. This experiment is conducted on the cube geometry, by still considering isotropic random diffusion defined by the same set of parameters defined in Section 7.1. Additional parameters with respect to the scar are listed here below:

- Scar center $\mathbf{x}_0 = [0.25, -0.5, 0.25]^T$.
- Scar width bounds $[\mu_{\min}, \mu_{\max}] = [0.05, 0.4]$.
- Scar conductivity bounds $[\sigma_{\min}, \sigma_{\max}] = [10^{-8}, 2.5 \cdot 10^{-4}]$.

Recall that the cube is given by $D = [-0.5, 0.5]^3$. In Figure 7.17, we show the evolution of the mean and the variance of the stochastic solution to the new

setting considering the scar, in ordering to underline its effect on the electrical propagation.



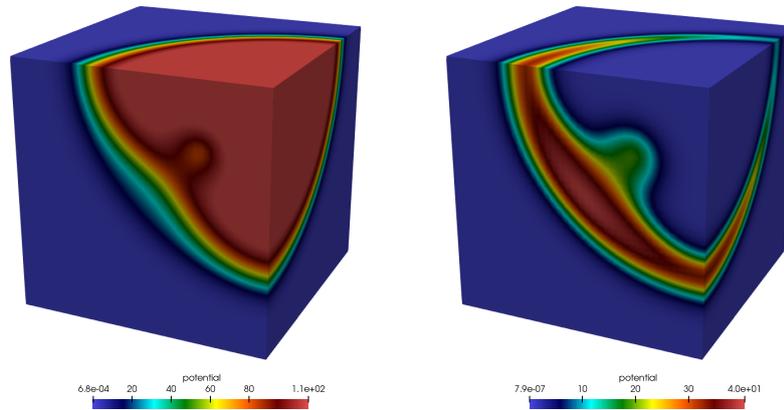


Figure 7.17. Mean (left) and variance (right) over time after the inclusion of a scar.

In Figures 7.18 and 7.19, we show the usual controlled convergence and asymptotical work graphs.

(i) **Controlled convergence of the general error:**

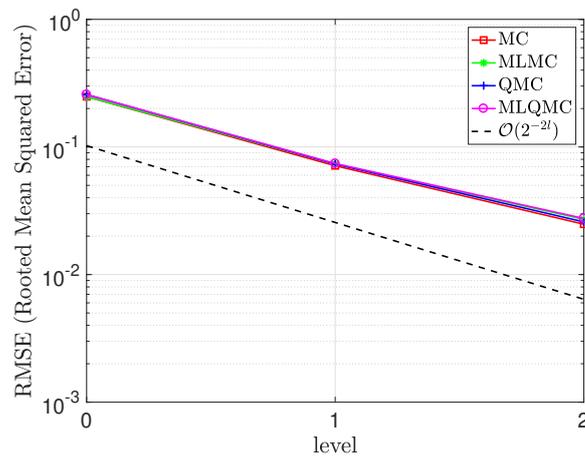


Figure 7.18. Controlled convergence graphs for the cube after inclusion of a scar in the L^2 -norm.

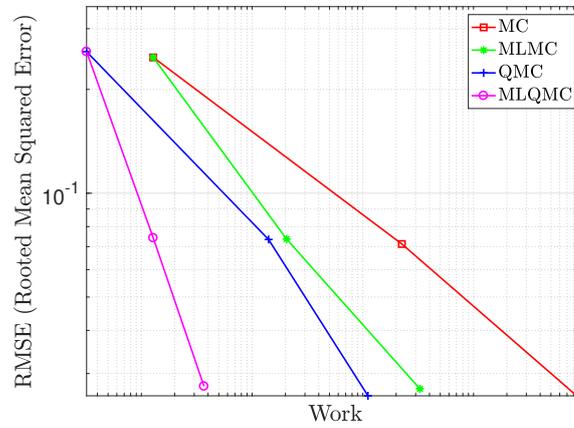
(ii) Asymptotical work behaviour:

Figure 7.19. Work comparison for MC/MLMC/QMC/MLQMC. Work is computed on the basis of execution time.

Chapter 8

Conclusion

Many sources of uncertainty can be encountered whenever considering parameter dependent PDEs. Those uncertainties are related to the impossibility, with the means at hand, of sharply estimating given parameters.

In electrophysiology, this may materialize in many aspects. As an example, obtaining a precise geometry of the fiber network within the myocardium is still out of reach even with state-of-the-art imaging techniques. Consequently, models describing the electrical propagation inside the cardiac muscle are most likely to be impacted by this uncertainty, in a way that severely complicates patient-specific virtual therapy.

With respect to the aforementioned challenges, in this thesis:

- We modeled the fibers' uncertainties with respect to the orientations and conductivities, by means of a KL expansion expressed as spatially correlated random perturbations on a given initial state.
- We employed the monodomain equation with the FHN ion channel model and evaluated the effect of random diffusion fields on the statistics of the resulting quantities-of-interest, i.e. the transmembrane potential, the activation time and the action potential for given sites.
- We applied sampling methods such as standard MC and QMC, as well as their multilevel pendants, i.e. MLMC and MLQMC (See Chapter 4). For a non-constraining usage of multilevel methods on complex geometries, we also employed a multilevel inverse construction (see Subsection 4.4.2).

- To render these methods easily applicable, we constructed an efficient solution framework that fully takes advantage of all possible concurrencies.
- We conducted extensive numerical studies on simple and complex geometries. We evaluated the convergence rates of the considered methods and compared their performances.
- We additionally included a realistic example of a scar fibrosis defined by an uncertainty on its width and conductivities.

We highlight the following results:

1. We have built an efficient computational framework for addressing large scale UQ problems, that is fully parallel in space, time and stochastics. In particular, we took advantage of three layers of parallelization: levels, batches of samples, and spatio-temporal grid. More precisely, we employed a space-time all-at-once approach and solved the arising systems with parallel space-time PGMRES.
2. We tackled the convergence problems with respect to Newton's method by means of multiple time blocks strategy and samples based Newton initial guess strategy, i.e. MTB, LNIG and GNIG (see Chapter 6). By doing so, we did not only improve the robustness of the iteration scheme but also drastically accelerated the convergence with a reduction of Newton iterations by an order of magnitude.
3. We evaluated through extensive numerical studies the convergence rates for the sampling methods at hand. In particular, the optimal convergence rate with respect to MLQMC is achieved, proving therefore that regularity requirements hold for this considered problem, and most likely the respective classes of problems. This validates the use of multilevel and deterministic UQ techniques for non-linear time dependent parabolic equations, with uncertainties in the diffusion coefficient.
4. The resulting MLQMC framework is capable of handling multilevel hierarchies of non-nested meshes, allowing for realistic geometries and reducing dramatically the cost for quantities-of-interest estimation, in comparison with standard MC.

Appendix A

Short review of normed vector spaces

In the upcoming definitions, X is always considered to be a nonempty set.

Definition A.1 (Metric space). A metric on X is a function $d : X \times X \rightarrow \mathbb{R}$ such that the following assumptions hold $\forall x, y, z \in X$:

(i) $d(x, y) \geq 0$ and $d(x, y) = 0$ if and only if $x = y$, (nonnegativity)

(ii) $d(x, y) = d(y, x)$, (symmetry)

(iii) $d(x, y) \leq d(x, z) + d(z, y)$. (triangular inequality)

(X, d) is then called a metric space and $d(x, y)$ the distance between x and y .

Definition A.2 (Convergence). A sequence of points $\{x_n\}_{n=1}^{\infty}$ in a metric space X is said to converge to an $x \in X$ if $\forall \epsilon > 0$, there exists an $N \in \mathbb{N}$ such that

$$d(x_n, x) < \epsilon, \quad \forall n \geq N.$$

Definition A.3 (Cauchy sequence). A sequence of points $\{x_n\}_{n=1}^{\infty}$ in a metric space X is said to be a Cauchy sequence if $\forall \epsilon > 0$, there exists an $N \in \mathbb{N}$ such that

$$d(x_n, x_m) < \epsilon, \quad \forall n, m \geq N.$$

The triangular inequality implies that a convergent sequence is necessarily a Cauchy sequence. The opposite does not necessarily hold, and characterizes so-called complete spaces.

Definition A.4 (Complete metric space). A metric space X is said to be complete if every Cauchy sequence converges to an element $x \in X$.

Definition A.5 (Dense subset). A subset S of a metric space (X, d) is said to be dense in X if for every $x \in X$ and every $\epsilon > 0$, we have that

$$\mathcal{B}(x, \epsilon) \cap S \neq \emptyset,$$

where $\mathcal{B}(x, \epsilon) = \{y \in X \text{ such that } d(x, y) < \epsilon\}$, that is the open ball centered in x of radius ϵ .

Remark. An example of a dense subset is provided by \mathbb{Q} with respect to \mathbb{R} .

In the following, we consider the notion of vector space to be known to the reader.

Definition A.6 (Norm). Let X be a vector space over \mathbb{R} . A norm on X is a function denoted as $\|\cdot\| : X \rightarrow \mathbb{R}$ such that the following holds $\forall x, y \in X$ and $c \in \mathbb{R}$:

- (i) $\|x\| \geq 0$ and $\|x\| = 0$ if and only if $x = 0$, (nonnegativity)
- (ii) $\|cx\| = |c| \|x\|$, (homogeneity)
- (iii) $\|x + y\| \leq \|x\| + \|y\|$. (triangular inequality)

The vector space X , together with its norm, i.e. $(X, \|\cdot\|)$, is then called a normed vector space.

Remark. The norm provides the notion of individual vector's length, while a metric describes the distance between given points in a set. The latter can therefore be defined for arbitrary sets, while norms are defined only for vector spaces.

Definition A.7 (Seminorm). A seminorm, usually denoted as $|\cdot|$, fulfills the same set of assumptions than that of a norm, except that if $|x| = 0$ for a given $x \in X$, we do not necessarily have $x = 0$.

Lemma A.1 (Induced metric). Let $(X, \|\cdot\|)$ be a normed space. The function $d : X \times X \rightarrow \mathbb{R}$

$$d(x, y) = \|x - y\|, \quad \forall x, y \in X,$$

defines a metric on X , called the induced metric.

Definition A.8 (Banach space). A normed vector space X is called a Banach space if it is complete, in accordance with the definition of completeness of a space provided by Definition A.4

Definition A.9 (Inner product). Let X be a vector space over \mathbb{R} . An inner product on X is a function $\langle \cdot, \cdot \rangle : X \times X \rightarrow \mathbb{R}$ such that the following holds $\forall x, y, z \in X$ and $\alpha, \beta \in \mathbb{R}$:

(i) $\langle x, x \rangle \geq 0$ and $\langle x, x \rangle = 0$ if and only if $x = 0$, (nonnegativity)

(ii) $\langle x, y \rangle = \langle y, x \rangle$, (symmetry)

(iii) $\langle \alpha x + \beta y, z \rangle = \alpha \langle x, z \rangle + \beta \langle y, z \rangle$. (linearity)

The vector space X is then called an inner product space.

Remark. The assumption (iii) describes the linearity only in the first term. The linearity on the second argument is automatically induced by symmetry, i.e. assumption (ii).

Lemma A.2 (Induced norm). If $\langle \cdot, \cdot \rangle$ is an inner product on a given vector space X then the function defined $\forall x \in X$ such that

$$\|x\| = \langle x, x \rangle^{1/2},$$

represents a norm on X , called the inner product induced norm on X .

Theorem A.1 (Cauchy–Schwarz inequality). Let X be an inner product space with corresponding inner product $\langle \cdot, \cdot \rangle$. The following inequality holds $\forall x, y \in X$:

$$|\langle x, y \rangle| \leq \|x\| \|y\|,$$

where $\|\cdot\|$ denotes the inner product induced norm.

Definition A.10 (Hilbert space). A Hilbert space H is an inner product space that is complete with respect to the inner product induced norm.

Definition A.11 (Separability). A Hilbert space is said to be separable if it contains a dense countable subset.

Remark. The set \mathbb{R} is separable for example, since \mathbb{Q} is countable.

Appendix B

Basic notions of measure theory

B.1 Measurable spaces

Definition B.1 (σ -algebra). Let X be a given set. A σ -algebra or σ -field \mathcal{F} is a collection of subsets of X satisfying the three following properties:

- (i) $\emptyset \in \mathcal{F}$,
- (ii) $A \in \mathcal{F} \implies A^c \in \mathcal{F}$, (closure under complementation)
- (iii) If $A_1, A_2, \dots, A_n, \dots \in \mathcal{F} \implies \bigcup_{n=1}^{\infty} A_n \in \mathcal{F}$. (closure under countable union)

The combination of (i) and (ii) implies that $X \in \mathcal{F}$. Furthermore, we have by definition that $\mathcal{F} \subseteq \mathcal{P}(X)$ where $\mathcal{P}(X)$ denotes the power set of X .

Proposition B.1. Suppose that we are given σ -algebras \mathcal{F}_i on X , for $i \in \mathcal{I}$ (index set). Then, $\bigcap_{i \in \mathcal{I}} \mathcal{F}_i$ is also a σ -algebra.

Proposition B.2. For $\mathcal{M} \subseteq \mathcal{P}(X)$, there exists a smallest σ -algebra that contains \mathcal{M} denoted $\sigma(\mathcal{M})$ such that:

$$\sigma(\mathcal{M}) := \bigcap_{\mathcal{M} \subseteq \mathcal{F}} \mathcal{F},$$

where the \mathcal{F} 's are σ -algebras. We call it the σ -algebra generated by \mathcal{M} .

Definition B.2 (Borel σ -algebra). Let X be a metric space with open sets corresponding to $\mathcal{B}(x, \epsilon)$, $\forall x \in X$ and $\forall \epsilon > 0$. We call the Borel σ -algebra, denoted by $\mathcal{B}(X)$, the σ -algebra generated by the open sets of X .

Definition B.3 (Measure function). Let X be a set and \mathcal{F} a σ -algebra over X . A measure μ on \mathcal{F} is a function $\mu : \mathcal{F} \rightarrow \mathbb{R}$ satisfying the following properties:

- (i) $\mu(A) \geq 0$, $\forall A \in \mathcal{F}$, (non-negativity)
- (ii) $\mu(\emptyset) = 0$, (null empty-set)
- (iii) For every countable collection of $\{A_n\}_{n=1}^{\infty}$, $A_n \in \mathcal{F}$, such that $A_i \cap A_j = \emptyset$, $\forall i \neq j$, we have that $\mu(\bigcup_{i=1}^{\infty} A_i) = \sum_{n=1}^{\infty} \mu(A_n)$. (countable additivity)

The pair (X, \mathcal{F}) is called a measurable space and $A \in \mathcal{F}$ is called a measurable set. The triple (X, \mathcal{F}, μ) is called a measure space.

Remark. Provided (iii), if a single set has finite measure, assumption (ii) holds automatically provided that

$$\mu(A) = \mu(A \cup \emptyset) = \mu(A) + \mu(\emptyset).$$

Definition B.4 (Measurable function). Let (X, \mathcal{F}_X) and (Y, \mathcal{F}_Y) be two measurable spaces. A function $f : X \rightarrow Y$ is said to be measurable if for every Y -measurable set $B \in \mathcal{F}_Y$, the inverse $A := f^{-1}(B)$ is X -measurable, i.e. $A \in \mathcal{F}_X$.

B.2 Probability and Bochner spaces

Definition B.5 (Probability measure). Let X be a set and \mathcal{F} a σ -algebra over X . A probability measure \mathbb{P} is a measure function such that $0 \leq \mathbb{P}(A) \leq 1$, $\forall A \in \mathcal{F}$ and $\mathbb{P}(X) = 1$.

Definition B.6 (Sample space). A sample space Ω is the collection of possible outcomes of a random experiment. It is also referred to as the universe of all possible outcomes.

Definition B.7 (Probability space). A probability space is a measurable space constituted by a triplet $(\Omega, \mathcal{F}, \mathbb{P})$, where Ω is a sample space, \mathcal{F} a σ -algebra on Ω and \mathbb{P} a probability measure. In this context, the σ -algebra \mathcal{F} is called an event space, and \mathbb{P} a probability function.

Definition B.8 (Random variable). Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. A function $X : \Omega \rightarrow \mathbb{R}$ is called a random variable if for $a \in \mathbb{R}$, we have

$$X([a, +\infty[)^{-1} = \{\omega \in \Omega, X(\omega) \geq a\} \in \mathcal{F}.$$

Definition B.9 (Borel measurability). Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and H a Hilbert space. A function $f : \Omega \rightarrow H$ is Borel measurable if f is measurable with respect to $B(H)$, that is the Borel σ -algebra induced by H . This is also referred to as the weak measurability.

Definition B.10 (Separably valued function). A function $f : \Omega \rightarrow H$ is said to be almost surely separably valued if there exists a subset $N \subset \Omega$ with $\mathbb{P}(N) = 0$ such that $f(\Omega \setminus N) \subset H$ is separable.

Definition B.11 (Bochner measurability). A function f is said to be Bochner measurable, or strongly measurable, if there exists a sequence of separably valued function $\{f_n\}_{n=1}^{\infty}$, such that:

$$f(\omega) = \lim_{n \rightarrow \infty} f_n(\omega),$$

for almost every $\omega \in \Omega$ and $f^{-1}(x)$ is measurable $\forall x \in H$.

Theorem B.1. A function $f : \Omega \rightarrow H$ is Bochner (strongly) measurable if and only if it is Borel (weakly) measurable almost surely separably valued.

Definition B.12. A Bochner measurable function $f : \Omega \rightarrow H$ is Bochner integrable if there exists a sequence $\{s_n\}_{n=1}^{\infty}$ of simple functions (i.e. $s_n = \sum_{i=1}^n \chi_{A_i} x_i$, with $A_i \in \mathcal{F}$ and $x_i \in H$) such that:

$$\int_{\Omega} \|f - s_n\|_H d\mathbb{P}(\omega) \xrightarrow{n \rightarrow \infty} 0.$$

Proposition B.3. A Bochner measurable function $f : \Omega \rightarrow H$ is Bochner integrable if and only if $\omega \mapsto \|f(\omega)\|_H \in \mathbb{R}$ is integrable with respect to \mathbb{P} (this is equivalent to write $f \in L^1(\Omega; H)$). We then have $\forall A \in \mathcal{F}$:

$$\left\| \int_A f(\omega) d\mathbb{P}(\omega) \right\|_H \leq \int_A \|f(\omega)\|_H d\mathbb{P}(\omega).$$

Therefore, if $f \in L^1(\Omega; H)$, we have that $\mathbb{E}[f] = \int f(\omega) d\mathbb{P}(\omega) \in H$ exists.

B.3 Lebesgue Measure and Integral on \mathbb{R}

Let E be a subset of \mathbb{R} ($E \subset \mathbb{R}$) and $(I_j)_{j \in \mathcal{J}}$ be a finite, or countable infinite, coverage of R with open intervals.

Definition B.13 (Outer measure of a set). *The outer measure of a E is defined by*

$$\mathcal{M}_{\text{out}}(E) = \inf_{(I_j)_{j \in \mathcal{J}}} \sum_j l(I_j),$$

where $l(I_j)$ denotes the length of the interval I_j .

Definition B.14 (Inner measure of a set). *Suppose that E has a finite outer measure. Let $\mathcal{F}(E)$ be the set of closed subsets of E . The inner measure of E is defined to be:*

$$\mathcal{M}_{\text{in}}(E) = \sup_{K \subset \mathcal{F}(E)} \mathcal{M}_{\text{out}}(K),$$

Definition B.15 (Measurability in the sense of Lebesgue). *If a set E is such that $\mathcal{M}_{\text{in}}(E) = \mathcal{M}_{\text{out}}(E)$, we say that it is measurable in the sense of Lebesgue and we set its Lebesgue measure to:*

$$\mathcal{M}(E) = \mathcal{M}_{\text{in}}(E) = \mathcal{M}_{\text{out}}(E).$$

Proposition B.4. *Every countable set of points has a zero measure.*

Definition B.16. *A property is said to hold almost everywhere if the set of points if it does not hold only on sets of measure zero.*

Definition B.17 (Lebesgue measurable functions). *Let $f : E \subset \mathbb{R} \rightarrow \mathbb{R}$. The function f is said to be Lebesgue measurable if $\forall a \in \mathbb{R}$, the set $\{x \in E \mid f(x) < a\}$ is a measurable set in the sense of Lebesgue.*

Definition B.18 (Simple functions). *Let $E \subset \mathbb{R}$. We denote with χ_E the characteristic function of the set E , that is:*

$$\chi_E(x) = \begin{cases} 1, & \text{if } x \in E \\ 0, & \text{otherwise.} \end{cases}$$

A simple function $f : E \rightarrow \mathbb{R}$ is a function such that it can be written as a linear combination of characteristic functions on measurable sets, i.e.

$$f(x) = \sum_{i=1}^n a_i \chi_{E_i}(x), \tag{B.1}$$

where $E_i \cap E_j = \emptyset$, $\forall i \neq j$, and $a_i \in \mathbb{R}$, $\forall i = 1, \dots, n$.

Definition B.19 (Integral on positive simple functions). *Let f be a positive simple function, i.e. $a_i > 0$ in (B.1). The Lebesgue integral of the function f is given by:*

$$\int_E f \, d\mu = \sum_{i=1}^n a_i \mu(E_i),$$

where we set $a_i \mu(E_i) = 0$ whenever $a_i = 0$ and $\mu(E_i) = \infty$.

Definition B.20. *Let $f : \mathbb{R} \rightarrow \mathbb{R}$ and A a measurable subset of \mathbb{R} . We have that*

$$\int_A f \, d\mu = \int_{\mathbb{R}} f \chi_A \, d\mu = \sum_{i=1}^n a_i \mu(E_i \cap A),$$

Definition B.21. *Let f be a positive Lebesgue measurable function. The integral of f is given by:*

$$\int f \, d\mu = \sup_{s \leq f} \int s \, d\mu,$$

where s represents a simple function. If $\int f \, d\mu$ is finite, f is said to be Lebesgue integrable (equivalently, integrable in the sense of Lebesgue).

Definition B.22. *Let f be an arbitrary, measurable function. We define:*

$$f^+(x) = \begin{cases} f(x), & \text{if } f(x) \geq 0 \\ 0, & \text{otherwise,} \end{cases}$$

and

$$f^-(x) = \begin{cases} -f(x), & \text{if } f(x) \leq 0 \\ 0, & \text{otherwise.} \end{cases}$$

in such a way that we have $f = f^+ - f^-$. The integral of f is defined with respect to the positive functions f^+ and f^- , by means of

$$\int f \, d\mu = \int f^+ \, d\mu - \int f^- \, d\mu.$$

Remark. • If $\int f^+ \, d\mu > \infty$ and $\int f^- \, d\mu > \infty$, then $\int f \, d\mu$ does not exist.

- If $\int f^+ \, d\mu > \infty$ (respectively f^-) and $\int f^- \, d\mu$ is finite (respectively f^+), then $\int f \, d\mu = +\infty$ (respectively $-\infty$).
- f is said to be Lebesgue integrable if $\int f \, d\mu < +\infty$.

Remark. All the definitions above can be generalized to functions $f : D \rightarrow \mathbb{R}$ with $D \subset \mathbb{R}^n$. The construction, mainly the sets measurability, is obtained by considering Borel sets in \mathbb{R}^n obtained by tensor product of Borel sets in \mathbb{R} , i.e.

$$\underbrace{\mathcal{B}(\mathbb{R}) \otimes \dots \otimes \mathcal{B}(\mathbb{R})}_{n \text{ times}} = \mathcal{B}(\mathbb{R}^n).$$

The product measure is considered on the basis of the usual Lebesgue measure. The theorem of Fubini–Tonelli (c.f. [Coh13]) relates the evaluation of integrals in \mathbb{R}^n , to that in \mathbb{R} .

Appendix C

Sobolev Spaces

In the following, we always consider $D \in \mathbb{R}^d$ to be a d -dimensional domain with a Lipschitz boundary.

Definition C.1. For $1 \leq p \leq \infty$, we define

$$L^p(D) = \{f : D \rightarrow \mathbb{R}, f \text{ measurable and } \|f\|_{L^p(D)} < \infty\},$$

where we define the norms:

$$\|f\|_{L^p(D)} = \begin{cases} \left(\int_D |f(\mathbf{x})|^p d\mathbf{x}\right)^{1/p}, & \text{for } 1 \leq p < \infty, \\ \text{ess sup}_{\mathbf{x} \in D} |f(\mathbf{x})|, & \text{for } p = \infty. \end{cases} \quad (\text{C.1})$$

We recall that the essential supremum of a function f is given by:

$$\text{ess sup}_{\mathbf{x} \in D} |f(\mathbf{x})| = \inf\{C \geq 0 \text{ such that } |f(\mathbf{x})| \leq C \text{ almost everywhere}\}.$$

In other terms, the $L^p(D)$ spaces define sets of equivalence classes of functions, that are integrable on D up to the power $1 \leq p < \infty$. For $p = \infty$ in turn, $L^\infty(D)$ represents the set of essentially bounded functions.

Remark. The equivalence class of a function is the set of functions that are defined to be equal almost everywhere, i.e. they may differ only on sets of measure zero. Provided that f is a representative of an equivalence class of interest, we write:

$$\lfloor f \rfloor = \{g : D \rightarrow \mathbb{R}, g(\mathbf{x}) = f(\mathbf{x}) \text{ almost everywhere}\}.$$

Proposition C.1. The spaces $L^p(D)$ are Banach spaces with respect to the norms defined in (C.1).

Proposition C.2. For $p = 2$, the spaces $L^2(D)$ is a Hilbert where the inner product inducing the norm is provided by:

$$(f, g) = \int_D f(\mathbf{x})g(\mathbf{x}) d\mathbf{x}.$$

We therefore also have the Cauchy–Schwarz inequality in $L^2(D)$:

$$(f, g) = \int_D f(\mathbf{x})g(\mathbf{x}) d\mathbf{x} \leq \|f\|_{L^2(D)} \|g\|_{L^2(D)} = \left(\int_D f^2(\mathbf{x}) d\mathbf{x} \right)^{1/2} \left(\int_D g^2(\mathbf{x}) d\mathbf{x} \right)^{1/2}$$

Proposition C.3 (L^p -embedding). For $1 \leq p < q \leq \infty$, we have the following inclusion:

$$L^q(D) \subset L^p(D).$$

In order to define a more general class of Sobolev spaces, we need to introduce the notion of weak derivative.

Definition C.2. We denote with:

$$C_0^\infty(D) = \{v : D \rightarrow \mathbb{R} : v \in C^\infty(D), \text{supp}(v) \subset D\},$$

the space of infinitely differentiable functions with compact support in D , where the support of a function is defined by means of:

$$\text{supp}(v) = \overline{\{\mathbf{x} \in D : v(\mathbf{x}) \neq 0\}}.$$

We also introduce the multi-index notation $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_n)$ of order $|\boldsymbol{\alpha}| = \sum_{k=1}^n \alpha_k$, for which we have:

$$D^{\boldsymbol{\alpha}} f = \frac{\partial^{|\boldsymbol{\alpha}|} f}{\partial x_1^{\alpha_1} \dots \partial x_n^{\alpha_n}},$$

corresponding to the α_k -th partial derivative of the function f in each coordinate x_k . Furthermore, we say that $f \in L_{\text{loc}}^1(D)$ if we have

$$\int_{D'} |f(\mathbf{x})| d\mathbf{x} < \infty, \forall D' \subset D.$$

Definition C.3 (Weak derivative). Let f and $F' \in L^1_{loc}(D)$. If for every $v \in C_0^\infty(D)$, we have

$$\int_D F(\mathbf{x})v(\mathbf{x}) \, d\mathbf{x} = (-1)^{|\alpha|} \int_D f(\mathbf{x})D^\alpha v(\mathbf{x}) \, d\mathbf{x},$$

then F is called the weak derivative of f corresponding to the multi-index α . We write $F(\mathbf{x}) = D^\alpha f(\mathbf{x})$.

Definition C.4 (H^k -spaces). The spaces $H^k(D)$ consist of all equivalence classes of functions in $L^2(D)$ such that all their weak derivatives of multi-index $|\alpha| \leq k$ exist and are square integrable, i.e.

$$H^k(D) = \{f \in L^2(D) : D^\alpha f \in L^2(D), \forall |\alpha| \leq k\}.$$

We furthermore equip the spaces $H^k(D)$ with the norm

$$\|f\|_{H^k(D)} = \left(\sum_{|\alpha| \leq k} \int_D |D^\alpha f(\mathbf{x})|^2 \, d\mathbf{x} \right)^{1/2},$$

induced by the inner product

$$(f, g)_{H^k(D)} = (f, g)_{L^2(D)} + \sum_{1 \leq |\alpha| \leq k} (D^\alpha f, D^\alpha g)_{L^2(D)}. \quad (\text{C.2})$$

Proposition C.4. The $H^k(D)$ spaces are complete with respect to the norm induced by the inner product (C.2), and therefore represent Hilbert spaces.

Remark. By definition of the spaces $H^k(D)$ and $L^2(D)$, the following embedding holds $\forall k$,

$$H^k(D) \subset L^2(D).$$

We furthermore recover $L^2(D)$ for $k = 0$, i.e. $H^0(D) = L^2(D)$.

Definition C.5 (General Sobolev spaces). The spaces $W^{k,p}(D)$ are generalization of the spaces $L^p(D)$ and $H^k(D)$, that is:

$$W^{k,p}(D) = \{f \in L^p(D) : D^\alpha f \in L^p(D), \forall |\alpha| \leq k\},$$

equipped with the norms:

$$\|f\|_{W^{k,p}(D)} = \begin{cases} \left(\sum_{|\alpha| \leq k} \int_D |D^\alpha f(\mathbf{x})|^p \, d\mathbf{x} \right)^{1/p}, & \text{for } 1 \leq p < \infty, \\ \sum_{|\alpha| \leq k} \text{ess sup}_{x \in D} |D^\alpha f(\mathbf{x})|^p, & \text{for } p = \infty. \end{cases} \quad (\text{C.3})$$

Proposition C.5. *The spaces $W^{k,p}(D)$ are complete with respect to the norms (C.3), and therefore Banach spaces.*

Remark. • For $k = 0$, we recover the L^p -spaces, i.e. $W^{0,p}(D) = L^p(D)$.

• For $p = 2$, we recover the H^k -spaces, i.e. $W^{k,2}(D) = H^k(D)$.

Lemma C.1 (Poincaré–Friedrich inequality). *Let us consider the case $k = 1$ and functions in $W_0^{1,p}(D)$, with*

$$W_0^{1,p}(D) = \{f \in W^{1,p}(D) \text{ such that } f|_{\partial D} = 0\},$$

i.e. functions in $W^{1,p}(D)$ with a zero trace at the boundary. We then have the following inequality:

$$\int_D |f(\mathbf{x})|^p d\mathbf{x} \leq C_D \int_D \|\nabla f(\mathbf{x})\|_2^p d\mathbf{x}, \quad (\text{C.4})$$

where $\|\cdot\|_2$ denotes the Euclidian norm and $C_D > 0$ is a constant depending on the domain D .

Remark. *The inequality defined in (C.4) implies that for functions in $W^{1,p}(D)$ with zero trace (i.e. $W_0^{1,p}(D)$), the norm (C.3) is equivalent to the semi-norm $|\cdot|_{W^{k,p}(D)}$ defined by*

$$|f|_{W^{k,p}(D)} = \begin{cases} \left(\sum_{|\alpha|=k} \int_D |D^\alpha f(\mathbf{x})|^p d\mathbf{x} \right)^{1/p}, & \text{for } 1 \leq p < \infty, \\ \sum_{|\alpha|=k} \text{ess sup}_{x \in D} |D^\alpha f(\mathbf{x})|^p, & \text{for } p = \infty. \end{cases}$$

Remark. *The classical Poincaré inequality is given for $p = 2$ and reads as*

$$\|f\|_{L^2(D)} \leq C_D \|\nabla f\|_{L^2(D)},$$

for every function $f \in H_0^1(D)$.

We finish by enunciate an important identity, very useful in variational formulations.

Theorem C.1 (Green first formula). *The following identity holds $\forall f \in H^2(D)$ and $v \in H^1(D)$:*

$$\int_D \nabla f(\mathbf{x}) \nabla v(\mathbf{x}) d\mathbf{x} = \int_{\partial D} \frac{\partial f}{\partial \mathbf{n}}(\mathbf{s}) v(\mathbf{s}) ds - \int_D \Delta f(\mathbf{x}) v(\mathbf{x}) d\mathbf{x},$$

where \mathbf{s} defines the surface variable, and \mathbf{n} its outer normal.

Bibliography

- [AL12] M. AlGhatrif and J. Lindsay. A brief review: history to understand fundamentals of electrocardiography. *Journal of community hospital internal medicine perspectives*, 2(1):14383, 2012. [2](#)
- [AM89] A. K. Aziz and P. Monk. Continuous finite elements in space and time for the heat equation. *Mathematics of Computation*, 52(186):255–274, 1989. [48](#)
- [BAA⁺16] S. Balay, S. Abhyankar, M. F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, K. Rupp, B. F. Smith, S. Zampini, H. Zhang, and H. Zhang. PETSc users manual. Technical Report ANL-95/11 - Revision 3.7, Argonne National Laboratory, 2016. [110](#)
- [Bab71] I. Babuška. Error-bounds for finite element method. *Numerische Mathematik*, 16(4):322–333, 1971. [34](#)
- [Bai69] C. Baiocchi. Régularité des solutions d’une équation différentielle abstraite. *Publications mathématiques et informatique de Rennes*, (1):1–26, 1969. [72](#)
- [Bat06] K.-J. Bathe. *Finite element procedures*. Klaus-Jurgen Bathe, 2006. [28](#)
- [BB15] S. Ben Bader. Inverse identification of infarcted areas inside the cardiac muscle through an optimal control formulation for cardiac electrophysiology equations. Technical report, 2015. [22](#), [101](#)
- [BBBQ⁺21] S. Ben Bader, P. Benedusi, A. Quaglino, P. Zulian, and R. Krause. Space-time multilevel monte carlo methods and their application to cardiac electrophysiology. *Journal of Computational Physics*, 433:110164, 2021. [109](#)

- [BBN14] F. Bonizzoni, A. Buffa, and F. Nobile. Moment equations for the mixed formulation of the hodge laplacian with stochastic loading term. *IMA Journal of Numerical Analysis*, 34(4):1328–1360, 2014. [52](#)
- [BBPT12] J. D. Bayer, R. C. Blake, G. Plank, and N. A. Trayanova. A novel rule-based algorithm for assigning myocardial fiber orientation to computational heart models. *Annals of biomedical engineering*, 40(10):2243–2254, 2012. [24](#), [128](#)
- [BC15] C. Bierig and A. Chernov. Convergence analysis of multilevel monte carlo variance estimators and application for random obstacle problems. *Numerische Mathematik*, 130(4):579–613, 2015.
- [BCF⁺10] M. Boulakia, S. Cazeau, M. A. Fernández, J.-F. Gerbeau, and N. Zenzemi. Mathematical modeling of electrocardiograms: a numerical study. *Annals of biomedical engineering*, 38(3):1071–1097, 2010. [21](#), [135](#)
- [bdF22] J. B. J. baron de Fourier. *Théorie analytique de la chaleur*. Firmin Didot, 1822. [27](#)
- [Ben20] P. Benedusi. *Parallel space-time multilevel methods with application to electrophysiology: theory and implementation*. PhD thesis, Università della Svizzera italiana, 2020. [29](#), [117](#)
- [BGK⁺18a] P. Benedusi, C. Garoni, R. Krause, X. Li, and S. Serra-Capizzano. Space-time fe-dg discretization of the anisotropic diffusion equation in any dimension: the spectral symbol. *SIAM Journal on Matrix Analysis and Applications*, 39(3):1383–1420, 2018.
- [BGK⁺18b] P. Benedusi, C. Garoni, R. Krause, X. Li, and S. Serra-Capizzano. Space-Time FE-DG Discretization of the Anisotropic Diffusion Equation in Any Dimension: The Spectral Symbol. *SIAM Journal on Matrix Analysis and Applications*, 39(3):1383–1420, 2018.
- [BGW15] J. Biehler, M. W. Gee, and W. A. Wall. Towards efficient uncertainty quantification in complex and large-scale biomechanical problems based on a bayesian multi-fidelity scheme. *Biomech. Model. Mechanobiol.*, 14(3):489–513, 2015.

- [BHK⁺21] S. B. Bader, H. Harbrecht, R. Krause, M. Multerer, A. Quaglino, and M. Schmidlin. Space-time multilevel quadrature methods and their application for cardiac electrophysiology. *arXiv preprint arXiv:2105.02007*, 2021.
- [BOCF08] A. Bueno-Orovio, E. M. Cherry, and F. H. Fenton. Minimal model for human ventricular action potentials in tissue. *Journal of theoretical biology*, 253(3):544–560, 2008. [14](#)
- [BPP02] T. J. Baker, P. Pebay, and J. Pousin. Dynamic meshing for finite element based segmentation of cardiac imagery. In *WCCM V-Fifth World Congress on Computational Mechanics, Vienna*, 2002.
- [BPP⁺18] J. Bayer, A. J. Prassl, A. Pashaei, J. F. Gomez, A. Frontera, A. Neic, G. Plank, and E. J. Vigmond. Universal ventricular coordinates: A generic framework for describing position within the heart and transferring data. *Medical image analysis*, 45:83–93, 2018. [24](#)
- [BR06] F. Bornemann and C. Rasch. Finite-element discretization of static Hamilton-Jacobi equations based on a local variational principle. *Computing and Visualization in Science*, 9(2):57–69, 2006.
- [BS84] R. Beyar and S. Sideman. A computer study of the left ventricular performance based on fiber structure, sarcomere dynamics, and transmural electrical propagation velocity. *Circ. Res.*, 55(3):358–375, 1984.
- [BSZ11] A. Barth, C. Schwab, and N. Zollinger. Multi-level monte carlo finite element method for elliptic pdes with stochastic coefficients. *Numerische Mathematik*, 119(1):123–161, 2011. [78](#), [82](#)
- [BW17] J. Biehler and W. Wall. The impact of personalized probabilistic wall thickness models on peak wall stress in abdominal aortic aneurysms. *International Journal for Numerical Methods in Biomedical Engineering*, 2017.
- [BZ14] I. Bilonis and N. Zabararas. Solution of inverse problems with limited forward solver evaluations: a Bayesian perspective. *Inverse Prob.*, 30(1):015004, January 2014.
- [CBC⁺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. *Prog. Biophys. Mol. Biol.*, 104(1-3):22–48, 2011.
- [CFPS14] P. Colli Franzone, L. F. Pavarino, and S. Scacchi. *Mathematical Cardiac Electrophysiology*. Springer, 2014.
- [CG93] P. Colli Franzone and L. Guerri. Spreading of excitation in 3-D models of the anisotropic cardiac tissue. I. validation of the eikonal model. *Math. Biosci.*, 113:145–209, 1993.
- [CGM15] C. Corrado, J.-F. Gerbeau, and P. Moireau. Identification of weakly coupled multiphysics problems. application to the inverse problem of electrocardiography. *Journal of Computational Physics*, 283:271–298, 2015.
- [Che08] S. Cheng. A crash course on the lebesgue integral and measure theory, 2008. [8](#)
- [CK13] M. Capinski and P. E. Kopp. *Measure, integral and probability*. Springer Science & Business Media, 2013. [8](#)
- [CMO18] J. Camaño, C. Muñoz, and R. Oyarzúa. Numerical analysis of a dual-mixed problem in non-standard Banach spaces. 2018. [72](#)
- [CO10] S. Conti and A. O’Hagan. Bayesian emulation of complex multi-output and dynamic computer models. *Journal of Statistical Planning and Inference*, 140(3):640–651, March 2010.
- [Coh13] D. L. Cohn. *Measure theory*. Springer, 2013. [8](#), [151](#)
- [CST13] J. Charrier, R. Scheichl, and A. L. Teckentrup. Finite element error analysis of elliptic pdes with random coefficients and its application to multilevel Monte Carlo methods. *SIAM Journal on Numerical Analysis*, 51(1):322–352, 2013.
- [CWH⁺16] R. Chabiniok, V. Y. Wang, M. Hadjicharalambous, L. Asner, J. Lee, M. Sermesant, E. Kuhl, A. A. Young, P. Moireau, M. P. Nash, D. Chapelle, and D. A. Nordsletten. Multiphysics and multiscale modelling, data–model fusion and integration of organ physiology in the clinic: ventricular cardiac mechanics. *Interface Focus*, 6(2):20150083, 2016.

- [CZC⁺17] C. Chavez, N. Zemzemi, Y. Coudière, F. Alonso-Atienza, and D. Alvarez. Inverse problem of electrocardiography: estimating the location of cardiac ischemia in a 3d geometry. In *Functional Imaging and modelling of the heart (FIMH2015)*, volume 9126. Springer International Publishing, 2017. [21](#), [22](#)
- [d'A47] J. L. R. d'Alembert. Recherches sur la courbe que forme une corde tendue mise en vibration. 1747. [27](#)
- [DB97] C. J. Davidson and R. O. Bonow. Cardiac catheterization. *Libby P*, 10, 1997. [18](#)
- [DBV11] H. Dierckx, O. Bernus, and H. Verschelde. Accurate eikonal-curvature relation for wave fronts in locally anisotropic reaction-diffusion systems. *Phys. Rev. Lett.*, 107(10):108101, 2011.
- [DGR⁺17] D. Drzisga, B. Gmeiner, U. Råde, R. Scheichl, and B. I. Wohlmuth. Scheduling massively parallel multigrid for multilevel Monte Carlo methods. *SIAM Journal on Scientific Computing*, 39(5):S873–S897, 2017. [111](#)
- [Dia88] P. Diaconis. Bayesian numerical analysis. *Statistical decision theory and related topics IV*, 1:163–175, 1988.
- [DJS96] J. E. Dennis Jr and R. B. Schnabel. *Numerical methods for unconstrained optimization and nonlinear equations*, volume 16. SIAM, 1996.
- [DM96] E. G. Dyakonov and S. McCormick. *Optimization in solving elliptic problems*. CRC Press, 1996.
- [DMDJ06] P. Del Moral, A. Doucet, and A. Jasra. Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society Series B-Statistical Methodology*, 68:411–436, 2006.
- [Ein95] W. Einthoven. Über die Form des menschlichen Elektrokardiogramms. *Pflugers Arch fd ges des Physiol*, 60:101, 1895. [2](#)
- [EvDB14] B. Erem, P. van Dam, and D. H. Brooks. Identifying model inaccuracies and solution uncertainties in noninvasive activation-based imaging of cardiac excitation using convex relaxation. *IEEE Transactions on Medical Imaging*, 33(4):902–912, apr 2014.

- [FFK⁺17] R. D. Falgout, S. Friedhoff, T. V. Kolev, S. P. MacLachlan, J. B. Schroder, and S. Vandewalle. Multigrid methods with space–time concurrency. *Computing and Visualization in Science*, 18(4):123–143, Aug 2017.
- [FG93] P. C. Franzone and L. Guerri. Spreading of excitation in 3-d models of the anisotropic cardiac tissue. i. validation of the eikonal model. *Mathematical biosciences*, 113(2):145–209, 1993.
- [FGPR18] S. R. Franco, F. J. Gaspar, M. A. V. Pinto, and C. Rodrigo. Multigrid method based on a space-time approach with standard coarsening for parabolic problems. *Applied Mathematics and Computation*, 317:25–34, 2018.
- [FGTV85] P. C. Franzone, L. Guerri, B. Taccardi, and C. Viganotti. Finite element approximation of regularized solutions of the inverse potential problem of electrocardiography and applications to experimental data. *Calcolo*, 22(1):91–186, 1985. [21](#)
- [Fis03] G. Fishman. *Monte Carlo: Concepts, Algorithms, and Applications*. Springer, New York, corrected edition edition, February 2003. [6](#)
- [Fit61] R. FitzHugh. Impulses and physiological states in theoretical models of nerve membrane. *Biophysical journal*, 1(6):445–466, 1961. [4](#), [15](#)
- [FK93] D. A. French and J. T. King. Analysis of a robust finite element approximation for a parabolic equation with rough boundary data. *Mathematics of computation*, 60(201):79–104, 1993.
- [FPS14] P. C. Franzone, L. F. Pavarino, and S. Scacchi. *Mathematical cardiac electrophysiology*, volume 13. Springer, 2014. [15](#)
- [FPT05] P. C. Franzone, L. Pavarino, and B. Taccardi. Simulating patterns of excitation, repolarization and action potential duration with cardiac bidomain and monodomain models. *Mathematical biosciences*, 197(1):35–66, 2005. [15](#)
- [FZ10] M. A. Fernández and N. Zemzemi. Decoupled time-marching schemes in computational cardiac electrophysiology and ecg numerical simulation. *Mathematical biosciences*, 226(1):58–75, 2010. [21](#), [135](#)

- [Gan15] M. J. Gander. 50 years of time parallel time integration. In *Multiple Shooting and Time Domain Decomposition Methods*, pages 69–113. Springer, 2015. [29](#)
- [GH13a] W. Gong and M. Hinze. Error estimates for parabolic optimal control problems with control and state constraints. *Computational Optimization and Applications*, 56(1):131–151, 2013.
- [GH13b] M. Griebel and H. Harbrecht. On the construction of sparse tensor product spaces. *Mathematics of computation*, 82(282):975–994, 2013. [79](#)
- [GH14] M. Griebel and H. Harbrecht. Approximation of bi-variate functions: singular value decomposition versus sparse grids. *IMA journal of numerical analysis*, 34(1):28–54, 2014. [59](#), [61](#)
- [GHM15] M. Griebel, H. Harbrecht, and M. Multerer. Multilevel quadrature for elliptic parametric partial differential equations on non-nested meshes. 2015. [80](#)
- [Gil08] M. Giles. Improved multilevel monte carlo convergence using the milstein scheme. In *Monte Carlo and Quasi-Monte Carlo Methods 2006*, pages 343–358. Springer, 2008. [78](#)
- [Gil13a] M. B. Giles. Multilevel monte carlo methods. In *Monte Carlo and Quasi-Monte Carlo Methods 2012*, pages 83–103. Springer, 2013. [78](#)
- [Gil13b] M. B. Giles. Multilevel Monte Carlo methods. In *Monte Carlo and Quasi-Monte Carlo Methods 2012*, pages 83–103. Springer, 2013. bibtex: giles2013multilevel.
- [Gil15] M. B. Giles. Multilevel monte carlo methods. *Acta Numerica*, 24:259–328, 2015. [6](#)
- [GKMP20] L. Gander, R. Krause, M. Multerer, and S. Pezzuto. Space-time shape uncertainties in the forward and inverse problem of electrocardiography. *arXiv preprint arXiv:2010.16104*, 2020. [24](#)
- [GKN⁺15] I. G. Graham, F. Y. Kuo, J. A. Nichols, R. Scheichl, C. Schwab, and I. H. Sloan. Quasi-monte carlo finite element methods for elliptic pdes with lognormal random coefficients. *Numerische Mathematik*, 131(2):329–368, 2015. [62](#)

- [GKZ18] M. J. Gander, F. Kwok, and H. Zhang. Multigrid interpretations of the parareal algorithm leading to an overlapping variant and mgrid. *Computing and Visualization in Science*, 19(3):59–74, 2018. [29](#)
- [GN16] M. J. Gander and M. Neumuller. Analysis of a new space-time parallel multigrid algorithm for parabolic problems. *SIAM Journal on Scientific Computing*, 38(4):A2173–A2208, 2016.
- [Hac85] W. Hackbusch. Parabolic multi-grid methods. In *Proc. Of the Sixth Int’L. Symposium on Computing Methods in Applied Sciences and Engineering, VI*, pages 189–197, Amsterdam, The Netherlands, The Netherlands, 1985. North-Holland Publishing Co.
- [HB06] C. Hesch and P. Betsch. A comparison of computational methods for large deformation contact problems of flexible bodies. *ZAMM - Journal of Applied Mathematics and Mechanics / Zeitschrift für Angewandte Mathematik und Mechanik*, 86(10):818–827, 2006. [94](#)
- [Hei01] S. Heinrich. Multilevel monte carlo methods. In *International Conference on Large-Scale Scientific Computing*, pages 58–67. Springer, 2001. [78](#)
- [Hen93] C. S. Henriquez. Simulating the electrical behavior of cardiac tissue using the bidomain model. *Critical reviews in biomedical engineering*, 21(1):1–77, 1993. [4](#), [16](#)
- [Hen14] C. S. Henriquez. A brief history of tissue models for cardiac electrophysiology. *IEEE Trans. Biomed. Eng.*, 61:1457–1465, 2014.
- [HH52] 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*, 117(4):500–544, 1952. [4](#), [13](#)
- [HH14] D. E. Hurtado and D. Henao. Gradient flows and variational principles for cardiac electrophysiology: toward efficient and robust numerical simulations of the electrical activity of the heart. *Computer Methods in Applied Mechanics and Engineering*, 273:238–254, 2014.
- [Hla61] E. Hlawka. Funktionen von Beschränkter Variatiou in der Theorie der Gleichverteilung. *Annali di Matematica Pura ed Applicata*, 54(1):325–333, 1961. [83](#)

- [HPS12a] H. Harbrecht, M. Peters, and R. Schneider. On the low-rank approximation by the pivoted Cholesky decomposition. *Applied numerical mathematics*, 62(4):428–440, 2012.
- [HPS12b] H. Harbrecht, M. Peters, and M. Siebenmorgen. On multilevel quadrature for elliptic stochastic partial differential equations. In *Sparse grids and applications*, pages 161–179. Springer, 2012. [78](#), [79](#)
- [HPS15] H. Harbrecht, M. Peters, and M. Siebenmorgen. Efficient approximation of random fields for numerical applications. *Numerical Linear Algebra with Applications*, 22(4):596–617, 2015. [59](#)
- [HPS16] H. Harbrecht, M. Peters, and M. Siebenmorgen. Analysis of the domain mapping method for elliptic diffusion problems on random domains. *Numerische Mathematik*, 134(4):823–856, 2016. [85](#)
- [HPS17a] H. Harbrecht, M. Peters, and M. Siebenmorgen. On the quasi-Monte Carlo method with Halton points for elliptic PDEs with log-normal diffusion. *Mathematics of Computation*, 86(304):771–797, 2017. [78](#), [84](#)
- [HPS17b] H. Harbrecht, M. D. Peters, and M. Schmidlin. Uncertainty quantification for pdes with anisotropic random diffusion. *SIAM Journal on Numerical Analysis*, 55(2):1002–1023, 2017. [71](#), [72](#), [84](#), [99](#)
- [HRS⁺16] J. S. Hesthaven, G. Rozza, B. Stamm, et al. *Certified reduced basis methods for parametrized partial differential equations*. Springer, 2016.
- [HV95] G. Horton and S. Vandewalle. A space-time multigrid method for parabolic partial differential equations. *SIAM Journal on Scientific Computing*, 16(4):848–864, 1995.
- [Joh19] J. Johnsen. Well-posed final value problems and Duhamel’s formula for coercive Lax–Milgram operators. *arXiv preprint arXiv:1906.03117*, 2019. [72](#)
- [JW08] W.-K. Jeong and R. T. Whitaker. A fast iterative method for eikonal equations. *SIAM J. Sci. Comput.*, 30(5):2512–2534, 2008.
- [Kat10] A. M. Katz. *Physiology of the Heart*. Lippincott Williams & Wilkins, 2010. [12](#)

- [KBMZ18] A. Karoui, L. Bear, P. Migerditichan, and N. Zemzemi. Evaluation of fifteen algorithms for the resolution of the electrocardiography imaging inverse problem using ex-vivo and in-silico data. *Frontiers in physiology*, 9:1708, 2018. [22](#)
- [KC71] K. T. Kavanagh and R. W. Clough. Finite element applications in the characterization of elastic solids. *International Journal of Solids and Structures*, 7(1):11–23, 1971. [28](#)
- [KK16] D. Krause and R. Krause. Enabling local time stepping in the parallel implicit solution of reaction–diffusion equations via space-time finite elements on shallow tree meshes. *Applied Mathematics and Computation*, 277:164–179, 2016.
- [Kla11] R. Klabunde. *Cardiovascular physiology concepts*. Lippincott Williams & Wilkins, 2011. [12](#)
- [Kou09] P-S. Koutsourelakis. Accurate uncertainty quantification using inaccurate computational models. *SIAM J. Sci. Comput.*, 31(5):3274–3300, 2009.
- [KPD⁺12] D. Krause, M. Potse, T. Dickopf, R. Krause, A. Auricchio, and F. W. Prinzen. Hybrid parallelization of a large-scale heart model. In R. Keller, D. Kramer, and J.-P. Weiss, editors, *Facing the Multicore-Challenge II*, volume 7174 of *Lecture Notes in Computer Science*, pages 120–132, Berlin, 2012. Springer.
- [Kre78] E. Kreyszig. *Introductory functional analysis with applications*, volume 1. Wiley, New York, 1978. [33](#)
- [Kre15] M. Kreuter. Sobolev spaces of vector-valued functions. *Ulm University Faculty of Mathematics and Economics*, 2015.
- [KS98] J. P. Keener and J. Sneyd. *Mathematical physiology*, volume 1. Springer, 1998. [13](#), [16](#)
- [KW08] M. H. Kalos and P. A. Whitlock. *Monte Carlo Methods*. John Wiley & Sons, October 2008.
- [KZ16] R. Krause and P. Zulian. A parallel approach to the variational transfer of discrete fields between arbitrarily distributed unstructured finite element meshes. *SIAM Journal on Scientific Computing*, 38(3):C307–C333, 2016.

- [LB13] M. G. Larson and F. Bengzon. *The finite element method: theory, implementation, and applications*, volume 10. Springer Science & Business Media, 2013. [28](#), [35](#)
- [LGT03] G. T. Lines, P. Grottum, and A. Tveito. Modeling the electrical activity of the heart: a bidomain model of the ventricles embedded in a torso. *Computing and Visualization in Science*, 5(4):195–213, 2003. [21](#)
- [LHS97] I. J. LeGrice, P. J. Hunter, and B. Smaill. Laminar structure of the heart: a mathematical model. *American Journal of Physiology-Heart and Circulatory Physiology*, 272(5):H2466–H2476, 1997. [24](#)
- [LK10] L. Leifsson and S. Koziel. Multi-fidelity design optimization of transonic airfoils using physics-based surrogate modeling and shape-preserving response prediction. *Journal of Computational Science*, 1(2):98–106, 2010.
- [LM12] J. L. Lions and E. Magenes. *Non-homogeneous boundary value problems and applications: Vol. 1*, volume 181. Springer Science & Business Media, 2012. [72](#)
- [LMQR13] T. Lassila, A. Manzoni, A. Quarteroni, and G. Rozza. A reduced computational and geometrical framework for inverse problems in hemodynamics. *International journal for numerical methods in biomedical engineering*, 29(7):741–776, 2013. [22](#)
- [LMT01] J. Lions, Y. Maday, and G. Turinici. A parareal in time discretization of PDEs. *Comptes Rendus de l'Académie des Sciences-Series I-Mathematics* 332, pages 661–668, 2001. [29](#)
- [Löh87] R. Löhner. An adaptive finite element scheme for transient problems in cfd. *Computer Methods in Applied Mechanics and Engineering*, 61(3):323–338, 1987. [28](#)
- [LP16] W. Lam and D. Pennell. Imaging of the heart: historical perspective and recent advances. *Postgraduate medical journal*, 92(1084):99–104, 2016. [3](#)
- [LSC⁺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*, 269(2):H571–H582, 1995. [16](#), [24](#)
- [LSK⁺13] P. Lamata, M. Sinclair, E. Kerfoot, A. Lee, A. Crozier, B. Blazevic, S. Land, A. J. Lewandowski, D. Barber, S. Niederer, and N. Smith. An automatic service for the personalization of ventricular cardiac meshes. *Journal of The Royal Society Interface*, 11(91):20131023, dec 2013.
- [Mat42] C. Matteucci. Sur un phenomene physiologique produit par les muscles en contraction. *Ann Chim Phys.*, 6:339–341, 1842. [2](#)
- [McD16] E. McDonald. *All-at-once solution of time-dependent PDE problems*. PhD thesis, University of Oxford, 2016. [29](#)
- [MG78] W. T. Miller and D. B. Geselowitz. Simulation studies of the electrocardiogram. i. the normal heart. *Circulation Research*, 43(2):301–315, 1978.
- [MG15] E. Madenci and I. Guven. *The finite element method and applications in engineering using ANSYS*. Springer, 2015. [28](#)
- [MPG⁺16] 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*, 594(23):6833–6847, 2016.
- [MS03] C. C. Mitchell and D. G. Schaeffer. A two-current model for the dynamics of cardiac membrane. *Bulletin of mathematical biology*, 65(5):767–793, 2003. [14](#)
- [MS12] S. Mishra and C. Schwab. Sparse tensor multi-level monte carlo finite volume methods for hyperbolic conservation laws with random initial data. *Mathematics of computation*, 81(280):1979–2018, 2012. [78](#)
- [MW16] E. McDonald and A. Wathen. A simple proposal for parallel computation over time of an evolutionary process with implicit time stepping. In *Numerical Mathematics and Advanced Applications ENUMATH 2015*, pages 285–293. Springer, 2016.
- [Nel04] M. E. Nelson. Electrophysiological models. *Databasing the brain: from data to knowledge*, pages 285–301, 2004. [14](#)

- [Nie64] J. Nievergelt. Parallel methods for integrating ordinary differential equations. *Communications of the ACM*, 7(12):731–733, 1964. [29](#)
- [Nie92] H. Niederreiter. *Random number generation and quasi-Monte Carlo methods*. SIAM, 1992. [84](#)
- [NK11] C. Nagaiah and K. Kunisch. Higher order optimization and adaptive numerical solution for optimal control of monodomain equations in cardiac electrophysiology. *Applied numerical mathematics*, 61(1):53–65, 2011. [22](#)
- [NKP13] C. Nagaiah, K. Kunisch, and G. Plank. Optimal control approach to termination of re-entry waves in cardiac electrophysiology. *Journal of Mathematical Biology*, 67(2):359–388, 2013. [22](#)
- [NMSP11a] S. Niederer, L. Mitchell, N. Smith, and G. Plank. Simulating human cardiac electrophysiology on clinical time-scales. *Frontiers in physiology*, 2:14, 2011.
- [NMSP11b] S. Niederer, L. Mitchell, N. Smith, and G. Plank. Simulating human cardiac electrophysiology on clinical time-scales. *Front. Physiol.*, 2, 2011.
- [NNN⁺11] D. Nordsletten, S. Niederer, M. Nash, P. Hunter, and N. Smith. Coupling multi-physics models to cardiac mechanics. *Progress in biophysics and molecular biology*, 104(1-3):77–88, 2011. [4](#)
- [NPR⁺15] U. C. Nguyễn, M. Potse, F. Regoli, M. L. Caputo, G. Conte, R. Murzilli, S. Muzzarelli, T. Moccetti, F. W. Prinzen, R. Krause, and A. Auricchio. An in-silico analysis of the effect of heart position and orientation on the ECG morphology and vectorcardiogram parameters in patients with heart failure and intraventricular conduction defects. *J. Electrocardiol.*, 48:617–625, May 2015.
- [NVKN98] 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. *The Canadian journal of cardiology*, 14(1):123–134, 1998. [14](#)
- [O’H87] A. O’Hagan. Monte Carlo is Fundamentally Unsound. *Journal of the Royal Statistical Society. Series D (The Statistician)*, 36(2/3):247–249, 1987.

- [O'H91] A. O'Hagan. Bayes-Hermite quadrature. *Journal of statistical planning and inference*, 29(3):245–260, 1991.
- [ost14] Multiphysics object oriented simulation environment (moose), 2014.
- [Pag17] S. Pagani. Reduced-order models for inverse problems and uncertainty quantification in cardiac electrophysiology. 2017. [24](#), [101](#)
- [PBC05] A. Pullan, M. L. Buist, and L. K. Cheng. *Mathematically modelling the electrical activity of the heart: from cell to body surface and back again*. World Scientific Publishing Company, 2005. [16](#), [17](#)
- [PDR⁺06a] M. Potse, B. Dubé, J. Richer, A. Vinet, and R. M. Gulrajani. A comparison of monodomain and bidomain reaction-diffusion models for action potential propagation in the human heart. *IEEE Transactions on Biomedical Engineering*, 53(12):2425–2435, 2006. [18](#), [128](#)
- [PDR⁺06b] M. Potse, B. Dubé, J. Richer, A. Vinet, and R. M. Gulrajani. A comparison of monodomain and bidomain reaction-diffusion models for action potential propagation in the human heart. *IEEE Trans. Biomed. Eng.*, 53(12):2425–2435, 2006. [25](#)
- [Pet14] M. Peters. *Numerical methods for boundary value problems on random domains*. PhD thesis, University of Basel, 2014. [51](#), [53](#), [54](#), [57](#), [59](#)
- [PHS16] S. Pezzuto, J. Hake, and J. Sundnes. Space-discretization error analysis and stabilization schemes for conduction velocity in cardiac electrophysiology. *Int. J. Numer. Meth. Biomed. Engng.*, 2016.
- [PKB⁺12] M. Potse, D. Krause, L. Bacharova, R. Krause, F. W. Prinzen, and A. Auricchio. Similarities and differences between electrocardiogram signs of left bundle-branch block and left-ventricular uncoupling. *Europace*, 14(suppl_5):v33–v39, 2012.
- [PKK⁺14] M. Potse, D. Krause, W. Kroon, R. Murzilli, S. Muzzarelli, F. Regoli, E. Caiani, F. W. Prinzen, R. Krause, and A. Auricchio. Patient-specific modelling of cardiac electrophysiology in heart-failure patients. *Europace*, 16:iv56–iv61, 2014.

- [PKN17] M. Pisaroni, S. Krumscheid, and F. Nobile. Quantifying uncertain system outputs via the multilevel monte carlo method—part i: Central moment estimation. Technical report, 2017.
- [PKP⁺17] S. Pezzuto, P. Kal’avský, M. Potse, F. W. Prinzen, A. Auricchio, and R. Krause. Evaluation of a rapid anisotropic model for ecg simulation. *Front. Physiol.*, 8:265, 2017.
- [Poz19] S. Pozzi. *Enhancing multi-scale cardiac simulations by coupling electrophysiology and mechanics: a flexible high performance approach to cardiac electromechanics*. PhD thesis, Università della Svizzera italiana, 2019. [14](#), [25](#)
- [PP02] A. Papoulis and S. U. Pillai. *Probability, Random Variables and Stochastic Processes*. McGraw-Hill Europe, Boston, Mass., 4th edition edition, January 2002.
- [PQP19] S. Pezzuto, A. Quaglino, and M. Potse. On sampling spatially-correlated random fields for complex geometries. In *International Conference on Functional Imaging and Modeling of the Heart*, pages 103–111. Springer, 2019. [133](#)
- [PT12] L. F. Pavarino and A. Toselli. *Recent developments in domain decomposition methods*, volume 23. Springer Science & Business Media, 2012. [28](#)
- [PTH02a] A. J. Pullan, K. A. Tomlinson, and P. J. Hunter. A finite element method for an eikonal equation model of myocardial excitation wavefront propagation. *SIAM Journal on Applied Mathematics*, 63(1):324–350, 2002. [116](#)
- [PTH02b] A. J. Pullan, K. A. Tomlinson, and P. J. Hunter. A finite element method for an eikonal equation model of myocardial excitation wavefront propagation. *SIAM J. Appl. Math.*, 63(1):324–350, 2002.
- [PVG⁺11] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

- [PWG16] B. Peherstorfer, K. Willcox, and M. Gunzburger. Optimal model management for multifidelity monte carlo estimation. *SIAM Journal on Scientific Computing*, 38(5):A3163–A3194, 2016.
- [PWG17] B. Peherstorfer, K. Willcox, and M. Gunzburger. Survey of multifidelity methods in uncertainty propagation, inference, and optimization. *SIAM Rev.*, 2017.
- [PWGW12] A. Popp, B. I. Wohlmuth, M. W. Gee, and W. A. Wall. Dual quadratic mortar finite element methods for 3d finite deformation contact. *SIAM Journal on Scientific Computing*, 34(4):B421–B446, 2012. [94](#)
- [QBBK17] A. Quaglino, S. Ben Bader, and R. Krause. Sloth: A Python library for UQ. software, 2017. [121](#)
- [QMV⁺19] A. Quarteroni, A. Manzoni, C. Vergara, et al. *Mathematical modelling of the human cardiovascular system: data, numerical approximation, clinical applications*, volume 33. Cambridge University Press, 2019. [11](#)
- [QPK⁺18] A. Quaglino, S. Pezzuto, P.-S. Koutsourelakis, A. Auricchio, and R. Krause. Fast uncertainty quantification of activation sequences in patient-specific cardiac electrophysiology meeting clinical time constraints. *International journal for numerical methods in biomedical engineering*, 34(7):e2985, 2018.
- [QPK19] A. Quaglino, S. Pezzuto, and R. Krause. High-dimensional and higher-order multifidelity Monte Carlo estimators. *Journal of Computational Physics*, 388:300–315, 2019.
- [QSS08] A. Quarteroni, R. Sacco, and F. Saleri. *Méthodes Numériques: Algorithmes, analyse et applications*. Springer Science & Business Media, 2008.
- [QSS10] A. Quarteroni, R. Sacco, and F. Saleri. *Numerical mathematics*, volume 37. Springer Science & Business Media, 2010. [48](#)
- [Qua15] A. Quarteroni. *Modeling the heart and the circulatory system*, volume 14. Springer, 2015. [4](#)
- [QV08] A. Quarteroni and A. Valli. *Numerical approximation of partial differential equations*, volume 23. Springer Science & Business Media, 2008. [48](#)

- [RA06] J. M. Reinert and G. E. Apostolakis. Including model uncertainty in risk-informed decision making. *Ann. Nucl. Energy*, 33(4):354–369, 2006.
- [Ras03] C. E. Rasmussen. Gaussian processes in machine learning. In *Summer school on machine learning*, pages 63–71. Springer, 2003. [62](#)
- [RC11] G. Richardson and S. J. Chapman. Derivation of the bidomain equations for a beating heart with a general microstructure. *SIAM J. Appl. Math.*, 71(3):657–675, 2011.
- [RCG⁺90] R. Rasmusson, J. Clark, W. Giles, K. Robinson, R. Clark, E. Shibata, and D. Campbell. A mathematical model of electrophysiological activity in a bullfrog atrial cell. *American Journal of Physiology-Heart and Circulatory Physiology*, 259(2):H370–H389, 1990. [14](#)
- [RG03] C. E. Rasmussen and Z. Ghahramani. Bayesian Monte Carlo. In *Advances in Neural Information Processing Systems 15*. MIT Press, 2003.
- [RGJ⁺04] C. Ramanathan, R. N. Ghanem, P. Jia, K. Ryu, and Y. Rudy. Noninvasive electrocardiographic imaging for cardiac electrophysiology and arrhythmia. *Nature medicine*, 10(4):422–428, 2004. [21](#)
- [RMRM02] R. Ram-Mohan and L. R. Ram-Mohan. *Finite element and boundary element applications in quantum mechanics*, volume 5. Oxford University Press on Demand, 2002. [28](#)
- [RV19] L. Ratti and M. Verani. A posteriori error estimates for the monodomain model in cardiac electrophysiology. *Calcolo*, 56(3):1–33, 2019. [49](#)
- [RW06] C. Rasmussen and C. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006. [63](#)
- [SDM⁺17] C. Sánchez, G. D’Ambrosio, F. Maffessanti, E. Caiani, F. Prinzen, R. Krause, A. Auricchio, and M. Potse. Sensitivity analysis of ventricular activation and electrocardiogram in tailored models of heart-failure patients. *Med Biol Eng Comput.*, 2017.
- [Sha18] I. Shames. *Energy and finite element methods in structural mechanics*. Routledge, 2018.

- [Sie15] M. Siebenmorgen. *Quadrature methods for elliptic PDEs with random diffusion*. PhD thesis, University_of_Basel, 2015. 95
- [SRE⁺14] R. Speck, D. Ruprecht, M. Emmett, M. Bolten, and R. Krause. A space-time parallel solver for the three-dimensional heat equation. *Parallel Computing: Accelerating Computational Science and Engineering (CSE)*, 25:263–272, 2014.
- [SSP⁺69] D. D. Streeter, H. M. Spotnitz, D. P. Patel, J. Ross, and E. H. Sonnenblick. Fiber orientation in the canine left ventricle during diastole and systole. *Circ. Res.*, 24(3):339–347, 1969.
- [ST06] C. Schwab and R. A. Todor. Karhunen–loève approximation of random fields by generalized fast multipole methods. *Journal of Computational Physics*, 217(1):100–122, 2006. 59
- [SW99] W. Strunk and E. B. White. *The Elements of Style*. Longman Publishers, 4th edition, 1899.
- [Tra11] N. A. Trayanova. Whole-heart modeling: applications to cardiac electrophysiology and electromechanics. *Circulation research*, 108(1):113–128, 2011. 4
- [TSGU13] A. L. Teckentrup, R. Scheichl, M. B. Giles, and E. Ullmann. Further analysis of multilevel Monte Carlo methods for elliptic pdes with random coefficients. *Numerische Mathematik*, 125(3):569–600, 2013.
- [tTNNP04] K. H. W. J. ten Tusscher, D. Noble, P. J. Noble, and A. V. Panfilov. A model for human ventricular tissue. *Am. J. Physiol. Heart Circ. Physiol.*, 286:H1573–H1589, 2004.
- [tTP06a] K. H. W. J. ten Tusscher and A. V. Panfilov. Alternans and spiral breakup in a human ventricular tissue model. *American Journal of Physiology-Heart and Circulatory Physiology*, 291(3):H1088–H1100, sep 2006.
- [TTP06b] K. H. Ten Tusscher and A. V. Panfilov. Alternans and spiral breakup in a human ventricular tissue model. *American Journal of Physiology-Heart and Circulatory Physiology*, 291(3):H1088–H1100, 2006. 14

- [Tun78] L. Tung. *A bi-domain model for describing ischemic myocardial dc potentials*. PhD thesis, Massachusetts Institute of Technology, 1978. [16](#), [17](#)
- [TW06] A. Toselli and O. Widlund. *Domain decomposition methods-algorithms and theory*, volume 34. Springer Science & Business Media, 2006. [28](#)
- [VKS⁺14] C. T. Villongco, D. E. Krummen, P. Stark, J. H. Omens, and A. D. McCulloch. Patient-specific modeling of ventricular activation pattern using surface ECG-derived vectorcardiogram in bundle branch block. *Progress in Biophysics and Molecular Biology*, 115(2-3):305–313, 2014.
- [vOOvD11] A. van Oosterom, T. F. Oostendorp, and P. M. van Dam. Potential applications of the new ECGSIM. *J. Electrocardiol.*, 44:577–583, 2011.
- [Wal87] A. D. Waller. A demonstration on man of electromotive changes accompanying the heart’s beat. *The Journal of physiology*, 8(5):229–234, 1887. [2](#)
- [Wan02] X. Wang. A constructive approach to strong tractability using quasi-monte carlo algorithms. *Journal of Complexity*, 18(3):683–701, 2002. [85](#)
- [Woh00] B. I. Wohlmuth. A mortar finite element method using dual spaces for the lagrange multiplier. *SIAM journal on numerical analysis*, 38(3):989–1012, 2000. [94](#)
- [XWZ15] Y. Xia, K. Wang, and H. Zhang. Parallel Optimization of 3D Cardiac Electrophysiological Model Using GPU. *Comput. Math. Methods Med.*, 2015, 2015.
- [Zar68] S. Zaremba. Some applications of multidimensional integration by parts. In *Annales Polonici Mathematici*, volume 1, pages 85–96, 1968. [83](#)
- [ZKN⁺16] P. Zulian, A. Kopaničáková, M. C. G. Nestola, A. Fink, N. Fadel, V. Magri, T. Schneider, E. Botter, J. Mankau, and R. Krause. Utopia: A C++ embedded domain specific language for scientific computing software, 2016. [121](#)

- [ZT00] O. C. Zienkiewicz and R. L. Taylor. *The finite element method, vol. 2*. Butterworth-Heinemann, 2000. [28](#)
- [ZZ88] J. Zhu and O. Zienkiewicz. Adaptive techniques in the finite element method. *Communications in applied numerical methods*, 4(2):197–204, 1988. [28](#)