Multi-fidelity modeling of atrial fibrillation

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
Lia Gander

under the supervision of
Prof. Rolf Krause and Prof. Simone Pezzuto

June 2023

Dissertation Committee

Prof. Cesare Alippi	Università della Svizzera Italiana, Switzerland
Prof. Yves Coudière	Université de Bordeaux, France
Prof. Michael Multerer	Università della Svizzera Italiana, Switzerland
Prof. Francisco Sahli Costabal	Pontificia Universidad Católica de Chile, Chile

Dissertation accepted on 27 June 2023

Research Advisor	Co-Advisor
Prof. Rolf Krause	**Prof. Simone Pezzuto**

PhD Program Director

The PhD program Director *pro tempore*

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.

Lia Gander
Lugano, 27 June 2023

Abstract

Clinical investigations on atrial fibrillation, the most common cardiac arrhythmia, are based on computational models of the cardiac electrophysiology. Accurate high fidelity models are costly and have limited applicability in the clinical setting, where patient-specific *in-silico* assessments need to meet the practical time constraints. To speed up the computations, less accurate but faster low fidelity models could be employed. In this work, we develop two low fidelity models for atrial fibrillation that approximate the monodomain high fidelity model. One low fidelity model is based on the eikonal model, that we adapt to handle the re-entries that characterize atrial fibrillation. This model includes the restitution properties computed from the high fidelity model and handles anisotropy. The other low fidelity model is based on a coarser discretization of the computational domain. In this model the coarsening of the underlying atrial model is properly defined and the conduction velocity is adjusted. We assess the similarity of the low fidelity approximations to the high fidelity results in quantitative and qualitative studies and we explain the discrepancies. The accuracy of the low fidelity models depends on the considered metric. Here the main focus is on the inducibility of atrial fibrillation. The characterization of the atrial regions where the arrhythmia can be induced provides information for the evaluation and the design of the ablation treatment. Personalized inducibility maps can be obtained with a multi-fidelity method. The multi-fidelity approach, in which the high fidelity model and a correlated low fidelity model are combined, leads to a speed-up compared to single-fidelity approaches.

The eikonal low fidelity model shows a poor agreement to the high fidelity model in terms of atrial fibrillation inducibility. However, its qualitative accuracy in simpler numerical experiments and its very low cost (potentially real-time) make it attractive. This motivates our future interest in optimizing the implementation and in reducing the discrepancies to the high fidelity model. Instead, the low fidelity model based on a coarser discretization highly agrees to the high fidelity model on the atrial fibrillation inducibility. Moreover, thanks to this high correlation, it performs well in the multi-fidelity framework.

Acknowledgements

I thank Prof. Rolf Krause for giving me the opportunity to do this PhD and for his supervision. Many thanks to Prof. Simone Pezzuto for his supervision and for always being available to help. During these four years, the discussions with him were always fruitful and I learned a lot from him. Thanks to Prof. Michael Multerer for his help during the first part of my PhD in the collaboration that lead to my first publication (Gander et al. [2021]), which is not part of this thesis. Thanks to Prof. Francisco Sahli Costabal for his help and collaboration in the second part of my PhD. Thanks also to all the other co-authors for their work. I thank all my colleagues for their advices and for the nice discussions.

Ringrazio i miei genitori Lucia e Marco per tutto quello che hanno fatto per me e per avermi sempre sostenuta. Ringrazio il mio compagno Cyrill per il suo sostegno e per avermi spesso ricordato che lui è un dottore mentre io non lo sono. Grazie agli amici Aline, Sofia, Luca e Paride per i bei momenti di svago. Infine, grazie agli amici del Tennis Gordola per tutti gli allenamenti, le sfide e, soprattutto, i momenti di divertimento.

Contents

Chapter 1

Introduction

The heart is an organ whose function is to pump the blood around the body. This is achieved by a cyclical alternation of contraction and relaxation movements. The haemodynamics and the mechanics of the heart are ruled by its electrophysiological activity, which is determined by a propagating action potential. The heart is composed of four chambers, two atria in the upper part and two ventricles in the lower part. In normal conditions of sinus rhythm, the action potential initiated by a group of pacemaker cells on the right atrium propagates to the left atrium and to the ventricles (Colli Franzone et al. [2014]). In pathological conditions of arrhythmias, the electrical activity of the heart is irregular. There are various cardiac arrhythmias and the range of complications is wide. In this work we focus on atrial fibrillation (AF), the most common cardiac arrhythmia and a significant contributor to morbidity and mortality (Tsao et al. [2022]) due to the consequent cardiac dysfunctions and the associated risk of causing strokes (Schotten et al. [2011]).

AF is characterized by a chaotic electrical activity of the atria, self-sustained by propagating re-entrant waves in the atrial tissue. AF is triggered by the so-called ectopic foci, i.e. cells that generate an abnormal electrical stimulation that hinders and perturbs the normal sinus rhythm. The occurrence of AF is related to the electric and structural remodeling of the atrial substrate. In particular, the presence of low-conducting fibrotic tissue is a risk factor. Moreover, AF itself increases the ectopic activity and enhances the remodeling, thus increasing the frequency and the duration of AF events (Schotten et al. [2011]). The most common treatment for AF is ablation, which consists of isolating some parts of the atrial tissue. If these isolated portions are pro-arrhythmic, ablation reduces the risk of occurrence of AF. In the early stages of AF, the pulmonary veins isolation is the standard treatment (Chen et al. [1999]) as the ectopic foci are located in the pulmonary veins (Haïssaguerre et al. [1998]). However, due to the progressive nature of AF, the efficacy of ablation is currently suboptimal in the later stages of AF. Indeed, there is no ablation strategy optimized to target the regions that reduce the most the risk of AF occurrence (Verma et al. [2015]) as the location of the ectopic foci shifts away from the pulmonary veins (Kawai et al. [2019]). Moreover, the inter-individual variability in atrial geometry and substrate makes the optimization of the ablation treatment even more challenging.

Computational models of the cardiac electrophysiology are widely used to perform *in-silico* patient-specific clinical investigations. The most accurate electrophysiological model is the bidomain model, which consists of two parabolic equations for the spatio-temporal evolution of the action potential, coupled with a system of ODEs modeling the ionic properties of the tissue. The propagation of the action potential is characterized by a steep and short upstroke that determines the activation of the cells. To capture this upstroke, the numerical solution of the bidomain equation requires small spatial and temporal discretizations. Moreover, the bidomain system is degenerate, thus it is not possible to solve it numerically with fully explicit schemes. These requirements on the discretization parameters and on the numerical method often lead to unbearable computational costs of the bidomain model (Colli Franzone et al. [2014]). It is therefore common to resort to reduced models such as the monodomain model. The monodomain model consists of one parabolic reaction-diffusion equation coupled with the ionic ODE system (Colli Franzone et al. [2014]). The monodomain system is non-degenerate and can be solved numerically with fully explicit methods. The monodomain model is computationally feasible and is widely used in AF studies, e.g. to understand its inducibility and mechanisms (Potse et al. [2018], Potse [2019], Gharaviri et al. [2020], Gharaviri et al. [2021b], Gharaviri et al. [2021a]) or to design and evaluate ablation strategies (McDowell et al. [2015], Boyle et al. [2019]). With appropriate spatial and temporal discretizations, the monodomain model is an accurate approximation of the bidomain model (Potse et al. [2006], Nagel et al. [2023]). Thus, in this work, we refer to the monodomain model with sufficiently small discretization parameters as the high fidelity model. However, in order to capture the upstroke in the action potential, the conditions on the discretization parameters remain. Therefore it is computationally demanding to employ the high fidelity model. The computational cost of the high fidelity model often affects its applicability to personalized clinical applications in which the results need to be delivered within practical time constraints.

Many ways to speed up the computations have been proposed. One option is to rely on a faster implementation of the high fidelity model, e.g. based on

GPUs (Kaboudian et al. [2019]). Alternatively, one could employ a low fidelity model. The low fidelity models are less accurate, but faster, than the high fidelity model. Several low fidelity models are available. First, possible low fidelity models are given by approximate models based on simplified physics, such as the eikonal models (Neic et al. [2017], Loewe et al. [2019]). The eikonal models describe the activation times on the computational domain. The pure eikonal model (Colli Franzone et al. [1990a]) does not capture the effects of the wavefront curvature on the action potential propagation. The eikonal-diffusion (Colli Franzone et al. [1990b]) and the eikonal-curvature (Keener [1991]) models are extensions of the pure eikonal model that include the curvature effects. Contrary to the eikonal-curvature model, the eikonal-diffusion model also captures the boundary and the front collision effects on the action potential propagation (Pullan et al. [2002]). The eikonal models are computationally cheap since there is no time component in the underlying space dependent equations. Moreover, as the eikonal equations describe the time when the upstroke in the action potential occurs instead of the action potential itself, the eikonal models allow for a larger spatial discretization compared to the monodomain model. However, the coarse discretization and the fact that the eikonal approximations do not model the ionic properties of the cardiac tissue might lead to imprecise simulations of the action potential propagation. The reaction-eikonal equations (Neic et al. [2017]) allow to include the ionic model into a system that can be solved numerically with less computational resources compared to the monodomain system. This is achieved by approximating the diffusive term of the monodomain equation with a term based on the eikonal activation times. The reaction-eikonal models are thus more accurate than the other eikonal models. However, since the reaction-eikonal equations involve a temporal component and the ionic model, the reaction-eikonal models are computationally more expensive than the other eikonal models. Second, another possible low fidelity model can be obtained by simply using a coarser discretization of the computational domain in the numerical solution of the high fidelity equations (Dhamala et al. [2020], Pagani et al. [2021]). When explicit numerical schemes are used to solve the high fidelity equations, the larger spatial resolution is complemented by a larger time step. The coarser discretization parameters might lead to an imprecise description of the action potential propagation, but allow to reduce the computational cost. Finally, another possibility for low fidelity models is provided by reduced-order modeling (Fresca et al. [2020], Cicci et al. [2022]). Model reduction is achieved by projecting the high fidelity model equations or solution onto a basis with lower dimension compared to the original basis.

In this work, the main focus is on the inducibility of AF. In computational studies, the AF inducibility is often evaluated from the outcome of a tissue stimulation by a pacing protocol (McDowell et al. [2015], Potse et al. [2018], Potse [2019], Boyle et al. [2019], Gharaviri et al. [2021b], Gharaviri et al. [2021a]). The pacing protocol consists of a train of stimuli delivered from a given a pacing site. Depending on the pacing site, the stimulation might induce AF or not. Specifically, if the pacing results in a vanishing wave, then the induction of AF is not successful. Instead, if the pacing results in a self-sustained activity, then the induction of AF is successful. Depending on the result of the stimulation (if AF is induced or not), one can associate a class to each location of the atria. This binary classification defines the inducibility map of the atria, i.e. a map associating the class labels to the points on the atrial domain. The inducibility map provides useful information for designing and evaluating ablation strategies. Indeed, given a patient-specific atrial model, the inducibility map shows the regions that trigger AF, which might be the target of a personalized ablation treatment. Moreover, the inducibility map can be reduced into the metric defined as the fraction of tissue where AF is inducible. Thus, in the case of several ablation therapies being available, the inducibility metric allows to determine which therapy is the one that mostly reduces the risk of occurrence of AF.

The exploration of the whole atrial domain to obtain the inducibility map is unfeasible. Therefore, the inducibility map must be obtained from few local evaluations of the inducibility. The classification output is computed only at the locations included in the so-called training set. For the pacing sites of the training set, the inducibility is evaluated from the outcome of the simulation of the electrical activity during and after the pacing protocol. For all the other locations of the atria, the class label is estimated from the training data with a classification method (Neal [1998], Rasmussen and Williams [2006], Nickisch and Rasmussen [2008]). In the patient-specific case, the inducibility map needs to be obtained within clinical time constraints and, therefore, the computational budget is restricted. As a consequence, the number of high fidelity simulations is limited and the accuracy of the classification in the whole atrial domain might be affected by the low amount of training data. Moreover, the selection of the pacing sites to include in the training set might also affect the accuracy of the classification (Boyle et al. [2021]). Compared to single-fidelity methods, multi-fidelity approaches can reach higher accuracy with a fixed computational cost or, equivalently, can reach a target accuracy with a lower computational cost (Kennedy and O'Hagan [2000], Perdikaris et al. [2016], Quaglino et al. [2018], Quaglino et al. [2019], Sahli Costabal et al. [2019]). The idea is to combine training data obtained from the high fidelity model and training data obtained from a low fidelity model. The speedup is achieved by offsetting most of the computational burden to the cheap low fidelity model, thus allowing a wider inspection of the atrial domain. The improved accuracy is achieved if the low fidelity model is sufficiently correlated to the high fidelity model. Note that our approach to obtain personalized inducibility maps is to train a classifier for each patient-specific atrial anatomy. An alternative, simulation-free, approach is to train a classifier that predicts the inducibility outputs from the patient-specific atrial anatomy itself (Sahli Costabal et al. [2023]).

Gaussian processes (GPs) perform well in predicting classification and regression outputs when limited data is available. Moreover, GPs are of Bayesian type and include the uncertainty in the predictions (Neal [1998], Rasmussen and Williams [2006]). GP classification is more challenging than GP regression since there is no closed expression of the likelihood and some approximations are needed to obtain predictions (Nickisch and Rasmussen [2008]). Another challenging aspect is the extension of the GP classification and regression from Euclidean spaces to Riemannian manifolds such as the atrial surface. Indeed, GPs depend on covariance functions that need to be positive semi-definite. The covariance functions are typically based on the concept of distance and naively accounting for the geodesic distance instead of the Euclidean distance does not guarantee the positive semi-definiteness property (Pezzuto et al. [2019]). The most used covariance function in the machine learning community is the Matérn kernel (Rasmussen and Williams [2006]). The extension of Matérn GPs to Riemannian manifolds has been introduced in Borovitskiy et al. [2020]. In the cardiac modeling literature, GP classification has been performed on Euclidean spaces (Sahli Costabal et al. [2019], Sahli Costabal et al. [2020]), whereas GP regression has also been performed on manifolds (Coveney et al. [2019], Coveney et al. [2020]). The extension of the GP prediction to the multi-fidelity framework has been introduced in Kennedy and O'Hagan [2000]. Multi-fidelity GP classification has been performed in Sahli Costabal et al. [2019] on Euclidean spaces. The methodology presented in Gander et al. [2022] extends the multi-fidelity GP classification to manifolds. This methodology is a novel procedure that can be applied to estimate atrial inducibility maps. This GP classification method takes advantage of the low fidelity models and incorporates the topological properties of personalized atrial anatomies.

In this work, the goal is to develop novel low fidelity models for atrial fibrillation, either based on the eikonal model or based on a coarser discretization of the computational domain. We assess the quality of the low fidelity approximations in quantitative studies and in numerical experiments. We also provide explanations for the discrepancies compared to the high fidelity results. As a novelty, the numerical experiments include an assessment of the correlation between the high and the low fidelity models when used to evaluate the AF inducibility. If the correlation is adequate, we employ the low fidelity model in the multi-fidelity GP classification method of Gander et al. [2022] to estimate inducibility maps. We then compare the resulting estimate to the estimates resulting from single-fidelity approaches in terms of accuracy and computational cost. The goal is to understand which strategies reach higher accuracies with given computational cost.

Among the eikonal models we choose the pure eikonal model. We exclude the reaction-eikonal model because it is computationally more expensive than the other eikonal models and here we want to focus on cheap low fidelity models. The other eikonal models describe the activation times of a single activation. During AF events the electric activity is characterized by re-entries that imply several activations. In particular, the cardiac cells alternate between the depolarized, the repolarized and the excitable states. The cells are initially in the excitable state, meaning that they can activate. If the activation occurs, the cells become depolarized. When the activation is over, the cells move to the repolarized state until they become excitable again. Therefore, to apply the eikonal models to the simulation of AF, we have to include the re-excitability of the tissue in order to model the activation times of multiple activations. The pure eikonal equation can be solved by algorithms that iteratively pass through the nodes of the computational mesh and update the corresponding activation times. These algorithmic solution methods are suitable for adaptations such as the inclusion of the re-excitability. The iterative algorithms could, in principle, be applied also to solve the eikonal-curvature equation. However, at each iteration, the solution would need an estimate of the wavefront curvature. Instead, the iterative algorithms can not be applied to solve the eikonal-diffusion equation, because of the global coupling due to the diffusion term. These facts determine our focus on the pure eikonal model. The most common iterative algorithms for the solution of the pure eikonal equation are the fast marching method (FMM) (Kimmel and Sethian [1998], Sethian and Vladimirsky [2000]) and the fast iterative method (FIM) (Fu et al. [2011]). Contrary to the FIM, the FMM is a single-pass algorithm that guarantees the monotone acceptance of the activation times. These properties are suitable for the adaptation to account for the re-excitability and determine our choice of the FMM. The FMM is a Dijkstra-like method that considers updates coming from the elements of the mesh, contrary to the standard Dijkstra's method (Dijkstra [1959]) that only allows updates from the edges. In this work, we adapt the FMM algorithm to take into account the re-excitability and we compare the resulting eikonal low fidelity model to the monodomain high fidelity model in some 2D numerical experiments and in terms of AF inducibility on a 3D atrial geometry. This task has many challenging aspects. First, the restitution properties describing the re-excitability need to be incorporated in the algorithm. These restitution properties include the description of the action potential duration (APD) and of the conduction velocity (CV) as a function of the diastolic interval (DI). The APD is the time between the activation and the repolarization. The DI is the time between the repolarization and the next activation. Here, in order to maximize the similarity to the high fidelity model, we use the restitution properties of the monodomain model, that we need to compute. Second, the influence of heterogeneities in the tissue conductivity on the eikonal approximation has to be analyzed, since the pure eikonal model does not account for the resulting diffusion currents. To this end, we perform a quantitative comparison between the monodomain and the eikonal propagations in presence of discontinuities in conductivity due to the presence of fibrosis. Third, in order to guarantee the convergence of the FMM to the viscosity solution of the eikonal equation, the computational mesh needs to satisfy an acuteness condition (Mirebeau [2012], Mirebeau [2014]). It is not straightforward to satisfy this requirement, mostly in the anisotropic case, where the acuteness is intended with respect to the metric defined by the fibers direction and the anisotropy ratio. Here we adapt the computational mesh to the metrics using the Mmg software (Dobrzynski and Frey [2008]). Fourth, if the computational mesh of the eikonal low fidelity model is coarser than the computational mesh of the monodomain high fidelity model, then a proper mapping of the anatomical features from the monodomain mesh to the eikonal mesh needs to be defined. Finally, a pacing protocol needs to be designed to evaluate the AF inducibility with the eikonal low fidelity model. Here we follow the idea introduced in Azzolin et al. [2021], which consists of stimulating as soon as the pacing site becomes re-excitable.

In the literature, fast simulations of re-entries have mostly been carried out using reaction-eikonal models or combinations of the monodomain model and an eikonal model. In Jacquemet [2010] and Jacquemet [2012], re-entries along prescribed pathways are generated by combining the monodomain model with a generalization of the eikonal-diffusion model that handles re-entries. In Corrado and Zemzemi [2018], a reaction-eikonal model is used to simulate a train of stimuli. The same reaction-eikonal model is used in Gassa et al. [2021] to simulate re-entries induced by S1-S2 stimulations, i.e. stimulations at two different locations. In Barrios Espinosa et al. [2022], the monodomain model and the pure eikonal model solved with the FIM are alternated to generate a re-entry on a ring with the S1-S2 stimulation. In few works the solution algorithm for the pure eikonal model has been modified to handle re-entries. Our approach is similar to the one presented in Sermesant et al. [2007] and in Pernod et al. [2011], where the FMM is combined with a time-stepping to account for the re-excitability. In Sermesant et al. [2007] the modified algorithm is only used to simulate a train of stimuli. In Pernod et al. [2011] an interactive simulator that allows to design and test ablation strategies is presented. There the authors also show an example of a re-entry generated by a periodic pacing in presence of scars and isthmus, i.e. respectively non-conducting and low-conducting tissue. A similar approach has also been presented in Corrado and Zemzemi [2018], but using the Dijkstra's method. In this work, contrary to Sermesant et al. [2007], Pernod et al. [2011] and Corrado and Zemzemi [2018], the modified algorithm also includes the restitution curves of the APD and the CV computed from the monodomain model combined with the ionic model of the human atria. Moreover, here we take care of the convergence of the FMM to the viscosity solution of the eikonal equation. Furthermore, in this work we present the first attempt to evaluate the AF inducibility from a pacing protocol using the pure eikonal model.

The comparison in terms of AF inducibility between the high fidelity model and the low fidelity model based on a coarser discretization of the computational domain is less complex. Indeed, the low fidelity model based on coarser discretization parameters includes the restitution properties, it deals with the conduction heterogeneity and the anisotropy in the same manner as the high fidelity model and it inherits the pacing protocol from the high fidelity model. However, some aspects require some care. The computational mesh of the monodomain model is coarser in the low fidelity case than in the high fidelity case. Therefore, we have to define a proper mapping of the anatomical features from the fine to the coarse mesh. Moreover, the coarser spatial discretization in the low fidelity model affects the CV (Pezzuto et al. [2016]). Thus, in order to maximize the correlation between the two models, we adjust the CV of the low fidelity model to match the CV of the high fidelity model.

This work is organized as follows. The main aspects of cardiac electrophysiology and AF are introduced in Chapter 2. The high fidelity model of AF is introduced in Chapter 3. The eikonal low fidelity model with re-excitability and several tests of its similarity to the high fidelity model are presented in Chapter 4. The low fidelity model based on coarser discretization parameters and a test of its correlation to the high fidelity model in terms of AF inducibility are presented in Chapter 5. The multi-fidelity classification method to predict AF inducibility maps is presented in Chapter 6, together with the study of its accuracy and cost compared to single-fidelity approaches. Conclusions are drawn in Chapter 7.

Some of the numerical results presented in this work are visualized in the videos available at `https://drive.google.com/drive/folders/1bAEHYhed8_2gKVPz8k5rLoqZP4YUWPEl?usp=sharing`.

Chapter 2

Cardiac electrophysiology and atrial fibrillation

The heart is an organ composed of four chambers, namely two atria in the upper part and two ventricles in the lower part. Figure 2.1 shows the cardiac anatomy. The cardiac function is to pump blood around the body. In particular, the right chambers pump the deoxygenated blood through the pulmonary artery to the lungs, where it gets oxygenated, and back to the heart through the pulmonary veins. The left chambers then pump the oxygenated blood through the aorta around the body, where it deoxygenates, and back to the heart through the ve-

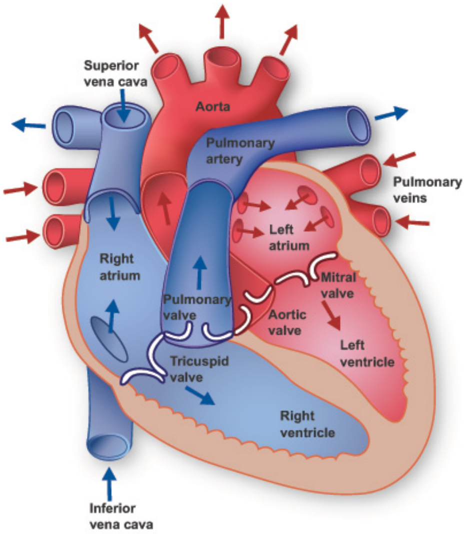

Figure 2.1. Cardiac anatomy and blood circulation. Source: `https://www.texasheart.org/heart-health/heart-information-center/topics/heart-anatomy`.

nae cavae (Colli Franzone et al. [2014]). The blood circulation in the heart is illustrated in Figure 2.1. The cardiac function is ruled by a cyclical alternation of contraction and relaxation movements. The haemodynamics and the mechanics of the heart are regulated by its electrophysiological activity (Colli Franzone et al. [2014]).

In this chapter we introduce the cardiac electrophysiology and the pathological condition of AF. The cardiac electrophysiology is presented in Section 2.1 and AF is presented in Section 2.2.

2.1 Cardiac electrophysiology

In this section we present the main aspects of the cardiac electrophysiology. At the cellular level, the electrical activity is ruled by ion fluxes between the intracellular and the extracellular spaces, across the ion channels located on the cellular membrane. The flow of ionic currents leads to an action potential, that is a strong variation in the transmembrane potential, i.e. the difference between the intracellular and the extracellular potentials (Colli Franzone et al. [2014]). The action potential consists of the five phases shown in Figure 2.2. The action

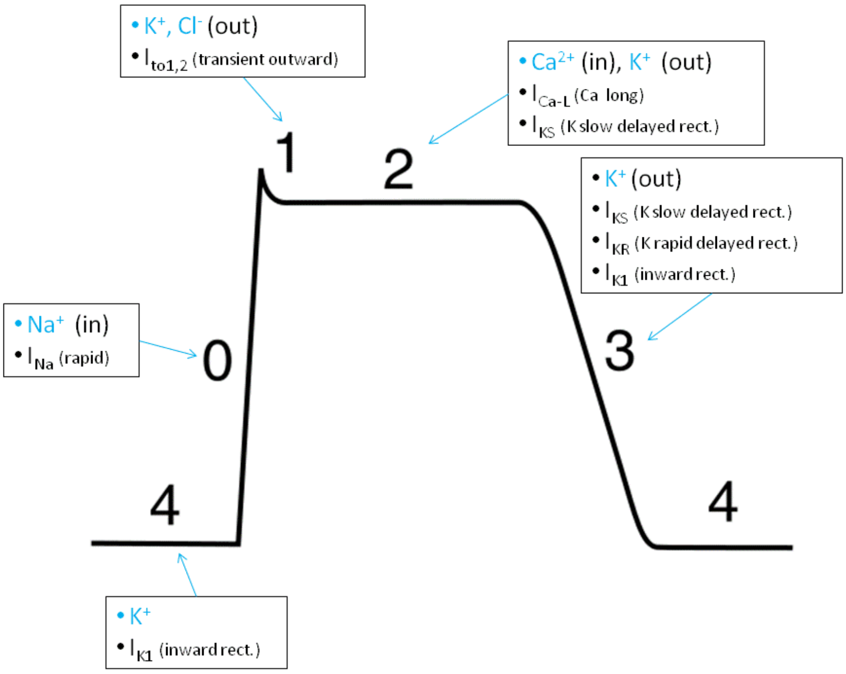

Figure 2.2. Action potential phases. Source: `https://en.wikipedia.org/wiki/Cardiac_action_potential`.

potential is determined by the Na^+, K^+, Cl^- and Ca^{2+} ions carried by several ionic currents. Phase 0 is the depolarization, during which the excitable cell activates and the transmembrane potential rapidly increases. The upstroke is due to the inflow of Na^+ ions carried by the rapid I_{Na} current. Phase 1 is a small downward deflection of the transmembrane potential. The deflection is due to the outflow of K^+ and Cl^- ions, respectively carried by the transient outward I_{to1} and I_{to2} currents. Phase 2 is the plateau, during which the transmembrane potential is roughly constant. The plateau is due to the balance between the inflow of Ca^{2+} ions, carried by the L-type I_{CaL} current, and the outflow of K^+ ions, carried by the slow delayed rectifier I_{KS} current. Phase 3 is the repolarization, during which the cell inactivates and the transmembrane potential rapidly decreases. The downstroke is due to a net outward current caused by the outflow of K^+ ions, carried by the slow and the rapid delayed rectifier I_{KS} and I_{KR} currents, and the inflow of K^+ ions, carried by the inward rectifier I_{K1} current. Phase 4 is the resting phase, during which the transmembrane potential takes its resting value and the cell becomes re-excitable. The resting value is maintained by the inward rectifier I_{K1} current. The mentioned ionic currents are those that mostly contribute to an action potential. Other currents also play a role in the cardiac electrophysiology. These include pump currents and exchanger currents (Colli Franzone et al. [2014]).

The excitable cells compose the cardiac tissue and their electric connections allow the propagation of the action potential. The electric connections are called gap junctions and connect the cells longitudinally or transversally. The action potential conduction is anisotropic, as it is faster along the longitudinal direction than along the transversal direction. Most connections are longitudinal and form the so-called cardiac fibers. The fibers show a laminar organization radially between the external part of the chamber walls (the epicardium) and the internal part (the endocardium) (Colli Franzone et al. [2014]). For more details on the cellular organization, we refer to Keener and Sneyd [2009b]. The cardiac conduction system is shown in Figure 2.3. In normal conditions of sinus rhythm, the sinoatrial node, which is located in the right atrium, is a natural pacemaker that initiates and regulates the action potential at each beat. The action potential that originates from the sinoatrial node then propagates through the right atrium and to the left atrium through the Bachmann's bundle. When the action potential reaches the atrioventricular node, which is located at the base of the atria, it propagates through the His bundle and the Purkinje network to the ventricles (Colli Franzone et al. [2014]).

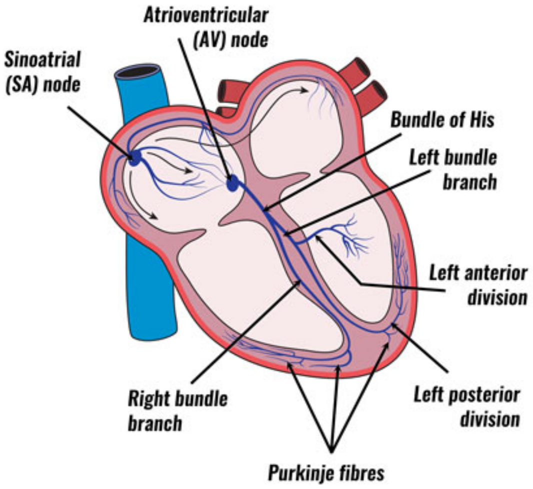

Figure 2.3. Cardiac conduction system. Source: `https://www.teachpe.com/anatomy-physiology/the-heart-conduction-system`.

2.2 Atrial fibrillation

In this section we present the main aspects of AF, the most common cardiac arrhythmia. AF is characterized by an abnormal electrical activity of the atria. We first introduce the atrial anatomy and some specific aspects of the atrial electrophysiology. Then, we move to the pathological condition of AF.

The atrial anatomy is shown in Figure 2.4. The left atrial geometry includes the pulmonary veins and the left atrial appendage (LAA). The right atrial geometry includes the venae cavae and the right atrial appendage (RAA). The coronary sinus and some interatrial connections are also part of the geometry. Moreover, there are some bundles, namely the Bachmann's bundle and the trabecular networks, which are shown in panels E and F of Figure 2.4. In the Bachmann's bundle, the CV along the fibers direction is higher compared to the other regions of the atria. The lines in Figure 2.4 show the fibers orientation. The epicardial fibers are shown in panels A and B, whereas the endocardial fibers are shown in panels C and D. Notice that, since in the region between the inferior and the superior vena cava the conduction is isotropic, the intercaval region has no fiber orientation (Gharaviri et al. [2020]).

The atria and the ventricles show some electrophysiological differences that are mainly due to differences in the density of some ionic currents and in the kinetics of some repolarizing currents. In particular, the plateau phase is less pronounced in the atria than in the ventricles, leading to a shorter APD, which

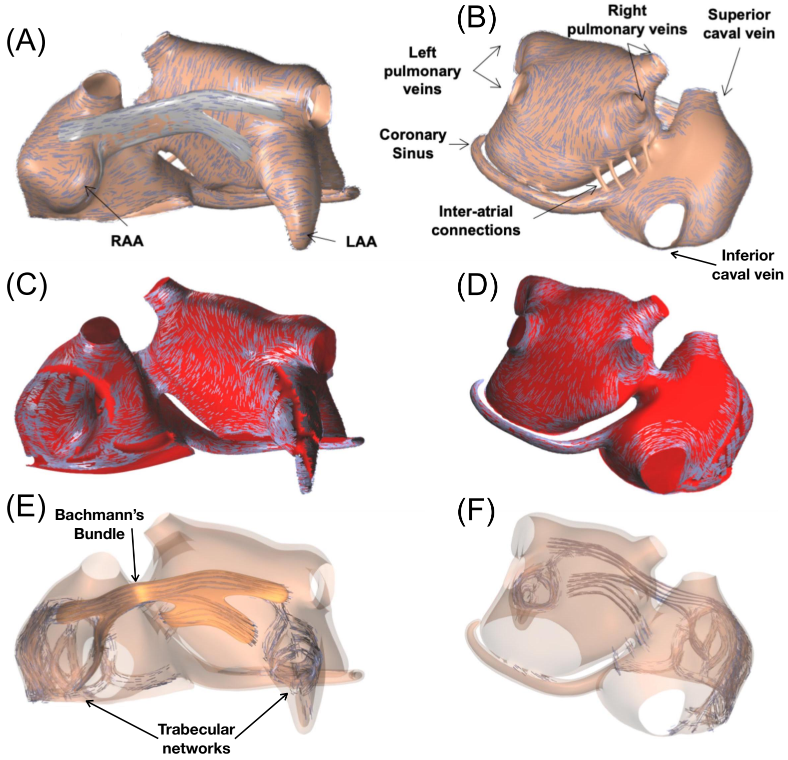

Figure 2.4. Atrial anatomy and fibers orientation. Left column: anterior view. Right column: posterior view. Panels A and B: epicardium. Panels C and D: endocardium. Panels E and F: bundles. Source: Gharaviri et al. [2020].

corresponds to the time between the activation and the repolarization. Moreover, the atrial cells are more excitable than the ventricular cells, which means that the critical size or strength of a stimulus needed to generate an action potential is lower in the atria than in the ventricles (Schotten et al. [2011]).

AF is characterized by a chaotic electrical activity of the atria. AF is triggered by abnormal electrical stimulations that occur away from the sinoatrial node, at the so-called ectopic foci. Under certain circumstances, this ectopic activity alters the normal sinus rhythm and induces a self-sustained electrical activity characterized by re-entrant waves that propagate in the atrial tissue. The ectopic activity is due to abnormal conditions that produce a cellular Ca^{2+} overload. The excess Ca^{2+} is then ejected by the I_{NaCa} exchanger current, which produces an inflow of Na^{+} that initiates an action potential (Schotten et al. [2011]). The induction of an AF event from the ectopic activity needs to be favored by a remodeled substrate. The remodeling is both electric and structural. The electric remodeling consists of a shortening of the effective refractory period (ERP), which is the time between activation and re-excitability. The structural remodeling consists of atrial tissue features that imply a slower and more heterogeneous conduction. The presence of endomysial fibrosis is the main feature of the structural remodeling. The fibrous tissue is located between the fibers, thus hindering the transversal propagation of the action potential and increasing the anisotropy. Another feature of the structural remodeling is the altered connexin expression, which consists of a reduction of the transversal gap junctions that increases the anisotropy. The structural remodeling can be caused by aging or by underlying heart diseases such as heart failure. An altered autonomic activity can be the cause of the electric remodeling and the Ca^{2+} overload (Schotten et al. [2011]).

AF causes the atrial contractile dysfunction, which affects the haemodynamical activity and increases the thromboembolic risk. If a clot forms and then gets pumped to the brain, it can cause a stroke. It is estimated that $20\% - 25\%$ of all strokes are due to AF (Schotten et al. [2011]). Therefore AF is a major contributor to mortality (Tsao et al. [2022]). AF also causes other heart diseases, such as atrial dilatation, cellular hypertrophy and atrial stretch, i.e. atrial pressure and volume overload (Schotten et al. [2011]). Hence AF is also a major contributor to morbidity (Tsao et al. [2022]). Moreover, AF promotes AF by increasing the ectopic activity and by enhancing the remodeling. Indeed, AF itself or its consequences cause cellular Ca^{2+} overload, shortening of the ERP and formation of fibrosis. As a consequence, in patients, the frequency and the duration of AF episodes increases over time. In the early stages of AF (paroxysmal AF) the episodes are rare and self-terminating, whereas in later stages of AF (persistent AF), the episodes are more frequent and sustained (Schotten et al. [2011]).

The most common treatment for AF is ablation, which consists of creating some isolation lines on the atrial tissue. The ablation lines have the effect of isolating some parts of tissue and, if these isolated portions are pro-arrhythmic, ablation reduces the risk of occurrence of AF. In paroxysmal AF, the ectopic foci are located in the pulmonary veins (Haïssaguerre et al. [1998]) and the pulmonary veins isolation is the standard treatment (Chen et al. [1999]). However, in persistent AF, the location of the ectopic foci shifts away from the pulmonary veins (Kawai et al. [2019]) and currently there is no ablation strategy optimized to target the regions that reduce the most the risk of an AF event (Verma et al. [2015]). The efficacy of ablation is thus suboptimal in persistent AF patients. Other treatments for AF such as pharmacological cardioversion and antiarrhythmic drugs also have limited efficacy (Schotten et al. [2011]). Instead, the anticoagulation therapy to reduce the thromboembolic risk is efficient, but it is potentially harmful and it does not prevent AF (Schotten et al. [2011]).

Chapter 3

High fidelity model of atrial fibrillation

In this chapter we present the high fidelity model of the AF electrophysiology. At the cellular level, the electrical activity is described by the ionic model. At the tissue level, the propagation of the action potential is described by the bidomain model, which considers the intracellular and the extracellular domains and models the intracellular and the extracellular potentials. The numerical solution of the equations involved in the bidomain model is excessively computationally demanding. Therefore we must resort to the monodomain model, which is a commonly used approximation of the bidomain model. The monodomain equation models the transmembrane potential, which is defined as the difference between the intracellular and the extracellular potentials (Colli Franzone et al. [2014]). Throughout this work we consider the monodomain model to be our high fidelity model.

In this work we focus on AF, so we present the high fidelity model on the atrial domain. The ionic model of AF is introduced in Section 3.1. The bidomain and the monodomain models are introduced in Sections 3.2 and 3.3, respectively. The pacing protocol to evaluate the AF inducibility is presented in Section 3.4. Our discretized atrial model is presented in Section 3.5 and the numerical solution of the monodomain equation is presented in Section 3.6.

3.1 Ionic model of atrial fibrillation

As described in Section 2.1, the cellular electrophysiological activity is driven by ion flows across the ion channels located on the membrane. In this section we introduce the ionic model that describes these ion flows. We consider the atrial tissue $\Omega \subset \mathbb{R}^3$ and the time interval $(0, T)$. The ionic model describes the ionic current $I_{\text{ion}}(v, \boldsymbol{w}, \boldsymbol{c})\colon \Omega \times (0, T) \to \mathbb{R}$ as a function of the transmembrane potential

$v\colon \Omega\times(0,T)\to\mathbb{R}$, the n_w gating variables $\boldsymbol{w}:\Omega\times(0,T)\to\mathbb{R}^{n_w}$ and the n_c ionic concentrations $\boldsymbol{c}\colon\Omega\times(0,T)\to\mathbb{R}^{n_c}$. The transmembrane potential is the potential difference across the membrane, the gating variables describe the proportion of open ion channels and the ionic concentrations refer to the intracellular domain. The gating variables and the ionic concentrations are modeled by an ODE system of the form

$$\begin{cases} \frac{\partial}{\partial t}\boldsymbol{w}(\boldsymbol{x},t)-\boldsymbol{G}_w(v,\boldsymbol{w})=0,\\ \frac{\partial}{\partial t}\boldsymbol{c}(\boldsymbol{x},t)-\boldsymbol{G}_c(v,\boldsymbol{w},\boldsymbol{c})=0.\end{cases} \tag{3.1}$$

We refer to Keener and Sneyd [2009a] and Colli Franzone et al. [2014] for more details on the ionic models.

Several ionic models exist and each model is specific for one membrane type. The ionic model for the human atria is the Courtemanche-Ramirez-Nattel model (Courtemanche et al. [1998]). This model describes the total ionic current I_{ion} as a sum of 12 currents which depend on $n_w=12$ gating variables and $n_c=3$ ionic concentration variables. In this work we consider this ionic model with some adaptations to the AF phenotype that capture the shortening of the ERP due to the electric remodeling. Additionally, the expressions for the rate coefficients of some gating variables are smoothed in order to increase the numerical stability (Potse [2019]).

The cell membrane separates the charges at the intracellular and extracellular domains, therefore it can be seen as a capacitor. The capacitance is defined as $C_{\text{m}}=Q/v$, where $Q\colon\Omega\times(0,T)\to\mathbb{R}$ is the charge across the membrane and v is the voltage potential needed to hold the charge. Assuming that the capacitance $C_{\text{m}}\in\mathbb{R}$ is constant, the capacitative current $I_{\text{cap}}\colon\Omega\times(0,T)\to\mathbb{R}$ is given by

$$I_{\text{cap}}(\boldsymbol{x},t)=\frac{\partial}{\partial t}Q(\boldsymbol{x},t)=C_{\text{m}}\frac{\partial}{\partial t}v(\boldsymbol{x},t).$$

The cell membrane can be modeled as an electrical circuit with a capacitor and a resistor (Keener and Sneyd [2009a], Colli Franzone et al. [2014]). The transmembrane current $I_{\text{m}}\colon\Omega\times(0,T)\to\mathbb{R}$ is thus the sum of the capacitative current I_{cap} and the ionic current I_{ion}, i.e.

$$I_{\text{m}}(\boldsymbol{x},t)=C_{\text{m}}\frac{\partial}{\partial t}v(\boldsymbol{x},t)+I_{\text{ion}}(v,\boldsymbol{w},\boldsymbol{c}). \tag{3.2}$$

3.2 Bidomain model

The action potential propagation on the tissue is modeled by the bidomain equation, that we derive in this section following Colli Franzone et al. [2014]. We consider the intracellular and extracellular domains Ω_i and Ω_e, respectively, of the atrial tissue $\Omega = \Omega_\mathrm{i} \cup \Omega_\mathrm{e}$. The intracellular and extracellular domains coexist everywhere on the tissue. We consider the average current densities $\boldsymbol{J}_\mathrm{i} \colon \Omega \times (0,T) \to \mathbb{R}^3$ and $\boldsymbol{J}_\mathrm{e} \colon \Omega \times (0,T) \to \mathbb{R}^3$ per unit area and the applied currents $I_{\mathrm{app,i}} \colon \Omega \times (0,T) \to \mathbb{R}$ and $I_{\mathrm{app,e}} \colon \Omega \times (0,T) \to \mathbb{R}$ in the intracellular and extracellular domains, respectively. The applied currents $I_{\mathrm{app,i}}$ and $I_{\mathrm{app,e}}$ and the transmembrane current I_m derived in Section 3.1 are meant per unit area. The intracellular and extracellular applied currents $i_{\mathrm{app,i}} \colon \Omega \times (0,T) \to \mathbb{R}$ and $i_{\mathrm{app,e}} \colon \Omega \times (0,T) \to \mathbb{R}$ and the transmembrane current $i_\mathrm{m} \colon \Omega \times (0,T) \to \mathbb{R}$ per unit volume are given by

$$i_{\mathrm{app,i}}(\boldsymbol{x},t) = \beta I_{\mathrm{app,i}}(\boldsymbol{x},t), \;\; i_{\mathrm{app,e}}(\boldsymbol{x},t) = \beta I_{\mathrm{app,e}}(\boldsymbol{x},t), \;\; i_\mathrm{m}(\boldsymbol{x},t) = \beta I_\mathrm{m}(\boldsymbol{x},t), \quad (3.3)$$

where $\beta \in \mathbb{R}$ is the surface-to-volume ratio. By the current conservation law on a generic volume V we have

$$\begin{cases} -\frac{1}{|V|}\int_{\partial V} \boldsymbol{J}_\mathrm{i}(\boldsymbol{x},t)\cdot \boldsymbol{n}(\boldsymbol{x}) \mathrm{d}\sigma_{\boldsymbol{x}} = \frac{1}{|V|}\int_V i_\mathrm{m}(\boldsymbol{x},t)\mathrm{d}\boldsymbol{x} - \frac{1}{|V|}\int_V i_{\mathrm{app,i}}(\boldsymbol{x},t)\mathrm{d}\boldsymbol{x}, \\ \frac{1}{|V|}\int_{\partial V} \boldsymbol{J}_\mathrm{e}(\boldsymbol{x},t)\cdot \boldsymbol{n}(\boldsymbol{x}) \mathrm{d}\sigma_{\boldsymbol{x}} = \frac{1}{|V|}\int_V i_\mathrm{m}(\boldsymbol{x},t)\mathrm{d}\boldsymbol{x} + \frac{1}{|V|}\int_V i_{\mathrm{app,e}}(\boldsymbol{x},t)\mathrm{d}\boldsymbol{x}, \end{cases}$$

where $|V|$ is the volume size of V, ∂V is the surface of V and $\boldsymbol{n}$ is the outward normal at ∂V. Note that the transmembrane current i_m is considered to flow from the intracellular to the extracellular domain. By taking the limit as $|V| \to 0$ we get

$$\begin{cases} \nabla \cdot \boldsymbol{J}_\mathrm{i}(\boldsymbol{x},t) = -i_\mathrm{m}(\boldsymbol{x},t) + i_{\mathrm{app,i}}(\boldsymbol{x},t), \\ \nabla \cdot \boldsymbol{J}_\mathrm{e}(\boldsymbol{x},t) = i_\mathrm{m}(\boldsymbol{x},t) + i_{\mathrm{app,e}}(\boldsymbol{x},t). \end{cases} \quad (3.4)$$

We also consider the intracellular and extracellular potentials $u_\mathrm{i} \colon \Omega \times (0,T) \to \mathbb{R}$ and $u_\mathrm{e} \colon \Omega \times (0,T) \to \mathbb{R}$, respectively. By Ohm's law we have

$$\begin{cases} \boldsymbol{J}_\mathrm{i}(\boldsymbol{x},t) = -\boldsymbol{D}_\mathrm{i}(\boldsymbol{x})\nabla u_\mathrm{i}(\boldsymbol{x},t), \\ \boldsymbol{J}_\mathrm{e}(\boldsymbol{x},t) = -\boldsymbol{D}_\mathrm{e}(\boldsymbol{x})\nabla u_\mathrm{e}(\boldsymbol{x},t), \end{cases} \quad (3.5)$$

where $\boldsymbol{D}_\mathrm{i} \colon \Omega \to \mathbb{R}^{3\times 3}$ and $\boldsymbol{D}_\mathrm{e} \colon \Omega \to \mathbb{R}^{3\times 3}$ are the intracellular and extracellular conductivity tensors, respectively. The tissue anisotropy due to the fibers orientation is encoded in the conductivity tensors, which are defined as

$$\boldsymbol{D}_{\mathrm{i,e}}(\boldsymbol{x}) = \sigma^l_{\mathrm{i,e}}(\boldsymbol{x})\boldsymbol{f}_l(\boldsymbol{x})\boldsymbol{f}_l^\intercal(\boldsymbol{x}) + \sigma^t_{\mathrm{i,e}}(\boldsymbol{x})\boldsymbol{f}_t(\boldsymbol{x})\boldsymbol{f}_t^\intercal(\boldsymbol{x}) + \sigma^n_{\mathrm{i,e}}(\boldsymbol{x})\boldsymbol{f}_n(\boldsymbol{x})\boldsymbol{f}_n^\intercal(\boldsymbol{x}), \quad (3.6)$$

where $\boldsymbol{f}_l \colon \Omega \to \mathbb{R}^3$ is the local fiber direction, $\boldsymbol{f}_t \colon \Omega \to \mathbb{R}^3$ is the direction transversal to the fiber and tangent to the radial lamina and $\boldsymbol{f}_n \colon \Omega \to \mathbb{R}^3$ is the direction orthogonal to both the fiber and the radial lamina. Moreover $\sigma^l_\mathrm{i} \colon \Omega \to \mathbb{R}$,

$\sigma_{\mathrm{i}}^{t}\colon \Omega \to \mathbb{R}$, $\sigma_{\mathrm{i}}^{n}\colon \Omega \to \mathbb{R}$ and $\sigma_{\mathrm{e}}^{l}\colon \Omega \to \mathbb{R}$, $\sigma_{\mathrm{e}}^{t}\colon \Omega \to \mathbb{R}$, $\sigma_{\mathrm{e}}^{n}\colon \Omega \to \mathbb{R}$ are the conductivities respectively along the directions $\boldsymbol{f}_l$, $\boldsymbol{f}_t$, $\boldsymbol{f}_n$ in the intracellular and extracellular domains, respectively. Note that $\boldsymbol{f}_l$, $\boldsymbol{f}_t$ and $\boldsymbol{f}_n$ are the orthonormal eigenvectors of the conductivity tensors and the conductivities are the corresponding eigenvalues.

The domain Ω is assumed to be insulated, therefore on the boundary $\partial\Omega$ the no current flow conditions

$$\begin{cases} \boldsymbol{n}(\boldsymbol{x})^{\mathsf{T}}\boldsymbol{J}_{\mathrm{i}}(\boldsymbol{x},t) = 0, \\ \boldsymbol{n}(\boldsymbol{x})^{\mathsf{T}}\boldsymbol{J}_{\mathrm{e}}(\boldsymbol{x},t) = 0 \end{cases} \tag{3.7}$$

must be assigned, where $\boldsymbol{n}\colon \partial\Omega \to \mathbb{R}^3$ is the unit outward normal to $\partial\Omega$. Additionally, initial conditions of the form

$$v(\boldsymbol{x},0) = v_0(\boldsymbol{x}), \quad \boldsymbol{w}(\boldsymbol{x},0) = \boldsymbol{w}_0(\boldsymbol{x}), \quad \boldsymbol{c}(\boldsymbol{x},0) = \boldsymbol{c}_0(\boldsymbol{x}) \tag{3.8}$$

must also be assigned, where $v_0\colon \Omega \to \mathbb{R}$, $\boldsymbol{w}_0\colon \Omega \to \mathbb{R}^{n_w}$ and $\boldsymbol{c}_0\colon \Omega \to \mathbb{R}^{n_c}$ represent the resting state of the system.

By combining the Equations (3.1)-(3.5), (3.7) and (3.8) we obtain the bidomain system

$$\begin{cases} \nabla\cdot(\boldsymbol{D}_{\mathrm{i}}\nabla u_{\mathrm{i}}) = \beta\left(C_{\mathrm{m}}\frac{\partial}{\partial t}v + I_{\mathrm{ion}}(v,\boldsymbol{w},\boldsymbol{c}) - I_{\mathrm{app,i}}\right) & \text{in } \Omega\times(0,T), \\ \nabla\cdot(\boldsymbol{D}_{\mathrm{e}}\nabla u_{\mathrm{e}}) = \beta\left(-C_{\mathrm{m}}\frac{\partial}{\partial t}v - I_{\mathrm{ion}}(v,\boldsymbol{w},\boldsymbol{c}) - I_{\mathrm{app,e}}\right) & \text{in } \Omega\times(0,T), \\ \frac{\partial}{\partial t}\boldsymbol{w} - \boldsymbol{G}_w(v,\boldsymbol{w}) = 0, \quad \frac{\partial}{\partial t}\boldsymbol{c} - \boldsymbol{G}_c(v,\boldsymbol{w},\boldsymbol{c}) = 0 & \text{in } \Omega\times(0,T), \\ \boldsymbol{n}^{\mathsf{T}}\boldsymbol{D}_{\mathrm{i}}\nabla u_{\mathrm{i}} = 0, \quad \boldsymbol{n}^{\mathsf{T}}\boldsymbol{D}_{\mathrm{e}}\nabla u_{\mathrm{e}} = 0 & \text{in } \partial\Omega\times(0,T), \\ v(\cdot,0) = v_0(\cdot), \quad \boldsymbol{w}(\cdot,0) = \boldsymbol{w}_0(\cdot), \quad \boldsymbol{c}(\cdot,0) = \boldsymbol{c}_0(\cdot) & \text{in } \Omega. \end{cases} \tag{3.9}$$

This system is composed of two parabolic equations, the ionic ODE system and the boundary and initial conditions. The bidomain system is degenerate, as the sum of the two parabolic equations does not depend on the time derivative of v. Therefore, fully explicit schemes can not be employed to solve the bidomain equation numerically. Moreover, the numerical solution of the bidomain equation requires a fine spatial discretization and a small time step to capture the steep and short upstroke that characterizes the transmembrane potential v during the activation. Under these constraints, the numerical computations are often unfeasible, so it is common to resort to approximated models. A common approximation is given by the monodomain model presented in Section 3.3 (Colli Franzone et al. [2014]).

3.3 Monodomain model

In this section we derive the monodomain approximation following Colli Franzone et al. [2014]. The monodomain model describes the transmembrane potential $v(\boldsymbol{x},t) = u_\mathrm{i}(\boldsymbol{x},t) - u_\mathrm{e}(\boldsymbol{x},t)$ under the assumption that the intracellular and the extracellular domains have the same anisotropy ratio. We consider the total current density $\boldsymbol{J}_\mathrm{tot}\colon \Omega \times (0,T) \to \mathbb{R}^3$ given by $\boldsymbol{J}_\mathrm{tot}(\boldsymbol{x},t) = \boldsymbol{J}_\mathrm{i}(\boldsymbol{x},t) + \boldsymbol{J}_\mathrm{e}(\boldsymbol{x},t)$. By Equation (3.5) and since $u_\mathrm{i}(\boldsymbol{x},t) = v(\boldsymbol{x},t) + u_\mathrm{e}(\boldsymbol{x},t)$, we have

$$\begin{aligned} \boldsymbol{J}_\mathrm{tot}(\boldsymbol{x},t) &= -\boldsymbol{D}_\mathrm{i}(\boldsymbol{x})\nabla u_\mathrm{i}(\boldsymbol{x},t) - \boldsymbol{D}_\mathrm{e}(\boldsymbol{x})\nabla u_\mathrm{e}(\boldsymbol{x},t) \\ &= -\boldsymbol{D}_\mathrm{i}(\boldsymbol{x})\nabla v(\boldsymbol{x},t) - \big(\boldsymbol{D}_\mathrm{i}(\boldsymbol{x}) + \boldsymbol{D}_\mathrm{e}(\boldsymbol{x})\big)\nabla u_\mathrm{e}(\boldsymbol{x},t), \end{aligned}$$

which is equivalent to

$$\begin{aligned} \boldsymbol{D}_\mathrm{e}(\boldsymbol{x})\nabla u_\mathrm{e}(\boldsymbol{x},t) = &- \boldsymbol{D}_\mathrm{e}(\boldsymbol{x})\big(\boldsymbol{D}_\mathrm{i}(\boldsymbol{x}) + \boldsymbol{D}_\mathrm{e}(\boldsymbol{x})\big)^{-1}\boldsymbol{D}_\mathrm{i}(\boldsymbol{x})\nabla v(\boldsymbol{x},t) \\ &- \boldsymbol{D}_\mathrm{e}(\boldsymbol{x})\big(\boldsymbol{D}_\mathrm{i}(\boldsymbol{x}) + \boldsymbol{D}_\mathrm{e}(\boldsymbol{x})\big)^{-1}\boldsymbol{J}_\mathrm{tot}(\boldsymbol{x},t). \end{aligned} \tag{3.10}$$

By Equation (3.6) and since $\boldsymbol{f}_l(\boldsymbol{x})\boldsymbol{f}_l(\boldsymbol{x})^\intercal + \boldsymbol{f}_t(\boldsymbol{x})\boldsymbol{f}_t(\boldsymbol{x})^\intercal + \boldsymbol{f}_n(\boldsymbol{x})\boldsymbol{f}_n(\boldsymbol{x})^\intercal = \boldsymbol{I}$, which is equivalent to $\boldsymbol{f}_l(\boldsymbol{x})\boldsymbol{f}_l(\boldsymbol{x})^\intercal = \boldsymbol{I} - \boldsymbol{f}_t(\boldsymbol{x})\boldsymbol{f}_t(\boldsymbol{x})^\intercal - \boldsymbol{f}_n(\boldsymbol{x})\boldsymbol{f}_n(\boldsymbol{x})^\intercal$, we have

$$\begin{aligned} \boldsymbol{D}_\mathrm{e}(\boldsymbol{x})\big(\boldsymbol{D}_\mathrm{i}(\boldsymbol{x}) + \boldsymbol{D}_\mathrm{e}(\boldsymbol{x})\big)^{-1} =& \frac{\sigma_\mathrm{e}^l(\boldsymbol{x})}{\sigma_\mathrm{i}^l(\boldsymbol{x}) + \sigma_\mathrm{e}^l(\boldsymbol{x})}\boldsymbol{f}_l(\boldsymbol{x})\boldsymbol{f}_l(\boldsymbol{x})^\intercal \\ &+ \frac{\sigma_\mathrm{e}^t(\boldsymbol{x})}{\sigma_\mathrm{i}^t(\boldsymbol{x}) + \sigma_\mathrm{e}^t(\boldsymbol{x})}\boldsymbol{f}_t(\boldsymbol{x})\boldsymbol{f}_t(\boldsymbol{x})^\intercal \\ &+ \frac{\sigma_\mathrm{e}^n(\boldsymbol{x})}{\sigma_\mathrm{i}^n(\boldsymbol{x}) + \sigma_\mathrm{e}^n(\boldsymbol{x})}\boldsymbol{f}_n(\boldsymbol{x})\boldsymbol{f}_n(\boldsymbol{x})^\intercal \\ =& \frac{\sigma_\mathrm{e}^l(\boldsymbol{x})}{\sigma_\mathrm{i}^l(\boldsymbol{x}) + \sigma_\mathrm{e}^l(\boldsymbol{x})}\boldsymbol{I} \\ &+ \left(\frac{\sigma_\mathrm{e}^t(\boldsymbol{x})}{\sigma_\mathrm{i}^t(\boldsymbol{x}) + \sigma_\mathrm{e}^t(\boldsymbol{x})} - \frac{\sigma_\mathrm{e}^l(\boldsymbol{x})}{\sigma_\mathrm{i}^l(\boldsymbol{x}) + \sigma_\mathrm{e}^l(\boldsymbol{x})}\right)\boldsymbol{f}_t(\boldsymbol{x})\boldsymbol{f}_t(\boldsymbol{x})^\intercal \\ &+ \left(\frac{\sigma_\mathrm{e}^n(\boldsymbol{x})}{\sigma_\mathrm{i}^n(\boldsymbol{x}) + \sigma_\mathrm{e}^n(\boldsymbol{x})} - \frac{\sigma_\mathrm{e}^l(\boldsymbol{x})}{\sigma_\mathrm{i}^l(\boldsymbol{x}) + \sigma_\mathrm{e}^l(\boldsymbol{x})}\right)\boldsymbol{f}_n(\boldsymbol{x})\boldsymbol{f}_n(\boldsymbol{x})^\intercal. \end{aligned}$$

Under the equal anisotropy ratio assumption, which states that

$$\frac{\sigma_\mathrm{i}^l(\boldsymbol{x})}{\sigma_\mathrm{e}^l(\boldsymbol{x})} = \frac{\sigma_\mathrm{i}^t(\boldsymbol{x})}{\sigma_\mathrm{e}^t(\boldsymbol{x})} = \frac{\sigma_\mathrm{i}^n(\boldsymbol{x})}{\sigma_\mathrm{e}^n(\boldsymbol{x})},$$

the latter two terms in the above equation vanish and we obtain

$$\boldsymbol{D}_{\mathrm{e}}(\boldsymbol{x})\big(\boldsymbol{D}_{\mathrm{i}}(\boldsymbol{x})+\boldsymbol{D}_{\mathrm{e}}(\boldsymbol{x})\big)^{-1}=\frac{\sigma_{\mathrm{e}}^{l}(\boldsymbol{x})}{\sigma_{\mathrm{i}}^{l}(\boldsymbol{x})+\sigma_{\mathrm{e}}^{l}(\boldsymbol{x})}\boldsymbol{I}.$$

By additionally defining the monodomain conductivity tensor $\boldsymbol{D}_{\mathrm{m}}\colon\Omega\to\mathbb{R}^{3\times 3}$ as

$$\boldsymbol{D}_{\mathrm{m}}(\boldsymbol{x}):=\boldsymbol{D}_{\mathrm{e}}(\boldsymbol{x})\big(\boldsymbol{D}_{\mathrm{i}}(\boldsymbol{x})+\boldsymbol{D}_{\mathrm{e}}(\boldsymbol{x})\big)^{-1}\boldsymbol{D}_{\mathrm{i}}(\boldsymbol{x}),$$

Equation (3.10) becomes

$$\boldsymbol{D}_{\mathrm{e}}(\boldsymbol{x})\nabla u_{\mathrm{e}}(\boldsymbol{x},t)=-\boldsymbol{D}_{\mathrm{m}}(\boldsymbol{x})\nabla v(\boldsymbol{x},t)-\frac{\sigma_{\mathrm{e}}^{l}(\boldsymbol{x})}{\sigma_{\mathrm{i}}^{l}(\boldsymbol{x})+\sigma_{\mathrm{e}}^{l}(\boldsymbol{x})}\boldsymbol{J}_{\mathrm{tot}}(\boldsymbol{x},t). \tag{3.11}$$

Equations (3.3) and (3.4) imply that

$$\nabla\cdot\boldsymbol{J}_{\mathrm{tot}}(\boldsymbol{x},t)=\nabla\cdot\boldsymbol{J}_{\mathrm{i}}(\boldsymbol{x},t)+\nabla\cdot\boldsymbol{J}_{\mathrm{e}}(\boldsymbol{x},t)=\beta\big(I_{\mathrm{app,i}}(\boldsymbol{x},t)+I_{\mathrm{app,e}}(\boldsymbol{x},t)\big)$$

and, consequently, from Equation (3.11) we obtain

$$\begin{aligned}\nabla\cdot\big(\boldsymbol{D}_{\mathrm{e}}(\boldsymbol{x})\nabla u_{\mathrm{e}}(\boldsymbol{x},t)\big)=&-\nabla\cdot\big(\boldsymbol{D}_{\mathrm{m}}(\boldsymbol{x})\nabla v(\boldsymbol{x},t)\big)\\&-\beta\frac{\sigma_{\mathrm{e}}^{l}(\boldsymbol{x})}{\sigma_{\mathrm{i}}^{l}(\boldsymbol{x})+\sigma_{\mathrm{e}}^{l}(\boldsymbol{x})}\big(I_{\mathrm{app,i}}(\boldsymbol{x},t)+I_{\mathrm{app,e}}(\boldsymbol{x},t)\big).\end{aligned} \tag{3.12}$$

By combining Equation (3.12) with the second equation of the bidomain system (3.9) and by defining the monodomain applied current $I_{\mathrm{app}}\colon\Omega\times(0,T)\to\mathbb{R}$ as

$$I_{\mathrm{app}}(\boldsymbol{x},t):=\frac{\sigma_{\mathrm{e}}^{l}(\boldsymbol{x})I_{\mathrm{app,i}}(\boldsymbol{x},t)-\sigma_{\mathrm{i}}^{l}(\boldsymbol{x})I_{\mathrm{app,e}}(\boldsymbol{x},t)}{\sigma_{\mathrm{i}}^{l}(\boldsymbol{x})+\sigma_{\mathrm{e}}^{l}(\boldsymbol{x})},$$

we get

$$\nabla\cdot\big(\boldsymbol{D}_{\mathrm{m}}(\boldsymbol{x})\nabla v(\boldsymbol{x},t)\big)=\beta\left(C_{\mathrm{m}}\frac{\partial}{\partial t}v(\boldsymbol{x},t)+I_{\mathrm{ion}}(v,\boldsymbol{w},\boldsymbol{c})-I_{\mathrm{app}}(\boldsymbol{x},t)\right). \tag{3.13}$$

Concerning the boundary condition, Equation (3.7) implies that

$$\boldsymbol{n}(\boldsymbol{x})^{\mathsf{T}}\boldsymbol{J}_{\mathrm{tot}}(\boldsymbol{x},t)=\boldsymbol{n}(\boldsymbol{x})^{\mathsf{T}}\boldsymbol{J}_{\mathrm{i}}(\boldsymbol{x},t)+\boldsymbol{n}(\boldsymbol{x})^{\mathsf{T}}\boldsymbol{J}_{\mathrm{e}}(\boldsymbol{x},t)=0$$

and, consequently, from Equation (3.11) and the fourth equation in the bidomain system (3.9) we obtain

$$\boldsymbol{n}(\boldsymbol{x})^{\mathsf{T}}\boldsymbol{D}_{\mathrm{m}}(\boldsymbol{x})\nabla v(\boldsymbol{x},t)=-\boldsymbol{n}(\boldsymbol{x})^{\mathsf{T}}\boldsymbol{D}_{\mathrm{e}}(\boldsymbol{x})\nabla u_{\mathrm{e}}(\boldsymbol{x},t)=0. \tag{3.14}$$

By combining the Equations (3.13) and (3.14) with the Equations (3.1) and (3.8) we obtain the monodomain system

$$\begin{cases} \nabla \cdot (\boldsymbol{D}_\mathrm{m} \nabla v) = \beta \left(C_\mathrm{m} \frac{\partial}{\partial t} v + I_\mathrm{ion}(v, \boldsymbol{w}, \boldsymbol{c}) - I_\mathrm{app} \right) & \text{in } \Omega \times (0, T), \\ \frac{\partial}{\partial t} \boldsymbol{w} - \boldsymbol{G}_w(v, \boldsymbol{w}) = 0, \quad \frac{\partial}{\partial t} \boldsymbol{c} - \boldsymbol{G}_c(v, \boldsymbol{w}, \boldsymbol{c}) = 0 & \text{in } \Omega \times (0, T), \\ \boldsymbol{n}^\intercal \boldsymbol{D}_\mathrm{m} \nabla v = 0 & \text{in } \partial\Omega \times (0, T), \\ v(\cdot, 0) = v_0(\cdot), \quad \boldsymbol{w}(\cdot, 0) = \boldsymbol{w}_0(\cdot), \quad \boldsymbol{c}(\cdot, 0) = \boldsymbol{c}_0(\cdot) & \text{in } \Omega. \end{cases} \tag{3.15}$$

This system is composed of a parabolic reaction-diffusion equation, the ionic ODE system and the boundary and initial conditions. In the parabolic equation, the reaction term is the ionic current I_ion and the diffusion term is $\nabla \cdot (\boldsymbol{D}_\mathrm{m} \nabla v)$. In this work we set the capacitance $C_\mathrm{m} = 1\ \mu\mathrm{F/cm}^2$ and the surface-to-volume ratio $\beta = 800\ \mathrm{cm}^{-1}$ (Gharaviri et al. [2020]). Moreover, note that the resting transmembrane potential is approximately −80 mV.

3.4 Pacing protocol

The main focus of this work is the inducibility of AF. The AF inducibility is typically evaluated from the outcome of a tissue stimulation with a pacing protocol. In this section we introduce the pacing protocol for the high fidelity model. The pacing protocol is encoded in the function I_app and consists of a sequence of N_stim stimuli of the form

$$I_\mathrm{app}(\boldsymbol{x}, t; \boldsymbol{x}_\mathrm{stim}, \boldsymbol{\tau}_\mathrm{stim}) = \begin{cases} I_\mathrm{max}, & (\boldsymbol{x}, t) \in B_r(\boldsymbol{x}_\mathrm{stim}) \times \cup_{j=1}^{N_\mathrm{stim}} [\tau_j, \tau_j + \Delta\tau], \\ 0, & \text{otherwise}, \end{cases}$$

where $\boldsymbol{x}_\mathrm{stim}$ is the pacing site, $\boldsymbol{\tau}_\mathrm{stim} = \{\tau_j\}_{j=1}^{N_\mathrm{stim}}$ are the stimulation times, I_max is the strength of the stimuli, $B_r(\boldsymbol{x}_\mathrm{stim}) = \{\boldsymbol{x} \in \Omega : \boldsymbol{x}_\mathrm{stim} \leq \boldsymbol{x} \leq \boldsymbol{x}_\mathrm{stim} + 2r\}$ is a neighborhood of the pacing site with radius r and $\Delta\tau$ is the duration of the stimuli. Our pacing protocol is tailored to the ionic model and has been tested in a single-cell preparation (Potse et al. [2018], Potse [2019], Gharaviri et al. [2021b], Gharaviri et al. [2021a]). The result is a pacing protocol consisting of $N_\mathrm{stim} = 14$ stimuli applied at the delivery times $\boldsymbol{\tau}_\mathrm{stim}$ reported in Figure 3.1, with decreasing temporal distance. Each stimulus has strength $I_\mathrm{max} = 800\ \mu\mathrm{A\ cm}^{-2}$, radius $r = 0.4$ cm and duration $\Delta\tau = 4$ ms.

After the pacing protocol is over, depending on the pacing site $\boldsymbol{x}_\mathrm{stim}$, AF is induced or not. If the pacing results in a vanishing wave, i.e. if the system asymptotically approaches the resting state, then the induction of AF is not successful. Instead, if the pacing results in a self-sustained activity, then the induction of AF is successful.

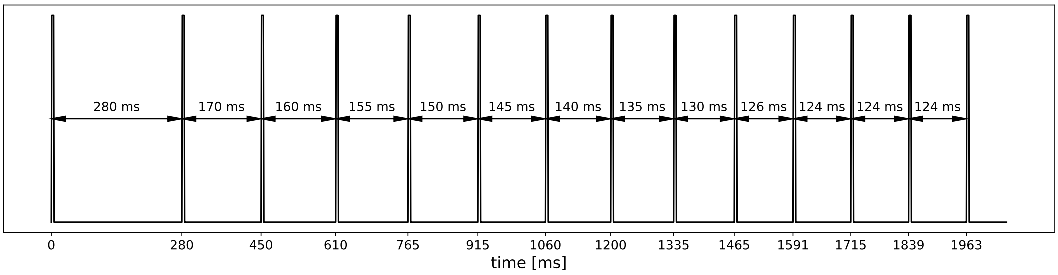

Figure 3.1. Pacing protocol encoded in I_{app}.

3.5 Atrial model

The atrial model on which we numerically solve the monodomain equation is described in Gharaviri et al. [2020]. In this section we summarize its main aspects. The atrial anatomy reproduces the features described in Section 2.2 and illustrated in Figure 2.4, including the bundles and the fibers. The atrial geometry Ω is obtained from imaging data, while the bundles and the fibers are added based on anatomical studies. The numerical solution of the monodomain system (3.15) is performed on a computational mesh of hexahedrons discretizing Ω with resolution h. Each hexahedral element is assigned to a tissue type and a fiber orientation.

The tissue type determines the longitudinal and transversal conductivities $\sigma^l_{\mathrm{i,e}}(\boldsymbol{x})$ and $\sigma^t_{\mathrm{i,e}}(\boldsymbol{x})$. Table 3.1 shows the intracellular and extracellular conductivities along and across the fiber direction associated to the tissue types. The conductivities are expressed in mS/cm. For the across-sheet conductivities, we assume $\sigma^n_{\mathrm{i,e}}(\boldsymbol{x}) = \sigma^t_{\mathrm{i,e}}(\boldsymbol{x})$, so that the conductivity tensors are given by

Type	σ^l_{i}	σ^t_{i}	σ^l_{e}	σ^t_{e}
atrium	3.0	0.3	3.0	1.2
bundle	9.0	0.3	9.0	1.2
intercaval	1.5	1.5	1.5	1.5
endomysial fibrosis	3.0	0.0	3.0	1.2
diffuse fibrosis	0.75	0.075	0.75	0.3

Table 3.1. Conductivities in mS/cm

$$\boldsymbol{D}_{\mathrm{i,e}}(\boldsymbol{x}) = \sigma^t_{\mathrm{i,e}}(\boldsymbol{x})\boldsymbol{I} + \big(\sigma^l_{\mathrm{i,e}}(\boldsymbol{x}) - \sigma^t_{\mathrm{i,e}}(\boldsymbol{x})\big)\boldsymbol{f}_l(\boldsymbol{x})\boldsymbol{f}_l(\boldsymbol{x})^\intercal.$$

Most of the healthy atrial tissue has type "atrium". The healthy tissue that is not assigned to the type "atrium" is located in the Bachmann's bundle, in the bundles connecting the coronary sinus to the left atrium or between the superior and the inferior vena cava. In the bundles the conduction along the fibers direction is increased. Therefore the tissue type "bundle" has increased longitudinal conductivities. In the intercaval region the conduction is isotropic. Thus the tissue type "intercaval" has equal longitudinal and transversal conductivities. The presence of fibrosis is modeled by a random field sampled on the atrial geometry (Pezzuto et al. [2019]). In particular, the random field is the average of a large-scale random field determining the patchyness and a small-scale random field. The hexahedral elements whose values of the random field are above a certain threshold are then marked as fibrotic. The threshold is selected such that a given fraction of the atrial tissue is marked as fibrotic, taking into account that the fibrosis does not extend to the bundles. The endomysial fibrosis associated to AF hinders the propagation across the fibers direction. Therefore the tissue type "endomysial fibrosis" has zero transversal intracellular conductivity. In this work we also consider another type of fibrosis, namely diffuse fibrosis, which reduces the conduction both along and across the fibers direction (Zahid et al. [2016]). Thus the tissue type "diffuse fibrosis" has reduced intracellular and extracellular conductivities. The presence of ablated tissue is simply modeled by marking the ablated elements as non-conductive.

The fiber orientation is encoded in the angles $\alpha(\boldsymbol{x})$, $\gamma(\boldsymbol{x})$ and $\theta(\boldsymbol{x})$ that are assigned to the hexahedral elements. These angles determine the entries

$$\begin{cases} a_{11}(\boldsymbol{x}) = \sin\big(\gamma(\boldsymbol{x})\big)\cos\big(\theta(\boldsymbol{x})\big), \\ a_{12}(\boldsymbol{x}) = \sin\big(\gamma(\boldsymbol{x})\big)\sin\big(\theta(\boldsymbol{x})\big), \\ a_{13}(\boldsymbol{x}) = \cos\big(\gamma(\boldsymbol{x})\big), \\ a_{21}(\boldsymbol{x}) = \sin\big(\alpha(\boldsymbol{x})\big)\sin\big(\theta(\boldsymbol{x})\big) - \cos\big(\alpha(\boldsymbol{x})\big)\cos\big(\gamma(\boldsymbol{x})\big)\cos\big(\theta(\boldsymbol{x})\big), \\ a_{22}(\boldsymbol{x}) = -\sin\big(\alpha(\boldsymbol{x})\big)\cos\big(\theta(\boldsymbol{x})\big) - \cos\big(\alpha(\boldsymbol{x})\big)\cos\big(\gamma(\boldsymbol{x})\big)\sin\big(\theta(\boldsymbol{x})\big), \\ a_{23}(\boldsymbol{x}) = \cos\big(\alpha(\boldsymbol{x})\big)\sin\big(\gamma(\boldsymbol{x})\big), \\ a_{31}(\boldsymbol{x}) = \cos\big(\alpha(\boldsymbol{x})\big)\sin\big(\theta(\boldsymbol{x})\big) + \sin\big(\alpha(\boldsymbol{x})\big)\cos\big(\gamma(\boldsymbol{x})\big)\cos\big(\theta(\boldsymbol{x})\big), \\ a_{32}(\boldsymbol{x}) = -\cos\big(\alpha(\boldsymbol{x})\big)\cos\big(\theta(\boldsymbol{x})\big) + \sin\big(\alpha(\boldsymbol{x})\big)\cos\big(\gamma(\boldsymbol{x})\big)\sin\big(\theta(\boldsymbol{x})\big), \\ a_{33}(\boldsymbol{x}) = -\sin\big(\alpha(\boldsymbol{x})\big)\sin\big(\gamma(\boldsymbol{x})\big) \end{cases}$$

of the matrix

$$A(\boldsymbol{x}) = \begin{bmatrix} a_{11}(\boldsymbol{x}) & a_{12}(\boldsymbol{x}) & a_{13}(\boldsymbol{x}) \\ a_{21}(\boldsymbol{x}) & a_{22}(\boldsymbol{x}) & a_{23}(\boldsymbol{x}) \\ a_{31}(\boldsymbol{x}) & a_{32}(\boldsymbol{x}) & a_{33}(\boldsymbol{x}) \end{bmatrix},$$

which, together with the conductivities, determines the conductivity tensors as

$$\boldsymbol{D}_{\mathrm{i,e}}(\boldsymbol{x}) = \boldsymbol{A}(\boldsymbol{x})^{\mathsf{T}} \begin{bmatrix} \sigma^n_{\mathrm{i,e}}(\boldsymbol{x}) & 0 & 0 \\ 0 & \sigma^t_{\mathrm{i,e}}(\boldsymbol{x}) & 0 \\ 0 & 0 & \sigma^l_{\mathrm{i,e}}(\boldsymbol{x}) \end{bmatrix} \boldsymbol{A}(\boldsymbol{x}).$$

3.6 Numerical solution of the monodomain equation

In this section we present the numerical methods that we employ to solve the monodomain system (3.15). We perform the simulations on the computational mesh of hexahedral elements with resolution h. We use a second-order finite difference method for the spatial discretization of the diffusion term $\nabla\cdot(\boldsymbol{D}_{\mathrm{m}}\nabla v)$. The finite difference scheme follows from the second-order Taylor expansion of the products between the coefficients of the monodomain conductivity tensor $\boldsymbol{D}_{\mathrm{m}}$ and the transmembrane potential v (Asencor and Panizo [1991], Saleheen and Ng [1997]). Notice that the transmembrane potential v is computed on the nodes of the computational mesh. The fact that the monodomain conductivity tensor $\boldsymbol{D}_{\mathrm{m}}$ is defined on the elements of the computational mesh introduces discontinuities in the conductivity coefficients at the boundaries between different elements. To handle this difficulty and the consequent singularities in the derivatives of the conductivity coefficients, we use the transition layer method (Panizo et al. [1977], Saleheen and Ng [1997]).

The monodomain system is non-degenerate and, as a consequence, explicit schemes can be employed for its time integration. We use the first-order explicit Euler method for the transmembrane potential v and the concentration variables $\boldsymbol{c}$ and we employ the Rush-Larsen method for the gating variables $\boldsymbol{w}$. The schemes are applied with the time step Δt. The partial derivatives of v and $\boldsymbol{c}$ with respect to time are thus approximated by

$$\frac{\partial}{\partial t} v(\boldsymbol{x},t) \approx \frac{v(\boldsymbol{x},t+\Delta t)-v(\boldsymbol{x},t)}{\Delta t} \quad \text{and} \quad \frac{\partial}{\partial t}\boldsymbol{c}(\boldsymbol{x},t) \approx \frac{\boldsymbol{c}(\boldsymbol{x},t+\Delta t)-\boldsymbol{c}(\boldsymbol{x},t)}{\Delta t}.$$

The Rush-Larsen method (Rush and Larsen [1978]) can be applied to gating variables $w\colon \Omega\times(0,T)\to\mathbb{R}$ that satisfy an ODE of the form

$$\frac{\partial}{\partial t} w(\boldsymbol{x},t) = a(v)\big(1-\boldsymbol{w}(\boldsymbol{x},t)\big) - b(v)\boldsymbol{w}(\boldsymbol{x},t), \tag{3.16}$$

where a and b are some voltage-dependent rate constants. By assuming that a and b are constant, the solution of (3.16) reads

$$w(\boldsymbol{x},t) = w(\boldsymbol{x},\infty) - \big(w(\boldsymbol{x},\infty)-w(\boldsymbol{x},0)\big)\exp\big(-(a+b)t\big),$$

where $w(\boldsymbol{x},\infty) = a/(a+b)$ is the steady state solution. The Rush-Larsen scheme for the update of the gating variable therefore is

$$w(\boldsymbol{x}, t+\Delta t) = w(\boldsymbol{x},\infty) - \big(w(\boldsymbol{x},\infty) - w(\boldsymbol{x},t)\big)\exp\big(-(a+b)\Delta t\big),$$

where the rate constants are computed from the transmembrane potential at time t, i.e. $a = a\big(v(\boldsymbol{x},t)\big)$ and $b = b\big(v(\boldsymbol{x},t)\big)$. The Rush-Larsen scheme is first-order explicit and it is more stable than the forward Euler scheme. The increased stability is necessary as $1/(a+b)$ is very small for some gating variables.

We perform the monodomain simulations using the Propag-5 software (Potse et al. [2006], Krause et al. [2012]) on CSCS (Swiss National Supercomputing Centre). Similarly to the bidomain equation, also the monodomain equation needs to be solved numerically with a fine spatial discretization and a small time step, in order to capture the steep and short upstroke that characterizes the transmembrane potential v during the activation. Even though the monodomain model is less accurate than the bidomain model, in this work our high fidelity model is the monodomain system (3.15) solved numerically on a hexahedral mesh with resolution $h = 0.02$ cm and with the time step $\Delta t = 0.01$ ms (Gharaviri et al. [2020]).

Chapter 4

Re-excitable eikonal low fidelity model

In this chapter we introduce the low fidelity model based on the algorithmic solution of the eikonal equation. The eikonal equation is an approximation of the monodomain equation that allows a fast computation of the activation times thanks to solution algorithms such as the FMM (Kimmel and Sethian [1998], Sethian and Vladimirsky [2000]). The eikonal equation models a single activation and does not take into account the re-excitability of the tissue. Since the electrical activity during AF is characterized by re-entrant wavelets that determine multiple activations, the re-excitability is a key feature that needs to be included in our low fidelity model. Therefore, our eikonal low fidelity model consists of an adaptation of the FMM to account for the re-excitability. Our goal is to maximize the similarity between the low fidelity model and the high fidelity model of Chapter 3. Therefore, the re-excitability included in the FMM is based on the restitution properties of the monodomain model. We test the accuracy of the eikonal low fidelity model in qualitative and quantitative studies. We first visually compare the 2D simulations of a re-entrant spiral wave obtained with the high and the low fidelity models. In particular, we study the effect of fibrosis and anisotropy. The effect of fibrosis is also investigated in a quantitative study in the non re-entrant case. Then, we compare the models in 2D numerical experiments where a re-entry is initiated by a S1-S2 stimulation in presence of scars. Finally, we move to 3D numerical experiments on the atrial geometry and to the comparison of the models in terms of AF inducibility after the stimulation by a pacing protocol.

The eikonal equation is introduced in Section 4.1. The restitution curves of the high fidelity model are shown in Section 4.2 and the eikonal algorithm with re-excitability is presented in Section 4.3. The effect of fibrosis is studied in Section 4.4 and the anisotropy is studied in Section 4.5. The 2D simulations of spiral waves are shown along Sections 4.3, 4.4 and 4.5. The 2D numerical experiments in presence of scars are shown in Section 4.6 and the 3D numerical experiments are presented in Section 4.7. A discussion follows in Section 4.8.

In this chapter, for the sake of simplicity, we consider the computational domain of the eikonal model to be a surface discretized by triangular elements, instead of a geometry discretized by tetrahedral elements. This approximation is motivated by the fact that the atrial wall is thin. Moreover, we consider the healthy tissue to be all of type "atrium". Additionally, we consider the fibrotic tissue type to be "diffuse fibrosis". The reason for this choice is that, with the tissue type "endomysial fibrosis", due to the zero transversal intracellular conductivity, the monodomain simulations might exhibit numerical artifacts that would hinder the comparison with the eikonal simulations.

4.1 Eikonal equation

In this section we derive the eikonal equation, we explain its link to the monodomain equation and we state its solution. The eikonal equation is derived in Subsection 4.1.1 and its relation to the monodomain equation is shown in Subsection 4.1.2. The solution is formulated in Subsection 4.1.3.

4.1.1 Derivation

In this subsection we derive the eikonal equation for the activation time $\phi : \Omega \to \mathbb{R}$ on the atrial tissue $\Omega \subset \mathbb{R}^3$. We consider the local fibers direction $\boldsymbol{f}_l : \Omega \to \mathbb{R}^3$, the transversal direction $\boldsymbol{f}_t : \Omega \to \mathbb{R}^3$ and the across-sheet direction $\boldsymbol{f}_n : \Omega \to \mathbb{R}^3$ introduced in Section 3.2, which form an orthonormal basis of $\mathbb{R}^3$. We assume that the conduction velocities along $\boldsymbol{f}_l$, $\boldsymbol{f}_t$ and $\boldsymbol{f}_n$ are known and are given by $CV_l : \Omega \to \mathbb{R}$, $CV_t : \Omega \to \mathbb{R}$ and $CV_n : \Omega \to \mathbb{R}$, respectively. We then consider an arbitrary unit direction $\boldsymbol{d} : \Omega \to \mathbb{R}^3$ given by $\boldsymbol{d}(\boldsymbol{x}) = d_l \boldsymbol{f}_l(\boldsymbol{x}) + d_t \boldsymbol{f}_t(\boldsymbol{x}) + d_n \boldsymbol{f}_n(\boldsymbol{x})$ with $\sqrt{d_l^2 + d_t^2 + d_n^2} = 1$. The conduction velocity $CV_{\boldsymbol{d}} : \Omega \to \mathbb{R}$ along the direction $\boldsymbol{d}$ is given by

$$\begin{aligned} CV_{\boldsymbol{d}}(\boldsymbol{x}) &= \sqrt{d_l^2 CV_l^2(\boldsymbol{x}) + d_t^2 CV_t^2(\boldsymbol{x}) + d_n^2 CV_n^2(\boldsymbol{x})} \\ &= CV_l(\boldsymbol{x}) \sqrt{\boldsymbol{d}(\boldsymbol{x})^\intercal \boldsymbol{D}(\boldsymbol{x}) \boldsymbol{d}(\boldsymbol{x})}, \end{aligned} \tag{4.1}$$

where the tensor $\boldsymbol{D} : \Omega \to \mathbb{R}^{3\times 3}$ describes the anisotropy and is given by

$$\boldsymbol{D}(\boldsymbol{x}) = \boldsymbol{f}_l(\boldsymbol{x})\boldsymbol{f}_l(\boldsymbol{x})^\intercal + \frac{CV_t(\boldsymbol{x})^2}{CV_l(\boldsymbol{x})^2}\boldsymbol{f}_t(\boldsymbol{x})\boldsymbol{f}_t(\boldsymbol{x})^\intercal + \frac{CV_n(\boldsymbol{x})^2}{CV_l(\boldsymbol{x})^2}\boldsymbol{f}_n(\boldsymbol{x})\boldsymbol{f}_n(\boldsymbol{x})^\intercal.$$

The conduction velocity $CV_{\boldsymbol{d}}$ is also given by the inverse of the directional derivative of ϕ along $\boldsymbol{d}$, i.e.

$$CV_{\boldsymbol{d}}(\boldsymbol{x}) = \left(\lim_{h\to 0} \frac{\phi\big(\boldsymbol{x} + h\boldsymbol{d}(\boldsymbol{x})\big) - \phi(\boldsymbol{x})}{h} \right)^{-1} = \frac{1}{\nabla_{\boldsymbol{d}}\phi(\boldsymbol{x})} = \frac{1}{\nabla\phi(\boldsymbol{x})^{\intercal}\boldsymbol{d}(\boldsymbol{x})}. \tag{4.2}$$

Since ϕ is the activation time, the unit direction of propagation is given by $\boldsymbol{d}(\boldsymbol{x}) = \nabla\phi(\boldsymbol{x})/\|\nabla\phi(\boldsymbol{x})\|$ (Colli Franzone et al. [1990a]). By substituting the propagation direction into Equation (4.1), we obtain

$$CV_{\boldsymbol{d}}(\boldsymbol{x}) = \frac{CV_l(\boldsymbol{x})}{\|\nabla\phi(\boldsymbol{x})\|}\sqrt{\nabla\phi(\boldsymbol{x})^{\intercal}\boldsymbol{D}(\boldsymbol{x})\nabla\phi(\boldsymbol{x})}. \tag{4.3}$$

Similarly, by substituting the propagation direction into Equation (4.2), we obtain

$$CV_{\boldsymbol{d}}(\boldsymbol{x}) = \frac{1}{\|\nabla\phi(\boldsymbol{x})\|}. \tag{4.4}$$

By combining the Equations (4.3) and (4.4), we obtain the eikonal equation

$$CV_l(\boldsymbol{x})\sqrt{\nabla\phi(\boldsymbol{x})^{\intercal}\boldsymbol{D}(\boldsymbol{x})\nabla\phi(\boldsymbol{x})} = 1. \tag{4.5}$$

4.1.2 Relation to the monodomain equation

In this subsection we show the relation between the eikonal equation (4.5) and the monodomain equation

$$\nabla\cdot\big(\boldsymbol{D}_{\mathrm{m}}(\boldsymbol{x})\nabla v(\boldsymbol{x},t)\big) = \beta\left(C_{\mathrm{m}}\frac{\partial}{\partial t}v(\boldsymbol{x},t) + I_{\mathrm{ion}}(v)\right). \tag{4.6}$$

Note that, compared to the monodomain equation (3.13), we assume here that the ionic current I_{ion} depends only on the transmembrane potential v and we neglect the applied current I_{app}. We look for solutions to Equation (4.6) taking the form

$$v(\boldsymbol{x},t) = u\big(t - \phi(\boldsymbol{x})\big) = u(\xi). \tag{4.7}$$

It holds

$$\boldsymbol{D}_{\mathrm{m}}(\boldsymbol{x})\nabla v(\boldsymbol{x},t) = -\boldsymbol{D}_{\mathrm{m}}(\boldsymbol{x})\nabla\phi(\boldsymbol{x})u'(\xi),$$

thus

$$\nabla\cdot\big(\boldsymbol{D}_{\mathrm{m}}(\boldsymbol{x})\nabla v(\boldsymbol{x},t)\big) = \nabla\phi(\boldsymbol{x})^{\intercal}\boldsymbol{D}_{\mathrm{m}}(\boldsymbol{x})\nabla\phi(\boldsymbol{x})u''(\xi) - \nabla\cdot\big(\boldsymbol{D}_{\mathrm{m}}(\boldsymbol{x})\nabla\phi(\boldsymbol{x})\big)u'(\xi).$$

Moreover, we have

$$\frac{\partial}{\partial t}v(\boldsymbol{x},t) = u'(\xi).$$

So, by plugging the solution (4.7) into the monodomain equation (4.6), we obtain

$$\beta^{-1}\nabla\phi(\boldsymbol{x})^{\intercal}\boldsymbol{D}_{\mathrm{m}}(\boldsymbol{x})\nabla\phi(\boldsymbol{x})u''(\xi)-\big(\beta^{-1}\nabla\cdot\big(\boldsymbol{D}_{\mathrm{m}}(\boldsymbol{x})\nabla\phi(\boldsymbol{x})\big)+C_{\mathrm{m}}\big)u'(\xi)-I_{\mathrm{ion}}(u)=0.$$

Using the change of variables

$$\hat{\xi}=\xi\beta^{-1/2}\sqrt{\nabla\phi(\boldsymbol{x})^{\intercal}\boldsymbol{D}_{\mathrm{m}}(\boldsymbol{x})\nabla\phi(\boldsymbol{x})},$$

we get

$$u''(\hat{\xi})-\frac{\beta^{-1}\nabla\cdot\big(\boldsymbol{D}_{\mathrm{m}}(\boldsymbol{x})\nabla\phi(\boldsymbol{x})\big)+C_{\mathrm{m}}}{\beta^{-1/2}\sqrt{\nabla\phi(\boldsymbol{x})^{\intercal}\boldsymbol{D}_{\mathrm{m}}(\boldsymbol{x})\nabla\phi(\boldsymbol{x})}}u'(\hat{\xi})-I_{\mathrm{ion}}(u)=0. \tag{4.8}$$

Under suitable assumptions, the eigenvalue problem

$$\begin{cases} u''(\xi)-\rho u'(\xi)-I_{\mathrm{ion}}\big(u(\xi)\big)=0, & \xi\in\mathbb{R},\\ I_{\mathrm{ion}}\big(u(\xi)\big)\to 0, & |\xi|\to\infty \end{cases} \tag{4.9}$$

has a unique solution (Colli Franzone et al. [2014]). By comparing the Equations (4.8) and (4.9), we deduce

$$\rho=\frac{\beta^{-1}\nabla\cdot\big(\boldsymbol{D}_{\mathrm{m}}(\boldsymbol{x})\nabla\phi(\boldsymbol{x})\big)+C_{\mathrm{m}}}{\beta^{-1/2}\sqrt{\nabla\phi(\boldsymbol{x})^{\intercal}\boldsymbol{D}_{\mathrm{m}}(\boldsymbol{x})\nabla\phi(\boldsymbol{x})}},$$

where $\rho\in\mathbb{R}$ is a parameter that depends on the ionic model. By rearranging we obtain the eikonal-diffusion equation

$$\frac{\rho}{C_{\mathrm{m}}\sqrt{\beta}}\sqrt{\nabla\phi(\boldsymbol{x})^{\intercal}\boldsymbol{D}_{\mathrm{m}}(\boldsymbol{x})\nabla\phi(\boldsymbol{x})}-\frac{1}{C_{\mathrm{m}}\beta}\nabla\cdot\big(\boldsymbol{D}_{\mathrm{m}}(\boldsymbol{x})\nabla\phi(\boldsymbol{x})\big)=1. \tag{4.10}$$

With reasonable values for the parameters ρ, C_{m}, β and for the conductivities, the advection term of the eikonal-diffusion equation (4.10) is dominant over the diffusion term. The pure eikonal approximation is obtained by disregarding the diffusion term. Thus the eikonal equation is

$$\frac{\rho}{C_{\mathrm{m}}\sqrt{\beta}}\sqrt{\nabla\phi(\boldsymbol{x})^{\intercal}\boldsymbol{D}_{\mathrm{m}}(\boldsymbol{x})\nabla\phi(\boldsymbol{x})}=1. \tag{4.11}$$

The pure eikonal equation (4.11) and the eikonal-diffusion equation (4.10) correspond when the monodomain conductivity tensor $\boldsymbol{D}_{\mathrm{m}}$ is constant and the front is planar, i.e. $\nabla^2\phi=0$. The pure eikonal model does not account for the diffusion currents due to heterogeneities in the conductivity or in the propagation.

We recall that the monodomain conductivity tensor $\boldsymbol{D}_{\mathrm{m}} \colon \Omega \to \mathbb{R}^{3\times 3}$ is

$$\boldsymbol{D}_{\mathrm{m}}(\boldsymbol{x}) = \sigma_{\mathrm{m}}^{l}(\boldsymbol{x})\boldsymbol{f}_{l}(\boldsymbol{x})\boldsymbol{f}_{l}^{\intercal}(\boldsymbol{x}) + \sigma_{\mathrm{m}}^{t}(\boldsymbol{x})\boldsymbol{f}_{t}(\boldsymbol{x})\boldsymbol{f}_{t}^{\intercal}(\boldsymbol{x}) + \sigma_{\mathrm{m}}^{n}(\boldsymbol{x})\boldsymbol{f}_{n}(\boldsymbol{x})\boldsymbol{f}_{n}^{\intercal}(\boldsymbol{x}),$$

where the monodomain longitudinal, transversal and across-sheet conductivities $\sigma_{\mathrm{m}}^{l,t,n} \colon \Omega \to \mathbb{R}$ are given by

$$\sigma_{\mathrm{m}}^{l,t,n}(\boldsymbol{x}) = \frac{\sigma_{\mathrm{e}}^{l,t,n}(\boldsymbol{x})\sigma_{\mathrm{i}}^{l,t,n}(\boldsymbol{x})}{\sigma_{\mathrm{e}}^{l,t,n}(\boldsymbol{x}) + \sigma_{\mathrm{i}}^{l,t,n}(\boldsymbol{x})}.$$

We again consider the unit propagation direction $\boldsymbol{d}(\boldsymbol{x}) = \nabla\phi(\boldsymbol{x})/\|\nabla\phi(\boldsymbol{x})\|$ and we write it as $\boldsymbol{d}(\boldsymbol{x}) = d_l\boldsymbol{f}_l(\boldsymbol{x}) + d_t\boldsymbol{f}_t(\boldsymbol{x}) + d_n\boldsymbol{f}_n(\boldsymbol{x})$ with $\sqrt{d_l^2 + d_t^2 + d_n^2} = 1$. By substituting $\nabla\phi(\boldsymbol{x}) = \|\nabla\phi(\boldsymbol{x})\|\boldsymbol{d}(\boldsymbol{x})$ into the eikonal equation (4.11) we obtain

$$\frac{\rho\|\nabla\phi(\boldsymbol{x})\|}{C_{\mathrm{m}}\sqrt{\beta}}\sqrt{d_l^2\sigma_{\mathrm{m}}^{l}(\boldsymbol{x}) + d_t^2\sigma_{\mathrm{m}}^{t}(\boldsymbol{x}) + d_n^2\sigma_{\mathrm{m}}^{n}(\boldsymbol{x})} = 1. \tag{4.12}$$

Similarly, by substituting $\nabla\phi(\boldsymbol{x}) = \|\nabla\phi(\boldsymbol{x})\|\boldsymbol{d}(\boldsymbol{x})$ into the eikonal equation (4.5) we obtain

$$\|\nabla\phi(\boldsymbol{x})\|\sqrt{d_l^2 CV_l^2(\boldsymbol{x}) + d_t^2 CV_t^2(\boldsymbol{x}) + d_n^2 CV_n^2(\boldsymbol{x})} = 1,$$

which, by Equation (4.1), is equivalent to

$$\|\nabla\phi(\boldsymbol{x})\| CV_{\boldsymbol{d}}(\boldsymbol{x}) = 1. \tag{4.13}$$

By combining the Equations (4.12) and (4.13), we find that the CV along the propagation direction is given by

$$CV_{\boldsymbol{d}}(\boldsymbol{x}) = \frac{\rho\sqrt{d_l^2\sigma_{\mathrm{m}}^{l}(\boldsymbol{x}) + d_t^2\sigma_{\mathrm{m}}^{t}(\boldsymbol{x}) + d_n^2\sigma_{\mathrm{m}}^{n}(\boldsymbol{x})}}{C_{\mathrm{m}}\sqrt{\beta}}. \tag{4.14}$$

Equation (4.14) allows to estimate the parameter ρ from the CV of the monodomain model. Since the monodomain CV is affected by a spatial discretization error (Pezzuto et al. [2016]), a good estimate needs to be computed at a fine spatial resolution h. We choose $h = 0.005$ cm and we use the implicit Euler method with the usual time step $\Delta t = 0.01$ ms for the time integration of the monodomain equation. We perform a 1D numerical experiment in a tissue line of type "atrium" to compute the conduction velocity along the fibers direction and obtain $CV_l = 68$ cm/s. Since in the tissue of type "atrium" we have $\sigma_{\mathrm{m}}^{l} = 1.5$ mS/cm, by Equation (4.14) we obtain the estimate $\rho = 1.57$.

The eikonal-diffusion equation (4.10) complemented by a boundary condition and an initial condition leads to the eikonal-diffusion system

$$\begin{cases} \frac{\rho}{C_\mathrm{m}\sqrt{\beta}}\sqrt{\nabla\phi(\boldsymbol{x})^\intercal \boldsymbol{D}_\mathrm{m}(\boldsymbol{x})\nabla\phi(\boldsymbol{x})} - \frac{1}{C_\mathrm{m}\beta}\nabla\cdot\big(\boldsymbol{D}_\mathrm{m}(\boldsymbol{x})\nabla\phi(\boldsymbol{x})\big) = 1, & \boldsymbol{x}\in\Omega, \\ \boldsymbol{n}(\boldsymbol{x})^\intercal \boldsymbol{D}_\mathrm{m}(\boldsymbol{x})\nabla\phi(\boldsymbol{x}) = 0, & \boldsymbol{x}\in\partial\Omega, \\ \phi(\boldsymbol{x}) = 0, & \boldsymbol{x}\in\Omega_0, \end{cases} \tag{4.15}$$

where $\boldsymbol{n}\colon \partial\Omega \to \mathbb{R}^3$ is the unit outward normal to $\partial\Omega$ and $\Omega_0 \subset \Omega$ denotes the region where the action potential is initiated. The eikonal-diffusion equation (4.10) accounts for the wavefront curvature effects that are not captured by the pure eikonal equation (4.11). Moreover, the eikonal-diffusion system (4.15) can be solved numerically by the finite element method. Contrary to the space-time monodomain equation (4.6), the eikonal-diffusion equation (4.10) only depends on the spatial component. Therefore, the eikonal-diffusion model is computationally less costly than the monodomain model. However, since in the eikonal-diffusion equation (4.10) the advection term is dominant over the diffusion term, the spatial resolution h of the computational mesh needs to be sufficiently small to guarantee the numerical stability. The pure eikonal equation (4.11) has no diffusion term, so its numerical solution is not subject to constraints on the spatial resolution h of the computational mesh. Moreover, the numerical solution of the pure eikonal equation (4.11) can be computed algorithmically following procedures that iteratively move through the nodes of the mesh and update the activation times. Several iterative procedures are available, so that there is no need to use the finite element method. Therefore, the computational cost of the pure eikonal model is lower than the cost of the eikonal-diffusion model. Furthermore, the eikonal-diffusion equation (4.10) describes the activation times of a single activation and is not appropriate to model re-entrant activities during which several activations occur. The pure eikonal equation (4.11) also describes a single activation, but its algorithmic numerical solution methods are potentially more suitable to be extended to model the re-excitability and multiple activations. Therefore, in this work we focus on the pure eikonal model. Our goal is to compare the eikonal low fidelity model to the monodomain high fidelity model of Chapter 3, where the spatial resolution is $h = 0.02$ cm. Since the spatial discretization affects the monodomain CV (Pezzuto et al. [2016]), the high fidelity model does not achieve the CV predicted by Equation (4.14). Therefore, we prefer the formulation (4.5) of the eikonal equation over the formulation (4.11), as it allows to set the CVs computed with the monodomain high fidelity model. The eikonal equation (4.5) complemented by an initial condition leads to the eikonal system

$$\begin{cases} CV_l(\boldsymbol{x})\sqrt{\nabla\phi(\boldsymbol{x})^\intercal \boldsymbol{D}(\boldsymbol{x})\nabla\phi(\boldsymbol{x})} = 1, & \boldsymbol{x} \in \Omega, \\ \phi(\boldsymbol{x}) = 0, & \boldsymbol{x} \in \Omega_0. \end{cases} \tag{4.16}$$

In this work, we consider the eikonal model on the atrial surface $\mathscr{S} \subset \Omega$. Thus we only consider the fibers direction $\boldsymbol{f}_l \colon \mathscr{S} \to \mathbb{R}^3$ and the transversal direction $\boldsymbol{f}_t \colon \mathscr{S} \to \mathbb{R}^3$, which span the tangent space of $\mathscr{S}$, so that the tensor $\boldsymbol{D} \colon \mathscr{S} \to \mathbb{R}^{3\times 3}$ describing the anisotropy is given by

$$\boldsymbol{D}(\boldsymbol{x}) = \boldsymbol{f}_l(\boldsymbol{x})\boldsymbol{f}_l(\boldsymbol{x})^\intercal + \frac{CV_t(\boldsymbol{x})^2}{CV_l(\boldsymbol{x})^2}\boldsymbol{f}_t(\boldsymbol{x})\boldsymbol{f}_t(\boldsymbol{x})^\intercal.$$

Moreover, the gradient of the activation time is the surface gradient

$$\nabla_{\mathscr{S}}\phi(\boldsymbol{x}) = \nabla\phi(\boldsymbol{x}) - \big(\nabla\phi(\boldsymbol{x})^\intercal \boldsymbol{n}(\boldsymbol{x})\big)\boldsymbol{n}(\boldsymbol{x}),$$

where $\boldsymbol{n} \colon \mathscr{S} \to \mathbb{R}^3$ is the unit normal to $\mathscr{S}$. The eikonal system formulated on the surface reads

$$\begin{cases} CV_l(\boldsymbol{x})\sqrt{\nabla_{\mathscr{S}}\phi(\boldsymbol{x})^\intercal \boldsymbol{D}(\boldsymbol{x})\nabla_{\mathscr{S}}\phi(\boldsymbol{x})} = 1, & \boldsymbol{x} \in \mathscr{S}, \\ \phi(\boldsymbol{x}) = 0, & \boldsymbol{x} \in \mathscr{S}_0, \end{cases} \tag{4.17}$$

where $\mathscr{S}_0 \subset \mathscr{S}$ denotes the region where the action potential is initiated.

4.1.3 Viscosity solution

In this subsection we formulate the solution of the eikonal equation (4.17). The eikonal equation (4.17) is a static Hamilton-Jacobi equation of the form

$$\begin{cases} H\big(\boldsymbol{x}, \nabla_{\mathscr{S}}\phi(\boldsymbol{x})\big) = 0, & \boldsymbol{x} \in \mathscr{S}, \\ \phi(\boldsymbol{x}) = \phi_0(\boldsymbol{x}), & \boldsymbol{x}_0 \in \mathscr{S}_0, \end{cases}$$

where H is the Hamiltonian and ϕ_0 is the Dirichlet data. The solutions of the Hamilton-Jacobi equations have been extensively studied in Lions [1982]. Since, in general, there are many solutions that satisfy the Hamilton-Jacobi equation almost everywhere, the concept of viscosity solution has been introduced. The viscosity solutions additionally satisfy an entropy condition where $\nabla_{\mathscr{S}}\phi$ does not exist and there is no initial condition. Under suitable assumptions on the Hamiltonian and the Dirichlet data, the viscosity solution exists and is unique (Lions [1982], Bornemann and Rasch [2006]). The viscosity solution is given by the Hopf-Lax formula (Lions [1982], Bornemann and Rasch [2006]), that, in the case of the eikonal equation (4.17), reads (Mirebeau [2012], Mirebeau [2014])

$$\phi(\boldsymbol{x}) = \inf_{\substack{\boldsymbol{\psi}\in C^1([0,1],\mathscr{S}),\\ \boldsymbol{\psi}(0)\in\mathscr{S}_0,\ \boldsymbol{\psi}(1)=\boldsymbol{x}}} \operatorname{length}(\boldsymbol{\psi}), \tag{4.18}$$

where the length of the path $\boldsymbol{\psi}$ connecting $\mathscr{S}_0$ to $\boldsymbol{x} \in \mathscr{S}$ is

$$\operatorname{length}(\boldsymbol{\psi}) = \int_0^1 CV_l^{-1}\big(\boldsymbol{\psi}(s)\big)\sqrt{\boldsymbol{\psi}'(s)^\intercal \boldsymbol{D}^{-1}\big(\boldsymbol{\psi}(s)\big)\boldsymbol{\psi}'(s)}\,\mathrm{d}s. \tag{4.19}$$

4.2 Restitution curves of the high fidelity model

AF is characterized by re-entrant wavelets that chaotically propagate on the atrial tissue. During an AF event, the atrial cells undergo several consecutive activations. The cells alternate between the active, the repolarized and the excitable states. The interval between consecutive activations is called cardiac cycle. The eikonal equation (4.17) describes the activation times of a single activation. Since we aim to simulate AF events using the eikonal model, we have to include the re-excitability to allow for multiple activations. Our goal is to maximize the similarity between the eikonal low fidelity model and the high fidelity model of Chapter 3. Therefore, the re-excitability properties that we include in the eikonal model are based on the monodomain model. In this section we numerically compute the required re-excitability properties of the high fidelity model.

The shape of the action potential is determined by the ionic model. Figure 4.1 shows the action potential obtained with the ionic model of Section 3.1. The transmembrane potential v resulting from two stimuli delivered with a cycle length of 300 ms is computed with the high fidelity model at a fixed location and plotted over time. One cardiac cycle starts with the activation and consists of the APD and the DI. In this work, we define the activation time as the time instant when the transmembrane potential reaches the −62 mV threshold and we define the APD as the time interval during which the transmembrane potential is above the −62 mV threshold. The APD is followed by the DI, which is the time interval until the next activation. After the action potential following an activation is over, the cell needs some resting time before becoming re-excitable. Therefore, for the next activation to be possible, the DI has to be larger than a threshold $\mathrm{DI}_{\mathrm{min}}$. Moreover, the APD and the CV of a stimulus depend on the DI of the previous cardiac cycle through the so-called restitution curves (Colli Franzone et al. [2014]).

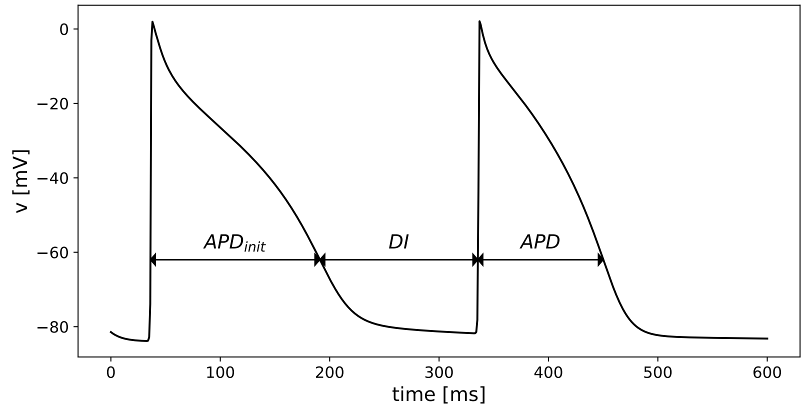

Figure 4.1. Action potential over time at a fixed location.

We compute the APD and CV restitution curves in 1D numerical experiments on a line of tissue of length 5 cm using the high fidelity model. Typically, the restitution curves are computed by delivering several stimuli with a fixed cycle length and then measuring the APD and the CV of the last stimulus (Xie et al. [2002], Cherry et al. [2008], Wilhelms et al. [2013]). The purpose of the sequence of stimuli is to reach the equilibrium values of the DI, APD and CV associated to the cycle length. To test this methodology, we measure the APD and the CV of 100 stimuli delivered with a cycle length of 300 ms on the tissue of type "atrium" and we report the results in the plots of Figure 4.2. The values of the APD do not stabilize, meaning that the APD and the DI do not reach an equilibrium value. Similarly, the values of the CV oscillate and do not reach an equilibrium as well. The reason for the instabilities is likely the drift of some ionic concentrations due to the lack of the charge conservation, which is caused by the stimulations

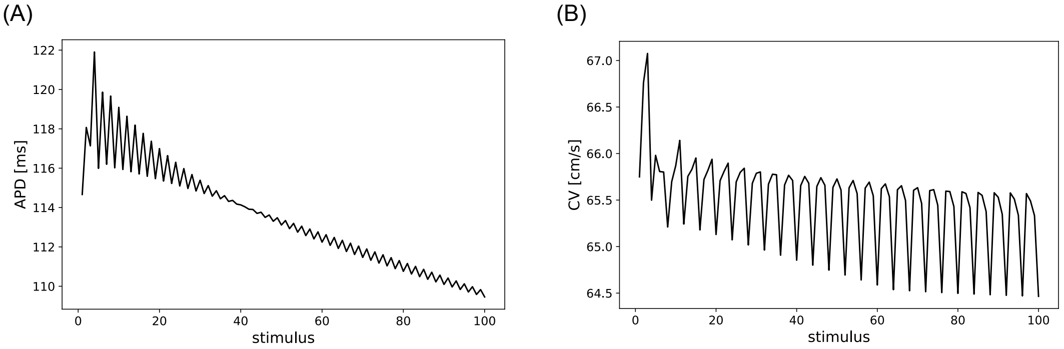

Figure 4.2. APDs (panel A) and CVs (panel B) of the stimuli delivered with a cycle length of 300 ms.

(Colli Franzone et al. [2014]). Specifically, the stimulations carry some charges that break the charge conservation property of the ionic model. It is possible to modify the ionic model equations to account for the applied current and to maintain the charge conservation (Colli Franzone et al. [2014]). Our current implementation of the Courtemanche model does not include this modification. An alternative is to compute the restitution curves on a ring with periodic boundary conditions, so that one tissue stimulation is sufficient to obtain multiple activations (Courtemanche et al. [1996]). In this work, we are interested in the short term comparison to the high fidelity model. Therefore, we compute the restitution curves by delivering only two stimuli and measuring the APD and the CV of the second stimulus. The two stimuli have size 0.2 cm and are delivered from one end of the tissue line.

In the numerical experiments of this chapter, we consider the healthy tissue of type "atrium" and the fibrotic tissue of type "diffuse fibrosis". Therefore, we compute the restitution curves on both tissue types. But first, we compute the initial CV on both tissues. To do so, we compute the activation times along the line following the first stimulation. Figure 4.3 shows the activation times obtained on the healthy and the fibrotic tissues. As expected, the activation times on the fibrotic tissue are higher than the activation times on the healthy tissue. The activation times determine the CVs. The initial conduction velocity on the healthy tissue is $CV_{l,\mathrm{init,h}} = 65$ cm/s. On the fibrotic tissue the initial conduction velocity is $CV_{l,\mathrm{init,f}} = 28$ cm/s. Notice that on the healthy tissue we have $\sigma^l_{\mathrm{m}} = \sigma_{\mathrm{h}} = 1.5$ mS/cm and on the fibrotic tissue we have $\sigma^l_{\mathrm{m}} = \sigma_{\mathrm{f}} = 0.375$ mS/cm $= 0.25 \cdot \sigma_{\mathrm{h}}$. By Equation (4.14), the ratio between the fibrotic and the healthy CVs should be $\sqrt{\sigma_{\mathrm{f}}/\sigma_{\mathrm{h}}} = 0.5$. However, our numerical ratio is $CV_{l,\mathrm{init,f}}/CV_{l,\mathrm{init,h}} = 0.43$. The discrepancy is due to the numerical error

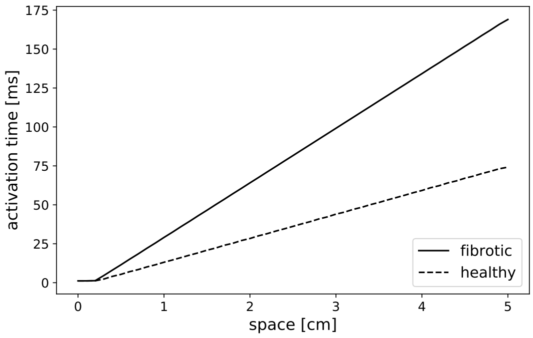

Figure 4.3. Activation times on the fibrotic (black line) and the healthy (dashed line) tissues.

introduced by the space discretization (Pezzuto et al. [2016]).

Now we focus on the numerical computation of the restitution curves. We perform 1D simulations with cycle lengths ranging between 150 ms and 3150 ms. The APD of the second stimulus is computed from the time instants when the potential is equal to −62 mV at a fixed location. The CV of the second stimulus is computed from the activation times at two locations. In the restitution curves the APD and the CV are plotted against the DI. The restitution curves of the APD are shown in Figures 4.4 and 4.5, panels A, for the healthy and the fibrotic tissues, respectively. The APD curves of the healthy and the fibrotic tissues do not match. Moreover, both APD curves show a discontinuity. To investigate whether the mismatch and the discontinuity are due to numerical discretization errors, we compute the restitution curves for the finer spatial resolutions $h = 0.01$ cm and $h = 0.005$ cm. In these two cases, the time integration of the monodomain equation is done implicitly with the usual time step $\Delta t = 0.01$ ms. The restitution curves obtained at the usual resolution $h = 0.02$ cm and at the finer resolutions $h = 0.01$ cm and $h = 0.005$ cm are shown in Figures 4.4 and 4.5 for the healthy and the fibrotic tissues, respectively. The panels A show the restitution curves of the APD and the panels B show the restitution curves of the CV. The size of the discontinuity in the APD restitution curve decreases as h decreases. This suggests that the APD restitution curve converges to a continuous function as $h \to 0$. Moreover, the mismatch between the APD restitution curves on the healthy and the fibrotic tissues disappears for $h = 0.005$ cm. The CV restitution curves show a similar shape. Only the CV restitution curve computed on the healthy tissue with $h = 0.005$ cm slightly differs from the others. However, the CV restitution curves are shifted because the finer space discretizations have the effect of increasing the CV (Pezzuto et al. [2016]).

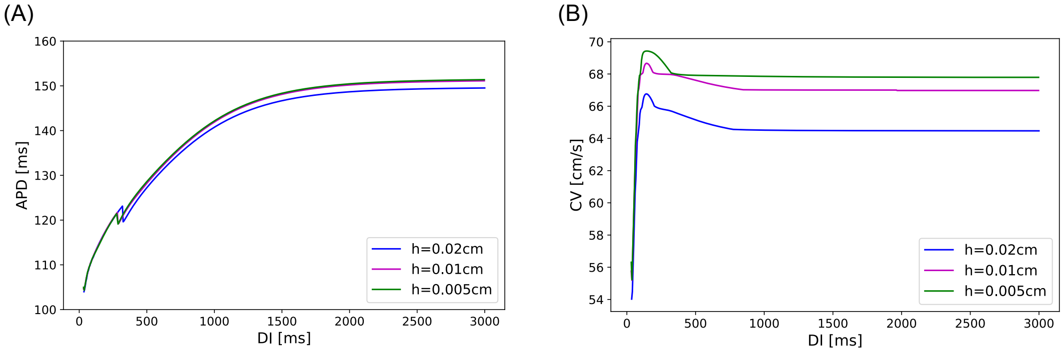

Figure 4.4. Restitution curves of the APD (panel A) and the CV (panel B) at various resolutions for the healthy tissue.

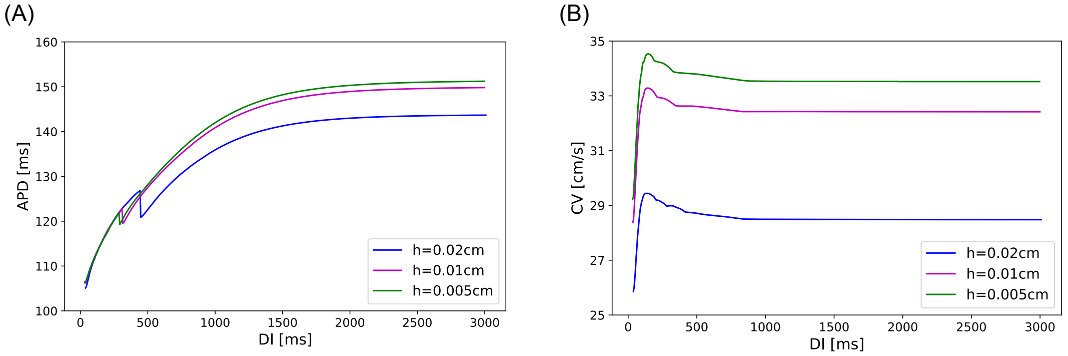

Figure 4.5. Restitution curves of the APD (panel A) and the CV (panel B) at various resolutions for the fibrotic tissue.

To summarize, the restitution curve of the APD is independent of the tissue type and is a smoothed version of the curves computed with $h = 0.005$ cm. Therefore, to obtain the restitution curve of the APD, we apply a Savitzky-Golay filter (Savitzky and Golay [1964]) to the curve of Figure 4.4, panel A, computed with $h = 0.005$ cm. The result is shown in Figure 4.6, panel A. Instead, the restitution curve of the CV depends on the tissue type, but only in terms of scaling, not in terms of shape. Therefore, to obtain the restitution curves of the CV, we rescale the curve of Figure 4.4, panel B, computed with $h = 0.005$ cm to match the initial conduction velocities $CV_{l,\text{init,h}} = 65$ cm/s and $CV_{l,\text{init,f}} = 28$ cm/s. The results are shown in Figure 4.6, panel B, for the healthy and the fibrotic tissues. Moreover, to determine the initial action potential duration APD_{init} and the threshold DI_{min}, we exploit the computations performed at the resolution $h = 0.005$ cm. The APD of the first stimulus is $\text{APD}_{\text{init}} = 158$ ms and the minimal DI for which the second stimulus propagates is $\text{DI}_{\text{min}} = 32$ ms.

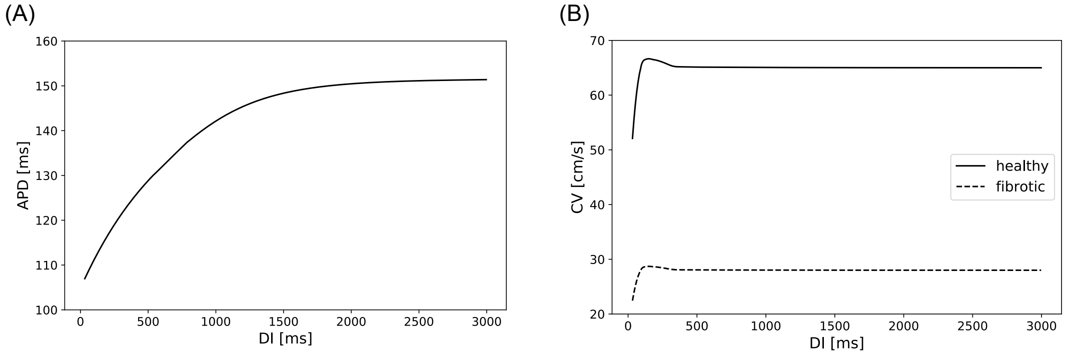

Figure 4.6. Restitution curve of the APD (panel A) and the CV (panel B).

4.3 Eikonal algorithm

The numerical solution of the eikonal equation (4.17) is computed on a triangulated mesh discretizing the atrial surface $\mathscr{S}$. Several algorithms have been developed for the solution of the eikonal equation. The common procedure is to iteratively pass through the nodes of the mesh until acceptance of the corresponding activation times. Our goal is to adapt an existing algorithm to account for the re-excitability of the tissue. In particular, a global time variable is needed in order to determine when the action potential on an activated node is over, to compute the DI and to determine when a previously activated node becomes re-excitable. Therefore, the underlying iterative procedure has to guarantee the monotonous increase of the global time variable and has to be single-pass, meaning that the activation time of each node is accepted only once per cardiac cycle. These requirements exclude the FIM (Fu et al. [2011]) and the method presented in Bornemann and Rasch [2006] based on an adaptive Gauss-Seidel iteration, and determine our choice of the FMM (Kimmel and Sethian [1998], Sethian and Vladimirsky [2000]). The FMM is a Dijkstra-like method that allows node updates coming from the triangles of the mesh, in contrast to the standard Dijkstra's method (Dijkstra [1959]) that only considers updates from the edges. Our adaptation of the FMM additionally needs to account for the variation of the APD and the CV determined by the restitution curves. To this end, a time-stepping needs to be introduced to perform the update of the DI. In this section we present the adaptation of the FMM.

The standard FMM is introduced in Subsection 4.3.1 and its convergence to the viscosity solution of the eikonal equation is analyzed in Subsection 4.3.2. The adaptation of the FMM to include the computation of the global time variable, the time-stepping and the restitution properties is presented in Subsection 4.3.3. The 2D simulations of a spiral wave in which we compare the eikonal model with re-excitability and the monodomain model in an homogeneous isotropic case are presented in Subsection 4.3.4.

4.3.1 Fast marching method

In this subsection we introduce the standard FMM. The FMM is an iterative algorithm to solve the eikonal equation (4.17) on a triangulated mesh discretizing $\mathscr{S}$. In the original formulation of the FMM, the conduction velocity CV_l along the fibers and the anisotropy tensor $\boldsymbol{D}$ are assigned to the nodes of the mesh. In this work, the CV_l and $\boldsymbol{D}$ are determined from the monodomain high fidelity model of Chapter 3. Specifically, the CV_l and $\boldsymbol{D}$ assigned to the eikonal computational mesh depend on the tissue type and the fibers direction assigned to the hexahedrons of the monodomain computational mesh. Therefore, in order to facilitate the mapping between the two meshes, here we slightly deviate from the original FMM and we assign the CV_l and $\boldsymbol{D}$ to the triangles of the mesh. The algorithm computes the activation times ϕ on the nodes of the mesh. We denote by $\boldsymbol{X}$ the set of nodes of the mesh and by $\boldsymbol{X}_0 \subset \boldsymbol{X}$ the set of nodes in $\mathscr{S}_0$. The iterative procedure starts from the nodes $\boldsymbol{x}_0 \in \boldsymbol{X}_0$ and moves through all the nodes $\boldsymbol{x} \in \boldsymbol{X}$. During the procedure three states are associated to the nodes, namely *Accepted*, *Considered* and *Far*. The *Accepted* nodes are those whose activation time is known. The *Considered* nodes have a temporary activation time that has not yet been accepted. The *Far* nodes have no activation time. At each iteration the node among the *Considered* with minimal temporary activation time is moved to the *Accepted* and its *Far* neighbors are moved to the *Considered*. Moreover, the temporary activation times of the *Considered* nodes are updated at each iteration. The pseudo-code of the FMM is reported in Algorithm 1. For a node $\boldsymbol{x} \in \boldsymbol{X}$, the set of neighbors is denoted by $N(\boldsymbol{x})$ and the set of triangles having $\boldsymbol{x}$ as vertex is denoted by $\mathscr{T}(\boldsymbol{x})$. The temporary activation time is denoted by $\tilde{\phi}$.

Algorithm 1: Fast marching method

```
# Initialization
for $\boldsymbol{x}_0 \in \boldsymbol{X}_0$ do
    tag $\boldsymbol{x}_0$ as Accepted
    set $\phi(\boldsymbol{x}_0) = 0$
for $\boldsymbol{x} \in \boldsymbol{X} \backslash \boldsymbol{X}_0$ do
    tag $\boldsymbol{x}$ as Far
    set $\phi(\boldsymbol{x}) = \infty$

# Update the neighbors
for $\boldsymbol{x}_n \in N(\boldsymbol{X}_0) \backslash Accepted$ do
    tag $\boldsymbol{x}_\mathrm{n}$ as Considered
    compute $\tilde{\phi}(\boldsymbol{x}_\mathrm{n}) = \mathrm{HL}\Big(\phi\big(N(\boldsymbol{x}_\mathrm{n})\big), CV_l\big(\mathscr{T}(\boldsymbol{x}_\mathrm{n})\big), \boldsymbol{D}\big(\mathscr{T}(\boldsymbol{x}_\mathrm{n})\big)\Big)$

# Iterate over the nodes
while $Considered \neq \emptyset$ do
    tag $\boldsymbol{x}_\mathrm{a} = \mathrm{argmin}_{\boldsymbol{x} \in Considered} \tilde{\phi}(\boldsymbol{x})$ as Accepted
    set $\phi(\boldsymbol{x}_\mathrm{a}) = \tilde{\phi}(\boldsymbol{x}_\mathrm{a})$
    for $\boldsymbol{x}_n \in N(\boldsymbol{x}_a) \backslash Accepted$ do
        tag $\boldsymbol{x}_\mathrm{n}$ as Considered
        compute $\tilde{\phi}(\boldsymbol{x}_\mathrm{n}) = \mathrm{HL}\Big(\phi\big(N(\boldsymbol{x}_\mathrm{n})\big), CV_l\big(\mathscr{T}(\boldsymbol{x}_\mathrm{n})\big), \boldsymbol{D}\big(\mathscr{T}(\boldsymbol{x}_\mathrm{n})\big)\Big)$
```

The HL function determining the temporary activation time of a node $\boldsymbol{x} \in \boldsymbol{X}$ is based on a local discretization of the Hopf-Lax formula (4.18), (4.19) given by (Mirebeau [2012], Mirebeau [2014])

$$\begin{aligned}\tilde{\phi}(\boldsymbol{x}) &= \mathrm{HL}\Big(\phi\big(N(\boldsymbol{x})\big), CV_l\big(\mathscr{T}(\boldsymbol{x})\big), \boldsymbol{D}\big(\mathscr{T}(\boldsymbol{x})\big)\Big) \\ &= \min_{\mathscr{T}\in\mathscr{T}(\boldsymbol{x})} \Big\{\min_{\boldsymbol{\lambda}} \Big\{\phi_{\boldsymbol{\lambda}} + CV_l(\mathscr{T})^{-1}\sqrt{(\boldsymbol{x}-\boldsymbol{x}_{\boldsymbol{\lambda}})^{\intercal}\boldsymbol{D}^{-1}(\mathscr{T})(\boldsymbol{x}-\boldsymbol{x}_{\boldsymbol{\lambda}})}\Big\}\Big\}, \qquad (4.20)\end{aligned}$$

where, denoting by $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ the other vertices of a triangle $\mathscr{T} \in \mathscr{T}(\boldsymbol{x})$, we have $\boldsymbol{x}_{\boldsymbol{\lambda}} = \lambda_1\boldsymbol{x}_1 + \lambda_2\boldsymbol{x}_2$ with $\lambda_1, \lambda_2 \geq 0$ and $\lambda_1 + \lambda_2 = 1$, and $\phi_{\boldsymbol{\lambda}}$ is a first order approximation of $\phi(\boldsymbol{x}_{\boldsymbol{\lambda}})$ given by $\phi_{\boldsymbol{\lambda}} = \lambda_1\phi(\boldsymbol{x}_1) + \lambda_2\phi(\boldsymbol{x}_2)$. The Hopf-Lax update (4.20) consists of two minimizations. One minimization is over the set $\mathscr{T}(\boldsymbol{x})$ of triangles having $\boldsymbol{x}$ as vertex. To solve this minimization problem, we simply compute the objective function for all triangles $\mathscr{T} \in \mathscr{T}(\boldsymbol{x})$ in order to find the minimizer. The other minimization is over the points $\boldsymbol{x}_{\boldsymbol{\lambda}}$ on the edge of the triangle $\mathscr{T}$ opposite to $\boldsymbol{x}$, which are determined by λ_1 and λ_2. Specifically, for each triangle $\mathscr{T} \in \mathscr{T}(\boldsymbol{x})$, we need to solve

$$\hat{\phi}_{\mathscr{T}}(\boldsymbol{x}) = \min_{\boldsymbol{\lambda}} \Big\{\phi_{\boldsymbol{\lambda}} + CV_l(\mathscr{T})^{-1}\sqrt{(\boldsymbol{x}-\boldsymbol{x}_{\boldsymbol{\lambda}})^{\intercal}\boldsymbol{D}^{-1}(\mathscr{T})(\boldsymbol{x}-\boldsymbol{x}_{\boldsymbol{\lambda}})}\Big\}. \qquad (4.21)$$

This objective function is convex, which implies that a unique minimum exists. To solve this minimization problem, we include the equality constraint in the objective function, which becomes

$$\mathscr{L}(\boldsymbol{\lambda},\mu) = \lambda_1\phi_1 + \lambda_2\phi_2 + CV_l(\mathscr{T})^{-1}\sqrt{(\boldsymbol{x}-\boldsymbol{x}_{\boldsymbol{\lambda}})^{\intercal}\boldsymbol{D}^{-1}(\mathscr{T})(\boldsymbol{x}-\boldsymbol{x}_{\boldsymbol{\lambda}})} + \mu(1-\lambda_1-\lambda_2),$$

where μ is the Lagrange multiplier and we use the notations $\phi_1 := \phi(\boldsymbol{x}_1)$ and $\phi_2 := \phi(\boldsymbol{x}_2)$. For $i = 1, 2$ and using the notation $\boldsymbol{p}_i := \boldsymbol{x} - \boldsymbol{x}_i$, the optimality conditions read

$$0 = \frac{\partial}{\partial\lambda_i}\mathscr{L}(\hat{\boldsymbol{\lambda}},\hat{\mu}) = \phi_i + CV_l(\mathscr{T})^{-1}\frac{\boldsymbol{p}_i^{\intercal}\boldsymbol{D}^{-1}(\mathscr{T})(\boldsymbol{x}-\boldsymbol{x}_{\hat{\boldsymbol{\lambda}}})}{\sqrt{(\boldsymbol{x}-\boldsymbol{x}_{\hat{\boldsymbol{\lambda}}})^{\intercal}\boldsymbol{D}^{-1}(\mathscr{T})(\boldsymbol{x}-\boldsymbol{x}_{\hat{\boldsymbol{\lambda}}})}} - \hat{\mu}.$$

These conditions are equivalent, for $i = 1, 2$, to

$$\hat{\mu} = \phi_i + CV_l(\mathscr{T})^{-1}\frac{\boldsymbol{p}_i^{\intercal}\boldsymbol{D}^{-1}(\mathscr{T})(\boldsymbol{x}-\boldsymbol{x}_{\hat{\boldsymbol{\lambda}}})}{\sqrt{(\boldsymbol{x}-\boldsymbol{x}_{\hat{\boldsymbol{\lambda}}})^{\intercal}\boldsymbol{D}^{-1}(\mathscr{T})(\boldsymbol{x}-\boldsymbol{x}_{\hat{\boldsymbol{\lambda}}})}}.$$

By combining these conditions with the equality constraint on $\hat{\lambda}_1$, $\hat{\lambda}_2$ and by the formulation of the minimization problem (4.21), we get

$$\begin{aligned}\hat{\mu} = (\hat{\lambda}_1 + \hat{\lambda}_2)\hat{\mu} &= \hat{\lambda}_1\phi_1 + \hat{\lambda}_2\phi_2 + CV_l(\mathscr{T})^{-1}\frac{(\hat{\lambda}_1\boldsymbol{p}_1 + \hat{\lambda}_2\boldsymbol{p}_2)^{\intercal}\boldsymbol{D}^{-1}(\mathscr{T})(\boldsymbol{x}-\boldsymbol{x}_{\hat{\boldsymbol{\lambda}}})}{\sqrt{(\boldsymbol{x}-\boldsymbol{x}_{\hat{\boldsymbol{\lambda}}})^{\intercal}\boldsymbol{D}^{-1}(\mathscr{T})(\boldsymbol{x}-\boldsymbol{x}_{\hat{\boldsymbol{\lambda}}})}} \\ &= \phi_{\hat{\boldsymbol{\lambda}}} + CV_l(\mathscr{T})^{-1}\sqrt{(\boldsymbol{x}-\boldsymbol{x}_{\hat{\boldsymbol{\lambda}}})^{\intercal}\boldsymbol{D}^{-1}(\mathscr{T})(\boldsymbol{x}-\boldsymbol{x}_{\hat{\boldsymbol{\lambda}}})} = \hat{\phi}_{\mathscr{T}}(\boldsymbol{x}).\end{aligned}$$

Therefore the optimality conditions are equivalent, for $i = 1, 2$, to

$$\hat{\phi}_{\mathcal{T}}(\boldsymbol{x}) - \phi_i = CV_l(\mathcal{T})^{-1} \frac{\boldsymbol{p}_i^{\mathsf{T}} \boldsymbol{D}^{-1}(\mathcal{T})(\boldsymbol{x} - \boldsymbol{x}_{\hat{\lambda}})}{\sqrt{(\boldsymbol{x} - \boldsymbol{x}_{\hat{\lambda}})^{\mathsf{T}} \boldsymbol{D}^{-1}(\mathcal{T})(\boldsymbol{x} - \boldsymbol{x}_{\hat{\lambda}})}}.$$

Denoting by $\boldsymbol{P}$ the matrix whose rows are $\boldsymbol{p}_1$, $\boldsymbol{p}_2$ and using the notations $\boldsymbol{\phi}_{1,2} := (\phi_1, \phi_2)^{\mathsf{T}}$ and $\mathbf{1} := (1, 1)^{\mathsf{T}}$, the optimality conditions lead to

$$\boldsymbol{P}^{-1}(\hat{\phi}_{\mathcal{T}}(\boldsymbol{x}) \cdot \mathbf{1} - \boldsymbol{\phi}_{1,2}) = CV_l(\mathcal{T})^{-1} \frac{\boldsymbol{D}^{-1}(\mathcal{T})(\boldsymbol{x} - \boldsymbol{x}_{\hat{\lambda}})}{\sqrt{(\boldsymbol{x} - \boldsymbol{x}_{\hat{\lambda}})^{\mathsf{T}} \boldsymbol{D}^{-1}(\mathcal{T})(\boldsymbol{x} - \boldsymbol{x}_{\hat{\lambda}})}}.$$

By squaring the expression and rearranging some terms, we obtain

$$(\hat{\phi}_{\mathcal{T}}(\boldsymbol{x}) \cdot \mathbf{1} - \boldsymbol{\phi}_{1,2})^{\mathsf{T}} \boldsymbol{P}^{-\mathsf{T}} \boldsymbol{D}(\mathcal{T}) \boldsymbol{P}^{-1} (\hat{\phi}_{\mathcal{T}}(\boldsymbol{x}) \cdot \mathbf{1} - \boldsymbol{\phi}_{1,2}) = CV_l(\mathcal{T})^{-2},$$

which is a quadratic equation for $\hat{\phi}_{\mathcal{T}}(\boldsymbol{x})$. Note that the quadratic equation corresponds to the discretization of the eikonal equation (4.5) on $\mathcal{T}$ with linear ϕ. The quadratic equation has either zero, one or two solutions. A solution $\hat{\phi}_{\mathcal{T}}(\boldsymbol{x})$ is valid only if the conditions $\hat{\lambda}_1, \hat{\lambda}_2 \geq 0$ are satisfied. With the notation $\hat{\boldsymbol{\lambda}} := (\hat{\lambda}_1, \hat{\lambda}_2)^{\mathsf{T}}$ and since $\boldsymbol{x} - \boldsymbol{x}_{\hat{\lambda}} = \boldsymbol{P}^{\mathsf{T}} \hat{\boldsymbol{\lambda}}$, we obtain

$$\hat{\phi}_{\mathcal{T}}(\boldsymbol{x}) \cdot \mathbf{1} - \boldsymbol{\phi}_{1,2} = CV_l(\mathcal{T})^{-1} \frac{\boldsymbol{P} \boldsymbol{D}^{-1}(\mathcal{T}) \boldsymbol{P}^{\mathsf{T}} \hat{\boldsymbol{\lambda}}}{\sqrt{(\boldsymbol{x} - \boldsymbol{x}_{\hat{\lambda}})^{\mathsf{T}} \boldsymbol{D}^{-1}(\mathcal{T})(\boldsymbol{x} - \boldsymbol{x}_{\hat{\lambda}})}}. \tag{4.22}$$

Therefore it is sufficient to check the sign of $\boldsymbol{P}^{-\mathsf{T}} \boldsymbol{D}(\mathcal{T}) \boldsymbol{P}^{-1} (\hat{\phi}_{\mathcal{T}}(\boldsymbol{x}) \cdot \mathbf{1} - \boldsymbol{\phi}_{1,2})$ to verify the condition $\hat{\boldsymbol{\lambda}} \geq 0$ and the validity of the solution $\hat{\phi}_{\mathcal{T}}(\boldsymbol{x})$. If both solutions of the quadratic equation are valid, then the solution to the optimization problem is the minimum among them. If the quadratic equation has no solution or none of the solutions is valid, then we can conclude that the update of the temporary activation time $\tilde{\phi}(\boldsymbol{x})$ does not come from the considered triangle $\mathcal{T} \in \mathcal{T}(\boldsymbol{x})$.

4.3.2 Convergence to the viscosity solution

In this subsection we analyze the convergence of the FMM. The FMM determines an ordering of the nodes and sequentially updates the temporary activation times exploiting the local discrete Hopf-Lax formula (4.20). The convergence of the solution obtained with the FMM to the viscosity solution of the eikonal equation given by the Hopf-Lax formula (4.18)-(4.19) depends on the ordering. It has been shown that the convergence is guaranteed if the causality property is satisfied (Kimmel and Sethian [1998], Sethian and Vladimirsky [2000], Mirebeau

[2012], Mirebeau [2014]). Locally, the causality property is satisfied if, assuming that the update for a node $\boldsymbol{x} \in \boldsymbol{X}$ comes from within the triangle $\mathscr{T} \in \mathscr{T}(\boldsymbol{x})$ with vertices $\boldsymbol{x}$, $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$, then $\tilde{\phi}(\boldsymbol{x}) > \max\{\phi(\boldsymbol{x}_1), \phi(\boldsymbol{x}_2)\}$.

By Equation (4.22), $\hat{\phi}_{\mathscr{T}}(\boldsymbol{x}) \cdot \boldsymbol{1} - \boldsymbol{\phi}_{1,2}$ has the same sign as

$$\boldsymbol{P}\boldsymbol{D}^{-1}(\mathscr{T})\boldsymbol{P}^{\intercal}\hat{\boldsymbol{\lambda}} = \begin{bmatrix} \boldsymbol{p}_1^{\intercal}\boldsymbol{D}^{-1}(\mathscr{T})(\hat{\lambda}_1\boldsymbol{p}_1 + \hat{\lambda}_2\boldsymbol{p}_2) \\ \boldsymbol{p}_2^{\intercal}\boldsymbol{D}^{-1}(\mathscr{T})(\hat{\lambda}_1\boldsymbol{p}_1 + \hat{\lambda}_2\boldsymbol{p}_2) \end{bmatrix}. \tag{4.23}$$

The fact that the update comes from within the triangle $\mathscr{T}$ means that $\hat{\lambda}_1, \hat{\lambda}_2 > 0$. Moreover, since the anisotropy tensor $\boldsymbol{D}(\mathscr{T})$ is symmetric positive definite, so is its inverse $\boldsymbol{D}^{-1}(\mathscr{T})$. Thus we have $\boldsymbol{p}_1^{\intercal}\boldsymbol{D}^{-1}(\mathscr{T})\boldsymbol{p}_1 > 0$ and $\boldsymbol{p}_2^{\intercal}\boldsymbol{D}^{-1}(\mathscr{T})\boldsymbol{p}_2 > 0$. As a consequence, a sufficient condition for the terms in Equation (4.23) to be positive is

$$\boldsymbol{p}_1^{\intercal}\boldsymbol{D}^{-1}(\mathscr{T})\boldsymbol{p}_2 = \boldsymbol{p}_2^{\intercal}\boldsymbol{D}^{-1}(\mathscr{T})\boldsymbol{p}_1 > 0. \tag{4.24}$$

Therefore, a sufficient condition for the causality property to be satisfied globally is that the condition (4.24) holds for all $\mathscr{T} \in \mathscr{T}(\boldsymbol{x})$, for all $\boldsymbol{x} \in \boldsymbol{X}$.

In the isotropic case we have $\boldsymbol{D}(\mathscr{T}) = \boldsymbol{I}$ for all triangles $\mathscr{T}$ of the mesh, so requiring the condition (4.24) to hold globally is equivalent to requiring that all the triangles of the mesh are acute. In the anisotropic case the acuteness condition is meant with respect to the metric defined by $\boldsymbol{D}^{-1}(\mathscr{T})$, for all triangles $\mathscr{T}$ of the mesh. In the numerical experiments of this chapter, we use the Mmg software (Dobrzynski and Frey [2008]) to create meshes in which the number of triangles satisfying the acuteness condition, with respect to the considered metrics, is maximized. To illustrate the dependence of the FMM solution on the triangulation, we consider an anisotropic square with side length of 1 cm and an action potential initiated on the lower left corner. The anisotropy is defined by the diagonal fibers direction, the longitudinal conduction velocity $CV_l = 65$ cm/s and the anisotropy ratio $CV_t/CV_l = 0.33$. We compute the activation times with the FMM method on a mesh adapted to the metric determined by the anisotropy tensor and on a non-adapted mesh. Both meshes have resolution $h = 0.1$ cm. The results are shown in Figure 4.7, panels A and B. Moreover, to illustrate the importance of considering updates coming from the triangles, in panel C we show the activation times computed with the edge Dijkstra's method on the adapted mesh. The solution obtained with the FMM method on the adapted mesh is more accurate than the one computed on the non-adapted mesh and the one obtained with the Dijkstra's method.

Alternatives to the mesh adaptation approach have been proposed in Sethian and Vladimirsky [2003] and Alton and Mitchell [2012]. These are ordered upwind methods that guarantee the convergence to the viscosity solution on any

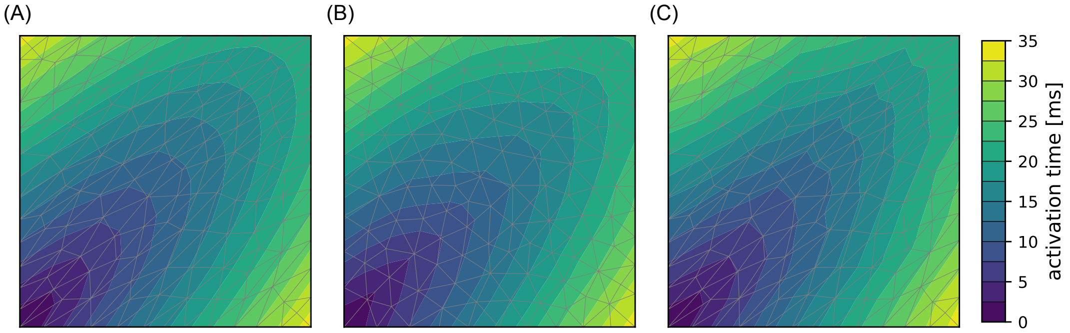

Figure 4.7. Contour plots of the activation times obtained with the FMM method on an adapted mesh (panel A), the FMM method on a non-adapted mesh (panel B) and the Dijkstra's method on an adapted mesh (panel C).

computational mesh. However, these algorithms are not suitable for our purpose of taking into account the re-excitability of the cardiac tissue. Indeed, the method from Sethian and Vladimirsky [2003] does not guarantee the monotonical acceptance of the nodes and the method from Alton and Mitchell [2012] is two-pass.

4.3.3 Algorithm with re-excitability

In this subsection we present the eikonal algorithm with re-excitability. The adaptation of the FMM of Subsection 4.3.1 to include the re-excitability requires several modifications. The activation time ϕ at the nodes of the mesh is now computed at each beat. The re-excitability is determined by the APD and the DI defined in Section 4.2, which are assigned to the nodes. When a node activates, the APD determines the time interval during which it is active and the DI determines when the node becomes re-excitable. The activation time ∞ is assigned between the beats, i.e. when the nodes are excitable. To facilitate the comparison between the eikonal low fidelity model and the monodomain high fidelity model, we define a pseudo-potential v, that we assign to the nodes. The variable v takes the value 1 when the node is active and the value 0 when the node is not active. The state *Accepted* is assigned to the non-excitable nodes. The state *Considered* is assigned to the excitable nodes whose activation times are being computed. Again, the *Considered* nodes have a temporary activation time $\tilde{\phi}$. The state *Far* is assigned to the excitable nodes that are not *Considered*. Again, the conduction velocity CV_l along the fibers and the anisotropy tensor $\boldsymbol{D}$ are assigned to the triangles of the mesh. The APD and the CV_l depend on the DI through the restitution curves. The function determining the APD from the DI

is denoted by a and is plotted in Figure 4.6, panel A. The function determining the CV_l from the tissue type and the DI is denoted by c. The function c is plotted in Figure 4.6, panel B, for the healthy tissue of type "atrium" and for the fibrotic tissue of type "diffuse fibrosis". The tissue type is assigned to the triangles and is denoted by TT.

The temporary activation time of a *Considered* node $\boldsymbol{x}$ is computed when a neighbor $\boldsymbol{x}'$ activates. When computing the temporary activation time of a *Considered* node $\boldsymbol{x}$, the CV_l of the triangles having $\boldsymbol{x}$ as vertex need to be known. Since the CV_l assigned to the triangles depends on the DI assigned to the nodes, a single triangle might take different CV_l values depending on which node is being considered. Therefore we use the notation $\mathscr{T}(\boldsymbol{x},\boldsymbol{x}')$ to denote the set of triangles having both $\boldsymbol{x}$ and $\boldsymbol{x}'$ as vertices and to refer to the propagation from $\boldsymbol{x}'$ to $\boldsymbol{x}$. The CV_l assigned to the triangles in $\mathscr{T}(\boldsymbol{x},\boldsymbol{x}')$ depends on the DI of $\boldsymbol{x}$ at the moment when $\boldsymbol{x}'$ activates. Note that $\mathscr{T}(\boldsymbol{x}',\boldsymbol{x})$ represents the same set of triangles as $\mathscr{T}(\boldsymbol{x},\boldsymbol{x}')$, but $\mathscr{T}(\boldsymbol{x}',\boldsymbol{x})$ refers to the propagation from $\boldsymbol{x}$ to $\boldsymbol{x}'$. Moreover, we use the notation $\mathscr{T}(\boldsymbol{x},:)$ to denote the set of triangles having $\boldsymbol{x}$ as vertex and to refer to the propagation to $\boldsymbol{x}$.

Since we are now taking into account the re-excitability, we can consider several action potential initiations. Therefore we consider N_{stim} stimuli. For $i = 1,...,N_{\text{stim}}$, the i-th stimulus is delivered from $\mathscr{S}_i \subset \mathscr{S}$ at time t_i. We denote by $\boldsymbol{X}_i \subset \boldsymbol{X}$ the set of nodes in $\mathscr{S}_i$, for $i = 1,...,N_{\text{stim}}$. Note that $t_1 < ... < t_{N_{\text{stim}}}$. Algorithm 2 provides the pseudo-code of our eikonal algorithm with re-excitability. The current time of the simulation is denoted by t, the time step is denoted by Δt and the final time of the simulation is denoted by T.

The FMM implementation relies on a min-heap structure to efficiently find the *Considered* node with minimal temporary activation time at each iteration. With such implementation, the computational complexity of the FMM is $O(N \log N)$, where N is the number of nodes in the mesh (Kimmel and Sethian [1998], Sethian and Vladimirsky [2000]). Indeed, the number of iterations needed to reach each node is $O(N)$ and, at each iteration, the heap is updated and reordered in $O(\log N)$ steps. Our implementation of Algorithms 1 and 2 is done in Python. The heap queue algorithm provided by the `heapq` module in Python is inappropriate for our purpose, as it does not allow to update the values in the heap. In the FMM, when a node becomes *Accepted*, its *Far* neighbors enter the heap and its *Considered* neighbors, which are already in the heap, get an updated temporary activation time. The update is allowed by the `heap` class in the `scikit-fmm` module, which we will use in our implementation. Notice that our Python implementation of Algorithm 1 is considerably less efficient than the implementation of the FMM provided by the `scikit-fmm` module, which is based

Algorithm 2: Eikonal algorithm with re-excitability

```
# Initialization
for $\boldsymbol{x} \in X$ do
    tag $\boldsymbol{x}$ as Far
    set $\phi(\boldsymbol{x}) = \infty$, $v(\boldsymbol{x}) = 0$
    set $DI(\boldsymbol{x}) = \infty$, $APD(\boldsymbol{x}) = APD_{\text{init}}$, $CV_l\big(\mathscr{T}(\boldsymbol{x},:)\big) = CV_{l,\text{init}}\big(\mathscr{T}(\boldsymbol{x})\big)$

set $t = t_1$
while $t<T$ do
    # Check for stimuli
    if $t \le t_i < t + \Delta t$ for some $i = 1, ..., N_{stim}$ then
        for $\boldsymbol{x}_i \in \boldsymbol{X}_i \backslash Accepted$ do
            tag $\boldsymbol{x}_i$ as Accepted
            set $\phi(\boldsymbol{x}_i) = t_i$, $v(\boldsymbol{x}_i) = 1$
            set $APD(\boldsymbol{x}_i) = a\big(DI(\boldsymbol{x}_i)\big)$, $DI(\boldsymbol{x}_i) = 0$, $CV_l\big(\mathscr{T}(\boldsymbol{x}_i,:)\big) = 0$
            for $\boldsymbol{x}_n \in N(\boldsymbol{x}_i) \backslash (\boldsymbol{X}_i \cup Accepted)$ do
                tag $\boldsymbol{x}_{\text{n}}$ as Considered
                set $CV_l\big(\mathscr{T}(\boldsymbol{x}_{\text{n}}, \boldsymbol{x}_i)\big) = c\Big(TT\big(\mathscr{T}(\boldsymbol{x}_{\text{n}}, \boldsymbol{x}_i)\big), DI(\boldsymbol{x}_{\text{n}})\Big)$
                compute $\tilde{\phi}(\boldsymbol{x}_{\text{n}}) = \text{HL}\Big(\phi\big(N(\boldsymbol{x}_{\text{n}})\big), CV_l\big(\mathscr{T}(\boldsymbol{x}_{\text{n}},:)\big), \boldsymbol{D}\big(\mathscr{T}(\boldsymbol{x}_{\text{n}},:)\big)\Big)$

    # Iterate within one time step
    while $\min_{\boldsymbol{x} \in Considered} \tilde{\phi}(\boldsymbol{x}) < t + \Delta t$ do
        tag $\boldsymbol{x}_{\text{a}} = \text{argmin}_{\boldsymbol{x} \in Considered} \tilde{\phi}(\boldsymbol{x})$ as Accepted
        set $\phi(\boldsymbol{x}_{\text{a}}) = \tilde{\phi}(\boldsymbol{x}_{\text{a}})$, $v(\boldsymbol{x}_{\text{a}}) = 1$
        set $APD(\boldsymbol{x}_{\text{a}}) = a\big(DI(\boldsymbol{x}_{\text{a}})\big)$, $DI(\boldsymbol{x}_{\text{a}}) = 0$, $CV_l\big(\mathscr{T}(\boldsymbol{x}_{\text{a}},:)\big) = 0$
        for $\boldsymbol{x}_n \in N(\boldsymbol{x}_a) \backslash Accepted$ do
            tag $\boldsymbol{x}_{\text{n}}$ as Considered
            set $CV_l\big(\mathscr{T}(\boldsymbol{x}_{\text{n}}, \boldsymbol{x}_{\text{a}}) = c\Big(TT\big(\mathscr{T}(\boldsymbol{x}_{\text{n}}, \boldsymbol{x}_{\text{a}})\big), DI(\boldsymbol{x}_{\text{n}})\Big)$
            compute $\tilde{\phi}(\boldsymbol{x}_{\text{n}}) = \text{HL}\Big(\phi\big(N(\boldsymbol{x}_{\text{n}})\big), CV_l\big(\mathscr{T}(\boldsymbol{x}_{\text{n}},:)\big), \boldsymbol{D}\big(\mathscr{T}(\boldsymbol{x}_{\text{n}},:)\big)\Big)$

    set $t = t + \Delta t$
    # Updates for the next time step
    for $\boldsymbol{x} \in Accepted$ such that $\phi(\boldsymbol{x}) + APD(\boldsymbol{x}) < t$ do
        set $v(\boldsymbol{x}) = 0$, $DI(\boldsymbol{x}) = t - \big(\phi(\boldsymbol{x}) + APD(\boldsymbol{x})\big)$

    for $\boldsymbol{x} \notin Accepted$ do
        set $DI(\boldsymbol{x}) = DI(\boldsymbol{x}) + \Delta t$

    for $\boldsymbol{x} \in Accepted$ such that $DI(\boldsymbol{x}) \ge DI_{min}$ do
        tag $\boldsymbol{x}$ as Far
        set $\phi(\boldsymbol{x}) = \infty$
```

on C++. The adaptations to the FMM required to include the re-excitability do not influence the computational complexity, which is $O(N \log N)$ also for Algorithm 2, but further reduce the efficiency of the implementation.

4.3.4 2D simulation of a spiral wave

In this subsection we perform a 2D numerical experiment to compare the eikonal model with re-excitability to the monodomain model of Chapter 3 in the case of a re-entrant spiral wave generated on an homogeneous isotropic tissue. In particular, we consider an isotropic square tissue slab of type "atrium" with 15 cm side length. To simulate the isotropic case with the monodomain model, we choose an arbitrary fiber direction and we set the transversal conductivities equal to the longitudinal conductivities, i.e. $\sigma^t_{\mathrm{i,e}}(\boldsymbol{x}) = \sigma^l_{\mathrm{i,e}}(\boldsymbol{x}) = 3.0$ mS/cm for all $\boldsymbol{x} \in \Omega$. The monodomain simulation is performed on a hexahedral mesh with one element thickness and with spatial resolution $h = 0.02$ cm. The time step for the monodomain simulation is $\Delta t = 0.01$ ms.

The eikonal simulation is performed on a triangulated mesh. In the eikonal model we set the transversal conduction velocity equal to the longitudinal conduction velocity, i.e. $CV_t(\mathscr{T}) = CV_l(\mathscr{T})$ for any triangle $\mathscr{T}$ of the mesh. As a consequence, the anisotropy tensor is equal to the identity, i.e. $\boldsymbol{D}(\mathscr{T}) = \boldsymbol{I}$ for any triangle $\mathscr{T}$ of the mesh. The initial conduction velocity is $CV_{l,\mathrm{init}}(\mathscr{T}) = 65$ cm/s and the restitution function c is the curve plotted in panel B of Figure 4.6 corresponding to the healthy tissue, for any triangle $\mathscr{T}$ of the mesh. Moreover, we set $\mathrm{APD}_{\mathrm{init}} = 158$ ms and the restitution function a of the APD is the curve plotted in panel A of Figure 4.6. Additionally, we set $\mathrm{DI}_{\mathrm{min}} = 32$ ms. For the eikonal simulation we use the coarser spatial resolution $h = 0.1$ cm, as in Pezzuto et al. [2017] and Quaglino et al. [2018]. We generate a triangulated mesh using the Mmg software. Only 0.18% of the triangles are non-acute. Furthermore, in the eikonal algorithm we set the time step $\Delta t = 1$ ms.

We first apply a stimulus to the left side at time $t = 0$ ms. Then we apply a second stimulus that generates a spiral wave. The second stimulus is applied to the lower left rectangle of size 6 cm × 12 cm at time $t = 275$ ms. Note that at time $t = 275$ ms part of the stimulated tissue might not be excitable. In the monodomain simulation, the stimulation in the non-excitable tissue results in a quick increase of the transmembrane potential, that then rapidly decreases back to its pre-stimulation value, without affecting the overall dynamics. In the eikonal simulation, because of the implementation of the eikonal algorithm, the non-excitable tissue is not activated. We perform the monodomain and the eikonal simulations with $T = 1000$ ms and we report the results in Figure 4.8. The figure illustrates the snapshots at five time instants. The active tissue is shown in red and the inactive tissue is shown in blue. Panel A shows the monodomain simulation, in which the tissue is active where the transmembrane potential v is above the -62 mV threshold. Note that the simulation is performed on a mesh with resolution $h = 0.02$ cm, but the result is visualized on a mesh with 0.1 cm resolution. Panel B shows the eikonal simulation, in which the tissue is active where the pseudo-potential v computed by the algorithm is 1. The result is visualized on the computational mesh with resolution $h = 0.1$ cm. At time $t = 275$ ms part of the tissue is still active due to the first stimulus and the second stimulus is applied. At time $t = 275$ ms the right part of the tissue is not excitable, therefore the second stimulus can only propagate upwards. Between $t = 275$ ms and $t = 390$ ms the right part of the tissue becomes re-excitable and the second stimulus can propagate in the right direction from above. The propagation then continues towards the right and downwards. At time $t = 390$ ms the portion activated by the second stimulus becomes inactive, but not yet re-excitable. Therefore the front can not yet propagate in the left direction. Around the time $t = 420$ ms the portion activated by the second stimulus becomes re-excitable and the front can propagate also towards the left, as we can see in the snapshots at time $t = 450$ ms. The front then starts a re-entrant propagation that induces a spiral wave around the time $t = 525$ ms. The re-entry self-sustains until the end of the simulation at time $t = T = 1000$ ms, when a spiral wave is still present.

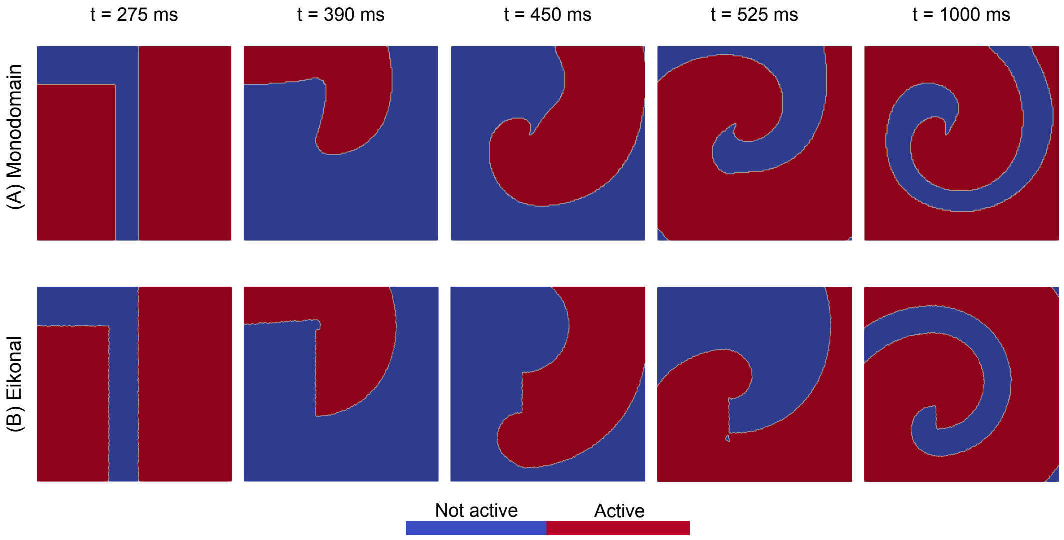

Figure 4.8. Monodomain (panel A) and eikonal (panel B) simulations of a spiral wave in the homogeneous isotropic case.

The CV of a front is not only determined by the tissue and the restitution properties. The curvature of the wavefront also affects the CV (Fast and Kléber [1997]). In particular, compared to a planar front, a convex front propagates with a reduced CV and a concave front propagates with an increased CV. This is due to the fact that, compared to a planar propagation, the current at the front diffuses to a larger tissue area during a convex propagation and to a smaller tissue area during a concave propagation. These curvature effects are captured by the monodomain model. In contrast, since the eikonal equation does not account for the diffusion currents, the eikonal model does not capture the wavefront curvature effects on the CV.

Hence, the curvature effects generate some differences between our monodomain and eikonal simulations. Since the first stimulus generates a planar front, the propagation and the following re-excitability occur in the same way in the two simulations. Therefore, at time $t = 275$ ms, when the second stimulus is applied, the snapshots correspond, except for the instantaneous effect of the stimulation of the non-excitable tissue in the monodomain simulation. The second stimulus generates a convex front, therefore after $t = 275$ ms the CV of the monodomain front is lower than the CV of the eikonal front. The faster propagation of the eikonal front can be observed in the snapshots at time $t = 390$ ms. As a consequence, compared to the monodomain re-entry, the eikonal re-entry is initiated lower in the tissue, as we can see in the snapshots at time $t = 450$ ms. This is the reason why the spiral wave initiated by the eikonal simulation is delayed compared to the one initiated by the monodomain simulation, as we can observe in the snapshots at time $t = 525$ ms. The eikonal re-entry remains shifted in space and time until the end of the simulation at time $t = T = 1000$ ms.

However, the spiral waves resulting from the two simulations have a similar shape. Moreover, compared to the monodomain result, the eikonal result is obtained with much lower computational resources. Indeed, the computing time of the monodomain simulation is 1 h 10min with Propag-5 on CSCS, whereas the computing time of the eikonal simulation is only 3 min with our Python implementation of Algorithm 2 on a laptop.

4.4 Effect of fibrosis

In this section we compare the eikonal and the monodomain models in presence of diffuse fibrosis. We first quantitatively compare the two models in the non re-entrant case. Then, we move to the re-entrant case by comparing the two models in 2D simulations of a spiral wave. The quantitative study in the non re-entrant case is presented in Subsection 4.4.1 and the 2D simulations of a re-entrant spiral wave are shown in Subsection 4.4.2.

In the 2D numerical experiments of this section we consider the isotropic case. In the eikonal model, the initial conduction velocity is $CV_{l,\mathrm{init}}(\mathscr{T}) = 65$ cm/s for the triangles $\mathscr{T}$ in the healthy regions and $CV_{l,\mathrm{init}}(\mathscr{T}) = 28$ cm/s for the triangles $\mathscr{T}$ in the fibrotic regions. Moreover, the anisotropy tensor is $\boldsymbol{D}(\mathscr{T}) = \boldsymbol{I}$ for any triangle $\mathscr{T}$ of the mesh, as the transversal conduction velocity is equal to the longitudinal conduction velocity, i.e. $CV_t(\mathscr{T}) = CV_l(\mathscr{T})$ for any triangle $\mathscr{T}$ of the mesh. In the monodomain model, we choose an arbitrary fibers direction and we set the transversal conductivities equal to the longitudinal conductivities. Thus we set $\sigma^t_{\mathrm{i,e}}(\boldsymbol{x}) = \sigma^l_{\mathrm{i,e}}(\boldsymbol{x}) = 3.0$ mS/cm for $\boldsymbol{x}$ in the healthy regions and $\sigma^t_{\mathrm{i,e}}(\boldsymbol{x}) = \sigma^l_{\mathrm{i,e}}(\boldsymbol{x}) = 0.75$ mS/cm for $\boldsymbol{x}$ in the fibrotic regions.

4.4.1 Quantitative study in the non re-entrant case

In this subsection we quantitatively compare the monodomain and the eikonal models in propagations in presence of diffuse fibrosis. Moreover, since we focus on the non re-entrant case, we also consider the eikonal-diffusion model. The work presented in this subsection has been done in collaboration with Prof. Rolf Krause, Prof. Martin Weiser, Prof. Francisco Sahli Costabal and Prof. Simone Pezzuto (Gander et al. [2023]). Some sentences are taken from Gander et al. [2023] *ad verbatim*.

In order to allow the comparison to the monodomain and the pure eikonal models, the CV of the eikonal-diffusion model needs to be adapted. To this end, we modify the parameter ρ according to Equation (4.14). In particular, we set the parameter ρ_{h} in the healthy regions and the parameter ρ_{f} in the fibrotic regions. These parameters are obtained from the longitudinal conduction velocities CV_l on the healthy and the fibrotic tissues, which are respectively 65 cm/s and 28 cm/s. Moreover, the longitudinal conductivity σ^l_{m} is $\sigma_{\mathrm{h}} = 1.5$ mS/cm in the healthy tissue and $\sigma_{\mathrm{f}} = 0.375$ mS/cm in the fibrotic tissue. Thus, by Equation (4.14), the resulting parameters are $\rho_{\mathrm{h}} = 1.51$ and $\rho_{\mathrm{f}} = 1.32$.

We first consider the 2D case of a square with 15 cm side length. The monodomain computational mesh consists of hexahedral elements with resolution $h = 0.02$ cm and has one element thickness. We sample the fibrosis on this hexahedral mesh following the procedure described in Section 3.5. The fraction of fibrotic tissue is 50% and the resulting fibrotic pattern is shown in Figure 4.9, panel A. The monodomain computational mesh can be seen as a 2D mesh of squares. The eikonal-diffusion and the pure eikonal equations are solved numerically on a triangulated mesh. We use the mesh obtained by splitting each square into two

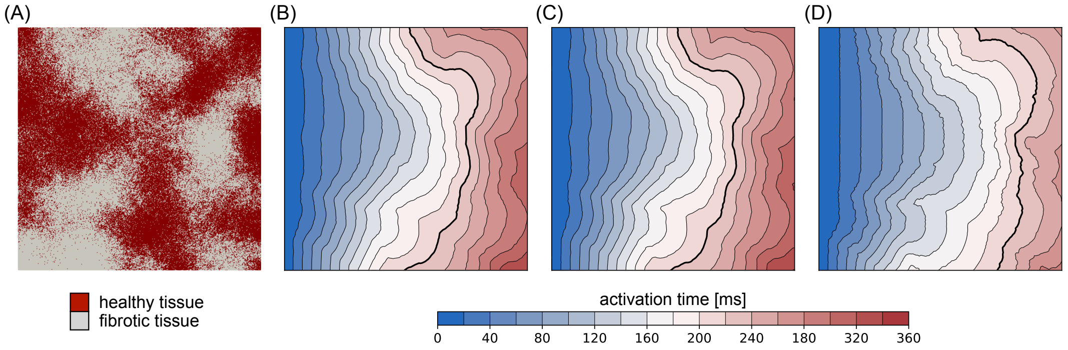

Figure 4.9. Panel A: fibrosis pattern. Panels B, C and D: contour plots of the activation times obtained with the monodomain, the eikonal-diffusion and the pure eikonal models, respectively.

triangles, so that the resolution $h = 0.02$ cm is the same as in the monodomain case and the fibrotic pattern is correctly captured. We consider a stimulus delivered from the left side of the square and we compute the activation times in the whole domain with the monodomain, the eikonal-diffusion and the pure eikonal models. The computing times of the simulations are 28 min with Propag-5 on CSCS for the monodomain model, 1 min with FEniCS on a laptop for the eikonal-diffusion model and 225 s with our Python implementation of Algorithm 1 on a laptop for the pure eikonal model. The computing time of the pure eikonal simulation is 0.25 s with the `scikit-fmm` module on a laptop. Since the maximum activation time is approximately 340 ms, the ratio between the computing and the simulation times is 250 ms/340 ms = 0.74 for the eikonal model implemented by the `scikit-fmm` module. Therefore, the eikonal model is (better than) real-time, but our Python implementation is slower by a factor of 225 s/0.25 s = 900. Note that, contrary to Algorithm 1, the FMM of the `scikit-fmm` module operates on a uniform grid where the CV is assigned to the nodes. The ratio between the computing and the simulation times is $6 \cdot 10^4$ ms/340 ms = 176 for the eikonal-diffusion model and $1.68 \cdot 10^6$ ms/340 ms = 4941 for the monodomain model. The results of the simulations are shown in panels B, C and D of Figure 4.9 for the monodomain, the eikonal-diffusion and the pure eikonal models, respectively. The activation times are reported in terms of isochrones with 20 ms spacing. The thick contour illustrates the position of the front after 220 ms. The pure eikonal activation times are smaller than the eikonal-diffusion activation times, which in turn are similar to the monodomain activation times. Indeed, the maximum difference between the eikonal-diffusion and the monodomain activation times is 7 ms, whereas the maximum difference between the pure eikonal and the monodomain activation times is 52 ms. The fact that the pure eikonal model does not account for the diffusion currents is the reason of the mismatch compared to the eikonal-diffusion and the monodomain models. The diffusion currents are due to heterogeneities in the propagation, e.g. because of the wavefront curvature, or in the conductivity, e.g. because of discontinuous coefficients. To quantify the effect of the discontinuity in the conductivity, due to the presence of fibrosis, on the eikonal approximation, we perform some 1D numerical experiments. In all the 1D cases we consider an action potential propagating on a tissue line from left to right and we compute the activation times of the monodomain, the eikonal-diffusion and the pure eikonal models.

First, we perform two numerical experiments on a tissue line of length 15 cm. We again consider the spatial resolution $h = 0.02$ cm, the healthy conductivity $\sigma_\mathrm{h} = 1.5$ mS/cm and the fibrotic conductivity $\sigma_\mathrm{f} = 0.375$ mS/cm. In the first numerical experiment, the tissue line is extracted from the 2D domain and consists of a random fibrotic pattern shown at the bottom of Figure 4.10, panel A, with 62.4% fibrotic tissue. The activation times of the first numerical experiment are reported in the plot of Figure 4.10, panel A. The activation times of the monodomain model are higher than the activation times of the eikonal-diffusion model, which in turn are slightly higher than the activation times of the pure eikonal model. This can be translated in terms of average CV, which is 31 cm/s for the monodomain model, 35 cm/s for the eikonal-diffusion model and 36 cm/s for the pure eikonal model. In the second numerical experiment, the tissue line consists of an ordered pattern in which the healthy tissue is on the left and the fibrotic tissue is on the right, as shown at the bottom of Figure 4.10, panel B. Note that the percentage of fibrotic tissue is the same as in the first numerical ex-

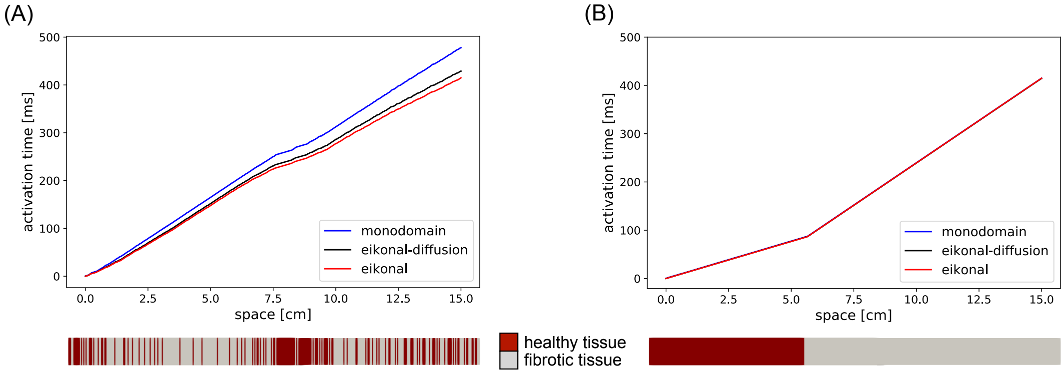

Figure 4.10. Activation times in 1D tissue with a random fibrotic pattern (panel A) and with an ordered fibrotic pattern (panel B).

periment. The activation times of the second numerical experiment are reported in the plot of Figure 4.10, panel B. The activation times of the monodomain, the eikonal-diffusion and the pure eikonal models are now very similar. As a consequence, the average CV is approximately 36 cm/s for the three models.

These observations suggest that the mismatch between the activation times of the three models is due to the discontinuities in the conductivity introduced at the boundaries between the healthy and the fibrotic tissue. Indeed, in the first numerical experiment the domain contains 151 discontinuities and the mismatch is significant, whereas in the second numerical experiment the domain contains only one discontinuity and the mismatch is negligible. To further investigate this aspect, we perform a numerical experiment in a tissue line of length 6 cm consisting of three portions of length 2 cm each. The first and the last portions are healthy tissue, whereas the middle portion is fibrotic tissue. We compare the activation times of the monodomain and the eikonal-diffusion models to the activation times of the pure eikonal model. Figure 4.11 shows the difference in the activation times compared to the pure eikonal model in this case where only two discontinuities are introduced. The plot shows two curves, one for the monodomain model and one for the eikonal-diffusion model. Note that we discard the results in the first and last portions of length 1 cm because they are affected by some boundary effects. Both curves show a negative jump as the

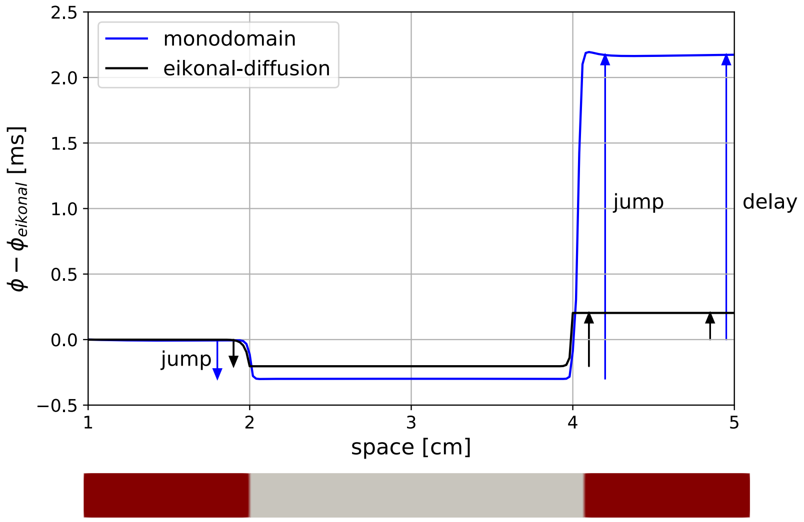

Figure 4.11. Difference in the activation times compared to the eikonal model. There is a negative jump when the front propagates from the healthy to the fibrotic tissue and a positive jump when the front propagates from the fibrotic to the healthy tissue. The sum of the two jumps is positive and is the delay.

action potential propagates from the healthy to the fibrotic tissue and a positive jump as the action potential propagates from the fibrotic to the healthy tissue. In both cases the two jumps are asymmetric, indeed the amplitude of the positive jump is higher than the amplitude of the negative jump. This asymmetry introduces a delay, which is the sum of the two jumps, in the monodomain and the eikonal-diffusion simulations compared to the eikonal simulation. When the conductivity presents many discontinuities, the total delay becomes significant, as in the numerical experiment of Figure 4.10, panel A. Moreover, the delay in the monodomain case is higher than in the eikonal-diffusion case. This explains why, in Figure 4.10, panel A, the total delay of the monodomain model is more apparent than the total delay of the eikonal-diffusion model.

To better understand the effect of the discontinuity in the conductivity on the propagations, we now consider various spatial resolutions h in the monodomain model and various fibrotic conductivities σ_{f} corresponding to different ratios of the healthy conductivity σ_{h}. In particular, we consider the coarse resolution $h = 0.04$ cm, the resolution $h = 0.02$ cm and the two finer resolutions $h = 0.01$ cm and 0.005 cm. At the resolution $h = 0.04$ cm, the time integration of the monodomain equation is done explicitly with time step $\Delta t = 0.02$ ms. Instead, at the resolutions $h = 0.01$ cm and $h = 0.005$ cm, the time integration is done implicitly with the usual time step $\Delta t = 0.01$ ms. We study 40 ratios $\sigma_{\mathrm{f}}/\sigma_{\mathrm{h}}$ ranging between 0.025 and 1 and we consider a domain consisting of two portions with different conductivities. We denote the conductivities on the left and right portions respectively by σ_{left} and σ_{right}. When setting the healthy portion to the left and the fibrotic portion to the right, the ratio $\sigma_{\mathrm{right}}/\sigma_{\mathrm{left}}$ takes our 40 values between 0.025 and 1. Instead, when setting the fibrotic portion to the left and the healthy portion to the right, the ratio $\sigma_{\mathrm{right}}/\sigma_{\mathrm{left}}$ takes 40 values between 1 and 40. For each of these 80 cases, we simulate the propagation with the monodomain and the eikonal-diffusion models and we compute the jump of the difference between the resulting activation times and the activation times of the pure eikonal model. The CVs in the fibrotic regions are computed from the monodomain simulations and are used in the pure eikonal model and for the adaptation of the parameter ρ_{f} in the eikonal-diffusion model. The jumps are plotted in Figure 4.12, panel A, against $\sqrt{\sigma_{\mathrm{right}}/\sigma_{\mathrm{left}}}$. The dashed lines are vertical asymptotes that indicate the propagation failure, which occurs with the monodomain model at all resolutions. As $h \to 0$, the monodomain curve converges towards the eikonal-diffusion curve computed with $h = 0.02$ cm, which is a straight line. Moreover, the jump is asymmetric, i.e. the amplitude of the jump at $\sqrt{\sigma_{\mathrm{h}}/\sigma_{\mathrm{f}}}$ is higher than the amplitude at $\sqrt{\sigma_{\mathrm{f}}/\sigma_{\mathrm{h}}}$. The delay introduced compared to the pure eikonal model is the sum of the jumps at $\sqrt{\sigma_{\mathrm{f}}/\sigma_{\mathrm{h}}}$ and at

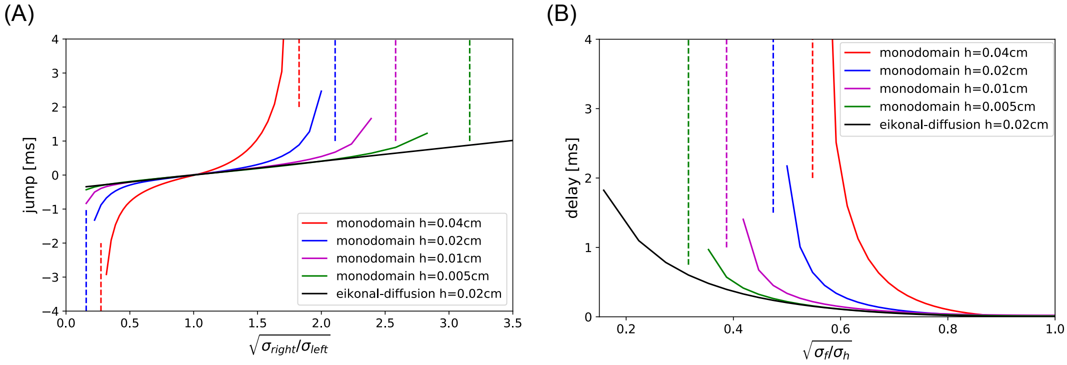

Figure 4.12. Panel A: jump of the difference in the activation times compared to the eikonal model for a front propagating from left to right. Panel B: delay introduced compared to the eikonal model.

$\sqrt{\sigma_\mathrm{h}/\sigma_\mathrm{f}}$. These delays are plotted in Figure 4.12, panel B, against $\sqrt{\sigma_\mathrm{f}/\sigma_\mathrm{h}}$. The dashed lines again illustrate the propagation failure. Again, there is convergence of the monodomain curve towards the eikonal-diffusion curve as $h \to 0$.

These numerical experiments show that the monodomain results converge towards the eikonal-diffusion results, therefore the eikonal-diffusion model is very accurate. The difference between the monodomain and the eikonal-diffusion delays depends on both the spatial resolution h and the value of $\sqrt{\sigma_\mathrm{f}/\sigma_\mathrm{h}}$, see Figure 4.12, panel B. The pure eikonal model does not capture the delays, therefore the error compared to the eikonal-diffusion model is given by the eikonal-diffusion delays, which depend on $\sqrt{\sigma_\mathrm{f}/\sigma_\mathrm{h}}$, see Figure 4.12, panel B. The fibrosis model is determined by the value of $\sqrt{\sigma_\mathrm{f}/\sigma_\mathrm{h}}$, which is inversely proportional to the difference between the healthy and the fibrotic CVs. Therefore, the accuracy of the monodomain and the pure eikonal models depends on the fibrosis model. Additionally, the accuracy of the monodomain model is also affected by the spatial resolution. To guarantee the accuracy of the monodomain model, the spatial resolution needs to be selected based on the fibrosis model. In contrast, there is no straightforward way to guarantee the accuracy of the pure eikonal model, as it only depends on the fibrosis model.

In the specific fibrosis model of this work we have $\sqrt{\sigma_\mathrm{f}/\sigma_\mathrm{h}} = 0.5$. Our monodomain spatial resolution $h = 0.02$ cm is sufficiently small to avoid artificial conduction blocks. However, the delay introduced by the monodomain model with $h = 0.02$ cm is 2.17 ms, which is significantly bigger than the delay of 0.20 ms introduced by the eikonal-diffusion model. Therefore, in 1D propagations, the pure eikonal model is more accurate than the monodomain model with $h = 0.02$ cm, as we can observe in Figure 4.10, panel A. However, in the

2D case, the monodomain model with $h = 0.02$ cm is more accurate than the pure eikonal model, as we can observe in Figure 4.9, panels B, C and D. This occurs because the effect of the wavefront curvature on the propagation is more significant than the effect of the delays introduced by the presence of fibrosis. The pure eikonal model does not capture these two effects. The monodomain model with $h = 0.02$ cm amplifies the delays but is accurate in describing the effect of the wavefront curvature.

We conclude this subsection with the derivation of a formula for the jump. We first consider the eikonal-diffusion equation (4.15) on $\Omega = [0, L]$ with $\Omega_0 = 0$ given by

$$\begin{cases} \frac{\rho}{C_\mathrm{m}\sqrt{\beta}}\sqrt{\sigma(x)}|\phi'(x)| - \frac{1}{C_\mathrm{m}\beta}\big(\sigma(x)\phi'(x)\big)' = 1, & x \in (0, L), \\ \phi(0) = 0, \quad \phi'(L) = 0 \end{cases} \tag{4.25}$$

We consider a discontinuous conductivity in $\ell \in (0, L)$, i.e.

$$\sigma(x) = \begin{cases} \sigma_\text{left}, & x \in (0, \ell), \\ \sigma_\text{right}, & x \in (\ell, L), \end{cases}$$

so that the eikonal-diffusion equation (4.25) is meant as a distributional ODE. By multiplying the eikonal-diffusion equation (4.25) by a test function $u \in H^1([0, L])$ with $u(0) = 0$ and by integrating we obtain

$$\frac{\rho}{C_\mathrm{m}\sqrt{\beta}}\int_0^L \sqrt{\sigma(x)}|\phi'(x)|u(x)\mathrm{d}x - \frac{1}{C_\mathrm{m}\beta}\int_0^L \big(\sigma(x)\phi'(x)\big)'u(x)\mathrm{d}x = \int_0^L u(x)\mathrm{d}x.$$

By integrating by parts the second term and by using the boundary conditions, we obtain

$$\begin{aligned} \int_0^L \big(\sigma(x)\phi'(x)\big)'u(x)\mathrm{d}x = & \int_0^L \sigma(x)\phi'(x)u'(x)\mathrm{d}x \\ = & -\sigma_\text{left}\int_0^\ell \phi'(x)u'(x)\mathrm{d}x - \sigma_\text{right}\int_\ell^L \phi'(x)u'(x)\mathrm{d}x \\ = & \sigma_\text{left}\int_0^\ell \phi''(x)u(x)\mathrm{d}x - \sigma_\text{left}\phi'(\ell^-)u(\ell^-) \\ & + \sigma_\text{right}\int_\ell^L \phi''(x)u(x)\mathrm{d}x + \sigma_\text{right}\phi'(\ell^+)u(\ell^+) \\ = & \int_0^L \big(\sigma(x)\phi'(x)\big)'u(x)\mathrm{d}x \\ & - \big(\sigma_\text{left}\phi'(\ell^-) - \sigma_\text{right}\phi'(\ell^+)\big)u(\ell), \end{aligned}$$

where in the last step we use the fact that the test function u is continuous. Since the equalities have to hold for every $u \in H^1([0,L])$, we deduce that $\sigma_{\mathrm{left}}\phi'(\ell^-) = \sigma_{\mathrm{right}}\phi'(\ell^+)$. The eikonal-diffusion equation (4.25) can be rewritten as

$$\begin{cases} \frac{\rho\sqrt{\sigma_{\mathrm{left}}}}{C_{\mathrm{m}}\sqrt{\beta}}|\phi'_{\mathrm{left}}(x)| - \frac{\sigma_{\mathrm{left}}}{C_{\mathrm{m}}\beta}\phi''_{\mathrm{left}}(x) = 1, & x \in (0,\ell), \\ \frac{\rho\sqrt{\sigma_{\mathrm{right}}}}{C_{\mathrm{m}}\sqrt{\beta}}|\phi'_{\mathrm{right}}(x)| - \frac{\sigma_{\mathrm{right}}}{C_{\mathrm{m}}\beta}\phi''_{\mathrm{right}}(x) = 1, & x \in (\ell,L), \\ \phi_{\mathrm{left}}(\ell) = \phi_{\mathrm{right}}(\ell), \quad \sigma_{\mathrm{left}}\phi'_{\mathrm{left}}(\ell) = \sigma_{\mathrm{right}}\phi'_{\mathrm{right}}(\ell), \\ \phi_{\mathrm{left}}(0) = 0, \quad \phi'_{\mathrm{right}}(L) = 0. \end{cases} \tag{4.26}$$

Using the change of variables

$$\hat{\phi}_{\mathrm{left/right}} = CV_{\mathrm{left/right}}\phi_{\mathrm{left/right}} = \frac{\rho\sqrt{\sigma_{\mathrm{left/right}}}}{C_{\mathrm{m}}\sqrt{\beta}}\phi_{\mathrm{left/right}},$$

obtained from Equation (4.14), and the notation

$$\varepsilon_{\mathrm{left/right}} := \frac{\sqrt{\sigma_{\mathrm{left/right}}}}{\rho\sqrt{\beta}},$$

Equation (4.26) becomes

$$\begin{cases} |\hat{\phi}'_{\mathrm{left}}(x)| - \varepsilon_{\mathrm{left}}\hat{\phi}''_{\mathrm{left}}(x) = 1, & x \in (0,\ell), \\ |\hat{\phi}'_{\mathrm{right}}(x)| - \varepsilon_{\mathrm{right}}\hat{\phi}''_{\mathrm{right}}(x) = 1, & x \in (\ell,L), \\ \frac{1}{\sqrt{\sigma_{\mathrm{left}}}}\hat{\phi}_{\mathrm{left}}(\ell) = \frac{1}{\sqrt{\sigma_{\mathrm{right}}}}\hat{\phi}_{\mathrm{right}}(\ell), & \sqrt{\sigma_{\mathrm{left}}}\hat{\phi}'_{\mathrm{left}}(\ell) = \sqrt{\sigma_{\mathrm{right}}}\hat{\phi}'_{\mathrm{right}}(\ell), \\ \hat{\phi}_{\mathrm{left}}(0) = 0, \quad \hat{\phi}'_{\mathrm{right}}(L) = 0. \end{cases} \tag{4.27}$$

Since we are considering a propagation from the left to the right of the interval $[0,L]$, the functions $\hat{\phi}'_{\mathrm{left/right}}$ are positive. Therefore, the two ODEs in Equation (4.27) have the form

$$\begin{cases} \hat{\phi}'(x) - \varepsilon\hat{\phi}''(x) = 1, & x \in (x_0, x_{\mathrm{end}}), \\ \hat{\phi}(x_0) = \hat{\phi}_0, \quad \hat{\phi}'(x_0) = \hat{\phi}'_0. \end{cases} \tag{4.28}$$

Note that ε is an estimate of the front thickness. The solution to Equation (4.28) reads

$$\hat{\phi}(x) = \hat{\phi}_0 + \varepsilon(1 - \hat{\phi}'_0) + x - x_0 + \varepsilon(\hat{\phi}'_0 - 1)\exp\big((x - x_0)/\varepsilon\big). \tag{4.29}$$

Since

$$\hat{\phi}'(x) = 1 + (\hat{\phi}'_0 - 1)\exp\big((x - x_0)/\varepsilon\big),$$

Equation (4.29) is equivalent to

$$\hat{\phi}(x) = \hat{\phi}_0 + x - x_0 + \varepsilon\big(\hat{\phi}'(x) - \hat{\phi}'_0\big). \tag{4.30}$$

The ODE for $\hat{\phi}_{\text{left}}$ in Equation (4.27) is of the form of Equation (4.28) with $\varepsilon = \varepsilon_{\text{left}}$, $x_0 = 0$ and $\hat{\phi}_0 = 0$. Moreover we have

$$\hat{\phi}'_0 = \hat{\phi}'_{\text{left}}(0) = CV_{\text{left}}\phi'_{\text{left}}(0) \approx CV_{\text{left}}\frac{1}{CV_{\text{left}}} = 1.$$

Therefore, by Equation (4.30) we get

$$\hat{\phi}_{\text{left}}(\ell) \approx \ell + \varepsilon_{\text{left}}\big(\hat{\phi}'_{\text{left}}(\ell) - 1\big). \tag{4.31}$$

Similarly, we have

$$\hat{\phi}'_{\text{right}}(\ell) = CV_{\text{right}}\phi'_{\text{right}}(\ell) \approx CV_{\text{right}}\frac{1}{CV_{\text{right}}} = 1.$$

Thus, from one of the boundary conditions in Equation (4.27), we obtain

$$\hat{\phi}'_{\text{left}}(\ell) = \sqrt{\frac{\sigma_{\text{right}}}{\sigma_{\text{left}}}}\hat{\phi}'_{\text{right}}(\ell) \approx \sqrt{\frac{\sigma_{\text{right}}}{\sigma_{\text{left}}}}.$$

By substituting this in Equation (4.31), we obtain

$$\hat{\phi}_{\text{left}}(\ell) \approx \ell + \varepsilon_{\text{left}}\left(\sqrt{\frac{\sigma_{\text{right}}}{\sigma_{\text{left}}}} - 1\right). \tag{4.32}$$

By rescaling Equation (4.32) back to the original quantity we get

$$\phi_{\text{left}}(\ell) \approx \frac{\ell}{CV_{\text{left}}} + \frac{\varepsilon_{\text{left}}}{CV_{\text{left}}}\left(\sqrt{\frac{\sigma_{\text{right}}}{\sigma_{\text{left}}}} - 1\right). \tag{4.33}$$

In order to obtain an estimate of the jump, we need the value of the activation time at ℓ obtained from the pure eikonal model. Therefore, we now consider the eikonal equation (4.16) on $\Omega = [0,\ell]$ with $\Omega_0 = 0$ given by

$$\begin{cases} CV_{\text{left}}|\phi'_{\text{eiko}}(x)| = 1, & x \in (0,\ell], \\ \phi_{\text{eiko}}(0) = 0. \end{cases} \tag{4.34}$$

Again, since we are considering a propagation from the left to the right of the interval $[0,\ell]$, the function ϕ'_{eiko} is positive. The solution to Equation (4.34) is thus

$$\phi_{\text{eiko}}(x) = \frac{x}{CV_{\text{left}}}. \tag{4.35}$$

By Equation (4.35) we get

$$\phi_{\text{eiko}}(\ell) = \frac{\ell}{CV_{\text{left}}}. \tag{4.36}$$

Finally, by combining Equations (4.33) and (4.36), we obtain the estimate of the jump given by

$$\text{jump} = \phi_{\text{left}}(\ell) - \phi_{\text{eiko}}(\ell) \approx \frac{\varepsilon_{\text{left}}}{CV_{\text{left}}}\left(\sqrt{\frac{\sigma_{\text{right}}}{\sigma_{\text{left}}}} - 1\right) = \frac{C_{\text{m}}}{\rho^2}\left(\sqrt{\frac{\sigma_{\text{right}}}{\sigma_{\text{left}}}} - 1\right). \tag{4.37}$$

Equation (4.37) highlights the asymmetry in the jump. Moreover, Equation (4.37) compares well to the numerical values for the jump obtained with the eikonal-diffusion model and shown in Figure 4.12, panel A. For example, for the case represented in Figure 4.11, the formula predicts the jumps −0.19 ms and 0.43 ms, while the numerical values are −0.20 ms and 0.41 ms.

4.4.2 2D simulation of a spiral wave

In this subsection we extend the 2D numerical experiment of Subsection 4.3.4 to the presence of diffuse fibrosis. We consider the fibrotic pattern of Figure 4.9, panel A, sampled on the hexahedral mesh with resolution $h = 0.02$ cm used for the monodomain simulation. For the eikonal simulation we use again a triangulated mesh with coarser resolution $h = 0.1$ cm. Since we are in the isotropic case, we use the mesh introduced in Subsection 4.3.4. The tissue type is assigned to the triangles based on the assignment to the hexahedral elements. Since the hexahedral mesh has one element thickness, it can also be seen as a 2D mesh of squares. We first map each square to the closest triangle. The distance is computed from the middle of the square to the centroid of the triangle. Then, for a fixed triangle, we consider all the squares mapped to it and we assign the tissue type based on a majority rule. We repeat this assignment procedure for all the mapped triangles. If there are some unmapped triangles, we assign the tissue type based on the assignment to the adjacent triangles. In particular, for each unmapped triangle, we consider all the adjacent triangles and, if at least one of those has been assigned to a tissue type, we use again a majority rule to assign the tissue type to the unmapped triangle. We repeat this procedure until all the triangles have been assigned to a tissue type. The resulting fibrotic pattern is shown in Figure 4.13, panel B, where the centroids of the triangles are colored based on the assigned tissue type. The panel A shows the original fibrotic pattern, i.e. the one of Figure 4.9, panel A. The large-scale features of the hexahedral mesh are captured by the triangulated mesh, whereas some small-scale features might not be present.

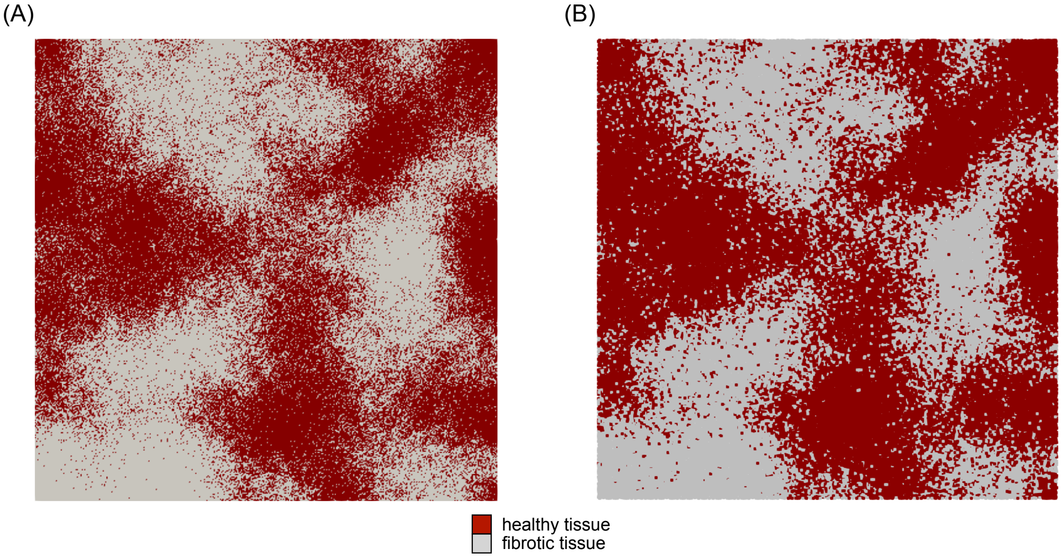

Figure 4.13. Fibrosis distribution in the hexahedral mesh with resolution 0.02 cm (panel A) and in the triangulated mesh with resolution 0.1 cm (panel B).

Since we simulate a re-entrant spiral wave, we need the restitution properties of the tissue. The restitution function c of the CV depends on the tissue type. Thus we consider the two restitution curves plotted in Figure 4.6, panel B, corresponding to the healthy and the fibrotic tissues.

Figure 4.14 shows the monodomain (panel A) and eikonal (panel B) simulations. The propagation of the first stimulus in the eikonal simulation is different from the propagation in the monodomain simulation, as we can see in the snapshots at the times $t = 275$ ms, $t = 390$ ms and $t = 450$ ms. This happens because of the observations of Subsection 4.4.1 and of the wavefront curvature affecting the monodomain propagation. Moreover, due to the presence of fibrosis, the propagation is slower than in the homogeneous case of Subsection 4.3.4. As a consequence, the portion of the tissue stimulated at time $t = 275$ ms that is excitable is smaller. Also the spiral wave generated by the second stimulus is affected by the heterogeneity of the tissue. As a consequence, since in the fibrotic regions the propagation is slower than in the healthy regions, the spiral wave is deformed, as we can observe in the snapshots at the times $t = 390$ ms, $t = 450$ ms, $t = 525$ ms and $t = T = 1000$ ms. The eikonal re-entry is shifted in space and time compared to the monodomain re-entry. This is due to the differences between the monodomain and the eikonal activation times in the boundaries of the tissue types studied in Subsection 4.4.1. Moreover, as in Subsection 4.3.4, the comparison is also affected by the effect of the wavefront cur-

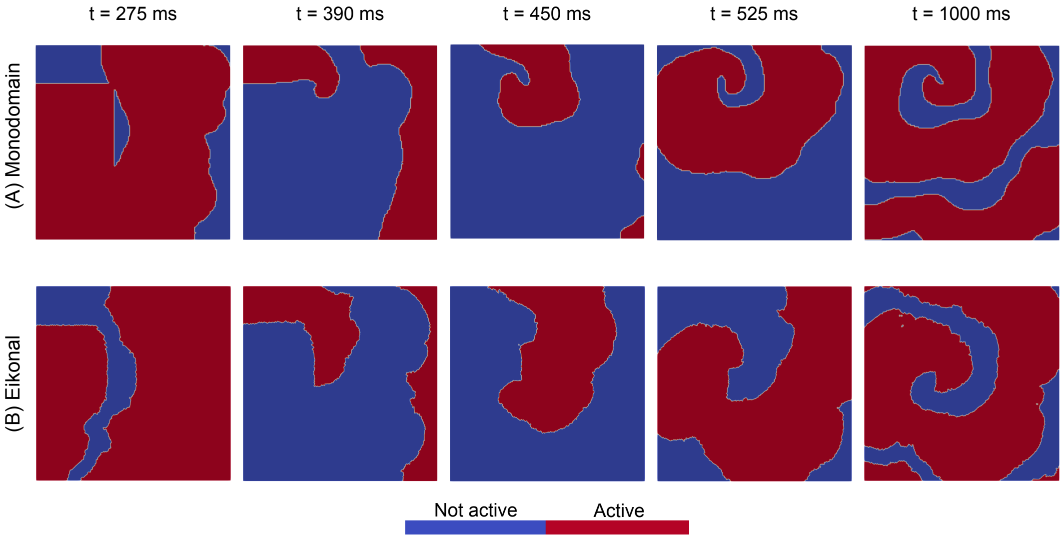

Figure 4.14. Monodomain (panel A) and eikonal (panel B) simulations of a spiral wave in the fibrotic isotropic case.

vature on the CV, which is present in the monodomain model and is not captured by the eikonal model. However, again, the shapes of the spiral waves resulting from the monodomain and the eikonal simulations are similar. Moreover, again, obtaining the the eikonal result has a much lower computational cost than obtaining the monodomain result.

4.5 Anisotropy

In this section we compare the eikonal model with re-excitability to the monodomain model in some anisotropic cases. Specifically, we extend the 2D simulation of a spiral wave shown in Subsection 4.3.4 to some anisotropic cases. In the anisotropic case, the anisotropy tensor of the eikonal model is no longer equal to the identity matrix. The anisotropy tensor $\boldsymbol{D}$ depends on the fiber direction $\boldsymbol{f}_l$ and on the anisotropy ratio CV_t/CV_l between the transversal and the longitudinal conduction velocities. Since the anisotropy ratio does not depend on the restitution, it is sufficient to know $CV_{t,\text{init}}/CV_{l,\text{init}}$. The initial transversal CV needs to be computed numerically from the monodomain model. Moreover, the triangulated mesh used in the eikonal simulation needs to be adapted to the metric defined by the anisotropy tensor, so that the acuteness condition is satisfied on the triangles. In the monodomain model the fiber orientation is encoded in the angles α, γ and θ defined in Section 3.5. The description of the anisotropy is completed by the conductivities $\sigma^l_{\mathrm{i,e}}$ and $\sigma^t_{\mathrm{i,e}}$ given in Table 3.1.

We first consider the homogeneous non-fibrotic case and then we consider the heterogeneous case in which the fibrosis is present. The simulations are presented in Subsection 4.5.1 for the homogeneous case and in Subsection 4.5.2 for the heterogeneous case.

4.5.1 2D simulation of a spiral wave in the homogeneous case

In this subsection we consider the homogeneous tissue of type "atrium" and two anisotropic cases. In the first case the fibers are horizontally oriented and in the second case the fibers are diagonally oriented. In other words, the fiber direction is $\boldsymbol{f}_l(\boldsymbol{x}) = [1, 0, 0]^\intercal$ in the first case and $\boldsymbol{f}_l(\boldsymbol{x}) = \left[\sqrt{2}/2, \sqrt{2}/2, 0\right]^\intercal$ in the second case, for all $\boldsymbol{x} \in \Omega$. To obtain these fiber directions, in the monodomain model we set $\gamma(\boldsymbol{x}) = -\pi$ and $\theta(\boldsymbol{x}) = -\pi/2$ for all $\boldsymbol{x} \in \Omega$. Moreover, we set $\alpha(\boldsymbol{x}) = 0$ in the first case and $\alpha(\boldsymbol{x}) = -\pi/4$ in the second case, for all $\boldsymbol{x} \in \Omega$. Furthermore, in the monodomain model the conductivities $\sigma^l_{\mathrm{i,e}}(\boldsymbol{x})$ and $\sigma^t_{\mathrm{i,e}}(\boldsymbol{x})$ are given by the values in Table 3.1 corresponding to the tissue type "atrium" for all $\boldsymbol{x} \in \Omega$.

In the eikonal model the fiber direction defines the anisotropy tensor together with the anisotropy ratio. The initial longitudinal conduction velocity is $CV_{l,\mathrm{init}}(\mathscr{T}) = 65$ cm/s for any triangle $\mathscr{T}$ of the mesh. The restitution function c of the CV is the curve plotted in panel B of Figure 4.6 corresponding to the healthy tissue. The initial transversal CV is obtained from a monodomain computation. We perform two 2D numerical experiments on our computational domains, i.e. the anisotropic square tissue slabs of type "atrium" with 15 cm side length. In the first numerical experiment the fibers are horizontally oriented and in the second numerical experiment the fibers are diagonally oriented. We deliver a stimulus of radius 0.1 cm from the middle of the domains that generates an action potential that propagates in all directions. We perform the simulations with the monodomain model and we compute the activation times along the fibers direction and in the transversal direction. In particular, in the first numerical experiment we compute the activation times along the horizontal and vertical directions, whereas in the second numerical experiment we compute the activation times along the two diagonal directions. Figure 4.15 shows the activation times along the longitudinal and transversal directions. Panel A shows the results of the numerical experiment with the horizontally oriented fibers and panel B shows the results of the numerical experiment with the diagonally oriented fibers. As expected, in both numerical experiments the activation times along the transversal direction are higher than the activation times along the longitudinal direction. The activation times determine the CVs. The CV along

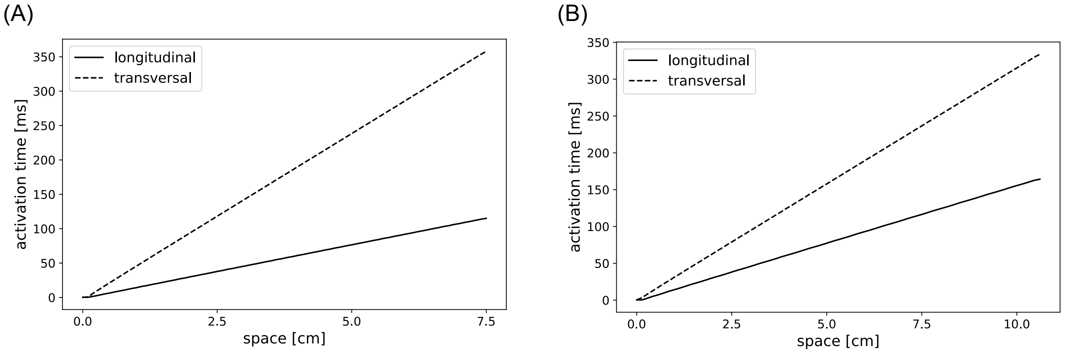

Figure 4.15. Activation times along the directions longitudinal and transversal to the fibers in a domain of type "atrium". Panel A: results for the case of horizontally oriented fibers. Panel B: results for the case of diagonally oriented fibers.

the transversal direction is different in the two numerical experiments. When the fibers are horizontally oriented, the transversal conduction velocity is 21 cm/s, so that $CV_{t,\text{init}}(\mathscr{T}) = 21$ cm/s for any triangle $\mathscr{T}$ of the mesh. As a consequence, the anisotropy ratio in the first case is $CV_{t,\text{init}}(\mathscr{T})/CV_{l,\text{init}}(\mathscr{T}) = 21/65 = 0.32$ for any triangle $\mathscr{T}$ of the mesh. When the fibers are diagonally oriented, the transversal conduction velocity is 32 cm/s, so that $CV_{t,\text{init}}(\mathscr{T}) = 32$ cm/s for any triangle $\mathscr{T}$ of the mesh. Consequently, in the second case the anisotropy ratio is $CV_{t,\text{init}}(\mathscr{T})/CV_{l,\text{init}}(\mathscr{T}) = 32/65 = 0.49$ for any triangle $\mathscr{T}$ of the mesh. The differences are due to the fact that in the second numerical experiment the fibers are not aligned with the computational grid (Pezzuto et al. [2016]). In both cases, the anisotropy ratio and the fiber direction define the anisotropy tensor $\boldsymbol{D}(\mathscr{T})$ for any triangle $\mathscr{T}$ of the mesh.

Since we change the anisotropy tensor, we have to ensure that the triangles of the mesh satisfy the acuteness condition. The mesh used in Subsection 4.3.4 is not appropriate for our anisotropic cases as the percentage of triangles not satisfying the acuteness condition is significant. Indeed, that percentage is 91.0% in the first case and 43.9% in the second case. Therefore we use the Mmg software to adapt the mesh to the metrics defined by the anisotropy tensors. This generates two adapted meshes that are appropriate for our anisotropic cases. Indeed, the adapted meshes have a less significant percentage of triangles not satisfying the acuteness condition. The percentage is only 1.1% for the first mesh and 0.9% for the second mesh. Figure 4.16 shows a detail of the meshes adapted to the cases of horizontal fibers (panel A) and diagonal fibers (panel B). Note that we ensure that, in terms of number of nodes, the two adapted meshes are

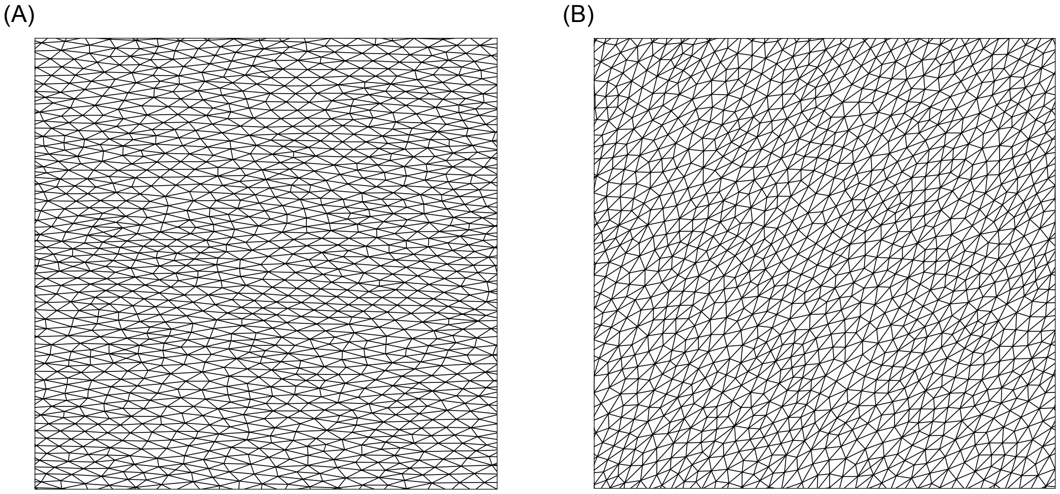

Figure 4.16. Details of the adapted meshes for the anisotropic cases with horizontally oriented fibers (panel A) and diagonally oriented fibers (panel B).

comparable to the mesh used in Subsection 4.3.4.

Figure 4.17 shows the monodomain (panel A) and the eikonal (panel B) simulations in the case of horizontally oriented fibers. The first stimulus propagates in the fibers direction, therefore, until the second stimulus is applied, the

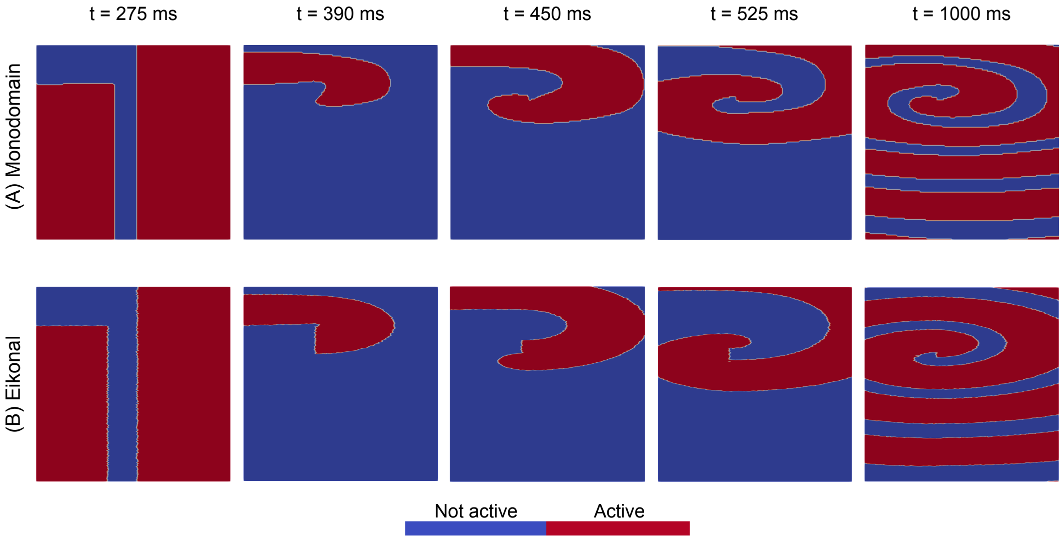

Figure 4.17. Monodomain (panel A) and eikonal (panel B) simulations of a spiral wave in the homogeneous anisotropic case with horizontally oriented fibers.

simulations are the same as those presented in Subsection 4.3.4. Indeed, the snapshots at time $t = 275$ ms correspond to the snapshots of Figure 4.8 at the same time. Instead, the second stimulus, which generates a spiral wave, propagates in all directions. In the anisotropic case the propagation in the directions non-longitudinal to the fibers is slower than in the isotropic case. Moreover, the velocity decreases as the angle between the propagation direction and the fibers direction increases. Therefore, the spiral wave generated in this anisotropic case is deformed horizontally, as we can see in the snapshots at the times $t = 390$ ms, $t = 450$ ms, $t = 525$ ms and $t = T = 1000$ ms.

Figure 4.18 shows the monodomain (panel A) and the eikonal (panel B) simulations in the case of diagonally oriented fibers. In this case the first stimulus does not propagate in the fibers direction. Therefore the propagation is slower than in the isotropic case of Subsection 4.3.4, mainly in the lower part of the domain. As a consequence, the portion of the tissue stimulated at time $t = 275$ ms that is excitable is smaller. The snapshots at the times $t = 275$ ms, $t = 390$ ms and $t = 450$ ms show a slight mismatch between the CV of the first stimulus in the comparison between the monodomain and the eikonal simulations. In particular, the monodomain propagation is slightly slower than the eikonal propagation. This is due to the wavefront curvature affecting the monodomain CV. Again, the second stimulus generates a spiral wave and propagates in all directions. Since

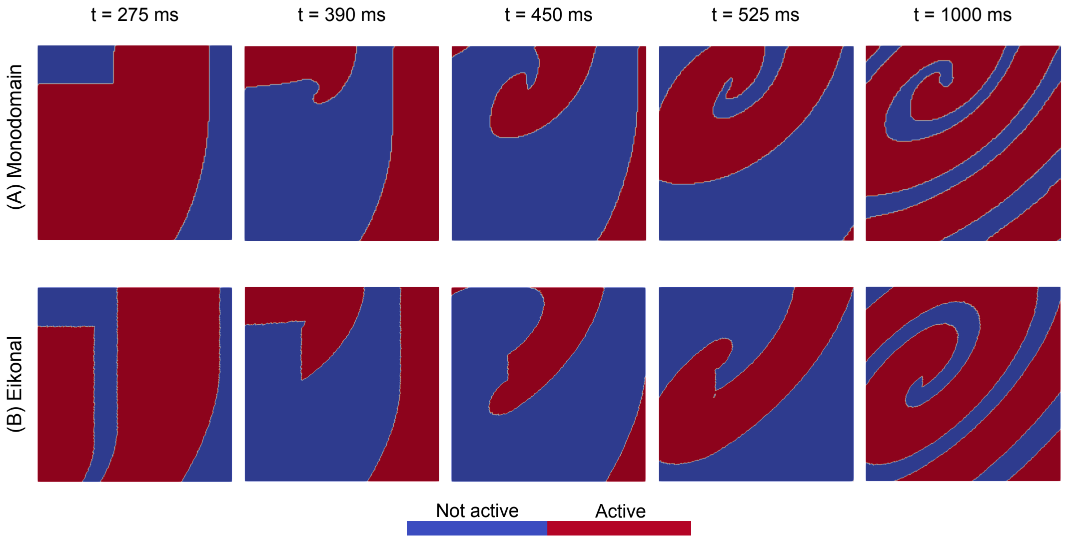

Figure 4.18. Monodomain (panel A) and eikonal (panel B) simulations of a spiral wave in the homogeneous anisotropic case with diagonally oriented fibers.

we are in an anisotropic case, the velocity decreases as the propagation direction deviates from the fibers direction. Therefore, the spiral wave generated in this case is deformed diagonally, as we can see in the snapshots at the times $t = 390$ ms, $t = 450$ ms, $t = 525$ ms and $t = T = 1000$ ms.

In both anisotropic cases, the eikonal re-entry is shifted in space and time compared to the monodomain re-entry. This is again mainly due to the effect of the wavefront curvature on the CV, which is captured by the monodomain model and is not present in the eikonal model. However, again, the spiral waves resulting from the eikonal simulations have a similar shape to those resulting from the monodomain simulations. Moreover, again, obtaining the eikonal results is computationally much less demanding than obtaining the monodomain results.

4.5.2 2D simulation of a spiral wave in the heterogeneous case

In this subsection we extend the two anisotropic cases of Subsection 4.5.1 to the presence of diffuse fibrosis. Thus we consider horizontally and diagonally oriented fibers, with the corresponding fiber direction $\boldsymbol{f}_l(\boldsymbol{x})$ and angles $\alpha(\boldsymbol{x})$, $\gamma(\boldsymbol{x})$, $\theta(\boldsymbol{x})$ given in Subsection 4.5.1 for all $\boldsymbol{x} \in \Omega$. Moreover, we consider the fibrotic pattern of Figure 4.9, panel A. In the monodomain model the heterogeneity of the tissue is encoded in the conductivities. For $\boldsymbol{x}$ in the healthy regions the conductivities $\sigma^l_{\mathrm{i,e}}(\boldsymbol{x})$ and $\sigma^t_{\mathrm{i,e}}(\boldsymbol{x})$ are given by the values in Table 3.1 corresponding to the tissue type "atrium", whereas for $\boldsymbol{x}$ in the fibrotic regions the conductivities $\sigma^l_{\mathrm{i,e}}(\boldsymbol{x})$ and $\sigma^t_{\mathrm{i,e}}(\boldsymbol{x})$ are given by the values in Table 3.1 corresponding to the tissue type "diffuse fibrosis".

In the eikonal model, the heterogeneity of the tissue is encoded in the CVs. The initial longitudinal conduction velocity is $CV_{l,\mathrm{init}}(\mathscr{T}) = 65$ cm/s for the triangles $\mathscr{T}$ in the healthy regions and $CV_{l,\mathrm{init}}(\mathscr{T}) = 28$ cm/s for the triangles $\mathscr{T}$ in the fibrotic regions. The restitution function c of the CV depends on the tissue type and is given by the two curves in Figure 4.6, panel B, corresponding to the healthy and the fibrotic tissues. For the triangles $\mathscr{T}$ in the healthy regions, the initial transversal conduction velocity is $CV_{t,\mathrm{init}}(\mathscr{T}) = 21$ cm/s when the fibers are horizontally oriented and $CV_{t,\mathrm{init}}(\mathscr{T}) = 32$ cm/s when the fibers are diagonally oriented. Thus, for the healthy triangles $\mathscr{T}$, the anisotropy ratio is $CV_{t,\mathrm{init}}(\mathscr{T})/CV_{l,\mathrm{init}}(\mathscr{T}) = 21/65 = 0.32$ for the case of horizontally oriented fibers and $CV_{t,\mathrm{init}}(\mathscr{T})/CV_{l,\mathrm{init}}(\mathscr{T}) = 32/65 = 0.49$ for the case of diagonally oriented fibers. The initial transversal CV in the fibrotic regions needs to be determined numerically from the monodomain model. We repeat the two numerical experiments of Subsection 4.5.1 in two anisotropic square tissue slabs of type "diffuse fibrosis". The obtained activation times along the longitudinal and transversal directions are shown in Figure 4.19. The results corresponding to the case of horizontally oriented fibers is shown in panel A and the results corresponding to the case of diagonally oriented fibers is shown in panel B. The activation times determine the CVs. When the fibers are horizontally oriented, the transversal conduction velocity is 6 cm/s, so that $CV_{t,\text{init}}(\mathscr{T}) = 6$ cm/s for the triangles $\mathscr{T}$ in the fibrotic regions. Consequently, in the first case the anisotropy ratio is $CV_{t,\text{init}}(\mathscr{T})/CV_{l,\text{init}}(\mathscr{T}) = 6/28 = 0.21$ for the fibrotic triangles $\mathscr{T}$. When the fibers are diagonally oriented, the transversal conduction velocity is 15 cm/s, so that $CV_{t,\text{init}}(\mathscr{T}) = 15$ cm/s for the triangles $\mathscr{T}$ in the fibrotic regions. As a consequence, in the second case the anisotropy ratio is $CV_{t,\text{init}}(\mathscr{T})/CV_{l,\text{init}}(\mathscr{T}) = 15/28 = 0.54$ for the fibrotic triangles $\mathscr{T}$. Again, the differences between the two anisotropic cases are due to the fact that in the second case the fibers are not aligned with the computational grid (Pezzuto et al. [2016]). In both anisotropic cases, the fiber direction and the anisotropy ratios define two anisotropy tensors $\boldsymbol{D}(\mathscr{T})$, one for the healthy triangles $\mathscr{T}$ and one for the fibrotic triangles $\mathscr{T}$.

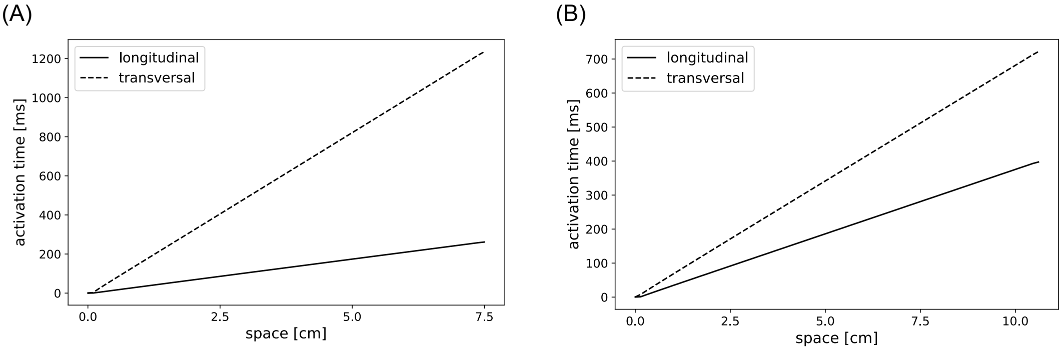

Figure 4.19. Activation times along the directions longitudinal and transversal to the fibers in a domain of type "diffuse fibrosis". Panel A: results for the case of horizontally oriented fibers. Panel B: results for the case of diagonally oriented fibers.

Since we are considering the anisotropic cases, we have to adapt the computational mesh to the metrics defined by the anisotropy tensors, in order to maximize the percentage of triangles satisfying the acuteness condition. Moreover, since in both anisotropic cases the anisotropy ratios of the healthy and the fibrotic tissue differ, we can not use the same meshes as in Subsection 4.5.1. Therefore, we repeat the process of generating two meshes from the mesh of Subsection 4.3.4 using the Mmg software. The percentage of triangles of the mesh of Subsection 4.3.4 not satisfying the acuteness condition is 93.5% for the case of the horizontal fiber direction and 38.6% for the case of the diagonal fiber direction. This percentage reduces to 4.7% for the adapted mesh of the case of horizontally oriented fibers and to 0.9% for the adapted mesh of the case of diagonally oriented fibers. The tissue type is assigned to the triangles of the adapted meshes following the procedure described in Subsection 4.4.2. Again, the triangulated meshes reproduce the large-scale features of the hexahedral mesh, whereas some small-scale features might not be captured.

The monodomain and the eikonal simulations are shown in Figures 4.20 and 4.21 for the cases of horizontally and diagonally oriented fibers, respectively. The monodomain simulations are shown in panels A and the eikonal simulation are shown in panels B. The figures show the snapshots at the times $t = 275$ ms, $t = 390$ ms, $t = 450$ ms, $t = 525$ ms and $t = T = 1000$ ms. The propagations are slower than in the homogeneous isotropic case of Subsection 4.3.4 because of both the anisotropy and the presence of fibrosis. Moreover, for the same reasons, the spiral waves are deformed. Again, there are differences between the monodomain and the eikonal simulations, both in the propagation of the first stimulus and in the re-entry generated by the second stimulus. These differences are due to two main reasons. First, the tissue heterogeneity due to the presence of fibrosis introduces a mismatch in the activation times, as observed in Subsection 4.4.1. Second, the wavefront curvature has an effect on the CV which is captured by the monodomain model and is not present in the eikonal model.

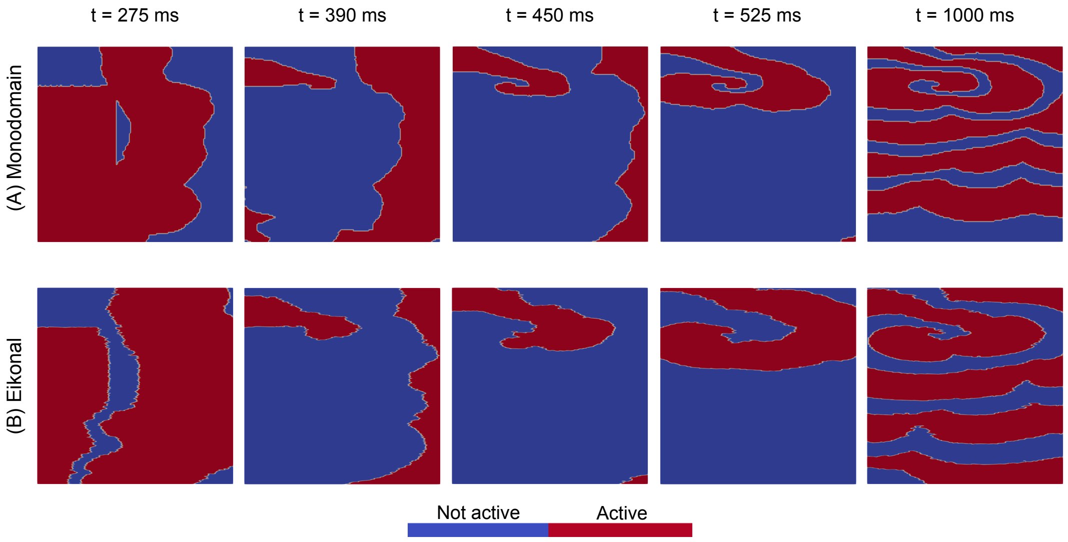

Figure 4.20. Monodomain (panel A) and eikonal (panel B) simulations of a spiral wave in the fibrotic anisotropic case with horizontally oriented fibers.

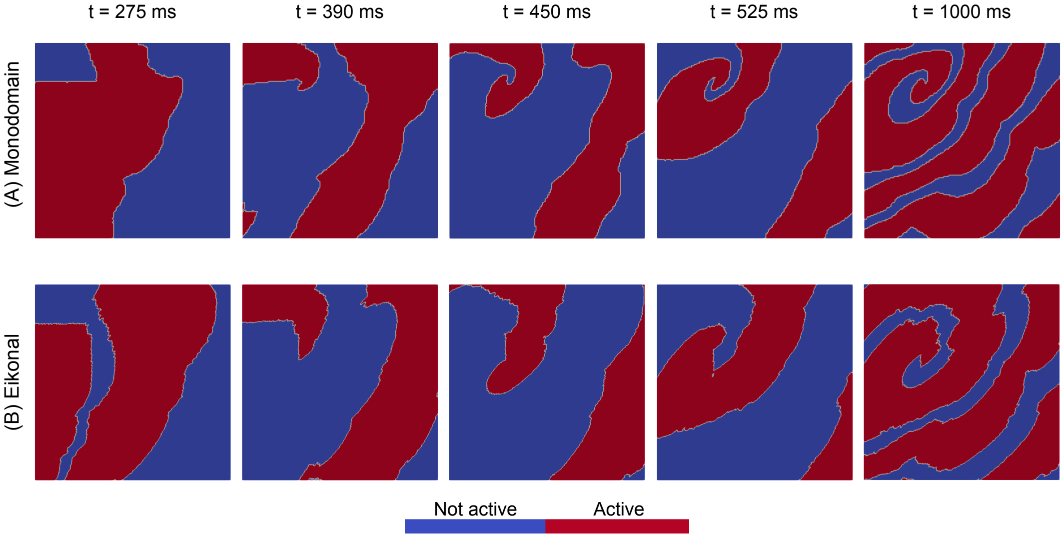

Figure 4.21. Monodomain (panel A) and eikonal (panel B) simulations of a spiral wave in the fibrotic anisotropic case with diagonally oriented fibers.

However, again, the results obtained with the cheap eikonal model are similar to the results obtained with the much more expensive monodomain model.

4.6 2D numerical experiments in presence of scars

In this section we perform some 2D numerical experiments in presence of non-conductive scars. In patients with ventricular tachycardia, the initiation of the arrhythmia is related to the presence of conductive channels between the scars (Pernod et al. [2011], Gionti et al. [2022]). We consider a square tissue slab with side length of 15 cm and with the two scars shown in Figure 4.22, panel A. We simulate the electrical activity during and after an S1-S2 stimulation. The first stimulus is applied at time $t = 0$ ms in the middle of the gap between the two scars, whereas the second stimulus is applied at the bottom of the gap between the two scars, see Figure 4.22, panel A. Both stimuli have radius 0.5 cm. The action potential generated by the first stimulus propagates in two directions around the scars. If the second stimulus is applied early, the surrounding tissue is not excitable and the generated action potential can not propagate. The result is a vanishing action potential, as illustrated in Figure 4.22, panel B. If the second stimulus is applied on time, only one side of the surrounding tissue is excitable and the generated action potential only propagates upwards. Subsequently, the action potential rotates around the scars and reaches again the

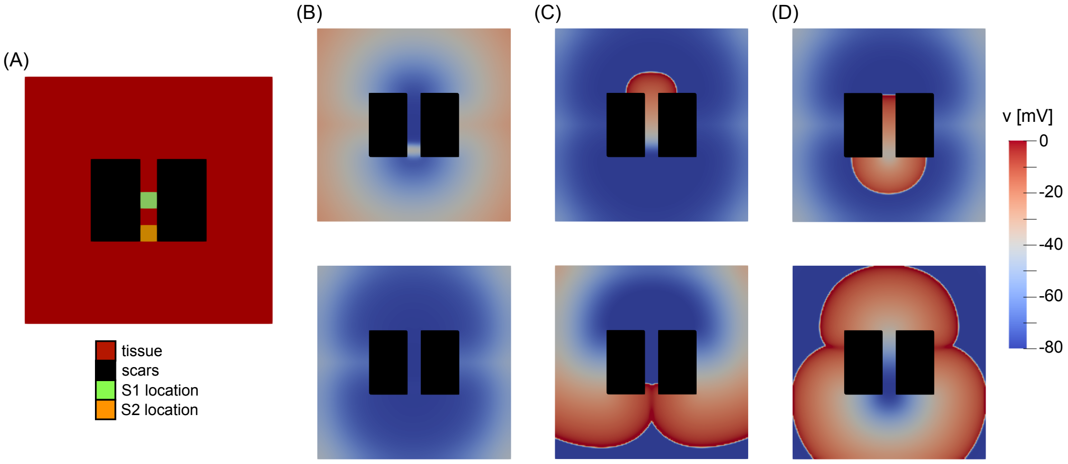

Figure 4.22. Panel A: tissue with two scars and locations of the two stimuli. Panel B: early S2 stimulation. Panel C: on time S2 stimulation. Panel D: late S2 stimulation. In panels B, C and D, two snapshots at different times show the result of the S2 stimulation.

location where the second stimulus was applied, as shown in Figure 4.22, panel C. A self-sustained re-entry is thus initiated that lasts until the end of the simulation. A similar dynamics underlies the initiation of ventricular tachycardia. If the second stimulus is applied late, the surrounding tissue is excitable and the generated action potential propagates in two directions around the scars. As a consequence, the two fronts collide, as illustrated in Figure 4.22, panel D, and the action potential vanishes. Therefore, there is a vulnerability interval during which the S2 stimulation induces a re-entry.

In this section we compare the vulnerability intervals and the dynamics of the re-entries obtained with the monodomain and the eikonal models. We consider the homogeneous non-fibrotic case and the heterogeneous case in which fibrosis is present, as well as the isotropic and the anisotropic cases. The homogeneous and the heterogeneous isotropic cases are presented in Subsections 4.6.1 and 4.6.2, respectively. The homogeneous and the heterogeneous anisotropic cases are presented in Subsections 4.6.3 and 4.6.4, respectively.

4.6.1 Homogeneous isotropic case

In this subsection we present the homogeneous isotropic case. We consider the same setting as in Subsection 4.3.4. The scars are modeled as non-conductive tissue for the monodomain simulations. In the eikonal algorithm the nodes in the scar regions are ignored.

The vulnerability interval obtained with the monodomain model is $[198, 216]$ ms, whereas the one obtained with the eikonal model is $[205, 220]$ ms. The length of the vulnerability interval is thus 18 ms in the monodomain case and 15 ms in the eikonal case. The length of the vulnerability interval roughly corresponds to the time needed by the action potential generated by the first stimulus to propagate through the region activated by the second stimulus, i.e. $\frac{1\ \text{cm}}{0.065\ \text{cm/ms}} \approx 15$ ms. Therefore, the lengths of the vulnerability intervals obtained with the two models are similar. The eikonal vulnerability interval is determined by the restitution properties of Section 4.2, which approximate the monodomain restitution curves as $h \to 0$. The monodomain restitution curves obtained with spatial resolution $h = 0.02$ cm differ from the restitution curves as $h \to 0$. This might be the reason why the eikonal vulnerability interval is shifted compared to the monodomain vulnerability interval. Moreover, in Section 4.2 the restitution properties are computed from 1D numerical experiments, but the monodomain restitution properties might be different in the 2D case in which the tissue area where the current is applied is larger.

To compare the re-entries induced with the monodomain and the eikonal models, we consider the case where the second stimulus is applied at time $t = 210$ ms. The monodomain and the eikonal simulations are shown in panels A and B of Figure 4.23, respectively. The snapshots at five times are shown. The final

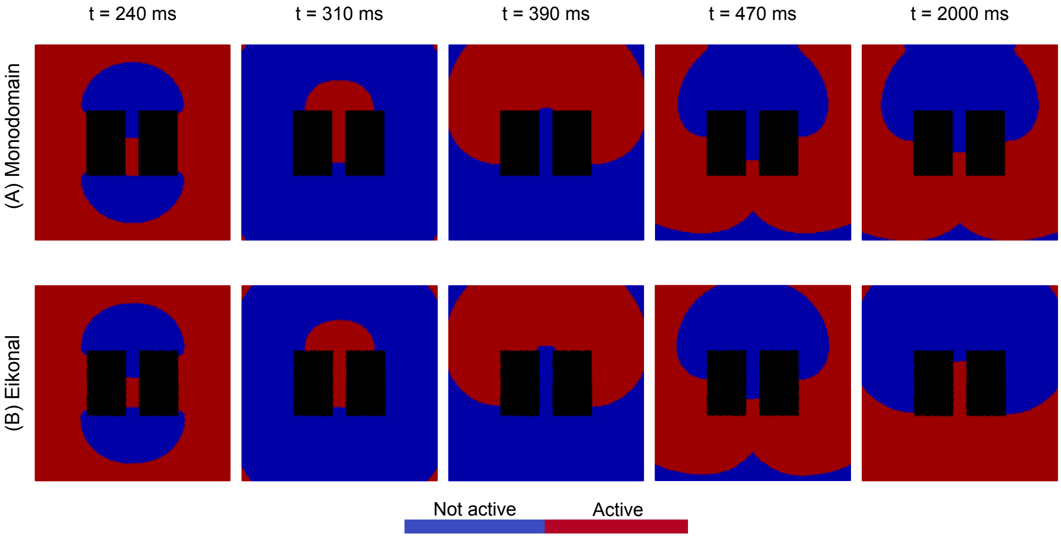

Figure 4.23. Monodomain (panel A) and eikonal (panel B) simulations of a re-entry induced by a S1-S2 stimulation in the homogeneous isotropic case.

time of the simulations is $T = 2000$ ms. After the second stimulus is delivered, the action potential propagates upwards between the two scars, as we can see in the snapshots at time $t = 240$ ms. The action potential then reaches the top of the gap between the two scars and rotates around them, as we can observe in the snapshots at the times $t = 310$ ms and $t = 390$ ms. Subsequently, the action potential reaches again the bottom of the gap between the two scars, as we can observe in the snapshots at time $t = 470$ ms. The action potential then keeps rotating around the scars until the end of the simulation. The monodomain and the eikonal simulations are very similar. The similarity is higher compared to the spiral wave simulation of Subsection 4.3.4. The reason is that here the rotation is anchored to the scars, whereas in the case of a spiral wave the tip of the rotor is free to move in the domain. The differences due to the fact that the eikonal model does not capture the wavefront curvature effects on the CV are limited. Indeed, these are visible only in the snapshots at time $t = T = 2000$ ms, which show that the eikonal propagation is slightly faster than the monodomain propagation. The computing time of the monodomain simulation is 1 h 30 min with Propag-5 on CSCS. The computing time of the eikonal simulation is much lower, indeed it is 7 min with our Python implementation of Algorithm 2 on a laptop.

4.6.2 Heterogeneous isotropic case

In this subsection we present the heterogeneous isotropic case. We consider the fibrosis pattern shown in Figure 4.13, panel A, and the setting of Subsection 4.4.2.

The vulnerability intervals obtained respectively with the monodomain and the eikonal models are $[202, 223]$ ms and $[206, 220]$ ms. Thus the lengths of the vulnerability intervals are 21 ms and 14 ms, respectively for the monodomain and the eikonal models. The monodomain vulnerability interval is shifted compared to the homogeneous case of Subsection 4.6.1 since the propagation of the action potential generated by the first stimulus is slower due to the presence of fibrosis. For the same reason, the monodomain vulnerability interval is longer than in the homogeneous case. The eikonal vulnerability interval is not affected in the same way because, in presence of fibrosis, the eikonal propagation is faster than the monodomain propagation, as observed in Subsection 4.4.1.

The monodomain and the eikonal simulations in which the second stimulus is applied at time $t = 210$ ms are shown in panels A and B of Figure 4.24, respectively. Because of the presence of fibrosis, the propagation is slower than in the homogeneous case of Subsection 4.6.1 and the waves are deformed. The

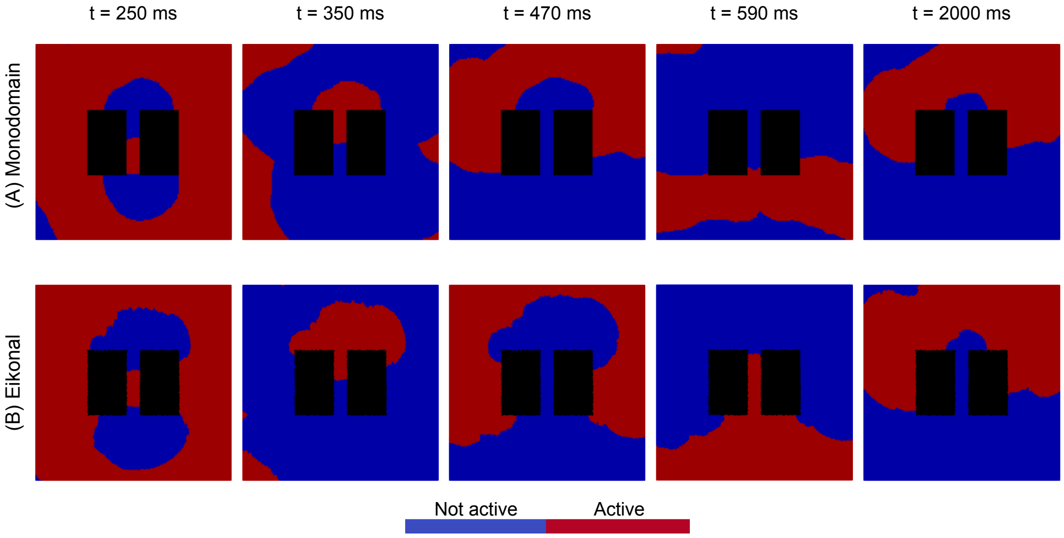

Figure 4.24. Monodomain (panel A) and eikonal (panel B) simulations of a re-entry induced by a S1-S2 stimulation in the heterogeneous isotropic case.

eikonal propagation is faster than the monodomain propagation, as we can see in the snapshots at the times $t = 250$ ms, $t = 350$ ms, $t = 470$ ms and $t = 590$ ms. This is a consequence of the tissue heterogeneity, as observed in Subsection 4.4.1.

4.6.3 Homogeneous anisotropic case

In this subsection we present the homogeneous anisotropic case. We consider the case of horizontally oriented fibers and the setting of the first numerical experiment of Subsection 4.5.1.

The vulnerability interval obtained with the monodomain model is $[211, 258]$ ms, whereas the one obtained with the eikonal model is $[239, 286]$ ms. There is agreement on the length of the vulnerability interval, which is 47 ms for both models. Since the action potential generated by the first stimulus initially propagates in the direction transversal to the fibers, its propagation is slower compared to the isotropic case of Subsection 4.6.1. As a consequence, both vulnerability intervals are shifted compared to the isotropic case. For the same reason, the vulnerability intervals are longer since, again, the length roughly corresponds to the time needed by the action potential generated by the first stimulus to propagate through the region activated by the second stimulus, i.e. $\frac{1 \text{ cm}}{0.021 \text{ cm/ms}} \approx 48$ ms. The eikonal vulnerability interval is again shifted compared to the monodomain vulnerability interval. The possible reason is again the way in which we computed the restitution properties in Section 4.2. Additionally, in the anisotropic case, the shift might be even more pronounced, as the 1D computation of the monodomain restitution properties is done only in the fibers direction.

We consider the case where the second stimulus is applied at time $t = 250$ ms to compare the re-entries induced with the monodomain and the eikonal models. The monodomain and the eikonal simulations are shown in panels A and B of Figure 4.25, respectively. The snapshots at the times $t = 350$ ms, $t = 470$ ms, $t = 690$ ms, $t = 820$ ms and $t = T = 2000$ ms are shown. Because of the anisotropy, the propagation is slower than in the isotropic case of Subsection 4.6.1. The monodomain and the eikonal simulations are very similar.

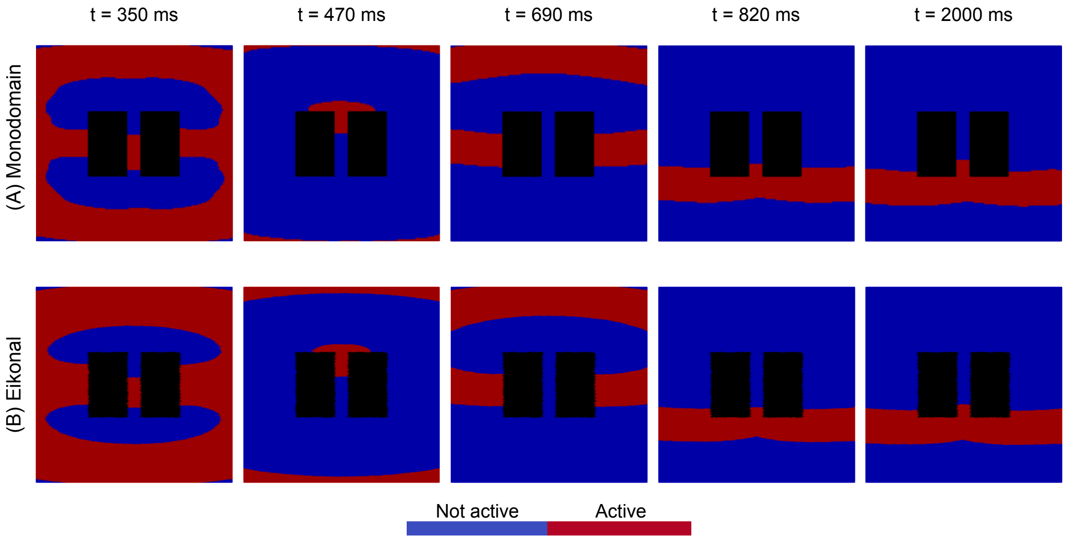

Figure 4.25. Monodomain (panel A) and eikonal (panel B) simulations of a re-entry induced by a S1-S2 stimulation in the homogeneous anisotropic case.

4.6.4 Heterogeneous anisotropic case

In this subsection we present the heterogeneous anisotropic case. We consider the fibrosis pattern shown in Figure 4.13, panel A, and the case of horizontally oriented fibers. The setting is thus the one of the first numerical experiment in Subsection 4.5.2.

The vulnerability intervals obtained respectively with the monodomain and the eikonal models are $[219, 276]$ ms and $[240, 288]$ ms. Thus the lengths of the vulnerability intervals are 57 ms and 48 ms, respectively for the monodomain and the eikonal models. Again, the propagation of the action potential generated by the first stimulus is slower than in the homogeneous case of Subsection 4.6.3 due to the presence of fibrosis. Therefore, compared to the homogeneous case, the monodomain vulnerability interval is shifted and longer. Again, the eikonal vulnerability interval is not affected in the same way because of the observations of Subsection 4.4.1.

The monodomain and the eikonal simulations in which the second stimulus is applied at time $t = 250$ ms are shown in panels A and B of Figure 4.26, respectively. Due to the presence of fibrosis, the propagation is slower than in the homogeneous case of Subsection 4.6.3 and the waves are deformed. Again, the eikonal propagation is faster than the monodomain propagation, as we can see in the snapshots at the times $t = 390$ ms, $t = 570$ ms, $t = 850$ ms and $t = 1080$ ms. This is again a consequence of the tissue heterogeneity, as observed in Subsection 4.4.1.

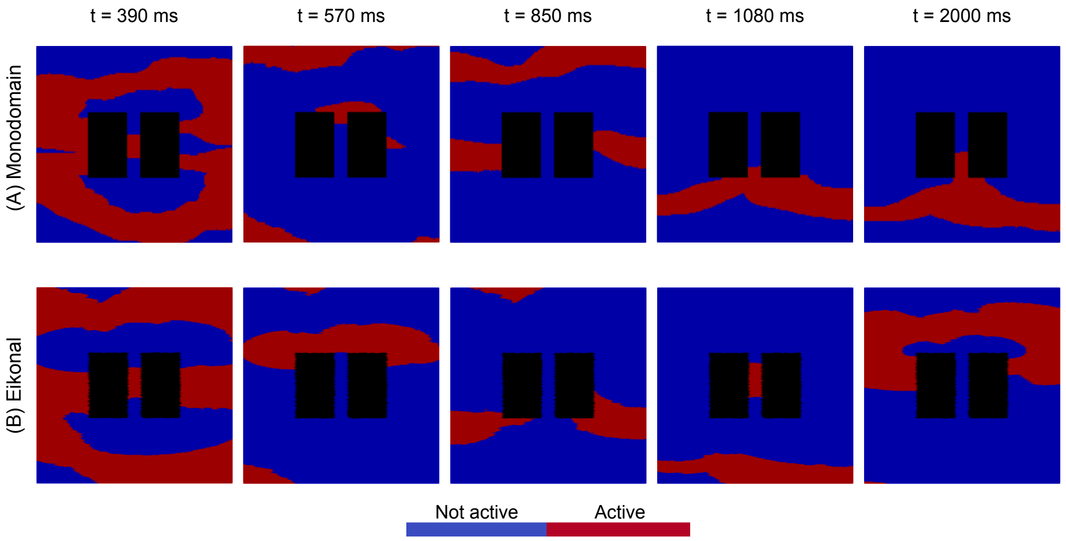

Figure 4.26. Monodomain (panel A) and eikonal (panel B) simulations of a re-entry induced by a S1-S2 stimulation in the heterogeneous anisotropic case.

4.7 3D numerical experiments

In this section we perform some 3D numerical experiments on the atrial geometry. We simulate the electrical activity during and after a pacing protocol with the eikonal model and we evaluate the AF inducibility. The pacing protocol is tailored to the restitution properties. The restitution curves included in the eikonal algorithm and computed in Section 4.2 result from two monodomain stimulations.

The restitution properties of the monodomain model depend on the number of stimulations, see Figure 4.2. Therefore, when several stimuli are applied, the restitution properties of the monodomain and the eikonal models do not match. As a consequence, the pacing protocol for the monodomain model introduced in Section 3.4 can not be applied to the eikonal model. To design an appropriate pacing protocol for the eikonal model, we resort to the idea introduced in Azzolin et al. [2021], i.e. pacing at the end of the ERP. The pacing protocol presented in Azzolin et al. [2021] consists of delivering a sequence of stimuli with interval between delivery times given by the ERP. The ERP is the sum of the APD and the $\mathrm{DI}_{\mathrm{min}}$. Here we consider $\mathrm{APD}_{\mathrm{init}} = 158$ ms, $\mathrm{DI}_{\mathrm{min}} = 32$ ms and the restitution function a of the APD given by the curve plotted in Figure 4.6, panel A. Therefore, in the pacing protocol for the eikonal model, the first interval between delivery times is $\mathrm{APD}_{\mathrm{init}} + \mathrm{DI}_{\mathrm{min}}$ and the next intervals are $a(\mathrm{DI}_{\mathrm{min}}) + \mathrm{DI}_{\mathrm{min}}$. The downside of this pacing protocol design is that there is no defined way to select the number of stimuli N_{stim}. However, we know that the number of stimuli required by the pacing protocol introduced in Azzolin et al. [2021] is lower than the number of stimuli required by a pacing protocol such as the one introduced in Section 3.4. Therefore, for the eikonal pacing protocol we decide to set $N_{\mathrm{stim}} = 14$ as in the monodomain case, since we know that additional stimuli would be superfluous. In some cases this choice of N_{stim} might lead to incorrect inducibility outcomes. Indeed, at some pacing sites AF could be induced after less than $N_{\mathrm{stim}} = 14$ stimuli and the successive stimuli could interfere with the induced AF dynamics and terminate it. However, we expect these cases to be rare. Moreover, we prefer to avoid that, at some pacing sites, AF is inducible but the inducibility outcome is incorrect because less than $N_{\mathrm{stim}} = 14$ stimuli are not sufficient, as we expect these cases to be more frequent. Note that we apply the pacing protocol of the eikonal model to pacing sites with radius 0.4 cm, as for the monodomain model. A direct comparison between the eikonal simulations and the monodomain simulations is not possible, as the pacing protocols for the two models are different. However, the similarity between the monodomain and the eikonal models in terms of AF inducibility can be investigated.

The computational domain is the epicardium of the atrial geometry described in Section 2.2 and illustrated in Figure 2.4, without coronary sinus. Notice that the domain does not include the trabecular networks, since these are attached to the endocardium. The computational mesh for the eikonal simulations consists of triangular elements and discretizes the epicardial surface. The computational mesh for the monodomain simulations discretizes the epicardium and consists of hexahedral elements. The presence of fibrosis is necessary for the pacing protocol to induce AF. Therefore, we consider the heterogeneous isotropic and anisotropic cases. The eikonal simulations are presented in Subsection 4.7.1 for the isotropic case and in Subsection 4.7.2 for the anisotropic case. The comparison between the monodomain and the eikonal models in terms of AF inducibility is shown in Subsection 4.7.3.

4.7.1 Isotropic case

In this subsection we present the eikonal simulations in the isotropic case. We set $CV_{l,\text{init}}(\mathscr{T}) = 65$ cm/s for the triangles $\mathscr{T}$ in the healthy regions and $CV_{l,\text{init}}(\mathscr{T}) = 28$ cm/s for the triangles $\mathscr{T}$ in the fibrotic regions. Moreover, the restitution function c of the CV depends on the tissue type and is given by the curves in Figure 4.6, panel B, for either the healthy or the fibrotic tissues. Since we consider the isotropic case, we set $CV_t(\mathscr{T}) = CV_l(\mathscr{T})$ for any triangle $\mathscr{T}$ of the mesh, so that the anisotropy tensor is equal to the identity, i.e. $\boldsymbol{D}(\mathscr{T}) = \boldsymbol{I}$ for any triangle $\mathscr{T}$ of the mesh. We use the Mmg software to generate an appropriate triangulated mesh for the eikonal simulations. The resolution $h = 0.1$ cm used in the 2D numerical experiments is not appropriate here because, in the resulting mesh, the percentage of triangles not satisfying the acuteness condition is 11.7%, which is significant. Therefore, we consider the finer resolution $h = 0.05$ cm, for which the percentage of triangles not satisfying the acuteness condition in the resulting mesh is 3.2%, which is acceptable. The fibrosis pattern is sampled on the hexahedral mesh with resolution $h = 0.04$ cm. We set the percentage of fibrotic tissue to 70% and we follow the sampling procedure described in Section 3.5, exploiting only the large-scale random field. The resulting fibrosis pattern is shown in Figure 4.27, panel A. The patchyness is needed for the monodomain

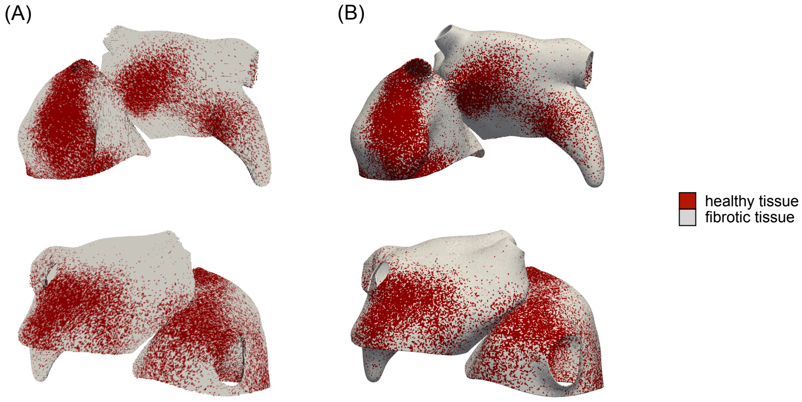

Figure 4.27. Fibrosis distribution in the hexahedral mesh with resolution 0.04 cm (panel A) and in the triangulated mesh with resolution 0.05 cm (panel B).

and the eikonal models to induce AF. To assign the tissue type to the triangles of the eikonal computational mesh, we follow a procedure that slightly differs from the one described in Subsection 4.4.2, as the hexahedral mesh is thicker than the triangulated mesh and because the resolutions are similar ($h = 0.04$ cm for the hexahedrons and $h = 0.05$ cm for the triangles). Specifically, the tissue type assigned to the triangles is the one of the closest hexahedron, where the distance is computed from the centroid of the triangle to the middle of the hexahedron. The result is shown in Figure 4.27, panel B, where the centroids of the triangles are colored based on their tissue type. The pattern is very similar to the original pattern of panel A.

We consider 20 pacing sites. The set of pacing sites is obtained by maximizing the spread on the epicardial surface. In particular, given our triangulated mesh, we first randomly choose a node and include it in the set. Then, we select the node that is further away from the starting node in terms of geodesic distance and add it to the set. We then repeat this procedure and at each iteration we add to the set the node that is further away from all the nodes already included in the set. We stop the procedure when the set reaches the requested size. For each pacing site we simulate the electrical activity during and after the pacing protocol with the eikonal model and we evaluate the AF inducibility. The computing time of one eikonal simulation with $T = 4000$ ms is 2 h 10 min with our Python implementation of Algorithm 2 on one CSCS node. The eikonal model induces AF in 9 cases, corresponding to 45% inducibility.

For one case in which the eikonal model induces AF, the simulation is shown in Figure 4.28. The figure shows ten snapshots of the electrical activity. The pacing site is indicated by the green star in the snapshot at time $t = 1890$ ms. The time $t = 1890$ ms is shortly after the delivery of the last stimulus of the pacing protocol. The front generated by the last stimulus collides with the front generated by the second-last stimulus, as we can see in the snapshot at time $t = 1910$ ms. As a consequence, the action potential generated by the last stimulus can not propagate in all directions and this dynamics creates a re-entrant path visible in the snapshots at the times $t = 1930$ ms, $t = 1980$ ms and $t = 2030$ ms. A re-entry is thus initiated, as we can observe in the snapshots at the times $t = 2070$ ms, $t = 2120$ ms, $t = 2170$ ms and $t = 2220$ ms. The self-sustained re-entrant activity then lasts until the end of the simulation at time $t = T = 4000$ ms.

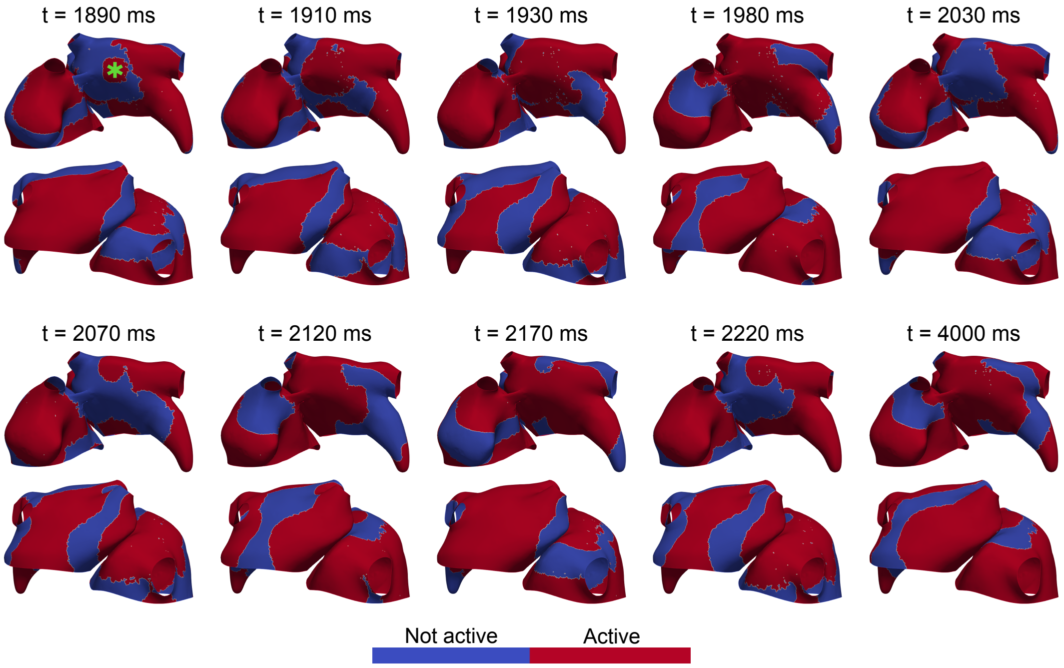

Figure 4.28. Eikonal simulation of a re-entry induced by a pacing protocol in the isotropic case. The green star illustrates the pacing site.

4.7.2 Anisotropic case

In this subsection we present the eikonal simulations in the anisotropic case. Again, we set $CV_{l,\mathrm{init}}(\mathscr{T}) = 65$ cm/s for the healthy triangles $\mathscr{T}$ and $CV_{l,\mathrm{init}}(\mathscr{T}) = 28$ cm/s for the fibrotic triangles $\mathscr{T}$. Moreover, again, the restitution function c of the CV is given by the curves in Figure 4.6, panel B, depending on the tissue type. To define the fibers orientation, we consider a simple vector field, obtained by projecting the horizontal vector on the epicardial surface. In particular, to assign the fibers orientation to the triangular elements of the mesh, we project the horizontal vector on the planes determined by the triangles. In Section 4.5 we have seen that the monodomain CV anisotropy ratio depends on both the tissue type and the fiber orientation. In this 3D setting, the fiber orientation varies over the epicardial domain. It is therefore unfeasible to set the anisotropy ratio of the eikonal model from the anisotropy ratio of the monodomain model. Thus, in order to set the anisotropy ratio for the eikonal simulations, we exploit Equation (4.14). By Equation (4.14), the anisotropy ratio CV_t/CV_l is equal to $\sqrt{\sigma_{\mathrm{m}}^t/\sigma_{\mathrm{m}}^l}$. The monodomain conductivities σ_{m}^l and σ_{m}^t are determined by the intracellular and extracellular conductivities $\sigma_{\mathrm{i,e}}^l$ and $\sigma_{\mathrm{i,e}}^t$ given in Table 3.1 and corresponding to the healthy tissue of type "atrium" and the fibrotic tissue of type "diffuse fibrosis". In both the healthy and the fibrotic tissues, we have $\sqrt{\sigma_{\mathrm{m}}^{t}/\sigma_{\mathrm{m}}^{l}} = 0.4$. Therefore, for the eikonal simulations, we set the anisotropy ratio $CV_t(\mathscr{T})/CV_l(\mathscr{T}) = 0.4$ for any triangle $\mathscr{T}$ of the mesh. Together with the fiber orientation, the anisotropy ratio determines the anisotropy tensor $\boldsymbol{D}(\mathscr{T})$ for any triangle $\mathscr{T}$ of the mesh. The computational mesh needs to be adapted to the metrics defined by the anisotropy tensors. Indeed, the mesh used in Subsection 4.7.1 is not suitable in this anisotropic case, since the percentage of triangles not satisfying the acuteness condition is 75.1%. Therefore, we use the Mmg software to adapt the mesh. In the resulting mesh, the percentage of triangles not satisfying the acuteness condition is 8.8%. The triangles that do not satisfy the acuteness condition are mostly located at the boundaries of the domain or in the regions where the surface curvature is high. Note that the adapted mesh is comparable to the mesh of Subsection 4.7.1 in terms of number of nodes. The assignment of the tissue type to the triangles of the adapted mesh with resolution $h = 0.05$ cm from the hexahedral mesh with resolution $h = 0.04$ cm follows the procedure described in Subsection 4.7.1. Again, the resulting fibrosis pattern is very similar to the original one.

We evaluate the AF inducibility of the eikonal model following the pacing protocol at the 20 pacing sites introduced in Subsection 4.7.1. The eikonal model induces AF in 14 cases, corresponding to 70% inducibility. The increased inducibility compared to the isotropic case is due to the increased variability in the action potential propagation caused by the anisotropy and is captured by the eikonal model.

The electrical activity of one case in which the eikonal model induces AF is shown in the ten snapshots of Figure 4.29. The pacing site is again indicated by the green star in the snapshot at time $t = 1890$ ms, which is shortly after the last stimulus of the pacing protocol is delivered. Again, the front generated by the last stimulus collides with the front generated by the second-last stimulus, as we can see in the snapshot at time $t = 1930$ ms. Consequently, a re-entrant path is created, as we can observe in the snapshots at the times $t = 1970$ ms, $t = 1990$ ms and $t = 2020$ ms. A re-entry is thus initiated, see the snapshots at the times $t = 2050$ ms, $t = 2110$ ms, $t = 2170$ ms and $t = 2200$ ms, and the self-sustained re-entrant activity again lasts until the end of the simulation at time $t = T = 4000$ ms.

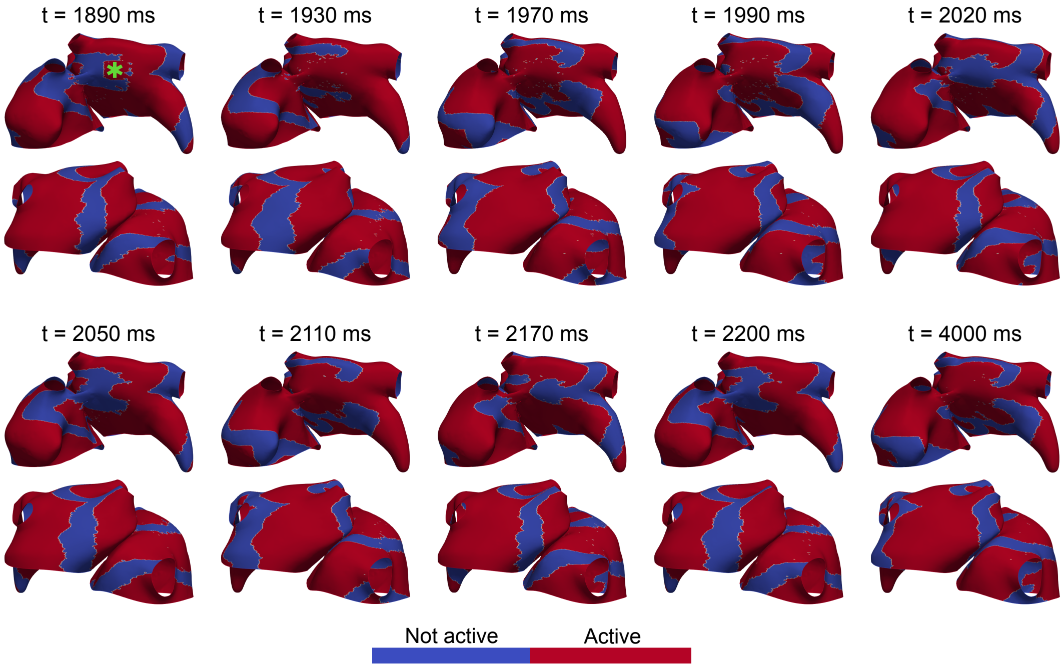

Figure 4.29. Eikonal simulation of a re-entry induced by a pacing protocol in the anisotropic case. The green star illustrates the pacing site.

4.7.3 Comparison to the high fidelity model

In this subsection we compare the eikonal model to the monodomain model in terms of AF inducibility. For the monodomain simulations, we set $\sigma^l_{\mathrm{i,e}}(\boldsymbol{x}) = 3.0\,\mathrm{mS/cm}$ for $\boldsymbol{x}$ in the healthy regions of type "atrium" and $\sigma^l_{\mathrm{i,e}}(\boldsymbol{x}) = 0.75\,\mathrm{mS/cm}$ for $\boldsymbol{x}$ in the fibrotic regions of type "diffuse fibrosis". For the sake of simplicity, we focus only on the isotropic case. Thus we consider an arbitrary fibers direction and we set $\sigma^t_{\mathrm{i,e}}(\boldsymbol{x}) = \sigma^l_{\mathrm{i,e}}(\boldsymbol{x})$ for all $\boldsymbol{x}$. The mesh used for the monodomain simulations has resolution $h = 0.02$ cm. The assignment of the tissue type to the hexahedral elements is based on the tissue type assigned to the coarse hexahedrons with resolution $h = 0.04$ cm that overlap with the fine elements. Thus the resulting fibrosis pattern corresponds to the original pattern of Figure 4.27, panel A.

We consider the 20 pacing sites introduced in Subsection 4.7.1 and shown in Figure 4.30, panel A. For each pacing site we simulate the electrical activity during and after the pacing protocol with the monodomain model and we evaluate the AF inducibility. The computing time of one monodomain simulation with $T = 4000$ ms is 1 h 40 min with Propag-5 on 8 CSCS nodes. The agree-

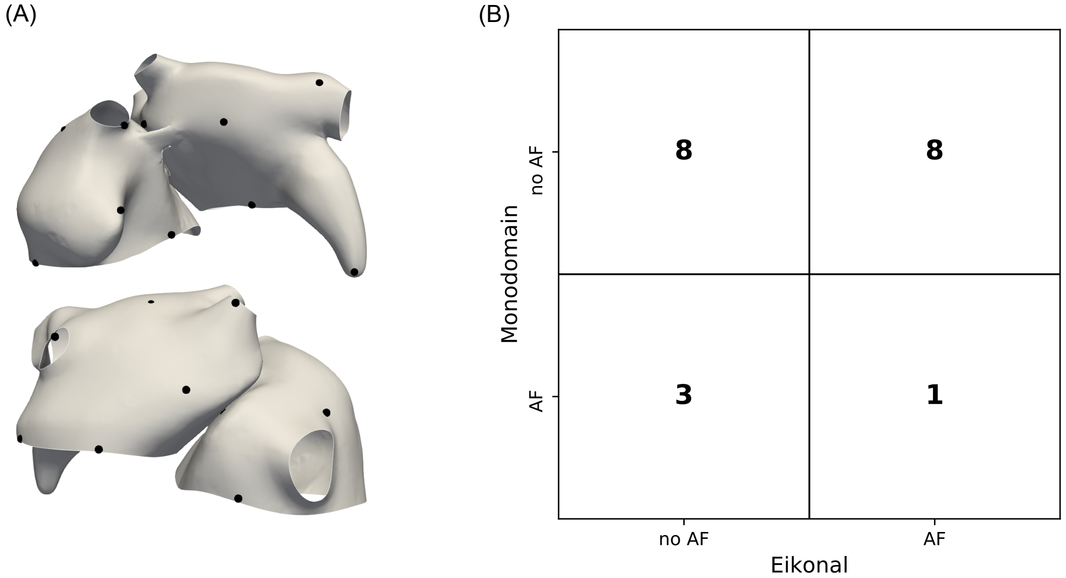

Figure 4.30. Panel A: 20 pacing sites. Panel B: agreement between the monodomain and the eikonal models.

ment between the inducibility outcomes obtained with the monodomain and the eikonal models is shown in Figure 4.30, panel B as a confusion matrix. The two models agree only in 9 cases, corresponding to 45% of the cases. In 3 cases (15%) the simulation performed with the monodomain model results in an AF event, whereas the one performed with the eikonal model does not. In 8 cases (40%) the AF induction is not successful with the monodomain model, whereas it is successful with the eikonal model. The monodomain model induces AF in 4 cases, corresponding to 20% inducibility, which is lower than the inducibility of the eikonal model.

4.8 Discussion

In this chapter we developed and tested a low fidelity model for atrial fibrillation based on the eikonal model, specifically on the FMM to solve the eikonal equation. The eikonal low fidelity model consists of an adaptation of the FMM to include the re-excitability and the restitution properties. The anisotropy is handled by adapting the computational mesh to the metric defined by the anisotropy tensor using the Mmg software. To maximize the similarity to the high fidelity model, we adjusted the CV in the eikonal model to the monodomain model. Moreover, we computed the restitution curves included in the eikonal algorithm from the monodomain model.

The eikonal model captures the main features of the re-entries. Indeed, it is very similar to the monodomain model in simulations of re-entries that rotate around scars in homogeneous tissue, see Subsections 4.6.1 and 4.6.3. However, the monodomain and the eikonal models are less similar in simulations of spiral waves whose rotor tip is free to move in the domain, see Subsections 4.3.4 and 4.5.1, or in presence of tissue heterogeneities, see Section 4.4 and Subsections 4.5.2, 4.6.2 and 4.6.4. Concerning the vulnerability interval of a S1-S2 simulation, the two models lead to different results, but similar in terms of interval length in the case of homogeneous tissue, see Section 4.6. Concerning the AF inducibility following a pacing protocol, the outputs computed with two models considerably differ, see Subsection 4.7.3. The main reasons of the discrepancies between the monodomain and the eikonal models are summarized as follows:

- Compared to the monodomain equation, the eikonal equation can be solved with a larger spatial resolution. To reduce the computational cost, the computational mesh used for the eikonal simulations is coarser than the mesh used for the monodomain simulations. As a consequence, a mapping needs to be defined to transfer the tissue information from the monodomain to the eikonal mesh. The mappings defined here allow the eikonal mesh to reproduce the large scale features of the computational domain, but some small scale features might not be captured. To avoid this issue, a finer mesh could be used, but the computational cost would increase.

- Our computation of the restitution curves does not guarantee that the restitution properties included in the eikonal algorithm always match the restitution properties of the monodomain model. Indeed, the restitution curves are computed from two monodomain stimulations. We have observed that the restitution properties computed with the monodomain model vary depending on the considered number of stimuli. Thus, the eikonal restitution properties differ from the monodomain restitution properties when several activations occur. This difference affects the evaluation of the AF inducibility, since, consequently, the pacing protocol used for the eikonal simulations is different from the pacing protocol used for the monodomain simulations. Moreover, the restitution curves are computed in 1D numerical experiments. Hence, the eikonal restitution properties might differ from the monodomain restitution properties in 2D and 3D numerical experiments. This difference might be amplified in presence of anisotropy. Furthermore, the monodomain restitution curves depend on the spatial resolution and on the tissue type. Instead, in the eikonal algorithm we use the restitution curves at convergence, which do not depend on the tissue type (up to a scaling factor for the CV). To overcome some of these issues, the ionic model could be modified to guarantee the charge conservation (Colli Franzone et al. [2014]). Alternatively, the restitution curves could be computed on a ring with periodic boundary conditions (Courtemanche et al. [1996]).

- The eikonal model does not capture the effect of the wavefront curvature on the CV. Indeed, the eikonal model does not account for the diffusion currents due to the heterogeneity in the propagation. In order to take into account the effect of the curvature, the eikonal-diffusion model could be employed. However, the inclusion of the re-excitability would require the solution of the eikonal-diffusion equation to be computed on the whole computational domain at each time step. Moreover, the eikonal-diffusion model would require a finer spatial resolution than the pure eikonal model. These constraints would lead to a higher computational cost. Alternatively, the eikonal-curvature model could be employed. The solution of the eikonal-curvature equation could, in principle, be obtained by iterative algorithms suitable to include the re-excitability, but an estimate of the curvature would be needed at each iteration.

- On the one hand, the eikonal model does not capture the jumps in the activation times and the delays introduced by the discontinuities in the conductivity in heterogeneous tissue. On the other hand, the discretized monodomain model amplifies these jumps and delays. This discrepancy is a novel aspect that had not been analyzed before, to the best of our knowledge. The eikonal model does not account for the diffusion currents due to the heterogeneity in the conductivity. In order to take into account the effect of discontinuous conductivity on the front propagation, the eikonal-diffusion model could be employed. However, again, the computational cost would increase.

- The convergence of the FMM to the viscosity solution of the eikonal equation is guaranteed if all the triangles of the computational mesh satisfy an acuteness condition. The mesh adaptation with the Mmg software minimizes the number of triangles that do not satisfy the acuteness condition. However, in some cases, the percentage of triangles not satisfying the acuteness condition might remain relevant. To overcome this issue, one could employ the FIM, which converges on all triangulations. However, since the FIM is not single-pass and does not guarantee the monotone acceptance of the activation times, the inclusion of the re-excitability would lead to a higher computational cost. Indeed, the eikonal equation would need to be solved on the whole computational domain at each time step.

The current speed-up obtained by the eikonal model varies depending on the application. The eikonal model is approximately 23 times faster than the monodomain model in the 2D simulations of spiral waves, approximately 13 times faster in the 2D simulations of re-entries in presence of scars and approximately 6 times faster in the 3D simulation of re-entries. The differences are due to the number of node updates performed in the various applications of the eikonal algorithm and to the resolution of the eikonal computational mesh, which is finer in the 3D numerical experiments than in the 2D cases. The current implementation of Algorithm 2 is not optimized. Indeed, the implementation of the FMM is done in Python instead of C++ and the adaptations to include the re-excitability are naively implemented. The eikonal algorithm can potentially be real-time, see Subsection 4.4.1.

The evaluation of the AF inducibility following a pacing protocol is challenging and the eikonal model is inaccurate. However, the eikonal model is qualitatively accurate in simpler numerical experiments and is potentially very cheap. Since the eikonal model is robust to changes in the ionic model (up to a re-computation of the restitution curves) and in the underlying anatomy (up to a re-computation of the CVs if the conductivity properties change), it might be employed in computational studies of cardiac arrhythmias characterized by simpler dynamics that occur on simpler anatomies, such as ventricular tachycardia. The qualitative accuracy and the low cost make the eikonal model attractive for *in-silico* studies where the computational resources are limited and motivate our interest in improving the model. Our future work includes the optimization of the algorithm implementation and the reduction of the discrepancies compared to the monodomain model. The improvement of the accuracy is likely to lead to an increase in the computational cost. Our future goal is to quantify this accuracy/cost trade-off.

Chapter 5

Low fidelity model based on coarser discretization parameters

In this chapter we introduce the low fidelity model based on coarser discretization parameters in the numerical solution of the monodomain system (3.15). Our goal is to maximize the correlation between the low fidelity model and the high fidelity model of Chapter 3 when evaluating the AF inducibility. To achieve this correlation, some aspects need to be taken into account. In particular, the coarsening of the atrial model of Section 3.5 needs to be appropriately defined. Moreover, since the coarsening of the computational mesh affects the CV (Pezzuto et al. [2016]), the CV of the low fidelity model has to be adapted to match the one of the high fidelity model.

The coarsening of the atrial model is presented in Section 5.1 and the adjustment of the CV is presented in Section 5.2. The numerical experiments in which we compare the high and the low fidelity models in terms of AF inducibility after the stimulation by a pacing protocol are shown in Section 5.3 and a discussion follows in Section 5.4.

5.1 Coarsening of the atrial model

In this section we present the coarsening of the atrial model. The atrial model of Section 3.5 consists of a mesh of hexahedral elements, in which each element is assigned to a tissue type and to three angles that determine the local fiber orientation. To use a larger spatial resolution h in the numerical solution of the monodomain equation, we need a coarser computational mesh. In the coarse mesh, the tissue type and the angles assigned to each hexahedral element depend on the assignments to the hexahedral elements of the fine mesh that overlap with the coarse element. In this work we double the spatial resolution and set it to $h = 0.04$ cm, so that eight fine elements overlap with each coarse element.

The assignment of the tissue type is based on a majority rule, i.e. the tissue type that is mostly represented among the fine elements is assigned to the coarse element. To illustrate the coarsening of the computational mesh, we consider an atrial model with 70% fibrotic tissue. The atrial anatomy is the one described in Section 3.5, which reproduces the features described in Section 2.2 and illustrated in Figure 2.4. Figure 5.1 shows the fine and the coarse meshes, where the hexahedral elements are colored according to their tissue type. The large-scale features of the fine mesh are present in the coarse mesh, whereas some small-scale features might not be captured by the coarse spatial discretization.

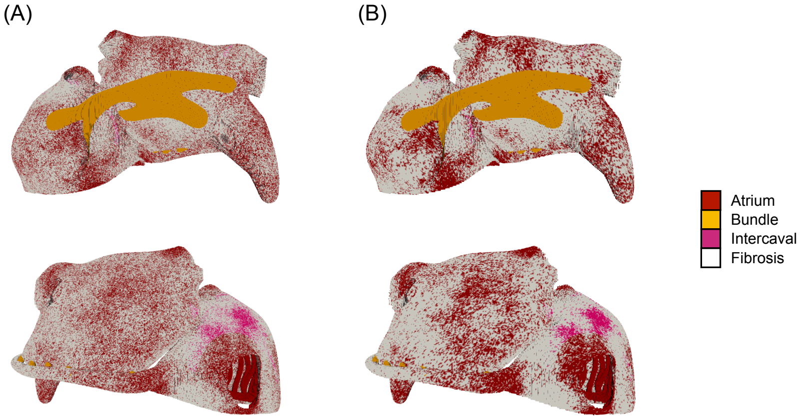

Figure 5.1. Tissue types assigned to the fine (panel A) and the coarse (panel B) computational meshes.

The majority rule is appropriate for variables taking few discrete values, as in the case of the tissue type. The angles take continuous values in $[-\pi, \pi]$. Therefore, the assignment of the angles is not based on a majority rule. Instead, the procedure sequentially inspects the fine elements and the first value in $[-\pi, \pi]$ found is assigned to the coarse element.

In this low fidelity model we also double the time step and set it to $\Delta t = 0.02$ ms. As a consequence, the low fidelity model is approximately 16 times faster than the high fidelity model.

5.2 Conduction velocity adjustment

In this section we present the CV adjustment. With finite differences, the coarser spatial discretization has the effect of reducing the CV (Pezzuto et al. [2016]).

We verify this fact with a 1D numerical experiment on a line of tissue of type "atrium" of length 5 cm. We deliver a stimulus of size 0.2 cm from one end that generates an action potential that propagates towards the other end. We perform one simulation with the fine discretization where the spatial resolution is $h = 0.02$ cm and the time step is $\Delta t = 0.01$ ms and another simulation with the coarse discretization where the spatial resolution is $h = 0.04$ cm and the time step is $\Delta t = 0.02$ ms. Figure 5.2 shows the activation times along the line obtained with the two simulations. The points on the line activate later with the coarse level simulation than with the fine level simulation, as a consequence of the reduced CV due to the coarser spatial discretization. Indeed, the CV measured on the fine level is 65 cm/s, whereas the one measured on the coarse level is 57 cm/s. This 12% difference is substantial and can not be neglected.

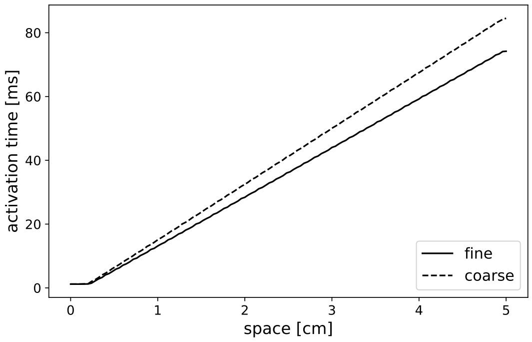

Figure 5.2. Activation times obtained with the fine (black line) and the coarse (dashed line) spatial discretizations.

The reduction in the CV affects the low fidelity model. In order to increase the inter-model correlation, the low fidelity model has to be adjusted to match the CV of the high fidelity model. To this end, we exploit Equation (4.14). Specifically, since the CV is inversely proportional to $\sqrt{\beta}$, we reduce the surface-to-volume ratio in the low fidelity model. By repeating the 1D simulation on the coarse level with $\beta = 450$ cm^{-1}, we obtain the same activation times and the same CV as obtained on the fine level with $\beta = 800$ cm^{-1}. Therefore, in the low fidelity model we set $\beta = 450$ cm^{-1}.

5.3 Numerical experiments

In this section we compare the high and the low fidelity models in terms of AF inducibility after the stimulation by a pacing protocol. Since the pacing protocol is tailored to the ionic model, the evaluation of the AF inducibility with the low fidelity model also exploits the pacing protocol of Section 3.4. In these numerical experiments we consider the atrial model of Figure 5.1 with endomysial fibrosis and the 20 pacing sites of Gharaviri et al. [2021a] and Gharaviri et al. [2021b], shown in Figure 5.3, panel A. For each pacing site we simulate the electrical activity during and after the pacing protocol and we evaluate the AF inducibility outcome with both the high and the low fidelity models. The agreement between the two models is shown in Figure 5.3, panel B as a confusion matrix. The two models agree on the inducibility result in 18 cases, corresponding to 90% of the cases. In one case (5%) the simulation performed with the high fidelity model results in an AF event, whereas the one performed with the low fidelity model does not. In one other case (5%) the AF induction is not successful with the high fidelity model, whereas it is successful with the low fidelity model. This determines the correlation between the AF inducibility results, which is 0.79. For both models, AF is induced in 12 cases, corresponding to 60% inducibility.

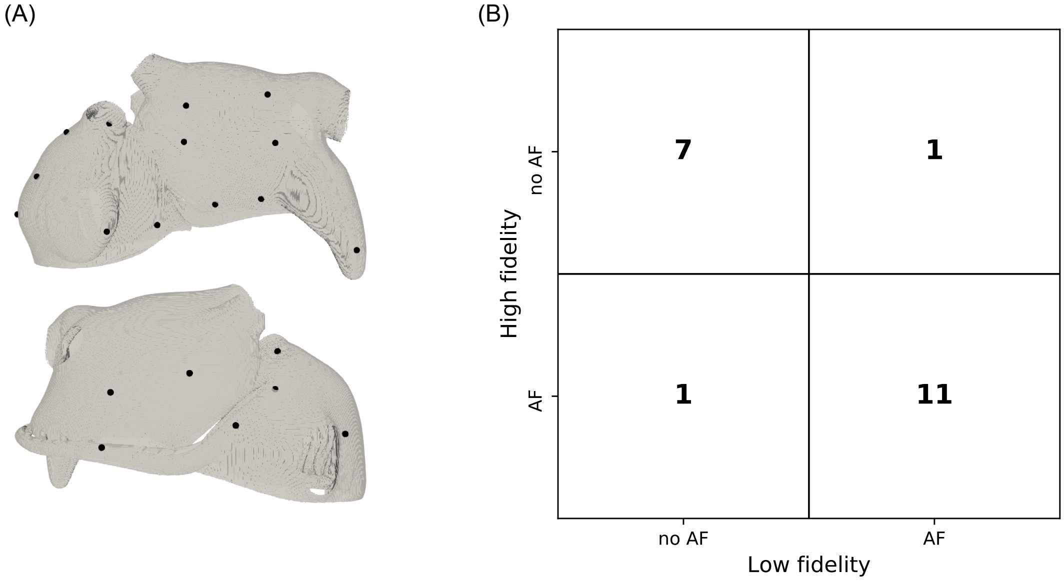

Figure 5.3. Panel A: 20 pacing sites. Panel B: agreement between the high and the low fidelity models.

For one case in which the two models agree on the AF induction, the electrical activity simulated with the high and the low fidelity models is shown in Figure 5.4. The figure shows five snapshots of the transmembrane potential v computed with the high (panel A) and the low (panel B) fidelity models. The final time of the simulations is $T = 4000$ ms. Both simulations are visualized on the coarse atrial mesh used for the low fidelity simulation. However, note that the high fidelity simulation is performed on the fine atrial mesh. The green star on the snapshots at time $t = 1965$ ms indicates the pacing site. The time

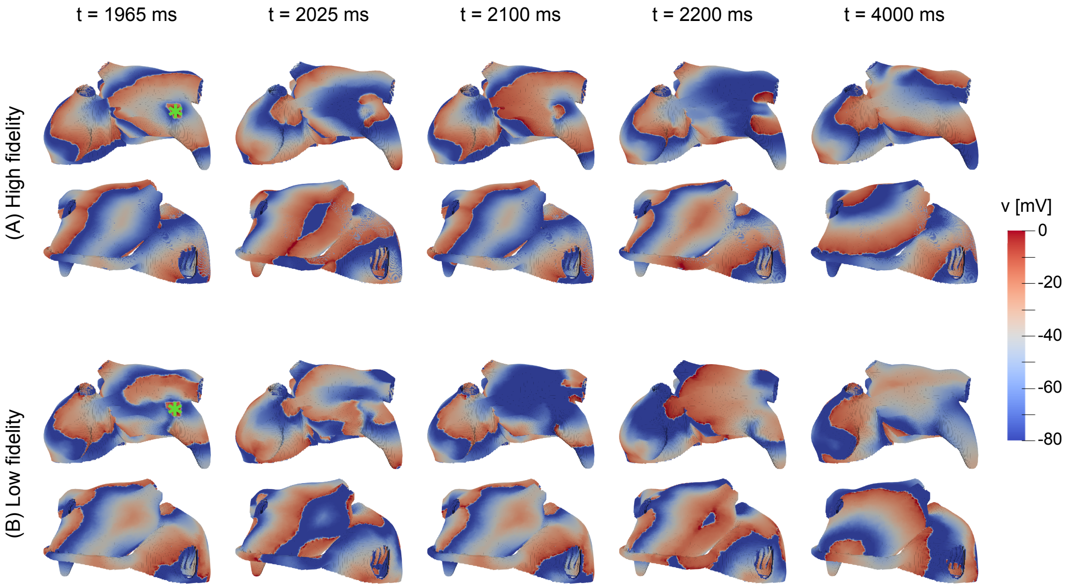

Figure 5.4. High (panel A) and low (panel B) fidelity simulations inducing AF. The green star illustrates the pacing site.

$t = 1965$ ms is shortly after the delivery of the last stimulus of the pacing protocol. In the high fidelity simulation, the fronts generated by the second-last and the last stimuli collide in a portion of the domain. As a consequence, the action potential generated by the last stimulus can not propagate in all directions, as we can see in the snapshot of panel A at time $t = 2025$ ms. This dynamics generates a re-entry, see the snapshot of panel A at time $t = 2100$ ms, that initiates a self-sustained activity, see the snapshot of panel A at time $t = 2200$ ms. The self-sustained activity lasts until the end of the simulation, as we can onserve in the snapshot of panel A at time $t = T = 4000$ ms. In the low fidelity simulation the dynamics is different, but again the collision between the fronts generated by the second-last and the last stimuli of the pacing protocol induces a re-entry that initiates a self-sustained activity, see the snapshots of panel B at the times $t = 2025$ ms, $t = 2100$ ms and $t = 2200$ ms. Again, the self-sustained activity lasts until the end of the simulation, as we can observe in the snapshot of panel B at time $t = T = 4000$ ms. Therefore, even though the dynamics of the low fidelity simulation is different from the dynamics of the high fidelity simulation, the low fidelity model agrees with the high fidelity model on the induction of

AF. Moreover, performing the low fidelity simulation requires less computational effort than performing the high fidelity simulation. Indeed, with Propag-5 on CSCS, the computing time of the high fidelity model is 1 h 40 min with 8 nodes, whereas the computing time of the low fidelity model is 14 min with 4 nodes.

5.4 Discussion

In this chapter we developed a low fidelity model based on a coarser discretization of the computational domain in the numerical solution of the monodomain equations. We tested the agreement to the high fidelity model in terms of AF inducibility in a numerical experiment. To maximize this agreement, we defined the coarsening of the underlying atrial model and we adjusted the CV in the low fidelity model to match the one of the high fidelity model.

The preliminary numerical results show that there is high correlation between the high and the low fidelity models when used to evaluate the AF inducibility. This high correlation is achieved since the low fidelity model inherits many properties from the high fidelity model and because the coarsening of the atrial model and the adjustment of the CV are properly addressed. Since there is high correlation, in the next chapter we combine the two models in a multi-fidelity framework to estimate atrial inducibility maps.

The high and the low fidelity models can be applied to any underlying atrial model. Indeed, the atrial geometry and the local conduction properties can be encoded in the fine mesh and can be transferred to the coarse mesh using the coarsening procedure. In this chapter we presented the numerical results on a single atrial model. In the next chapter we consider various atrial models obtained by combining different fibrosis patterns and different ablation strategies.

Chapter 6

Multi-fidelity characterization of the atrial inducible regions

In this chapter we combine the high fidelity model of Chapter 3 and the low fidelity model based on coarser discretization parameters of Chapter 5 in a multi-fidelity framework to characterize the atrial inducible regions. We characterize the atrial regions depending on the outcome of a stimulation by the pacing protocol of Section 3.4. The two possible outcomes, i.e. whether AF is induced or not, lead to our definition of the inducibility map of the atria. The problem of obtaining the inducibility map is thus a binary classification problem. To solve the problem, we design a multi-fidelity GP classifier that operates on the atrial surface. We perform some numerical experiments to evaluate the performance of the multi-fidelity GP classifier by comparing it to two single-fidelity classifiers, namely the nearest neighbor and the standard GP classifiers. The work presented in this chapter has been done in collaboration with Prof. Simone Pezzuto, Dr. Ali Gharaviri, Prof. Rolf Krause, Prof. Paris Perdikaris and Prof. Francisco Sahli Costabal (Gander et al. [2022]). Some sentences are taken from Gander et al. [2022] *ad verbatim*.

The classification problem is defined in Section 6.1. The two single-fidelity and the one multi-fidelity classifiers are presented in Section 6.2. The numerical experiments are shown in Section 6.3 and a discussion follows in Section 6.4.

6.1 The classification problem

In this section we define the classification problem, i.e. obtaining an AF inducibility map of the atria. We evaluate the AF inducibility from the outcome of the tissue stimulation with the pacing protocol of Section 3.4. The success of the pacing protocol in inducing AF depends on the pacing site. Indeed, the stimulation from some pacing sites leads to an AF event, whereas stimulating from some other pacing sites does not lead to an AF event. These two different outcomes allow us to define the inducibility map of the atria as a map showing the pacing sites that induce AF if stimulated by the pacing protocol.

In this work, we simulate the electrical activity on the atrial geometry Ω during and after pacing from various locations $\boldsymbol{x}_{\text{stim}}$. Since the atrial wall is thin, we restrict the pacing sites to be on a mid-wall atrial surface $\mathscr{S} \subset \Omega$. Our goal is to determine the set $\mathscr{A} \subset \mathscr{S}$ such that, if $\boldsymbol{x}_{\text{stim}} \in \mathscr{A}$, then AF is induced. In other words, we want to learn the indicator function $F \colon \mathscr{S} \to \{0, 1\}$ such that $F^{-1}(1) = \mathscr{A}$. The function F is thus a map showing the regions where AF is inducible. The problem of determining the inducibility map F is a binary classification problem.

In some cases, it might be useful to reduce the inducibility map into one metric. This metric is the inducibility $I \in \mathbb{R}$, which is defined as the fraction of tissue where AF is inducible, i.e.

$$I = \frac{|\mathscr{A}|}{|\mathscr{S}|} = \frac{1}{|\mathscr{S}|} \int_{\mathscr{S}} F(\boldsymbol{x}) \mathrm{d}\boldsymbol{x}.$$

6.2 Classifiers

In this section we present three classification methods to learn the inducibility map from the limited data contained in a training set. The first two methods, namely the nearest neighbor and the standard GP classifiers, are single-fidelity. The third method is the multi-fidelity GP classifier. In the single-fidelity classification, the training set contains only high fidelity data. In the multi-fidelity classification, the low fidelity data are also included. The classifiers allow to predict the class output at locations that are not contained in the training set. Therefore, by making predictions at a sufficiently large set of unknown inputs, it is possible to obtain an inducibility map of the atria.

The nearest neighbor classifier is introduced in Subsection 6.2.1. The single-fidelity and the multi-fidelity GP classifiers are presented in Subsections 6.2.2 and 6.2.3, respectively.

6.2.1 Nearest neighbor classifier

In this subsection we present the nearest neighbor classifier. The single-fidelity training set consists of N high fidelity input/output pairs of the form $\{(\boldsymbol{x}_i, y_i)\}_{i=1}^{N}$, where, for $i = 1, ..., N$, the inputs $\boldsymbol{x}_i \in \mathscr{S}$ are some pacing sites and the outputs $y_i \in \{0, 1\}$ are the corresponding class labels. The nearest neighbor classifier simply assigns to unknown locations the class label of the point in the training set that is closer in terms of geodesic distance. In other words, the prediction y^* for a new pacing site $\boldsymbol{x}^* \in \mathscr{S}$ is given by

$$y^* = y_{i^*} \quad \text{with} \quad i^* = \underset{i=1,\dots,N}{\arg\min}\, d_{\mathscr{S}}(\boldsymbol{x}^*, \boldsymbol{x}_i),$$

where $d_{\mathscr{S}}$ denotes the geodesic distance on the atrial surface $\mathscr{S}$. To compute the geodesic distance on the manifold, we use the heat method (Crane et al. [2013]).

6.2.2 Single-fidelity Gaussian process classifier

In this subsection we present the single-fideity GP classifier. GP classification is based on an intermediate variable $f(\boldsymbol{x})$ at the location $\boldsymbol{x} \in \mathscr{S}$ computed from a latent function f with GP prior (Rasmussen and Williams [2006]). By assuming, without loss of generality, standardized datasets, the zero-mean GP prior assigned to the latent function takes the form

$$f \sim \mathscr{GP}\big(0, k(\boldsymbol{x}, \boldsymbol{x}'; \boldsymbol{\theta})\big),$$

where k is the covariance kernel, which depends on the parameters $\boldsymbol{\theta}$. In contrast to nearest neighbor classification, GP classification includes the uncertainty in the prediction, by returning a class probability. This is done by passing the latent function f through a non-linear warping function $\sigma : \mathbb{R} \to [0, 1]$. The output is then the class probability $\mathbb{P}[y = 1] = \sigma\big(f(\boldsymbol{x})\big)$ (Nickisch and Rasmussen [2008]). Our choice for the warping function σ is the commonly used logistic sigmoid function $\sigma(f) = \big(1 + \exp(-f)\big)^{-1}$.

The covariance kernel is a fundamental element in GP classification. A common choice is the Matérn kernel, which, in a Euclidean space, reads

$$k(\boldsymbol{x}, \boldsymbol{x}'; \boldsymbol{\theta}) = \eta^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu} \frac{\|\boldsymbol{x} - \boldsymbol{x}'\|}{\ell} \right)^{\nu} K_\nu \left(\sqrt{2\nu} \frac{\|\boldsymbol{x} - \boldsymbol{x}'\|}{\ell} \right),$$

where Γ is the gamma function and K_ν is the modified Bessel function of the second kind. The parameters are $\boldsymbol{\theta} = (\eta, \ell)$, which respectively control the variance of the GP and the spatial correlation length-scale. The value ν controls the regularity of the latent function f, so this kernel allows to include smoothness assumptions (Rasmussen and Williams [2006]). However, the kernel can not be straightforwardly applied to our case, as we are working on the atrial surface $\mathscr{S}$. The definition of the kernel in the Euclidean space can not be extended to manifolds simply by replacing the Euclidean distance with the geodesic distance. Indeed, this approach can not guarantee the positive semi-definiteness of the covariance matrix resulting from the evaluation of k at input points, which is a necessary property for a kernel (Pezzuto et al. [2019], Borovitskiy et al. [2020]). So we need to follow an alternative approach based on the solution of a stochastic partial differential equation (Lindgren et al. [2011]). This approach guarantees the positive semi-definiteness of the covariance matrix when working on manifolds (Borovitskiy et al. [2020]). Moreover, the kernel can be expressed explicitly. The formulation of the kernel is based on the eigenpairs $\{(\lambda_i, \psi_i)\}_{i=1}^{\infty}$ of the Laplace-Beltrami operator $-\Delta$ with zero Neumann boundary condition, satisfying

$$\begin{cases} -\Delta\psi_i(\boldsymbol{x}) = \lambda_i\psi_i(\boldsymbol{x}), & \boldsymbol{x} \in \mathscr{S}, \\ -\boldsymbol{n}(\boldsymbol{x}) \cdot \nabla\psi_i(\boldsymbol{x}) = 0, & \boldsymbol{x} \in \partial\mathscr{S}, \end{cases}$$

for $i \in \mathbb{N}$. The Matérn kernel on the manifold is then given by (Borovitskiy et al. [2020])

$$k(\boldsymbol{x}, \boldsymbol{x}'; \boldsymbol{\theta}) = \frac{\eta^2}{C} \sum_{i=1}^{\infty} \left(\frac{1}{\ell^2} + \lambda_i \right)^{-\nu - \frac{d}{2}} \psi_i(\boldsymbol{x})\psi_i(\boldsymbol{x}'),$$

where C is a normalizing constant, d is the dimension of the manifold, ν is the smoothness value and $\boldsymbol{\theta} = (\eta, \ell)$ are the parameters. For the practical computation of the kernel, we use a triangulated mesh to discretize the atrial surface and we use finite elements shape functions to solve the eigenvalue problem. Moreover, we truncate the eigendecomposition to a finite number of eigenpairs, starting from the pair with the smallest eigenvalue. Notice that the eigendecomposition needs to be computed only once.

Now we can describe the single-fidelity GP classification in more detail. The training set of N high fidelity input/output pairs $\{(\boldsymbol{x}_i, y_i)\}_{i=1}^{N} = \{\boldsymbol{X}, \boldsymbol{y}\}$ is defined in Subsection 6.2.1. In the Bayesian approach, prior distributions are assigned to the kernel parameters (Neal [1998]). Our assumption for the priors over the parameters $\boldsymbol{\theta} = (\eta, \ell)$ is

$$\begin{aligned} \eta &\sim \text{HalfNormal}(\sigma = 1000), \\ \ell &\sim \text{Gamma}(\alpha = 1, \beta = 1). \end{aligned}$$

The posterior distributions over $\boldsymbol{\theta}$ can be estimated based on the training set. For this purpose, we use the No-U-Turn sampler (Hoffman and Gelman [2014]) to perform the Bayesian inference. The posterior distributions over $\boldsymbol{\theta}$ determine the posterior distribution of f, which allows to predict the classification output at a new pacing site $\boldsymbol{x}^* \in \mathscr{S}$ as follows. The posterior distribution of the predictive latent function f^* is obtained by conditioning on the training data (Rasmussen and Williams [2006]). Thus the predictive latent function follows a normal distribution $f^*(\boldsymbol{x}^*) \sim \mathcal{N}\big(\mu(\boldsymbol{x}^*), \Sigma(\boldsymbol{x}^*)\big)$ with mean and variance given by

$$\begin{aligned}
&\mu(\boldsymbol{x}^*) = k(\boldsymbol{x}^*, \boldsymbol{X}; \boldsymbol{\theta})k(\boldsymbol{X}, \boldsymbol{X}; \boldsymbol{\theta})^{-1} f(\boldsymbol{X}),\\
&\Sigma(\boldsymbol{x}^*) = k(\boldsymbol{x}^*, \boldsymbol{x}^*; \boldsymbol{\theta}) - k(\boldsymbol{x}^*, \boldsymbol{X}; \boldsymbol{\theta})k(\boldsymbol{X}, \boldsymbol{X}; \boldsymbol{\theta})^{-1} k(\boldsymbol{X}, \boldsymbol{x}^*; \boldsymbol{\theta}).
\end{aligned}$$

The mean and the variance can be sampled using $\boldsymbol{\theta}$ and f drawn from the posterior distributions. The resulting N_{s} samples μ_i and Σ_i, for $i = 1, ..., N_{\mathrm{s}}$, can be averaged to obtain

$$\hat{\mu} = \frac{1}{N_{\mathrm{s}}} \sum_{i=1}^{N_{\mathrm{s}}} \mu_i \quad \text{and} \quad \hat{\Sigma} = \frac{1}{N_{\mathrm{s}}} \sum_{i=1}^{N_{\mathrm{s}}} \Sigma_i.$$

Finally, the predictive latent function is drawn from $\hat{f}^* \sim \mathcal{N}(\hat{\mu}, \hat{\Sigma})$ and the prediction at the new location is given by $y^* = \sigma(\hat{f}^*)$.

6.2.3 Multi-fidelity Gaussian process classifier

In this subsection we present the multi-fidelity GP classifier. Multi-fidelity classification combines models with multiple levels of fidelity by handling a training set that contains data generated at the various levels. In this work, we focus on the two-level case, where the high fidelity model is combined with the low fidelity model of Chapter 5. The idea is to offset most of the computational burden to the fast low fidelity model, in order to gain in efficiency. To maintain the desirable accuracy, the multi-fidelity approach exploits the correlation between the high and the low fidelity model.

Considering two sources of information, the training set is composed of N_{L} low fidelity and N_{H} high fidelity input/output pairs $\{(\boldsymbol{x}_{\mathrm{L},i}, y_{\mathrm{L},i})_{i=1}^{N_{\mathrm{L}}}, (\boldsymbol{x}_{\mathrm{H},i}, y_{\mathrm{H},i})_{i=1}^{N_{\mathrm{H}}}\} = \{(\boldsymbol{X}_{\mathrm{L}}, \boldsymbol{y}_{\mathrm{L}}), (\boldsymbol{X}_{\mathrm{H}}, \boldsymbol{y}_{\mathrm{H}})\}$, where the subscripts L and H refer to the low and the high fidelity data, respectively. Multi-fidelity GP classification is a generalization of the single-fidelity framework of Subsection 6.2.2. In the multi-fidelity framework, we introduce the low and high fidelity latent functions f_{L} and f_{H}. The two latent functions are assumed to be related through the auto-regressive prior (Kennedy and O'Hagan [2000])

$$f_{\mathrm{H}}(\boldsymbol{x}) = \rho f_{\mathrm{L}}(\boldsymbol{x}) + \delta(\boldsymbol{x}),$$

where ρ is a scaling parameter and δ is another latent function, independent of f_{L}, that describes the difference between the two levels of fidelity. The functions

f_L and δ have zero-mean GP priors of the form (Kennedy and O'Hagan [2000])

$$\begin{aligned} f_\mathrm{L} &\sim \mathscr{GP}\big(0, k_\mathrm{L}(\boldsymbol{x},\boldsymbol{x}';\boldsymbol{\theta}_\mathrm{L})\big),\\ \delta &\sim \mathscr{GP}\big(0, k_\mathrm{H}(\boldsymbol{x},\boldsymbol{x}';\boldsymbol{\theta}_\mathrm{H})\big),\end{aligned}$$

where the covariance kernels k_L and k_H, respectively depending on the parameters $\boldsymbol{\theta}_\mathrm{L}$ and $\boldsymbol{\theta}_\mathrm{H}$, are associated to the low and the high fidelity levels, respectively. As a consequence, the joint prior distribution over the latent functions $f_\mathrm{L}(\boldsymbol{X}_\mathrm{L})$ and $f_\mathrm{H}(\boldsymbol{X}_\mathrm{H})$ takes the form

$$\begin{bmatrix} f_\mathrm{L}(\boldsymbol{X}_\mathrm{L})\\ f_\mathrm{H}(\boldsymbol{X}_\mathrm{H})\end{bmatrix} \sim \mathscr{N}\left(\begin{bmatrix}\mathbf{0}\\ \mathbf{0}\end{bmatrix}, \begin{bmatrix} k_\mathrm{LL}(\boldsymbol{X}_\mathrm{L},\boldsymbol{X}_\mathrm{L};\boldsymbol{\theta}_\mathrm{L}) & k_\mathrm{LH}(\boldsymbol{X}_\mathrm{L},\boldsymbol{X}_\mathrm{H};\boldsymbol{\theta}_\mathrm{L},\rho)\\ k_\mathrm{LH}(\boldsymbol{X}_\mathrm{H},\boldsymbol{X}_\mathrm{L};\boldsymbol{\theta}_\mathrm{L},\rho) & k_\mathrm{HH}(\boldsymbol{X}_\mathrm{H},\boldsymbol{X}_\mathrm{H};\boldsymbol{\theta}_\mathrm{L},\boldsymbol{\theta}_\mathrm{H},\rho)\end{bmatrix}\right),$$

with

$$\begin{aligned} k_\mathrm{LL}(\boldsymbol{X}_\mathrm{L},\boldsymbol{X}_\mathrm{L};\boldsymbol{\theta}_\mathrm{L}) &= k_\mathrm{L}(\boldsymbol{X}_\mathrm{L},\boldsymbol{X}_\mathrm{L};\boldsymbol{\theta}_\mathrm{L}),\\ k_\mathrm{LH}(\boldsymbol{X}_\mathrm{L},\boldsymbol{X}_\mathrm{H};\boldsymbol{\theta}_\mathrm{L},\rho) &= \rho k_\mathrm{L}(\boldsymbol{X}_\mathrm{L},\boldsymbol{X}_\mathrm{H};\boldsymbol{\theta}_\mathrm{L}),\\ k_\mathrm{HH}(\boldsymbol{X}_\mathrm{H},\boldsymbol{X}_\mathrm{H};\boldsymbol{\theta}_\mathrm{L},\boldsymbol{\theta}_\mathrm{H},\rho) &= \rho^2 k_\mathrm{L}(\boldsymbol{X}_\mathrm{H},\boldsymbol{X}_\mathrm{H};\boldsymbol{\theta}_\mathrm{L}) + k_\mathrm{H}(\boldsymbol{X}_\mathrm{H},\boldsymbol{X}_\mathrm{H};\boldsymbol{\theta}_\mathrm{H}).\end{aligned}$$

In the following, we will use the notation

$$K(\boldsymbol{X}_\mathrm{L},\boldsymbol{X}_\mathrm{H};\boldsymbol{\theta}_\mathrm{L},\boldsymbol{\theta}_\mathrm{H},\rho) := \begin{bmatrix} k_\mathrm{LL}(\boldsymbol{X}_\mathrm{L},\boldsymbol{X}_\mathrm{L};\boldsymbol{\theta}_\mathrm{L}) & k_\mathrm{LH}(\boldsymbol{X}_\mathrm{L},\boldsymbol{X}_\mathrm{H};\boldsymbol{\theta}_\mathrm{L},\rho)\\ k_\mathrm{LH}(\boldsymbol{X}_\mathrm{H},\boldsymbol{X}_\mathrm{L};\boldsymbol{\theta}_\mathrm{L},\rho) & k_\mathrm{HH}(\boldsymbol{X}_\mathrm{H},\boldsymbol{X}_\mathrm{H};\boldsymbol{\theta}_\mathrm{L},\boldsymbol{\theta}_\mathrm{H},\rho)\end{bmatrix}$$

for the joint covariance matrix. The components k_LL and k_HH describe the spatial correlation of the low and the high fidelity levels, respectively, whereas the component k_LH describes the cross-correlation between the two levels of fidelity. The presence of the latter component is the reason why the multi-fidelity GP classifier is able to exploit the inter-model correlation. In this work, we set both kernels k_L and k_H to be of Matérn type, hence the parameters are $\boldsymbol{\theta}_\mathrm{L} = (\eta_\mathrm{L}, \ell_\mathrm{L})$ and $\boldsymbol{\theta}_\mathrm{H} = (\eta_\mathrm{H}, \ell_\mathrm{H})$. Similarly to Subsection 6.2.2, we prescribe prior distributions over the kernel parameters $\boldsymbol{\theta}_\mathrm{L}$ and $\boldsymbol{\theta}_\mathrm{H}$ and over the scaling parameter ρ. Our assumption for the priors is

$$\begin{aligned} \eta_\mathrm{L}, \eta_\mathrm{H} &\sim \mathrm{HalfNormal}(\sigma = 1000),\\ \ell_\mathrm{L}, \ell_\mathrm{H} &\sim \mathrm{Gamma}(\alpha = 2, \beta = 2),\\ \rho &\sim \mathrm{Normal}(\mu = 0, \sigma = 10).\end{aligned}$$

The posterior distributions over $\boldsymbol{\theta}_\mathrm{L}$, $\boldsymbol{\theta}_\mathrm{H}$ and ρ, which in turn determine the posterior distribution of the latent functions, can be inferred based on the training set as in Subsection 6.2.2. The prediction of the classification output at a new pacing site $\boldsymbol{x}^* \in \mathscr{S}$ is again performed by conditioning on the training data. The posterior distribution of the predictive latent function is $f_{\mathrm{H}}^*(\boldsymbol{x}^*) \sim \mathcal{N}\big(\mu(\boldsymbol{x}^*), \Sigma(\boldsymbol{x}^*)\big)$ with

$$\begin{aligned}
\mu(\boldsymbol{x}^*) &= \tilde{k}(\boldsymbol{X}_{\mathrm{L}}, \boldsymbol{X}_{\mathrm{H}}, \boldsymbol{x}^*; \boldsymbol{\theta}_{\mathrm{L}}, \boldsymbol{\theta}_{\mathrm{H}}, \rho)^{\intercal} K(\boldsymbol{X}_{\mathrm{L}}, \boldsymbol{X}_{\mathrm{H}}; \boldsymbol{\theta}_{\mathrm{L}}, \boldsymbol{\theta}_{\mathrm{H}}, \rho)^{-1} \begin{bmatrix} f_{\mathrm{L}}(\boldsymbol{X}_{\mathrm{L}}) \\ f_{\mathrm{H}}(\boldsymbol{X}_{\mathrm{H}}) \end{bmatrix}, \\
\Sigma(\boldsymbol{x}^*) &= k_{\mathrm{HH}}(\boldsymbol{x}^*, \boldsymbol{x}^*; \boldsymbol{\theta}_{\mathrm{L}}, \boldsymbol{\theta}_{\mathrm{H}}, \rho) \\
&\quad - \tilde{k}(\boldsymbol{X}_{\mathrm{L}}, \boldsymbol{X}_{\mathrm{H}}, \boldsymbol{x}^*; \boldsymbol{\theta}_{\mathrm{L}}, \boldsymbol{\theta}_{\mathrm{H}}, \rho)^{\intercal} K(\boldsymbol{X}_{\mathrm{L}}, \boldsymbol{X}_{\mathrm{H}}; \boldsymbol{\theta}_{\mathrm{L}}, \boldsymbol{\theta}_{\mathrm{H}}, \rho)^{-1} \tilde{k}(\boldsymbol{X}_{\mathrm{L}}, \boldsymbol{X}_{\mathrm{H}}, \boldsymbol{x}^*; \boldsymbol{\theta}_{\mathrm{L}}, \boldsymbol{\theta}_{\mathrm{H}}, \rho),
\end{aligned}$$

where

$$\tilde{k}(\boldsymbol{X}_{\mathrm{L}}, \boldsymbol{X}_{\mathrm{H}}, \boldsymbol{x}^*; \boldsymbol{\theta}_{\mathrm{L}}, \boldsymbol{\theta}_{\mathrm{H}}, \rho) = \begin{bmatrix} k_{\mathrm{LH}}(\boldsymbol{X}_{\mathrm{L}}, \boldsymbol{x}^*; \boldsymbol{\theta}_{\mathrm{L}}, \rho) \\ k_{\mathrm{HH}}(\boldsymbol{X}_{\mathrm{H}}, \boldsymbol{x}^*; \boldsymbol{\theta}_{\mathrm{L}}, \boldsymbol{\theta}_{\mathrm{H}}, \rho) \end{bmatrix}.$$

Similarly to Subsection 6.2.2, the mean μ and the variance Σ can be sampled using $\boldsymbol{\theta}_{\mathrm{L}}$, $\boldsymbol{\theta}_{\mathrm{H}}$, ρ, f_{L} and f_{H} drawn from the posterior distributions. The samples are then averaged and the predictive latent function $\hat{f}_{\mathrm{H}}^*$ is drawn from the posterior distribution. Finally, the prediction at the new location is given by $y_{\mathrm{H}}^* = \sigma(\hat{f}_{\mathrm{H}}^*)$.

6.3 Numerical experiments

In this section we present some numerical experiments to evaluate the performance of the classifiers introduced in Section 6.2. We consider nine atrial models on a single atrial geometry, for which we estimate the AF inducibility map. The atrial anatomy is the one described in Section 3.5, which reproduces the features described in Section 2.2 and illustrated in Figure 2.4. We thus have a computational domain Ω that can be discretized on the fine and the coarse levels. The atrial model properties set on the fine level can be transferred to the coarse level following the procedure of Section 5.1. The high and the low fidelity simulations are performed on Ω. We also have an atrial surface $\mathscr{S}$, which is discretized by a triangulated mesh with 3′298 nodes. The classification is performed on $\mathscr{S}$. For this purpose, we center the nodes around the origin and we normalize the geometry by the largest standard deviation of one of its coordinates. The nine atrial models are obtained by combining the three fibrosis patterns and the three ablation treatments illustrated in Figure 6.1. Concerning the fibrosis, we consider one pattern with 50% fibrotic tissue, corresponding to moderate fibrosis, and two different patterns with 70% fibrotic tissue, corresponding to severe fibrosis. In all cases we consider endomysial fibrosis. Concerning the ablation treatments, we consider pulmonary veins isolation (PVI), which consists of creating ablation lines around the pulmonary veins, and PVI with roof lines (PVI+BOX), which

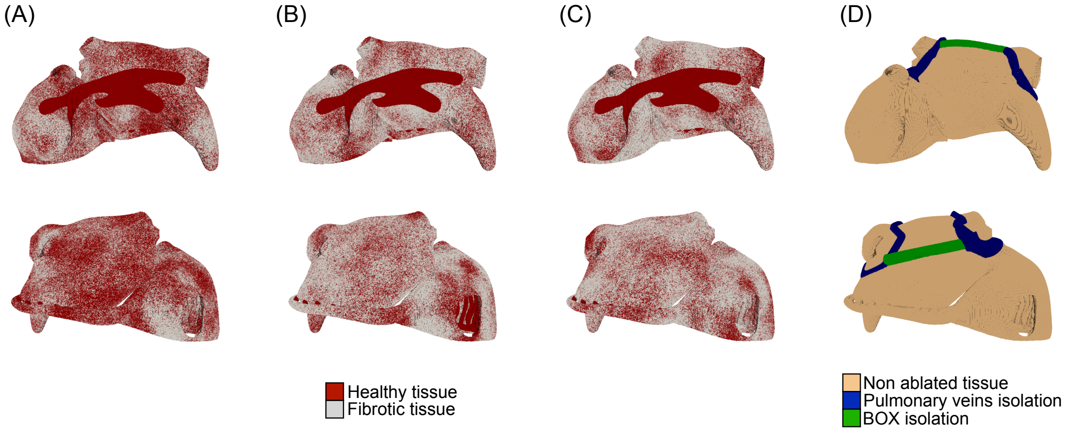

Figure 6.1. Atrial models. Panel A: moderate fibrosis pattern. Panel B: first pattern of severe fibrosis. Panel C: second pattern of severe fibrosis. Panel D: ablation lines.

consists of creating two ablation lines connecting the already present PVI lines. These are two standard-of-care ablation strategies employed in clinical practice. Additionally, we also consider the case of no ablation. We remark that the first case of severe fibrosis without ablation corresponds to the atrial model illustrated in Figure 5.1 and employed in Section 5.3.

We compare the performances of the nearest neighbor, the single-fidelity GP and the multi-fidelity GP classifiers. The classifiers exploit the information acquired on the training set to make predictions at all the nodes of the mesh discretizing the atrial surface $\mathscr{S}$. The performance of the classifiers is then measured by the balanced accuracy on a test set. This score is defined as

$$\text{balanced accuracy} = \frac{1}{2}\left(\frac{\#\text{ of predicted positives}}{\#\text{ of real positives}} + \frac{\#\text{ of predicted negatives}}{\#\text{ of real negatives}}\right).$$

This metric is suitable as it can handle imbalanced datasets. In particular, it detects if a classifier predominately predicts the class that is mostly represented in a dataset.

Our training and test sets contain 100 pacing sites and are obtained following the procedure described in Subsection 4.7.1, thus maximizing the spread on the atrial surface $\mathscr{S}$. The resulting training and test sets are shown in Figure 6.2, panel A. For each of our nine atrial models, we run the high fidelity model on the test set and both the high and the low fidelity models on the training set. This results in 1’800 high fidelity simulations and 900 low fidelity simulations, for a total cost of around 25’000 node-hours. We compare the inducibility outputs of the high and the low fidelity models on the training set. In Figure 6.2,

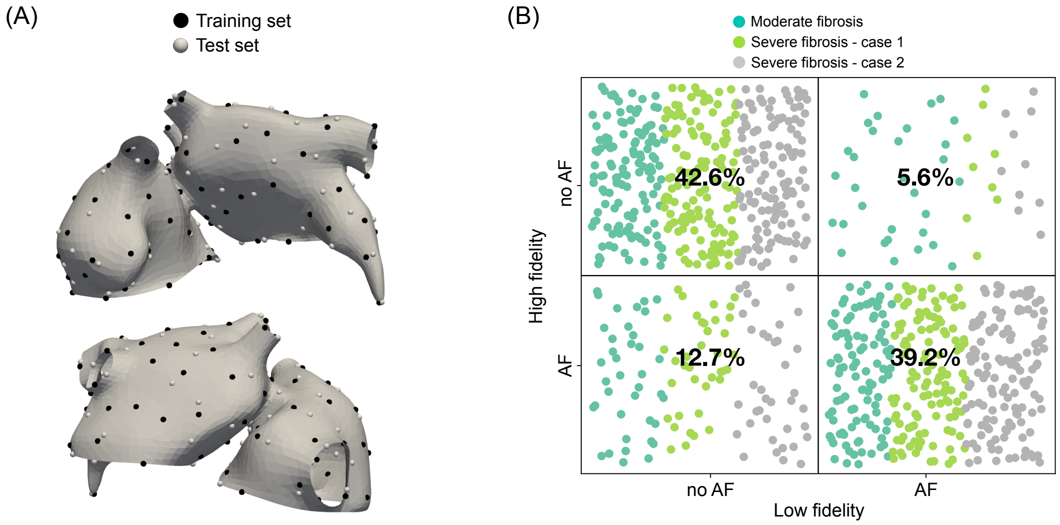

Figure 6.2. Panel A: training and test sets. Panel B: agreement between the high and the low fidelity models on the training set.

panel B, we show the confusion matrix of the agreement of the two models over 900 simulations per fidelity. The outputs of our nine atrial models are colored according to the level of fibrotic tissue. For more detailed results, we refer to Table 6.1, where we report the high and the low fidelity inducibility and the agreement separately for the nine atrial models. Here the inducibility is defined as the fraction of the simulations on the training set that result in an AF event. There is high agreement between the two models, with 81.8% of the low fidelity outputs matching the high fidelity outputs. However, the low fidelity inducibility is lower than the high fidelity inducibility, meaning that the low fidelity model is biased towards not yielding the AF event when the high fidelity model yields it. As expected, both models agree on the increase in inducibility as the degree of fibrotic tissue increases and on the decrease in inducibility when more ablation lines are present.

To analyze the performance of the classifiers, we consider an increasing size of the training set. In particular, the classifiers are trained with high fidelity input/output pairs ranging from 20 to 100. The multi-fidelity GP classifier additionally exploits 100 low fidelity training input/output pairs. The cost of the 100 low fidelity simulations is comparable to the cost of only $100/16 = 6.25$ high fidelity simulations. Additionally to the nearest neighbor, the single-fidelity GP and the multi-fidelity GP classifiers, here we also consider the low fidelity GP classifier trained following the procedure of Section 6.2.2 with 100 low fidelity input/output pairs. The cost of the low fidelity GP classifier is fixed and is

		Inducibility [%]		
Fibrosis	Ablation	High fidelity	Low fidelity	Agreement [%]
Moderate	-	58	55	77
Moderate	PVI	43	41	76
Moderate	PVI+BOX	40	38	78
Severe - case 1	-	62	52	82
Severe - case 1	PVI	50	42	86
Severe - case 1	PVI+BOX	47	36	85
Severe - case 2	-	65	57	84
Severe - case 2	PVI	54	44	84
Severe - case 2	PVI+BOX	48	38	84

Table 6.1. Inducibility and agreement on the training set.

comparable to the cost of 6.25 high fidelity simulations. The accuracy results for our nine atrial models are shown in Figure 6.3. The balanced accuracies of the nearest neighbor, the single-fidelity GP and the multi-fidelity GP classifiers are plotted against the cost expressed in terms of high fidelity simulations. The line corresponding to the multi-fidelity GP classifier is therefore shifted to the right by the cost of 100 low fidelity simulations. The balanced accuracy of the low fidelity GP classifier is represented by an horizontal line. As expected, for all the classifiers the accuracy tends to increase as the size of the training set increases. In general, for small training sets the multi-fidelity GP classifier reaches higher accuracies than the nearest neighbor and the single-fidelity GP classifiers, whereas for large training sets the accuracy of the classifiers is comparable. Interestingly, for low costs, the accuracies of the nearest neighbor and the single-fidelity GP classifiers are even below the accuracy of the low fidelity GP classifier.

Clinical studies attempting to design optimal ablation strategies are commonly based on the simulations resulting from the stimulation at 40 pacing sites (Boyle et al. [2019]). Therefore, in the following we focus on the classifiers that have the cost of 40 high fidelity simulations. The nearest neighbor and the single-fidelity GP classifiers trained with 40 high fidelity input/output pairs have such cost. The multi-fidelity alternative with equivalent cost is the multi-fidelity GP classifier trained with 33 high fidelity and 100 low fidelity input/output pairs. Additionally, we also consider the low fidelity GP classifier trained with 100 low fidelity input/output pairs, that has lower cost. The balanced accuracies of the

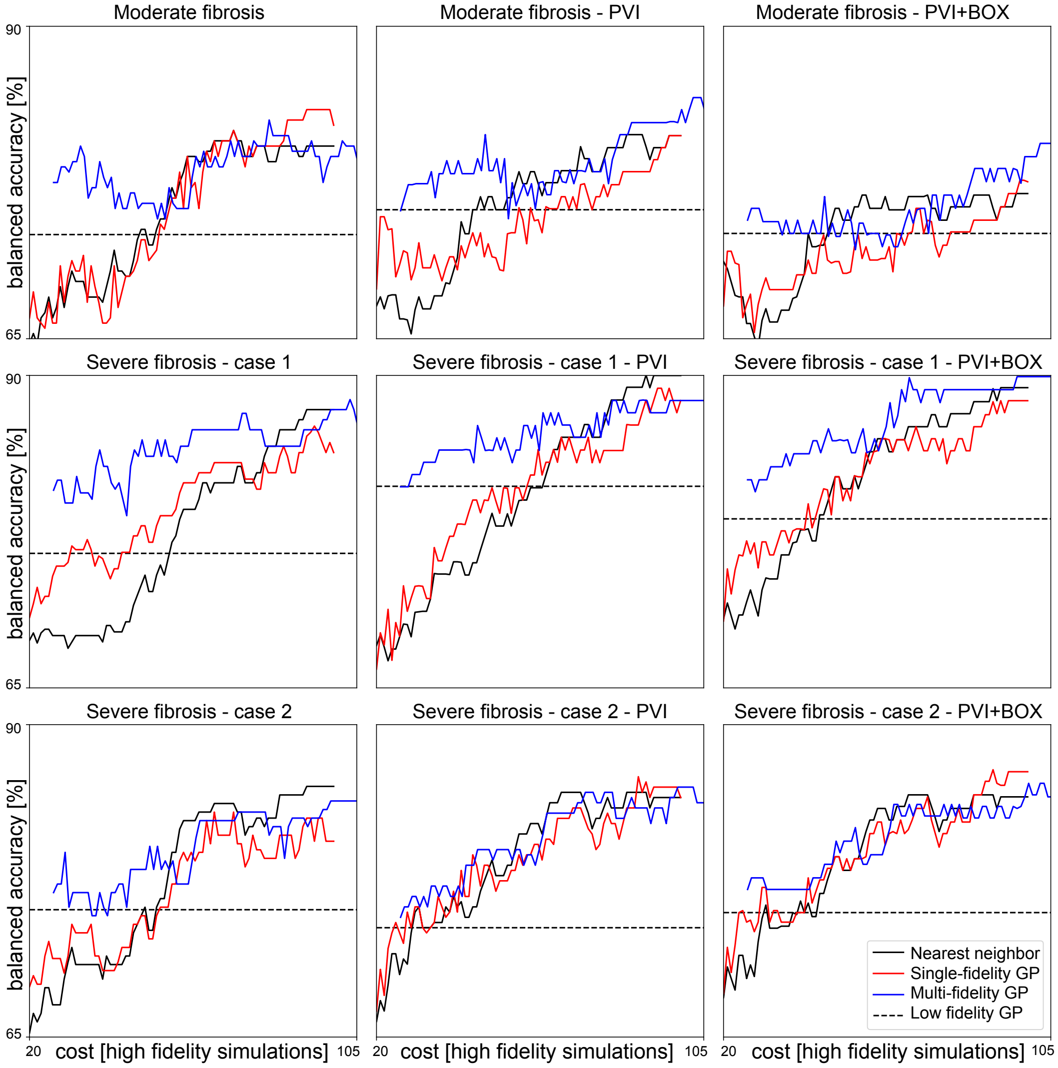

Figure 6.3. Balanced accuracy vs. cost of the nearest neighbor, the single-fidelity GP and the multi-fidelity GP classifiers for the nine atrial models, together with the balanced accuracy of the low fidelity GP classifier.

four classifiers are illustrated in Figure 6.4. The lines connect the accuracies that refer to the same atrial model. We use the Mann Whitney U test (Mann and Whitney [1947], Sahli Costabal et al. [2019]) with 5% significance level to evaluate whether the multi-fidelity GP and the low fidelity GP classifiers are more accurate than the single-fidelity GP and the nearest neighbor classifiers. The multi-fidelity GP classifier is significantly more accurate than the nearest neighbor (p-value=0.06) and the single-fidelity GP (p-value=0.020) classifiers. On average, the gain in accuracy of the multi-fidelity GP classifier is 6% compared to the nearest neighbor classifier and 4.9% compared to the single-fidelity GP classifier. Instead, the gain in accuracy of the low fidelity GP classifier on average is 2.8% compared to the nearest neighbor classifier and 1.7% compared to the single-fidelity GP classifier. However, the low fidelity GP classifier is not significantly more accurate than the nearest neighbor (p-value=0.091) and the single-fidelity GP (p-value=0.252) classifiers. These nearest neighbor, single-fidelity GP, multi-fidelity GP and low fidelity GP classifiers are shown in Figures 6.5, 6.6 and 6.7 for our nine atrial models. The predicted classification outputs are shown on the atrial surface. On top of the atrial surface, we show the classification outputs of the high fidelity simulations on the test set, which allow to compute the balanced accuracy. In general, the nearest neighbor and the single-fidelity GP classifiers show similar patterns and have comparable accuracies. Instead, the multi-fidelity GP classifier shares some features with the low fidelity GP classifier

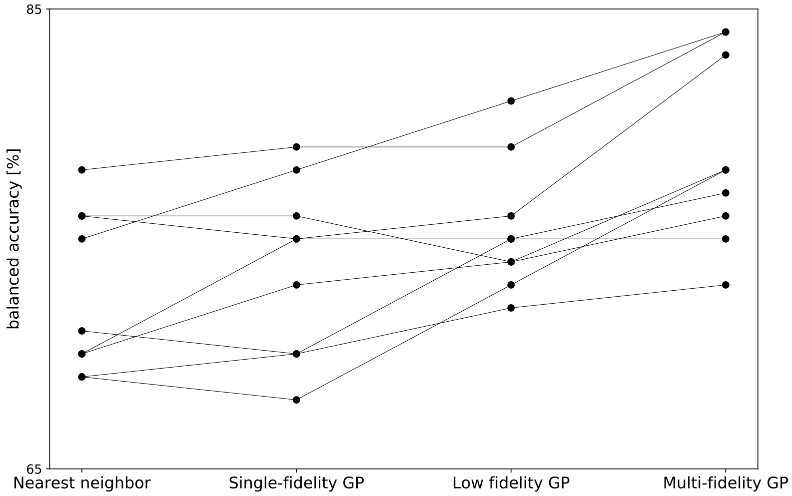

Figure 6.4. Balanced accuracies of the nearest neighbor, the single-fidelity GP, the low fidelity GP and the multi-fidelity GP classifiers for the nine atrial models.

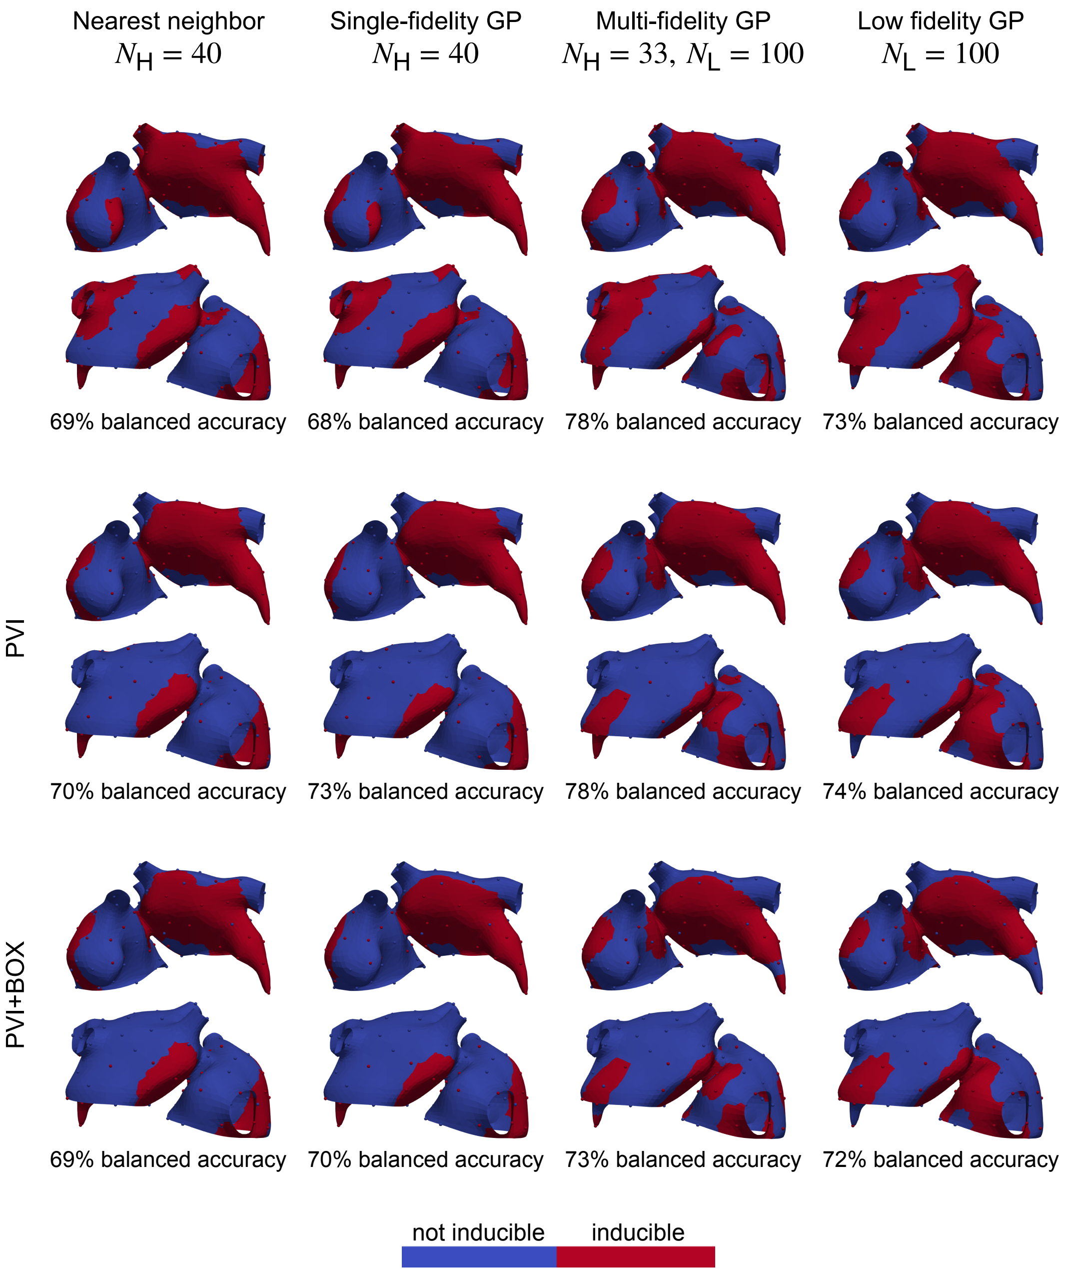

Figure 6.5. Classifiers for the atrial models with moderate fibrosis.

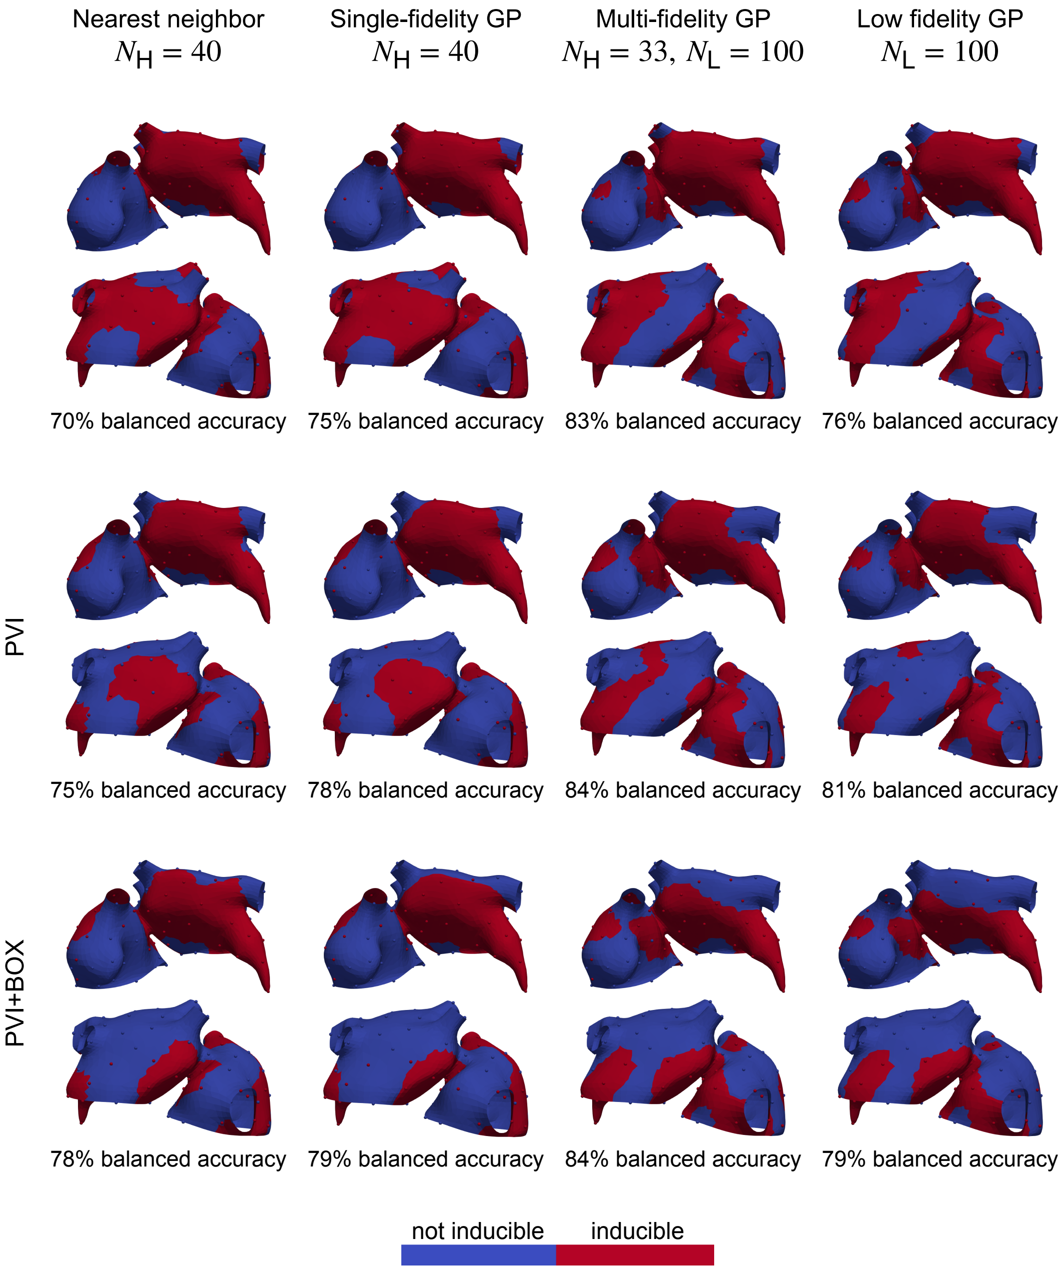

Figure 6.6. Classifiers for the atrial models with the first pattern of severe fibrosis.

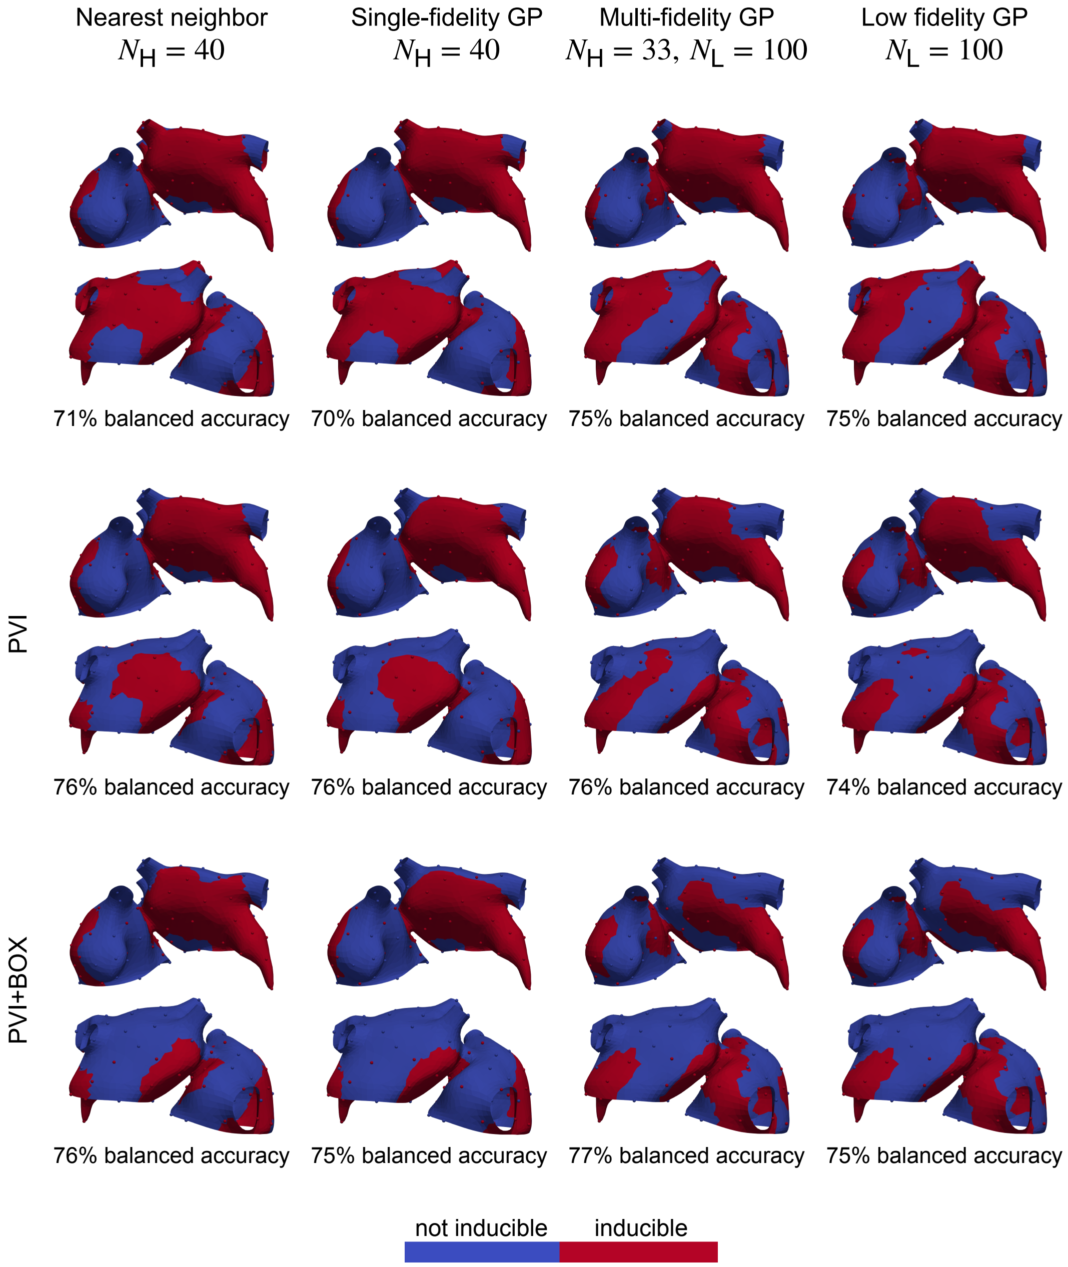

Figure 6.7. Classifiers for the atrial models with the second pattern of severe fibrosis.

that, in general, allow to reach a higher accuracy. In many cases, even the accuracy of the low fidelity GP classifier is higher than the accuracies of the nearest neighbor and the single-fidelity GP classifiers.

We conclude this section by remarking that the accuracy of the classifiers depends on the length scale of the inducibility map. In particular, the accuracy decreases as the length scale decreases. The shorter the length scale, the larger the training set must be in order to capture the classification regions with small characteristic size. An estimate of the length scale of the inducibility map can be obtained by training a single-fidelity classifier from the dataset of input/output pairs given by the high fidelity simulations run on both the training and test sets. The resulting average length scale over our nine atrial models is 0.28 cm. This small length scale of the inducibility maps indicates that pacing sites that are close on the atrial surface might lead to opposite AF inducibility outcomes. The average geodesic distance between the 100 pacing sites in the training set is 0.39 cm, which is larger than the average length scale of the inducibility maps. This may explain why the balanced accuracies of the classifiers are lower than 90%. However, the inducibility score is not strongly affected by the length scale of the inducibility map, as the overall fraction of tissue where AF is inducible is only slightly affected by the number of disconnected inducibility regions and by the smoothness of their boundaries.

6.4 Discussion

In this chapter we introduced a multi-fidelity GP classifier for the estimation of AF inducibility maps of the atria. We combined the high fidelity model and the low fidelity model based on coarser discretization parameters in this multi-fidelity framework. We evaluated the performance of the multi-fidelity method by comparing it to two single-fidelity methods, namely the single-fidelity GP and the nearest neighbor classifiers. The comparison was done in terms of accuracy and computational cost.

The GP classification is efficient, as the cost of training and predicting with the GP classifier is negligible compared to the cost of running the high fidelity simulations and is comparable to the cost of the nearest neighbor classification. Indeed, the training and the prediction phases take less than 5 min on a laptop. Moreover, the GP classification framework is non-intrusive and the atrial model is a black box. In this chapter we considered nine atrial models on a single geometry, but the same approach can be applied to any atrial model. In particular, thanks to the standardization and the normalization of the atrial geometry, the same prior distributions can be assumed on generic anatomies. Additionally, the topological properties of the anatomy are captured by the classifier, as it operates directly on the atrial surface. The methodology is also robust to changes in the local conduction properties of the tissue.

The low fidelity model is biased towards not inducing AF when the high fidelity model induces it. The reduced inducibility might be due to the fact that the low fidelity model captures less fine scale details than the high fidelity model, as it is based on a coarser discretization of the computational domain. The fine computational mesh allows to detect more tissue heterogeneities, which increase the risk of occurrence of an AF event. Another reason for the reduced inducibility might be the increase of the CV in the low fidelity model to match the high fidelity CV. This change, achieved by decreasing the surface-to-volume ratio, may be anti-arrhythmic. However, this change is required to reach higher correlation between the high and the low fidelity models and, therefore, to enhance the performance of the multi-fidelity GP classifier.

Indeed, compared to the single-fidelity classifiers, the multi-fidelity GP classifier can achieve higher accuracy with a fixed computational cost or, equivalently, it can achieve a target accuracy with lower computational cost. In particular, the multi-fidelity GP classifier performs better than the nearest neighbor and the single-fidelity GP classifiers for the typical cost of 40 high fidelity simulations. For low costs, e.g. when only 20 high fidelity simulations are feasible (Gharaviri et al. [2021a], Gharaviri et al. [2021b]), the nearest neighbor and the single-fidelity GP classifiers do not achieve sufficient accuracy, whereas the multi-fidelity GP classifier can already achieve a good estimate of the inducibility map. When the computational budget is even more limited, the cheap low fidelity GP classifier is a valid choice to get an insight into the inducible regions. Therefore it is worth taking advantage of the cheap low fidelity model, either in a single-fidelity or in a multi-fidelity framework. Moreover, the low fidelity model based on coarser discretization parameters is easy to implement and it does not need a training phase, so it does not require additional cost.

Chapter 7

Conclusion

In this work we first presented the monodomain high fidelity model for atrial fibrillation. We then developed two low fidelity models, one based on the eikonal model and one based on a coarser discretization of the computational domain in the numerical solution of the monodomain equation. We tested the quality of the low fidelity approximations compared to the high fidelity results in numerical studies and we explained the discrepancies. The numerical experiments included an evaluation of the AF inducibility. We finally employed the low fidelity model based on coarser discretization parameters in a multi-fidelity framework to estimate atrial inducibility maps. We compared the accuracy of the multi-fidelity estimates to the accuracy of estimates obtained from single-fidelity approaches with equivalent cost.

Personalized clinical investigations on AF, such as the evaluation and the design of the ablation treatment, are based on *in-silico* studies. The starting point is the patient-specific atrial anatomy (the geometry and the local conduction properties), which can be acquired through imaging techniques (Oakes et al. [2009], Siebermair et al. [2017]). The computational models then try to reproduce the realistic electrophysiology of the patients, for which the amount of measured data is very limited. The noise in the acquired anatomical data affects both the monodomain (Pathmanathan et al. [2019]) and the eikonal (Quaglino et al. [2018]) models. The effect of the variability in the anatomical data on the electrophysiological outcomes has been analyzed in the uncertainty quantification studies collected in Clayton et al. [2020]. In this work, we did not focus on the acquisition of patient-specific atrial anatomies. Instead, we focused on the electrophysiological models and on their applicability to personalized investigations. The accurate high fidelity model is computationally expensive and often fails to meet the clinical time constraints. Therefore, less accurate but faster low fidelity models need to be developed.

The eikonal low fidelity model is very attractive as it is potentially real-time. This model consists of an adaptation of the FMM to include the re-excitability. The model is qualitatively accurate in simple numerical experiments. However, the accuracy decreases in presence of heterogeneities in conduction. Moreover, when used to evaluate the AF inducibility, the model shows poor agreement with the high fidelity results. These discrepancies occur despite the adjustment of the CV and the use of the restitution properties computed from the monodomain model. The main reasons for the discrepancies compared to the high fidelity model are the fact that the eikonal model does not take into account the diffusion currents, either due to heterogeneities in the propagation or in the conductivity, and the fact that the restitution properties included in the eikonal algorithm do not always match the restitution properties of the monodomain model. Moreover, the eikonal computational mesh is coarser than the monodomain computational mesh and not all its elements satisfy the acuteness condition required by the FMM. The improvement of the accuracy of the eikonal low fidelity model by reducing the discrepancies to the high fidelity model is part of our future work. The computation of the restitution curves could be improved following the approaches described in Colli Franzone et al. [2014] or in Courtemanche et al. [1996]. Other improvements could be achieved by using the FIM to solve the eikonal equation, by employing the eikonal-diffusion or the eikonal-curvature models, or by refining the computational mesh. These improvements would lead to an increase in the computational cost. The quantification of this accuracy/cost trade-off is our future interest.

A valid alternative is the low fidelity model based on coarser discretization parameters. This model is approximately 16 times faster than the high fidelity model. Thanks to the adjustment of the CV and to the proper procedure to coarsen the underlying atrial model, the model is sufficiently accurate. In particular, when used to evaluate the AF inducibility, the model highly agrees to the high fidelity model. This high correlation can be exploited in the multi-fidelity framework to efficiently estimate atrial inducibility maps. Indeed, at low computational costs, the multi-fidelity estimates achieve higher accuracies than the single-fidelity estimates. Moreover, when the computational budget is limited and the high fidelity simulations are not feasible, the low fidelity model can be employed in the single-fidelity setting to obtain cheap estimates of the inducibility maps.

Bibliography

Ken Alton and Ian M. Mitchell. An ordered upwind method with precomputed stencil and monotone node acceptance for solving static convex Hamilton-Jacobi equations. *Journal of Scientific Computing*, 51:313–348, 2012. doi: 10.1007/s10915-011-9512-4.

Francisco J. Asencor and Manuel Panizo. Finite-difference operators in anisotropic inhomogeneous dielectrics: general case. *Journal of Computational Physics*, 95(2):387–399, 1991. doi: 10.1016/0021-9991(91)90282-P.

Luca Azzolin, Steffen Schuler, Olaf Dössel, and Axel Loewe. A reproducible protocol to assess arrhythmia vulnerability in silico: pacing at the end of the effective refractory period. *Frontiers in Physiology*, 12:656411, 2021. doi: 10.3389/fphys.2021.656411.

Cristian Barrios Espinosa, Jorge Sánchez, Olaf Dössel, and Axel Loewe. Diffusion reaction eikonal alternant model: towards fast simulations of complex cardiac arrhythmias. In *Computing in Cardiology (CinC)*, volume 498, 2022. doi: 10.22489/CinC.2022.054.

Folkmar Bornemann and Christian Rasch. Finite-element discretization of static Hamilton-Jacobi equations based on a local variational principle. *Computing and Visualization in Science*, 9:57–69, 2006. doi: 10.1007/s00791-006-0016-y.

Viacheslav Borovitskiy, Alexander Terenin, Peter Mostowsky, and Marc Deisenroth. Matérn Gaussian processes on Riemannian manifolds. In *Advances in Neural Information Processing Systems*, volume 33, pages 12426–12437, 2020.

Patrick M. Boyle, Tarek Zghaib, Sohail Zahid, Rheeda L. Ali, Dongdong Deng, William H. Franceschi, Joe B. Hakim, Michael J. Murphy, Adityo Prakosa, Stefan L. Zimmerman, et al. Computationally guided personalized targeted ablation of persistent atrial fibrillation. *Nature Biomedical Engineering*, 3:870–879, 2019. doi: 10.1038/s41551-019-0437-9.

Patrick M. Boyle, Alexander R. Ochs, Rheeda L. Ali, Nikhil Paliwal, and Natalia A. Trayanova. Characterizing the arrhythmogenic substrate in personalized models of atrial fibrillation: sensitivity to mesh resolution and pacing protocol in AF models. *EP Europace*, 23(Supplement_1):i3–i11, 2021. doi: 10.1093/europace/euaa385.

Shih-Ann Chen, Ming-Hsiung Hsieh, Ching-Tai Tai, Chin-Feng Tsai, V. S. Prakash, Wen-Chung Yu, Tsui-Lieh Hsu, Yu-An Ding, and Mau-Song Chang. Initiation of atrial fibrillation by ectopic beats originating from the pulmonary veins: electrophysiological characteristics, pharmacological responses, and effects of radiofrequency ablation. *Circulation*, 100(18):1879–1886, 1999. doi: 10.1161/01.CIR.100.18.1879.

Elizabeth M. Cherry, Harold M. Hastings, and Steven J. Evans. Dynamics of human atrial cell models: restitution, memory, and intracellular calcium dynamics in single cells. *Progress in Biophysics and Molecular Biology*, 98(1):24–37, 2008. doi: 10.1016/j.pbiomolbio.2008.05.002.

Ludovica Cicci, Stefania Fresca, Stefano Pagani, Andrea Manzoni, and Alfio Quarteroni. Projection-based reduced order models for parameterized nonlinear time-dependent problems arising in cardiac mechanics. *Mathematics in Engineering*, 5(2):1–38, 2022. doi: 10.3934/mine.2023026.

Richard H. Clayton, Yasser Aboelkassem, Chris D. Cantwell, Cesare Corrado, Tammo Delhaas, Wouter Huberts, Chon Lok Lei, Haibo Ni, Alexander V. Panfilov, Caroline Roney, et al. An audit of uncertainty in multi-scale cardiac electrophysiology models. *Philosophical Transactions of the Royal Society A - Mathematical, Physical and Engineering Sciences*, 378(2173), 2020. doi: 10.1098/rsta.2019.0335.

Piero Colli Franzone, Luciano Guerri, and Sergio Rovida. Wavefront propagation in an activation model of the anisotropic cardiac tissue: asymptotic analysis and numerical simulations. *Journal of Mathematical Biology*, 28:121–176, 1990a. doi: 10.1007/BF00163143.

Piero Colli Franzone, Luciano Guerri, and Stefania Tentoni. Mathematical modeling of the excitation process in myocardial tissue: influence of fiber rotation on wavefront propagation and potential field. *Mathematical Biosciences*, 101 (2):155–235, 1990b. doi: 10.1016/0025-5564(90)90020-Y.

Piero Colli Franzone, Luca F. Pavarino, and Simone Scacchi. *Mathematical cardiac electrophysiology*, volume 13. Springer, 2014.

Cesare Corrado and Nejib Zemzemi. A conduction velocity adapted eikonal model for electrophysiology problems with re-excitability evaluation. *Medical Image Analysis*, 43:186–197, 2018. doi: 10.1016/j.media.2017.11.002.

Marc Courtemanche, James P. Keener, and Leon Glass. A delay equation representation of pulse circulation on a ring in excitable media. *SIAM Journal on Applied Mathematics*, 56(1):119–142, 1996. doi: 10.1137/S0036139993259337.

Marc Courtemanche, Rafael J. Ramirez, and Stanley Nattel. Ionic mechanisms underlying human atrial action potential properties: insights from a mathematical model. *American Journal of Physiology - Heart and Circulatory Physiology*, 275(1):H301–H321, 1998. doi: 10.1152/ajpheart.1998.275.1.H301.

Sam Coveney, Cesare Corrado, Caroline H. Roney, Richard D. Wilkinson, Jeremy E. Oakley, Finn Lindgren, Steven E. Williams, Mark D. O'Neill, Steven A. Niederer, and Richard H. Clayton. Probabilistic interpolation of uncertain local activation times on human atrial manifolds. *IEEE Transactions on Biomedical Engineering*, 67(1):99–109, 2019. doi: 10.1109/TBME.2019.2908486.

Sam Coveney, Cesare Corrado, Caroline H. Roney, Daniel O'Hare, Steven E. Williams, Mark D. O'Neill, Steven A. Niederer, Richard H. Clayton, Jeremy E. Oakley, and Richard D. Wilkinson. Gaussian process manifold interpolation for probabilistic atrial activation maps and uncertain conduction velocity. *Philosophical Transactions of the Royal Society A*, 378(2173):20190345, 2020. doi: 10.1098/rsta.2019.0345.

Keenan Crane, Clarisse Weischedel, and Max Wardetzky. Geodesics in heat: a new approach to computing distance based on heat flow. *ACM Transactions on Graphics*, 32(5), 2013. doi: 10.1145/2516971.2516977.

Jwala Dhamala, Pradeep Bajracharya, Hermenegild J. Arevalo, John L. Sapp, B. Milan Horácek, Katherine C. Wu, Natalia A. Trayanova, and Linwei Wang. Embedding high-dimensional Bayesian optimization via generative modeling: parameter personalization of cardiac electrophysiological models. *Medical Image Analysis*, 62:101670, 2020. doi: 10.1016/j.media.2020.101670.

Edsger W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271, 1959. doi: 10.1007/BF01386390.

Cécile Dobrzynski and Pascal Frey. Anisotropic Delaunay mesh adaptation for unsteady simulations. In *Proceedings of the 17th International Meshing Roundtable*, pages 177–194. Springer, 2008. doi: 10.1007/978-3-540-87921-3_11.

Vladimir G. Fast and André G. Kléber. Role of wavefront curvature in propagation of cardiac impulse. *Cardiovascular Research*, 33(2):258–271, 1997. doi: 10.1016/S0008-6363(96)00216-7.

Stefania Fresca, Andrea Manzoni, Luca Dedè, and Alfio Quarteroni. Deep learning-based reduced order models in cardiac electrophysiology. *PLoS ONE*, 15(10):e0239416, 2020. doi: 10.1371/journal.pone.0239416.

Zhisong Fu, Won-Ki Jeong, Yongsheng Pan, Robert M. Kirby, and Ross T. Whitaker. A fast iterative method for solving the eikonal equation on triangulated surfaces. *SIAM Journal on Scientific Computing*, 33(5):2468–2488, 2011. doi: 10.1137/100788951.

Lia Gander, Rolf Krause, Michael Multerer, and Simone Pezzuto. Space-time shape uncertainties in the forward and inverse problem of electrocardiography. *International Journal for Numerical Methods in Biomedical Engineering*, 37(10): e3522, 2021. doi: 10.1002/cnm.3522.

Lia Gander, Simone Pezzuto, Ali Gharaviri, Rolf Krause, Paris Perdikaris, and Francisco Sahli Costabal. Fast characterization of inducible regions of atrial fibrillation models with multi-fidelity Gaussian process classification. *Frontiers in Physiology*, 13(757159), 2022. doi: 10.3389/fphys.2022.757159.

Lia Gander, Rolf Krause, Martin Weiser, Francisco Sahli Costabal, and Simone Pezzuto. On the accuracy of eikonal approximations in cardiac electrophysiology in the presence of fibrosis. In *International Conference on Functional Imaging and Modeling of the Heart*, pages 137–146, 2023. doi: 10.1007/978-3-031-35302-4_14.

Narimane Gassa, Nejib Zemzemi, Cesare Corrado, and Yves Coudière. Spiral waves generation using an eikonal-reaction cardiac electrophysiology model. In *International Conference on Functional Imaging and Modeling of the Heart*, pages 523–530, 2021. doi: 10.1007/978-3-030-78710-3_50.

Ali Gharaviri, Elham Bidar, Mark Potse, Stef Zeemering, Sander Verheule, Simone Pezzuto, Rolf Krause, Jos G. Maessen, Angelo Auricchio, and Ulrich Schotten.

Epicardial fibrosis explains increased endo-epicardial dissociation and epicardial breakthroughs in human atrial fibrillation. *Frontiers in physiology*, 11(68), 2020. doi: 10.3389/fphys.2020.00068.

Ali Gharaviri, Simone Pezzuto, Mark Potse, Giulio Conte, Stef Zeemering, Vladimír Sobota, Sander Verheule, Rolf Krause, Angelo Auricchio, and Ulrich Schotten. Synergistic antiarrhythmic effect of inward rectifier current inhibition and pulmonary vein isolation in a 3D computer model for atrial fibrillation. *EP Europace*, 23:i161–i168, 2021a. doi: 10.1093/europace/euaa413.

Ali Gharaviri, Simone Pezzuto, Mark Potse, Sander Verheule, Giulio Conte, Rolf Krause, Ulrich Schotten, and Angelo Auricchio. Left atrial appendage electrical isolation reduces atrial fibrillation recurrences: a simulation study. *Circulation: Arrhythmia and Electrophysiology*, 14(1):e009230, 2021b. doi: 10.1161/CIRCEP.120.009230.

Vincenzo Gionti, Simone Scacchi, Piero Colli Franzone, Luca F. Pavarino, Roberto Dore, and Cesare Storti. Role of scar and border zone geometry on the genesis and maintenance of re-entrant ventricular tachycardia in patients with previous myocardial infarction. *Frontiers in Physiology*, 13(834747), 2022. doi: 10.3389/fphys.2022.834747.

Michel Haïssaguerre, Pierre Jaïs, Dipen C. Shah, Atsushi Takahashi, Mélèze Hocini, Gilles Quiniou, Stéphane Garrigue, Alain Le Mouroux, Philippe Le Métayer, and Jacques Clémenty. Spontaneous initiation of atrial fibrillation by ectopic beats originating in the pulmonary veins. *New England Journal of Medicine*, 339(10):659–666, 1998. doi: 10.1056/NEJM199809033391003.

Matthew D. Hoffman and Andrew Gelman. The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15:1593–1623, 2014.

Vincent Jacquemet. An eikonal approach for the initiation of reentrant cardiac propagation in reaction-diffusion models. *IEEE Transactions on Biomedical Engineering*, 57(9):2090–2098, 2010. doi: 10.1109/TBME.2010.2051156.

Vincent Jacquemet. An eikonal-diffusion solver and its application to the interpolation and the simulation of reentrant cardiac activations. *Computer Methods and Programs in Biomedicine*, 108(2):548–558, 2012. doi: 10.1016/j.cmpb.2011.05.003.

Abouzar Kaboudian, Elizabeth M. Cherry, and Flavio H. Fenton. Real-time interactive simulations of large-scale systems on personal computers and cell phones: toward patient-specific heart modeling and other applications. *Science Advances*, 5(3):eaav6019, 2019. doi: 10.1126/sciadv.aav6019.

Shunsuke Kawai, Yasushi Mukai, Shujiro Inoue, Daisuke Yakabe, Kazuhiro Nagaoka, Kazuo Sakamoto, Susumu Takase, Akiko Chishaki, and Hiroyuki Tsutsui. Non-pulmonary vein triggers of atrial fibrillation are likely to arise from low-voltage areas in the left atrium. *Scientific Reports*, 9(12271), 2019. doi: 10.1038/s41598-019-48669-1.

James Keener and James Sneyd. *Mathematical physiology I: cellular physiology*, volume 8. Springer, 2009a.

James Keener and James Sneyd. *Mathematical physiology II: systems physiology*, volume 8. Springer, 2009b.

James P. Keener. An eikonal-curvature equation for action potential propagation in myocardium. *Journal of Mathematical Biology*, 29:629–651, 1991. doi: 10.1007/BF00163916.

Marc C. Kennedy and Anthony O'Hagan. Predicting the output from a complex computer code when fast approximations are available. *Biometrika*, 87(1): 1–13, 2000. doi: 10.1093/biomet/87.1.1.

Ron Kimmel and James A. Sethian. Computing geodesic paths on manifolds. *Proceedings of the National Academy of Sciences*, 95(15):8431–8435, 1998. doi: 10.1073/pnas.95.15.8431.

Dorian Krause, Mark Potse, Thomas Dickopf, Rolf Krause, Angelo Auricchio, and Frits Prinzen. Hybrid parallelization of a large-scale heart model. In *Facing the Multicore - Challenge II*, pages 120–132. Springer, 2012. doi: 10.1007/978-3-642-30397-5_11.

Finn Lindgren, Håvard Rue, and Johan Lindström. An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(4):423–498, 2011. doi: 10.1111/j.1467-9868.2011.00777.x.

Pierre-Louis Lions. *Generalized solutions of Hamilton-Jacobi equations*. Pitman, 1982.

Axel Loewe, Emanuel Poremba, Tobias Oesterlein, Armin Luik, Claus Schmitt, Gunnar Seemann, and Olaf Dössel. Patient-specific identification of atrial flutter vulnerability - a computational approach to reveal latent reentry pathways. *Frontiers in Physiology*, 9(1910), 2019. doi: 10.3389/fphys.2018.01910.

Henry B. Mann and Donald R. Whitney. On a test of whether one of two random variables is stochastically larger than the other. *The Annals of Mathematical Statistics*, 18(1):50–60, 1947.

Kathleen S. McDowell, Sohail Zahid, Fijoy Vadakkumpadan, Joshua Blauer, Rob S. MacLeod, and Natalia A. Trayanova. Virtual electrophysiological study of atrial fibrillation in fibrotic remodeling. *PLoS one*, 10(2):e0117110, 2015. doi: 10.1371/journal.pone.0117110.

Jean-Marie Mirebeau. On the accuracy of anisotropic fast marching. *arXiv:1205.5300*, 2012.

Jean-Marie Mirebeau. Anisotropic fast-marching on Cartesian grids using lattice basis reduction. *SIAM Journal on Numerical Analysis*, 52(4), 2014. doi: 10.1137/120861667.

Claudia Nagel, Cristian Barrios Espinosa, Karli Gillette, Matthias A.F. Gsell, Jorge Sánchez, Gernot Plank, Olaf Dössel, and Axel Loewe. Comparison of propagation models and forward calculation methods on cellular, tissue and organ scale atrial electrophysiology. *IEEE Transactions on Biomedical Engineering*, 70 (2):511–522, 2023. doi: 10.1109/TBME.2022.3196144.

Radford M. Neal. Regression and classification using Gaussian process priors. *Bayesian Statistics*, 6, 1998.

Aurel Neic, Fernando O. Campos, Anton J. Prassl, Steven A. Niederer, Martin J. Bishop, Edward J. Vigmond, and Gernot Plank. Efficient computation of electrograms and ECGs in human whole heart simulations using a reaction-eikonal model. *Journal of Computational Physics*, 346:191–211, 2017. doi: 10.1016/j.jcp.2017.06.020.

Hannes Nickisch and Carl E. Rasmussen. Approximations for binary Gaussian process classification. *Journal of Machine Learning Research*, 9:2035–2078, 2008.

Robert S. Oakes, Troy J. Badger, Eugene G. Kholmovski, Nazem Akoum, Nathan S. Burgon, Eric N. Fish, Joshua J.E. Blauer, Swati N. Rao, Edward V.R.

DiBella, Nathan M. Segerson, et al. Detection and quantification of left atrial structural remodeling with delayed-enhancement magnetic resonance imaging in patients with atrial fibrillation. *Circulation*, 119(13):1758–1767, 2009. doi: 10.1161/CIRCULATIONAHA.108.811877.

Stefano Pagani, Luca Dedè, Antonio Frontera, Matteo Salvador, Luca R. Limite, Andrea Manzoni, Felicia Lipartiti, Georgios Tsitsinakis, Alexios Hadjis, Paolo Della Bella, et al. A computational study of the electrophysiological substrate in patients suffering from atrial fibrillation. *Frontiers in Physiology*, 12(673612), 2021. doi: 10.3389/fphys.2021.673612.

Manuel Panizo, Antonio Castellanos, and José Rivas. Finite-difference operators in inhomogeneous anisotropic media. *Journal of Applied Physics*, 48(3):1054–1057, 1977. doi: 10.1063/1.323779.

Pras Pathmanathan, Jonathan M. Cordeiro, and Richard A. Gray. Comprehensive uncertainty quantification and sensitivity analysis for cardiac action potential models. *Frontiers in Physiology*, 10(721), 2019. doi: 10.3389/fphys.2019.00721.

Paris Perdikaris, Daniele Venturi, and George E. Karniadakis. Multifidelity information fusion algorithms for high-dimensional systems and massive data sets. *SIAM Journal on Scientific Computing*, 38(4):B521–B538, 2016. doi: 10.1137/15M1055164.

Erik Pernod, Maxime Sermesant, Ender Konukoglu, Jatin Relan, Hervé Delingette, and Nicholas Ayache. A multi-front eikonal model of cardiac electrophysiology for interactive simulation of radio-frequency ablation. *Computers & Graphics*, 35(2):431–440, 2011. doi: 10.1016/j.cag.2011.01.008.

Simone Pezzuto, Johan Hake, and Joakim Sundnes. Space-discretization error analysis and stabilization schemes for conduction velocity in cardiac electrophysiology. *International Journal for Numerical Methods in Biomedical Engineering*, 32(10):e02762, 2016. doi: 10.1002/cnm.2762.

Simone Pezzuto, Peter Kal’avskỳ, Mark Potse, Frits W. Prinzen, Angelo Auricchio, and Rolf Krause. Evaluation of a rapid anisotropic model for ECG simulation. *Frontiers in Physiology*, 8(265), 2017. doi: 10.3389/fphys.2017.00265.

Simone Pezzuto, Alessio Quaglino, and Mark Potse. On sampling spatially-correlated random fields for complex geometries. In *International Conference* *on Functional Imaging and Modeling of the Heart*, pages 103–111, 2019. doi: 10.1007/978-3-030-21949-9_12.

Mark Potse. Inducibility of atrial fibrillation depends chaotically on ionic model parameters. In *Computing in Cardiology (CinC)*, volume 46, 2019. doi: 10.22489/CinC.2019.410.

Mark Potse, Bruno Dubé, Jacques Richer, Alain Vinet, and Ramesh 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. doi: 10.1109/TBME.2006.880875.

Mark Potse, Ali Gharaviri, Simone Pezzuto, Angelo Auricchio, Rolf Krause, Sander Verheule, and Ulrich Schotten. Anatomically-induced fibrillation in a 3D model of the human atria. In *Computing in Cardiology (CinC)*, volume 45, 2018. doi: 10.22489/CinC.2018.366.

Andrew J. Pullan, Karl A. Tomlinson, and Peter 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. doi: 10.1137/S0036139901389513.

Alessio Quaglino, Simone Pezzuto, Phaedon-Stelios Koutsourelakis, Angelo Auricchio, and Rolf 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. doi: 10.1002/cnm.2985.

Alessio Quaglino, Simone Pezzuto, and Rolf Krause. High-dimensional and higher-order multifidelity Monte Carlo estimators. *Journal of Computational Physics*, 388:300–315, 2019. doi: 10.1016/j.jcp.2019.03.026.

Carl E. Rasmussen and Christopher K. I. Williams. *Gaussian processes for machine learning*. MIT press Cambridge, 2006.

Stanley Rush and Hugh Larsen. A practical algorithm for solving dynamic membrane equations. *IEEE Transactions on Biomedical Engineering*, BME-25(4): 389–392, 1978. doi: 10.1109/TBME.1978.326270.

Francisco Sahli Costabal, Paris Perdikaris, Ellen Kuhl, and Daniel E. Hurtado. Multi-fidelity classification using Gaussian processes: accelerating the prediction of large-scale computational models. *Computer Methods in Applied*

Mechanics and Engineering, 357:112602, 2019. doi: 10.1016/j.cma.2019.112602.

Francisco Sahli Costabal, Kinya Seo, Euan Ashley, and Ellen Kuhl. Classifying drugs by their arrhythmogenic risk using machine learning. *Biophysical Journal*, 118(5):1165–1176, 2020. doi: 10.1016/j.bpj.2020.01.012.

Francisco Sahli Costabal, Tomás Banduc, Lia Gander, and Simone Pezzuto. The fibrotic kernel signature: simulation-free prediction of atrial fibrillation. In *International Conference on Functional Imaging and Modeling of the Heart*, pages 87–96, 2023. doi: 10.1007/978-3-031-35302-4_9.

Hasan I. Saleheen and Kwong T. Ng. New finite difference formulations for general inhomogeneous anisotropic bioelectric problems. *IEEE Transactions on Biomedical Engineering*, 44(9):800–809, 1997. doi: 10.1109/10.623049.

Abraham Savitzky and Marcel J.E. Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry*, 36(8):1627–1639, 1964. doi: 10.1021/ac60214a047.

Ulrich Schotten, Sander Verheule, Paulus Kirchhof, and Andreas Goette. Pathophysiological mechanisms of atrial fibrillation: a translational appraisal. *Physiological Reviews*, 91(1):265–325, 2011. doi: 10.1152/physrev.00031.2009.

Maxime Sermesant, Ender Konukoglu, Hervé Delingette, Yves Coudière, Phani Chinchapatnam, Kawal S. Rhode, Reza Razavi, and Nicholas Ayache. An anisotropic multi-front fast marching method for real-time simulation of cardiac electrophysiology. In *International Conference on Functional Imaging and Modeling of the Heart*, pages 160–169, 2007. doi: 10.1007/978-3-540-72907-5_17.

James A. Sethian and Alexander Vladimirsky. Fast methods for the eikonal and related Hamilton-Jacobi equations on unstructured meshes. *Proceedings of the National Academy of Sciences*, 97(11):5699–5703, 2000. doi: 10.1073/pnas.090060097.

James A. Sethian and Alexander Vladimirsky. Ordered upwind methods for static Hamilton-Jacobi equations: theory and algorithms. *SIAM Journal on Numerical Analysis*, 41(1):325–363, 2003. doi: 10.1137/S0036142901392742.

Johannes Siebermair, Eugene G. Kholmovski, and Nassir Marrouche. Assessment of left atrial fibrosis by late gadolinium enhancement magnetic resonance imaging: methodology and clinical implications. *JACC: Clinical Electrophysiology*, 3(8):791–802, 2017. doi: 10.1016/j.jacep.2017.07.004.

Connie W. Tsao, Aaron W. Aday, Zaid I. Almarzooq, Alvaro Alonso, Andrea Z. Beaton, Marcio S. Bittencourt, Amelia K. Boehme, Alfred E. Buxton, April P. Carson, Yvonne Commodore-Mensah, et al. Heart disease and stroke statistics - 2022 update: a report from the American Heart Association. *Circulation*, 145 (8):e153–e639, 2022. doi: 10.1161/CIR.0000000000001052.

Atul Verma, Chen-yang Jiang, Timothy R. Betts, Jian Chen, Isabel Deisenhofer, Roberto Mantovan, Laurent Macle, Carlos A. Morillo, Wilhelm Haverkamp, Rukshen Weerasooriya, et al. Approaches to catheter ablation for persistent atrial fibrillation. *New England Journal of Medicine*, 372(19):1812–1822, 2015. doi: 10.1056/NEJMoa1408288.

Mathias Wilhelms, Hanne Hettmann, Mary M. Maleckar, Jussi T. Koivumäki, Olaf Dössel, and Gunnar Seemann. Benchmarking electrophysiological models of human atrial myocytes. *Frontiers in physiology*, 3(487), 2013. doi: 10.3389/fphys.2012.00487.

Fagen Xie, Zhilin Qu, Alan Garfinkel, and James N. Weiss. Electrical refractory period restitution and spiral wave reentry in simulated cardiac tissue. *American Journal of Physiology - Heart and Circulatory Physiology*, 283(1):H448–H460, 2002. doi: 10.1152/ajpheart.00898.2001.

Sohail Zahid, Hubert Cochet, Patrick M. Boyle, Erica L. Schwarz, Kaitlyn N. Whyte, Edward J. Vigmond, Rémi Dubois, Mélèze Hocini, Michel Haïssaguerre, Pierre Jaïs, et al. Patient-derived models link re-entrant driver localization in atrial fibrillation to fibrosis spatial pattern. *Cardiovascular Research*, 110(3):443–454, 2016. doi: 10.1093/cvr/cvw073.