



UNIVERSIDAD DE CHILE
FACULTAD DE CIENCIAS FÍSICAS Y MATEMÁTICAS
DEPARTAMENTO DE INGENIERÍA MATEMÁTICA

**THE FIBROTIC KERNEL SIGNATURE: SIMULATION-FREE PREDICTION
OF ATRIAL FIBRILLATION INDUCIBILITY**

TESIS PARA OPTAR AL GRADO DE MAGÍSTER EN CIENCIAS DE LA INGENIERÍA,
MENCIÓN MATEMÁTICAS APLICADAS

MEMORIA PARA OPTAR AL TÍTULO DE INGENIERO CIVIL MATEMÁTICO

TOMÁS BANDUC MORENO

PROFESOR GUÍA:
FRANCISCO SAHLI COSTABAL

PROFESOR CO-GUÍA:
AXEL OSSES ALVARADO

MIEMBROS DE LA COMISIÓN:
HANNE VAN DEN BOSCH
FRANCISCO FÖRSTER BURÓN

Este trabajo ha sido parcialmente financiado por:
FONDECYT INICIACIÓN 11220816; ANID ERAPERMED 134
CMM ANID BASAL FB210005

SANTIAGO DE CHILE

2024

RESUMEN DE LA TESIS PARA OPTAR AL GRADO DE
MAGÍSTER EN CIENCIAS DE LA INGENIERÍA,
MENCION MATEMÁTICAS APLICADAS
MEMORIA PARA OPTAR AL TÍTULO DE
INGENIERO CIVIL MATEMÁTICO
POR: TOMÁS BANDUC MORENO
FECHA: 2024
PROF. GUÍA: FRANCISCO SAHLI COSTABAL
PROF. CO-GUÍA: AXEL OSSES ALVARADO

LA FIRMA DEL NÚCLEO DE CALOR INFORMADA CON FIBROSIS: PREDICCIÓN DE INDUCIBILIDAD DE FIBRILACIÓN AURICULAR LIBRE DE SIMULACIONES

La fibrilación auricular (AF) es el tipo de arritmia más común en seres humanos, llegando a ser prevalente en la población a escala epidémica. Los modelos computacionales pueden contribuir a la mejora de las tasas de éxito para tratamientos terapéuticos de la AF, como la ablación cardíaca, donde partes del tejido son cauterizadas con el fin de bloquear puntos de iniciación de AF. Sin embargo, evaluar la eficacia de distintas estrategias de ablación requiere realizar costosas simulaciones donde se estimulan eléctricamente diferentes puntos y se evalúa si un episodio de fibrilación se indujo o no, obstaculizando la implementación clínica de estos modelos. En el presente trabajo de tesis se propone un método de clasificación eficiente capaz de predecir confiablemente la inducibilidad de AF en modelos cardíacos personalizados sin la necesidad de correr simulaciones adicionales. La metodología no requiere de re-entrenamiento cuando se introducen cambios en la anatomía cardíaca, distribución de fibrosis o líneas de ablación. Para lograr esto, se desarrolla un conjunto de atributos dados por una variante anisotrópica de la firma del núcleo de calor (*heat kernel signature* - hks) que incorpora la información de fibrosis y la orientación de fibras en el corazón: La firma del núcleo de calor con información fibrótica (*fibrotic kernel signature* - fks). La fks es más rápida de calcular que una simulación puntual de AF y puede predecir mapas de inducibilidad en la totalidad del dominio auricular cuando se emplea en conjunto con algoritmos de clasificación. Para evaluar la relación entre la fks y la inducibilidad de arritmia se realizaron dos estudios: uno para probar la capacidad de la firma en capturar el efecto de ablación sobre un modelo auricular simple, y otro para demostrar la generalidad del método propuesto en varias anatomías cardíacas sujetas a múltiples distribuciones de fibrosis. Los estimadores basados en la fks predicen la inducibilidad de AF con alta precisión y aproximan el porcentaje inducible de mejor manera que métodos alternativos. El flujo de trabajo de la fks puede acelerar significativamente los cálculos necesarios para testear eventos de arritmia, lo que es crucial para optimizar terapias de AF dentro de las restricciones del contexto clínico.

THESIS ABSTRACT FOR THE
MASTER'S DEGREE IN ENGINEERING SCIENCES,
MINOR APPLIED MATHEMATICS
THESIS FOR THE DEGREE OF
MATHEMATICAL ENGINEER
BY: TOMÁS BANDUC MORENO
DATE: 2024
GUIDE PROF.: FRANCISCO SAHLI COSTABAL
CO-GUIDE PROF.: AXEL OSSES ALVARADO

THE FIBROTIC KERNEL SIGNATURE: SIMULATION-FREE PREDICTION OF ATRIAL FIBRILLATION INDUCIBILITY

Atrial fibrillation (AF) is the most common type of cardiac arrhythmia, being prevalent at epidemic scale. Computational models can help improve success rates for therapeutic treatments of AF, such as ablation procedures, where regions of tissue are destroyed in a controlled fashion to block AF initiation points. However, evaluating the efficacy of different ablation strategies requires performing costly simulations by pacing at different points and checking whether AF has been induced, hindering the clinical application of these models. For this thesis work, an efficient classification method that can reliably predict AF inducibility in patient-specific cardiac models without running additional simulations is proposed. The methodology does not require re-training when changing atrial anatomy, fibrosis distribution or ablation lines. To achieve this, a set of features given by an anisotropic variant of the heat kernel signature (hks) that incorporates fibrosis information and fiber orientation is developed: the fibrotic kernel signature (fks). The fks is faster to compute than a single AF simulation, and it can predict AF inducibility in the entire atrial domain when paired with machine learning classifiers. To assess the relationship between the FKS and AF inducibility two studies were conducted: one study to evaluate the capability of the signature to capture the effect of ablation lines in a single atrial model, and another study to demonstrate the generality of the method across several anatomies subject to various fibrosis scenarios. fks-based estimators predicted AF inducibility with high accuracy and could approximate overall inducibility better than alternative methods. The fks pipeline can significantly speed-up the calculations required for AF testing, which is crucial to optimize therapies for AF within the time constraints of the clinical setting.

*A mis hermanos,
Matías y Benjamín.*

Agradecimientos

Quiero agradecer a mi familia por el apoyo y compañía durante todos estos años de estudio. Gracias a Carmen, mi madre, por su amor incondicional, su receptividad y preocupación por asegurarme un entorno de formación ideal, que nunca me faltara nada y que siempre al llegar a casa tuviera un plato servido para disfrutar; a Fernando, mi padre, por su sensibilidad, por enseñarme a amar sin miedo, a esforzarme, pero cuidando mi propio bienestar, y por impulsar sin restricciones mis sueños. Gracias a Daniel, pareja de mi padre, por compartirme sus experiencias de vida y sustentar mis ideas y principios; y a Francisco, marido de mi madre, por enseñarme a ser paciente y que siempre hay lugar para seguir aprendiendo y para mejorar.

Agradezco a mis hermanos por ser mis cómplices, por escucharme y comprenderme siempre que lo requerí. Gracias a Matías, mi hermano mellizo, por recordarme que la empatía y la generosidad son fundamentales para tener relaciones significativas, por hacerme ver que la imaginación nos mantiene apasionados en lo que cotidianamente hacemos y reafirmar el valor que tengo como profesional; y a mi hermano mayor, Benjamín, por toda la energía que pone en mantenernos unidos, por mostrarme la importancia de las acciones cuando se hacen con emoción y por darme la confianza de expresarme y ser una persona con iniciativa.

Gracias al profesor Francisco Sahli por tenerme bajo su tutela, por darme la oportunidad de participar en sus proyectos y compartirme su fascinación por la investigación; al profesor Axel Osses por todas las instancias de aprendizaje que me ha facilitado, por involucrarme en el ambiente académico y por impulsarme a tener un espíritu creativo y lógico. Agradezco a ambos por sus consejos para descubrir mis intereses vocacionales, por incentivar me a apuntar alto y por darme su apoyo para continuar con mis estudios. Gracias también a los académicos, académicas, funcionarios y funcionarias del departamento, a los equipos de coordinación y a la profesora Natacha Astromujoff por todas las experiencias y aprendizaje invaluable que me entregaron durante esta etapa.

Agradezco a Sebastián, mi pareja y mejor amigo, por ser mi razón de comenzar el día, por alentarme a ser crítico, a desarrollar mis opiniones y mis ideas, por respaldarme en mis proyectos y darme la motivación de completar este proceso. Gracias a Yeniffer, mi amiga y confidente, por mantenerme a flote durante estos años de estudio, por siempre estar atenta a nuestra amistad y por hacerme sentir acompañado en los momentos que más lo necesité.

Finalmente, quiero dar gracias a todas las personas que se vieron involucradas en mi vida durante el trayecto de la carrera, por cada experiencia que compartimos y por llevarme a ser quien soy el día de hoy. Gracias a *les autres c.s.*, mi familia durante especialidad, y a Kathryn Cuneo, quien me ha dado todos los ánimos desde la distancia. Espero seguir compartiendo y formando recuerdos con cada uno de ustedes. Los guardo en mi corazón con mucho cariño.

Table of Contents

Introduction	1
1. Preliminaries	5
1.1. Notation	5
1.2. Sobolev Spaces	6
1.3. Heat Semigroups	7
1.4. Kernel Spectral Decomposition	8
1.5. Heat Kernels of Second Order Elliptic Operators	9
2. Related Work	11
2.1. Surrogate Models of Atrial Fibrillation	11
2.2. ML-based Geometrical Features Approach	13
2.3. Multi-Fidelity Gaussian Process Classification	15
3. General Methodology	18
3.1. Atrial Fibrillation Modeling	18
3.1.1. Microscopic Bidomain Model	18
3.1.2. Macroscopic Bidomain System	20
3.1.3. Macroscopic Monodomain System	21
3.2. Pacing Protocols	23
3.2.1. Phase Singularity Distribution	24
3.2.2. Rapid Pacing	25
3.2.3. Pacing at the End of the Effective Refractory Period	26
3.3. The Fibrotic Kernel Signature	27
3.3.1. The Heat Kernel Signature	27
3.3.2. Fibrosis Information Inclusion	29
3.3.3. Signature Computation with Finite Elements	30
3.3.4. Method Workflow for AF Prediction	31
4. Study 1	36
4.1. Introduction	36
4.2. Methods	36
4.2.1. Atrial Modeling	36
4.2.2. Fibrotic Kernel Signature in the Atria	38
4.2.3. Prediction of Atrial Fibrillation	38
4.3. Results	39
4.4. Discussion and Conclusions	42

5. Study 2	44
5.1. Introduction	44
5.2. Methods	45
5.2.1. Atrial Modeling	45
5.2.2. Data	46
5.2.3. Fibrotic Kernel Signature in the Atria	49
5.2.4. Prediction of Atrial Fibrillation	50
5.3. Results	51
5.3.1. Atrial Fibrillation Simulations	51
5.3.2. Prediction of Inducibility	52
5.4. Discussion and Conclusions	56
6. General Discussion	58
7. Conclusion	61
Bibliography	63
Annex	72
.A. Performance with Standard Metrics	72

List of Tables

- 5.1. Summary of fiber (f) and transverse (t) conductivity and the resulting conduction velocities measured for planar wavefront propagation on a computational grid of average spatial discretization of $\Delta\mathbf{x} \approx 425\mu\text{m}$ for normal atrial tissue, Bachmann’s bundle, inferior/superior Vena Cava and electrically remodeled tissue (Banduc et al., 2024). Table reprinted from Banduc et al. (2024). 45
- 5.2. Summary statistics for all the simulations performed. Active points refers to the pacing points that were not passive tissue. Re-entries correspond to the sum of AF and flutter cases (Banduc et al., 2024). Table adapted from Banduc et al. (2024). 47
- 5.3. Mesh information. Table reprinted from Banduc et al. (2024). 48
- 5.4. Computation time for fks eigenproblem on each model and conductivity option (Banduc et al., 2024). 50

List of Figures

2.1.	Comparative of simulated rotor spiral waves using a CA and a biophysical model in a slab of tissue. Simulations are shown for five consecutive times, 120[ms] after the rotor stabilized. The CA simulation is color-coded by time in milliseconds (ms) from the activation (red tones) to the resting state (blue tones). The biophysical simulation shows the equivalent phenomenon, where colors represent transmembrane potential in millivolts (mV) (Serra et al., 2022). Figure adapted from Serra et al. (2022).	12
2.2.	Simulation time instants (ms) for CA and biophysical model. A ventricular tachycardia (VT) episode resulted from applying an S1–S2 protocol at the location labeled as ‘Stm’. First row shows the endocardial layer (light gray) and the activity in the border zone (BZ) of scar region in the CA model (the core zone, ‘CZ’, and healthy tissue are transparent), while second row shows the equivalent in the biophysical model with all tissue types visible (CZ in dark gray color). Colors in the CA simulation represent time (ms), and in the biophysical simulation represent potential (mV). A conduction block can be observed a few milliseconds after the S2 (white cross). VT is subsequently sustained across the slow conduction channel (SCC). White arrows show the direction of the wavefront (Serra et al., 2022). Figure adapted from Serra et al. (2022).	13
2.3.	Feature extraction for a potentially arrhythmogenic non-conductive tissue region, including the area, perimeter, surrounding fibrosis, and proximity to nearest non-conductive tissue (Bifulco et al., 2023). Figure adapted from Bifulco et al. (2023).	14
2.4.	A, dependence plot for perimeter. B, dependence plot for area. C, dependence plot of surrounding residual fibrosis percentage. D, dependence plot for proximity index. SHAP indicates Shapley additive explanations (Bifulco et al., 2023). Figure adapted from Bifulco et al. (2023).	15
2.5.	A, average balanced accuracy change for increasing length scale when the classifiers are trained with 100 samples. The dashed vertical line represents the average geodesic distance between training points of the fixed design. B, balanced accuracy comparison for the nearest neighbor, single-fidelity, and multi-fidelity classifiers, for all nine model scenarios and with a fixed budget of 40 high-fidelity simulations (Gander et al., 2022). Figure adapted from Gander et al. (2022).	17
3.1.	Illustration of cellular connectivity from slice of cardiac tissue. The structure presented repeats periodically to form extended layers of heart muscle.	19
3.2.	Illustration of muscle composition in ventricular tissue with orthonormal vectors aligned with laminae and fiber direction. Figure based on illustration from Kocica et al. (2006).	22

3.3.	Phase singularity below left pulmonary veins on atrial model with counterclockwise spiral rotation from electrical potential recordings (mV).	24
3.4.	Illustration of converted EGM signal to phase maps with Sawtooth and Hilbert Transform methods. Figure adapted from Lootens et al. (2024).	24
3.5.	A, electrogram of single cardiomyocyte under RP stimulation protocol. B, electrogram of cardiomyocyte with PEERP pacing strategy. The colored vertical lines correspond to the instants of stimulation. Note that the coupling interval from the PEERP protocol depends on the effective refractory period.	25
3.6.	ERP representation for myocyte action potential and binary search strategy for end of ERP estimation. The colored bands indicate the pulse application and their width the state of iteration. The color green indicates that the delivered stimulus successfully spread in neighboring cells, and red otherwise.	27
3.7.	Example of inducibility criterion $\mathbf{1}_{AF}$ for PEERP protocol and monodomain model.	32
3.8.	fks method pipeline. In the offline phase (top), classifiers are trained with the fks features to predict monodomain inducibility tags. Then, in the online phase (bottom), the fks is computed for a new geometry and passed through the classifier to generate an inducibility map. Figure adapted from Banduc et al. (2024).	35
4.1.	Atrial models considered in first study. Different shades of gray are used to mark healthy tissue, fibrosis and ablation scars.	37
4.2.	Instance of leave-one-out by case technique. A sample of size 100 is taken from each configuration (adding up to 800 samples), excluding the case with s_1 fibrosis distribution and PVI ablation, which is set aside for testing.	39
4.3.	fks over time for 20 randomly selected points for 9 different cases: moderate fibrosis (m), severe fibrosis - case 1 (s_1) and severe fibrosis - case 2 (s_2). The red curves represent points where AF was induced and the blue curves are points where AF was not induced. Figure reprinted from Sahli Costabal, Banduc, Gander, and Pezzuto (2023).	40
4.4.	Overall inducibility prediction for leave-one-out by case cross-validation with RF. Figure reprinted from Sahli Costabal et al. (2023).	40
4.5.	Three time steps for fks map in moderate fibrosis case under three ablation patterns: no ablation (NA), pulmonary vein isolation (PVI) and PVI with box enclosure (PVI+BOX).	41
4.6.	Data efficiency for balanced accuracy across all cases following a leave-on-out by case strategy. The solid lines show the performance of k-nearest neighbors, random forest and gradient boosting as the training data increases. The dashed line corresponds to the naïve approach. Figure reprinted from Sahli Costabal et al. (2023).	42
5.1.	Train geometries used in second study, the first was comprised of 3 fibrosis configurations with 3 severity patterns each, while the second and third had 2 fibrosis patterns with 3 severity scenarios each (Banduc et al., 2024).	48
5.2.	Test geometries used in second study. Each case corresponded to a different severity scenario (Banduc et al., 2024).	48
5.3.	fks extraction from aniso- (top) and isotropic (bottom) options. The first case considers the monodomain conductivity tensor to account for fiber direction, and the second an heterogeneous conductivity function modeling conduction variations due to fibrosis (Banduc et al., 2024).	49

5.4.	Convergence of inducibility. Changes in computed inducibility are shown as more simulations are considered. Cases are grouped by fibrosis severity. Each of the three leftmost panels show 8 different cases, while the rightmost sub-plot shows the evolution of percent error when considering the inducibility computed with all simulations as ground truth (Banduc et al., 2024). Figure reprinted from Banduc et al. (2024).	52
5.5.	Isotropic fks map evolution across the three test cases for the 1 st , the 20 th and the 50 th time step computed. Figure reprinted from Banduc et al. (2024). . . .	53
5.6.	Anisotropic fks map evolution across the three test cases for the 1 st , the 20 th and the 50 th time step computed. Figure reprinted from Banduc et al. (2024).	53
5.7.	fks samples from test anatomies for isotropic (top) and anisotropic (bottom) options of the signature. For each case, 20 points are randomly selected from the pacing sites and their corresponding fks evolution in time is shown. A curve is colored red if the pacing point induced AF, and blue if not. Figure reprinted from Banduc et al. (2024).	54
5.8.	Data efficiency for F1 score across test cases for a random forest (blue) and a multi-layer perceptron classifier (mlp, red). The solid line represents the median F1 score, while the shaded area represents the range between minimum and maximum F1 score over 10 runs with different random seeds (Banduc et al., 2024). Figure reprinted from Banduc et al. (2024).	54
5.9.	Overall inducibility predicted by the fks-trained RF and the naïve classifier. The error bars for the fks classifier represent the standard deviation over 10 runs with different random seeds. For reference, the inducibility computed using only 40 simulations is also included. The ground truth inducibility was computed using all available simulations (Banduc et al., 2024). Figure reprinted from Banduc et al. (2024).	55
5.10.	Inducibility maps and pacing locations for test cases. In the top row, fibrotic patterns are shown in gray/black, inducible sites are colored in red and non-inducible points are colored in blue. In the bottom row, predicted inducibility maps from the RF classifier are shown and colored under the same criterion used for inducible points (Banduc et al., 2024). Figure reprinted from Banduc et al. (2024).	55
.A.1.	Data efficiency for accuracy, sensitivity and specificity across test cases for a random forest (blue) and a multi-layer perceptron classifier (mlp, red). The solid lines represent the median values, while the shaded area represents the range between minimum and maximum scores over 10 runs with different random seeds (Banduc et al., 2024).	72

Introduction

Atrial fibrillation (AF) is the most common type of cardiac arrhythmia in humans (Tsao et al., 2023; as cited in Banduc et al., 2024). AF is described as an aberrant and self-sustained electrical activation in the atrial region of the heart (Schotten et al., 2011; as cited in Banduc et al., 2024). Research for its pathophysiological mechanisms and treatment is motivated by its incidence at epidemic scale (Kornej et al., 2020; as cited in Banduc et al., 2024) and co-occurrence alongside clinical conditions that contribute to its onset and perpetuation, such as age, hypertension, diabetes, thyroid disease, cardiomyopathy, heart failure and valvular disease (Schotten et al., 2011). It is expected that AF prevalence may increase in the future as longevity grows within the world population (Trohman et al., 2023).

About 25% of ischemic strokes are estimated to be caused by AF (Marini et al., 2005; as cited in Banduc et al., 2024), as thrombosis often occurs during arrhythmia events due to blood stasis. This disease not only increases the chances of ischemia but worsens prognosis in patients with preexisting heart conditions (Kornej et al., 2020; as cited in Banduc et al., 2024). Overall, independently of previous cardiac disease, death rates double in patients with AF (Schotten et al., 2011; as cited in Banduc et al., 2024).

AF is progressive. Initially, episode duration is at most 7 days (Joglar et al., 2024), which occur between extended periods of sinus rhythm (paroxysmal AF). The intermediate state is characterized by frequent and prolonged episodes of fibrillation, lasting longer than a week to terminate (persistent AF). Eventually, the disease becomes chronic (permanent AF). This progression is mainly explained by its self-generative nature: AF causes gradual processes of remodeling in the functionality of ionic channels of cardiac myocytes and fibrosis deposition in the atria (Schotten et al., 2011; as cited in Banduc et al., 2024), while fibrosis increases abnormalities in electrical conductivity throughout cardiac tissue and perpetuates AF (Pezzuto et al., 2018; as cited in Banduc et al., 2024).

Cardiac fibrosis is one of the main drivers of AF (Allessie et al., 2002; as cited in Banduc et al., 2024). It is defined as a pathological scarring event caused by collagen accumulation and fibroblast activation in the cardiac muscle (Hinderer & Schenke-Layland, 2019; as cited in Banduc et al., 2024). Fibrosis distribution in the atria is patient-specific and unfolds with AF in a vicious systemic loop. During fibrillation events, cellular hypertrophy leads to scar formation and atrial stretching. This increases tissue heterogeneity and anisotropy in conduction in the cardiac muscle, creating pathways for electrical re-entry and altering the extracellular matrix, which enhances AF inducibility and complexity (Gharaviri et al., 2020; Sahli Costabal et al., 2023; Schotten et al., 2011). In this sense, the more fibrosis is present, the more AF events are likely to occur, which triggers more fibrosis deposition.

There is a demonstrated positive correlation between fibrotic tissue distribution and point localization for AF initiation. McDowell et al. (2015) showed that unique fibrosis patterns determine persistent reentry events during the propagation of voltage that are independent of the stimulation points (known as *pacing points*) selected for testing AF inducibility and that mechanisms for atrial sites to initiate sustained episodes of AF given a standard stimulation protocol act in a distance-to-fibrosis manner. Likewise, recent evidence suggests that surrounding residual fibrosis from non-conductive cauterized tissue may generate a strong arrhythmogenic substrate (Bifulco et al., 2023). These findings suggest that fibrosis allocation in the human atria is informative as to determine whether stimuli in certain pacing locations can potentially induce AF (Gander et al., 2022; as cited in Banduc et al., 2024).

Therapeutic approaches are widely used for the treatment of AF as they require less resource management and cause less fibrillation recurrence than pharmacological cardioversion methods (Blomström-Lundqvist et al., 2019; Saad-Omer et al., 2020; Yi et al., 2019). Within this framework, catheter ablation is a technique used in patients with different states of progression of the disease. Ablation is a non-invasive surgical procedure in which sections of cardiac tissue are destroyed with radio-frequencies in a controlled fashion, to block AF initiation points (Joglar et al., 2024; as cited in Banduc et al., 2024). In patients with paroxysmal AF, this treatment induces optimal results when performed around pulmonary veins (PVs) due to the electrical blockage of ectopic pacemaker cells in the PV area (Haissaguerre et al., 1998; Chen et al., 1999; as cited in Banduc et al., 2024). However, for advanced stages of arrhythmia, the complexity in fibrosis patterns and cardiac structural remodeling lead to a shift in electrical abnormalities (Boyle et al., 2019), causing PVI to underperform in the reduction of inducible surface (Verma et al., 2015; Kawai et al., 2019; Gharaviri et al., 2021; as cited in Banduc et al., 2024), prompting the need to design ablation strategies fitted to each patient’s cardiac geometry and heart electrical properties.

As shown in recent retrospective and prospective studies, therapeutic approaches to AF, such as catheter ablation, tend to lead to better outcomes when tailored to patient fibrosis distribution and anatomy (Boyle et al., 2019; Loewe et al., 2019; McDowell et al., 2015; as cited in Banduc et al., 2024). These findings imply that it is possible to use patient-specific *in silico* models to estimate AF inducibility for various ablation scenarios, and then select the best treatment for the patient.

Patient-specific models of AF are typically based on the monodomain system. The model consists of a single reaction-diffusion equation coupled with a possibly large system of ordinary differential equations accounting for ionic dynamics through the cellular membrane and within each cardiomyocyte (Colli Franzone et al., 2014; as cited in Banduc et al., 2024). Despite considering intracellular and extracellular spaces superimposed and homogenized (Rosilho de Souza et al., 2024; as cited in Banduc et al., 2024), the stiffness in the temporal dynamic caused by the rapid upstroke of action potentials raises the requirement for high spatial resolution of the monodomain model to obtain accurate approximations of the conduction velocity (Pezzuto et al., 2016; Niederer, Kerfoot, et al., 2011; as cited in Banduc et al., 2024), being usually below 500[μm] for normal conduction and much higher for slow conduction velocities (Bishop & Plank, 2024; as cited in Banduc et al., 2024).

Pacing at a fixed number of well-distributed atrial locations is a standard method to assess

AF inducibility (Boyle et al., 2019; as cited in Banduc et al., 2024). However, each protocol needs to be repeated when ablation lines are added, fibrosis distribution is changed, or a new cardiac anatomy is considered. Monodomain solvers designed for running on many-core CPU or GPU systems have sped-up inducibility simulations (Niederer, Mitchell, et al., 2011; Neic et al., 2012; Kaboudian et al., 2019), but testing if a stimulation protocol in the atria can degrade into a sustained episode of arrhythmia remains costly both in terms of time and computational resources, hindering the clinical feasibility of this approach. The need to decrease costs to assure availability of personalized therapy has led some authors to propose adaptive pacing protocols (Azzolin et al., 2021), surrogate models of AF (Serra et al., 2022) and analysis of geometrical features (Bifulco et al., 2023) to pinpoint regions that could induce fibrillation. Methods integrating image-based modeling and deep learning have also been reported to successfully determine personalized ablation strategies without the burden of simulation, showing great potential to clinically translate as efficient tools for non-invasive treatment (Muffoletto et al., 2021; as cited in Banduc et al., 2024).

To control the computational burden of AF simulations with the monodomain equations, testing AF inducibility can be recast as a classification problem (Gander et al., 2022; as cited in Banduc et al., 2024), where a set of customized features encoding atrial geometry, fibrosis patterns and electrical parameters are used through a machine learning classifier trained to determine whether pacing from a given location produces a stable AF event. This thesis aims to offer a solution to the limitations of conventional AF simulations by proposing a general simulation-free classification method for AF inducibility: *The fibrotic kernel signature* (fks). The fks retains full anatomical properties and the structural complexity of a conventional simulation, including fiber direction and scar distribution. Instead of having to re-run the monodomain solver when fibrosis is modified, ablation strategies are adjusted or the atrial geometry is changed, the electrical and architectural information is stored in this set of features, which is based on the *heat kernel signature* (hks) (Sun et al., 2009). The hks is a time series that describes multi-scale geometrical and topological properties of compact domains via homogeneous diffusion within them. The hks is extensively used in shape analysis, particularly shape matching. Formally, the hks is defined as a point restriction of the heat kernel function, whose values can be efficiently computed using the Laplace operator spectrum. Intuitively, it can be explained as instantaneously applying an infinitely intense heat source at a point for an infinitely small window of time, and then measuring the temperature evolution at the same point. At first, the temperature changes will heavily depend locally on the geometry around the point, and ultimately will be influenced by the overall shape.

In the context of cardiac electrophysiology, the fks complements the previous method by considering heterogeneous and anisotropic processes of electrical propagation, hence including information not only from the atrial geometry but also the specific fibrotic distribution and fiber orientation. The fks can be interpreted as measuring the temporal evolution of the transmembrane voltage after applying a Dirac-delta potential at a point in the atria. The fks is paired with classic machine learning classifiers to predict if AF can be induced when pacing from a particular point.

In this work, a large dataset was created comprising different anatomies, a diversity of fibrotic levels, ablation strategies and multiple pacing locations to assess the relationship between the fks and AF inducibility. In this way, once a classifier was trained, only the

fks computations were required for studying inducibility in new cases, which was done at a fraction of the cost of a standard AF simulation. The mathematical basis, work related to AF inducibility characterization, computational methods, results and discussion are reviewed in this thesis with the following structure:

- *Chapter 1 - Preliminaries:* This chapter shows the theoretical results that evidence the functionality of the fks method. Sobolev spaces are defined, results from semigroup theory regarding parabolic equations and heat kernels are shown, and a spectral decomposition theorem is presented.
- *Chapter 2 - Related Work:* In this chapter, recent advancements in AF inducibility characterization are summarized. The reviewed bibliography consists of fast biophysical cardiac emulators (Serra et al., 2022), the use of a multi-fidelity Gaussian process for accelerated inducibility classification (Gander et al., 2022) and AF substrate description with geometric properties of non-conductive atrial tissue through an ML-based approach (Bifulco et al., 2023).
- *Chapter 3 - General Methodology:* This chapter encompasses the mathematical model used for the monodomain simulations, fks computations and dataset creation. Here, the hks features are introduced, its main properties are described and the concept of diffusion for point description is extended to a cardiac electrophysiology scenario, obtaining the fks time series. In the chapters that follow, the proposed method is validated with atrial geometries from real patients in two application studies: overall inducibility reduction from cardiac ablation and AF inducibility estimation for different fibrosis levels.
- *Chapter 4 - Study 1:* This study exhibits preliminary results for the capability of the fks method to efficiently predict the effect of distinct cardiac ablation strategies over AF inducibility for a single atrial model in various fibrosis scenarios, with common ML classifiers.
- *Chapter 5 - Study 2:* In this part, a larger dataset is used to test the generality of the introduced method, where the efficiency and accuracy of the fks features are evaluated for the task of identifying AF initiation points and estimate overall inducibility in unobserved cardiac anatomies with classifiers trained by several atrial models subject to diverse fibrosis levels.
- *Chapter 6 - General Discussion:* This chapter globally discusses the results of both studies regarding the reduction of time complexity for inducibility prediction with the fks, the behavior of the signature along varying atrial architectures, scar severity and fibrosis patterns, its proficiency to detect AF substrate for multiple designs of conductivity across the atria, and limitations to sufficiently characterize cardiac pro-arrhythmic structures.
- *Chapter 7 - Conclusion:* This section details the main features of the fks, highlights the competence of the signature to reliably and efficiently determine if pacing from a given location could degenerate into fibrillation, and concludes about the use of the proposed method in therapy planning for the clinical setting.

Chapter 1

Preliminaries

1.1. Notation

- $\mathcal{L}(X, Y)$: Space of linear bounded operators from X to Y .
- $\mathcal{C}(X, Y)$: Space of continuous functions from X to Y .
- $\mathcal{C}^k(X, Y)$: Space of k -continuously differentiable functions from X to Y .
- $\mathcal{C}^\infty(X, Y)$: Space of infinitely differentiable functions from X to Y .
- $\mathcal{D}(X, Y)$: Space of compactly supported functions in $\mathcal{C}^\infty(X, Y)$.
- $\mathcal{M}(X, Y)$: Space of measurable functions from X to Y .
- $L^p(X, Y)$: Space of p -integrable functions from X to Y .
- $L^\infty(X, Y)$: Space of essentially bounded functions from X to Y .
- $H^k(X, Y)$: Space of $L^2(X, Y)$ functions with 2-integrable (weak) derivatives of order k .
- $H_0^k(X, Y)$: Closure of $\mathcal{D}(X, Y)$ with respect to the norm $\|\cdot\|_{H^k}$.
- X^* : Dual space of X .
- \overline{X} : Closure of X .
- $D(L)$: Domain of operator L .
- $(\cdot, \cdot)_X$: Inner product of X .
- $\|\cdot\|_X$: Norm of X .
- $\|\cdot\|_p$: L_p -norm.
- \mathbb{P}_k : Space of polynomials of degree k .
- Ω_H : Cardiac domain, assumed bounded with Lipschitz boundary.
- Ω_h : Polyhedral approximation of cardiac domain Ω_H .
- $\mathbf{1}_{AF}(\mathbf{p}, \mathcal{M})$: Inducibility indicator function for protocol \mathbf{p} and model \mathcal{M} .

1.2. Sobolev Spaces

Definition 1.1 (*L^p Space - Cazenave & Haraux, 1999*) Let (X, \mathcal{A}, μ) be a measure space, Y be a Banach space and $p \in [1, \infty)$. The space $L^p(X, Y)$ is defined as the equivalence class of p -integrable functions from X to Y given by the a.e. equality relation

$$L^p(X, Y) := \mathcal{L}^p / \sim; \quad \mathcal{L}^p := \left\{ u \in \mathcal{M}(X, Y) : \int \|u\|_Y^p d\mu < \infty \right\}; \quad u \sim v \Leftrightarrow u = v \text{ a.e.}$$

Additionally, the space of locally p -integrable functions is defined as

$$L^p_{loc}(X, Y) := \{u \in \mathcal{M}(X, Y) : u|_K \in L^p(K, Y), \forall K \subseteq X, K \text{ compact}\}.$$

Definition 1.2 (*L^∞ Space - Cazenave & Haraux, 1999*) Let (X, \mathcal{A}, μ) be a measure space and Y a Banach space. The space $L^\infty(X, Y)$ is defined as the equivalence class of essentially bounded functions from X to Y given by the a.e. equality relation

$$L^\infty(X, Y) := \mathcal{L}^\infty / \sim; \quad \mathcal{L}^\infty := \{u \in \mathcal{M}(X, Y) : \exists \alpha > 0 \text{ s.t. } \|u\|_Y \leq \alpha \text{ a.e.}\}; \quad u \sim v \Leftrightarrow u = v \text{ a.e.}$$

Proposition 1.1 (**Hytönen et al., 2016**) The space $L^p(X, Y)$ is a Banach space provided the norm

$$\|u\|_p := \begin{cases} \left(\int \|u\|_Y^p d\mu \right)^{1/p}, & \text{if } p \in [1, \infty); \\ \inf\{\alpha : \|u\|_Y \leq \alpha \text{ a.e.}\}, & \text{if } p = \infty. \end{cases}$$

Corollary 1.1 The space $L^2(X, \mathbb{R})$ is a Hilbert space, with the following inner product:

$$(u, v)_{L^2} := \int_X uv dx.$$

Definition 1.3 (**Distributional Derivative - Hytönen et al., 2016**) Let X be an open set of \mathbb{R}^d equipped with the Lebesgue measure and Y a Banach space. For any multi-index α , $v \in L^1_{loc}(X, Y)$ is said to be the distributional (or weak) derivative of order α of $u \in L^1_{loc}(X, Y)$, written $D^\alpha u = v$, provided

$$\int_X u \partial^\alpha \phi dx = (-1)^{|\alpha|} \int_X v \phi dx, \quad \forall \phi \in \mathcal{D}(X, \mathbb{R}).$$

Definition 1.4 (**Sobolev Space - Hytönen et al., 2016**) Let $X \subseteq \mathbb{R}^d$ be an open set and Y a Banach space. For $1 \leq p \leq \infty$, the Sobolev space of order k and exponent p is defined as

$$W^{k,p}(X, Y) := \{u \in L^p(X, Y) : D^\alpha u \text{ exists and } D^\alpha u \in L^p(X, Y), \forall \alpha \text{ multi-index with } |\alpha| \leq k\}.$$

When $p = 2$, $W^{k,p}(X, Y)$ is denoted by $H^k(X, Y)$.

Proposition 1.2 (**Hytönen et al., 2016**) The space $W^{k,p}(X, Y)$ is a Banach space endowed

with the norm

$$\|u\|_{W^{k,p}} := \left(\sum_{|\alpha| \leq k} \|D^\alpha u\|_p^p \right)^{1/p}.$$

Corollary 1.2 *The space $H^2(X, \mathbb{R})$ is a Hilbert space with inner product*

$$(u, v)_{H^2} := \sum_{|\alpha| \leq k} \int_X (D^\alpha u)(D^\alpha v) dx.$$

1.3. Heat Semigroups

Definition 1.5 (Contraction Semigroups - Cazenave & Haraux, 1999) *Let X be a Banach space. A semigroup is a one-parameter family $\{T(t)\}_{t \geq 0} \subseteq \mathcal{L}(X, X)$ such that*

- $T(0) = I$;
- $T(s+t) = T(s)T(t), \forall s, t \in (0, \infty)$;
- $T(\cdot)x \in \mathcal{C}([0, \infty), X), \forall x \in X$.

Additionally, $\{T(t)\}_{t \geq 0}$ is said to be a contraction semigroup if $\|T(t)\|_{\mathcal{L}} \leq 1, \forall t \geq 0$.

Definition 1.6 (Semigroup Generator - Cazenave & Haraux, 1999) *The generator of a semigroup $\{T(t)\}_{t \geq 0} \subseteq \mathcal{L}(X, X)$ is a linear operator L defined by*

$$D(L) = \left\{ x \in X : \frac{T(t)x - x}{t} \text{ has limit in } X \text{ as } t \rightarrow 0^+ \right\}, \quad Lx = \lim_{t \rightarrow 0^+} \frac{T(t)x - x}{t}.$$

Definition 1.7 (Maximally Dissipative Operator - Cazenave & Haraux, 1999) *Let X be a Banach space. A linear operator $L : D(L) \subseteq X \rightarrow X$ is dissipative if*

$$\|x - \lambda Lx\|_X \geq \|x\|_X, \quad \forall x \in D(L), \forall \lambda > 0.$$

Moreover, if a dissipative operator L is such that for every $\lambda > 0$ and every $f \in X$, there exists $x \in D(L)$ such that $x - \lambda Lx = f$, then L is said to be m -dissipative.

Proposition 1.3 (Cazenave & Haraux, 1999) *Let H be a Hilbert space and $L : D(L) \rightarrow H$ be a linear operator in H . Then L is dissipative if and only if $(Lu, u)_H \leq 0, \forall u \in D(L)$.*

Proposition 1.4 (Cazenave & Haraux, 1999) *Let H be a Hilbert space and $L : D(L) \rightarrow H$ be a linear operator in H . If L is self-adjoint and dissipative, then L is m -dissipative.*

Theorem 1.1 (Hille-Yosida-Phillips Theorem - Arendt, 2006) *Let $L : D(L) \rightarrow H$ be a self-adjoint and dissipative linear operator acting on a Hilbert space H . Then, L generates a contraction semigroup $\{T(t)\}_{t \geq 0} \subseteq \mathcal{L}(H, H)$ of self-adjoint operators. Moreover, if L has compact resolvent, then $T(t)$ is compact for all $t > 0$.*

Theorem 1.2 (Heat Semigroup Generation - Cazenave & Haraux, 1999) *Let $L :$*

$D(L) \rightarrow H$ be a self-adjoint and dissipative linear operator acting on a Hilbert space H . Then, for all $u_0 \in D(L)$, the contraction semigroup $\{T(t)\}_{t \geq 0}$ generated by L guarantees that $u = T(\cdot)u_0$ satisfies the following heat equation:

$$(H) \quad \begin{cases} u \in \mathcal{C}([0, \infty), H) \cap \mathcal{C}((0, \infty), D(L)) \cap \mathcal{C}^1((0, \infty), H); \\ \frac{\partial u}{\partial t} - Lu = 0, \quad \forall t > 0; \\ u(0) = u_0. \end{cases}$$

The family $\{T(t)\}_{t \geq 0}$ is called the heat semigroup of L , and is typically denoted by $\{e^{tL}\}_{t \geq 0}$.

1.4. Kernel Spectral Decomposition

Definition 1.8 (Ultra-contractivity - E. Davies, 1987) Consider the spaces $H := L^2(\Omega, \mathbb{R})$ and $Y := L^\infty(\Omega, \mathbb{R})$ for some domain $\Omega \subseteq \mathbb{R}^d$. Let $\{T(t)\}_{t \geq 0} \subseteq \mathcal{L}(H, H)$ be a contraction semigroup. $\{T(t)\}_{t \geq 0}$ is said to be ultra-contractive if $T(t)H \subseteq Y$.

Theorem 1.3 (Kernel Existence - Arendt, 2006) Consider some domain $\Omega \subseteq \mathbb{R}^d$ and let $\{T(t)\}_{t \geq 0}$ be an ultra-contractive semigroup in $L^2(\Omega, \mathbb{R})$. Then, there exists a unique function $K : (0, \infty) \times \Omega \times \Omega \rightarrow \mathbb{R}$ such that $K(t, \cdot, \cdot) \in L^\infty(\Omega \times \Omega, \mathbb{R}), \forall t > 0$, and

$$(Tu)(x) = \int_{\Omega} K(t, x, y)u(y)dy \quad a.e.$$

The function K is called the heat kernel.

Theorem 1.4 (Spectral Theorem for Operators with Compact Resolvent - Arendt, 2006) Let $L : D(L) \rightarrow H$ be a linear operator over an infinite-dimensional separable Hilbert space H . If L is self-adjoint, dissipative and has a compact resolvent, then there exists an orthonormal basis $\{\varphi_i\}_{i=0}^\infty \subseteq D(L)$ of H and an increasing sequence of non-negative real numbers $\{\lambda_i\}_{i=0}^\infty$ such that $-L\varphi_i = \lambda_i\varphi_i$ and $\lim_{i \rightarrow \infty} \lambda_i = \infty$. Moreover, L is given by

$$D(L) = \left\{ u \in H : \sum_{i=0}^{\infty} |\lambda_i(u, \varphi_i)_H|^2 < \infty \right\}; \quad -Lu = \sum_{i=0}^{\infty} \lambda_i(u, \varphi_i)_H \varphi_i.$$

Theorem 1.5 (Mercer's Theorem - Dodziuk, 1981) Let Ω be a bounded domain and $K \in L^2(\Omega \times \Omega, \mathbb{R})$ be an integral kernel of an operator T in $L^2(\Omega, \mathbb{R})$. Assume further that $K \in \mathcal{C}(\overline{\Omega} \times \overline{\Omega}, \mathbb{R})$ and that K is symmetric. If T is compact, self-adjoint and $(Tu, u) \geq 0, \forall u \in L^2(\Omega, \mathbb{R})$, then there exists an eigenbasis $\{\psi_i\}_{i \in \mathbb{N}}$ of $L^2(\Omega, \mathbb{R})$ with non-increasing eigenvalues $\{\mu_i\}_{i=0}^\infty \subseteq [0, \infty)$ approaching 0 such that

$$K(x, y) = \sum_{i=0}^{\infty} \mu_i \psi_i(x) \psi_i(y)$$

1.5. Heat Kernels of Second Order Elliptic Operators

Definition 1.9 (Elliptic Tensor - Allaire, 2007) Let Ω be a measurable set in \mathbb{R}^n and let $A : \Omega \rightarrow \mathbb{R}^{d \times d}$ be an operator with tensorial coefficients, that is $\mathbf{A}(\mathbf{x}) = (a_{ij}(\mathbf{x}))_{1 \leq i, j \leq d}$. \mathbf{A} is elliptic (or uniformly positive, or coercive) if there exists $\alpha > 0$ such that, a.e. in Ω

$$(\mathbf{A}(\mathbf{x})\xi, \xi)_{\ell^2} \geq \alpha \cdot \|\xi\|_{\ell^2}^2, \quad \forall \xi \in \mathbb{R}^d.$$

Definition 1.10 (Second Order Elliptic Operator in Divergence Form - Arendt, 2002) Let Ω be an open set in \mathbb{R}^d and let $\mathbf{A} : \Omega \rightarrow \mathbb{R}^{d \times d}$ be an elliptic tensor. Assume that \mathbf{A} has bounded tensorial coefficients, that is $\mathbf{A}(\mathbf{x}) = (a_{ij}(\mathbf{x}))_{1 \leq i, j \leq d}$, with $a_{ij} \in L^\infty(\Omega, \mathbb{R})$. A second order elliptic operator in divergence form is a mapping $L_{\mathbf{A}} : H_{loc}^1(\Omega, \mathbb{R}) \rightarrow \mathcal{D}'(\Omega, \mathbb{R})^*$ given by

$$L_{\mathbf{A}}u = -\nabla \cdot (\mathbf{A}\nabla u) = -\sum_{i,j=1}^d \frac{\partial}{\partial x_i} \left(a_{ij} \frac{\partial u}{\partial x_j} \right)$$

A realization of $L_{\mathbf{A}}$ in $L^2(\Omega, \mathbb{R})$ is defined as the operator $\mathcal{L}_{\mathbf{A}} : D(\mathcal{L}_{\mathbf{A}}) \rightarrow L^2(\Omega, \mathbb{R})$ given by

$$D(\mathcal{L}_{\mathbf{A}}) := \left\{ u \in V : \exists f \in L^2(\Omega, \mathbb{R}) \text{ s.t. } a(u, v)_{L^2} = (f, v)_{L^2}, \forall v \in V \right\}; \quad \mathcal{L}_{\mathbf{A}}u := L_{\mathbf{A}}u;$$

where $a(u, v) : V \times V \rightarrow \mathbb{R}$ corresponds to the bilinear form associated to $L_{\mathbf{A}}$:

$$a(u, v) = \int_{\Omega} \left(\sum_{i,j=1}^d a_{ij} \frac{\partial u}{\partial x_i} \frac{\partial v}{\partial x_j} \right) d\mathbf{x};$$

and V is a closed subspace of $H^1(\Omega)$ encoding a boundary condition. In the pure Dirichlet case, $V = H_0^1(\Omega)$, and in the pure Neumann case, $V = H^1(\Omega)$. It can be noticed that $a(\cdot, \cdot)$ is a bounded and coercive form, and $a(\cdot, \cdot)$ is symmetric whenever $\mathbf{A}(\mathbf{x}) = \mathbf{A}(\mathbf{x})^T$.

Proposition 1.5 (Arendt, 1999/2000; Glück & Mui, 2024) Let Ω be a bounded domain with Lipschitz boundary and \mathbf{A} be a symmetric and elliptic tensor with bounded coefficients. Then, the operator $-\mathcal{L}_{\mathbf{A}}$ (equipped with Neumann or Dirichlet boundary conditions) is dissipative, self-adjoint and has a compact resolvent.

Theorem 1.6 (Existence of Heat Kernel - Arendt & ter Elst, 1997) The operator $-\mathcal{L}_{\mathbf{A}}$ (equipped with Neumann or Dirichlet boundary conditions) generates a heat semigroup $\{e^{-t\mathcal{L}_{\mathbf{A}}}\}_{t \geq 0}$ on $L^2(\Omega, \mathbb{R})$. The semigroup is self-adjoint, compact, non-negative and given by a kernel K defined on $(0, \infty) \times \Omega \times \Omega$.

Proposition 1.6 (E. B. Davies, 1989; Grigor'yan, 2006; Griepentrog et al., 2001) The heat kernel K induced by $-\mathcal{L}_{\mathbf{A}}$ satisfies the following properties:

- K is a fundamental solution of the heat equation (H);
- $K \in \mathcal{C}((0, \infty) \times \overline{\Omega} \times \overline{\Omega}, \mathbb{R})$;
- K is symmetric, i.e., $K(t, x, y) = K(t, y, x)$, $\forall t > 0, \forall x, y \in \overline{\Omega}$;

- K is strictly positive;
- K satisfies the semigroup identity:

$$K(t + s, x, y) = \int_{\Omega} K(t, x, z)K(s, z, y)dz, \quad \forall t, s > 0, \forall x, y \in \bar{\Omega}.$$

Chapter 2

Related Work

This chapter reviews three articles proposing novel methods to assess AF vulnerability in patient-specific cardiac models. The studies presented are driven by the compelling need to reduce the computational burden of monodomain simulations, so as to enable clinical translation of *in silico* models to the treatment of arrhythmia, allowing non-invasive and personalized therapy, which could potentially reduce prevalence and recurrence of this disease. Full mathematical background and detailed methodology are not thoroughly described on this chapter, but the essential ideas, basic methods and key results of each work are exhibited to set a precedent of the current state-of-the-art techniques to computationally guide treatment for AF.

2.1. Surrogate Models of Atrial Fibrillation

This section reviews the methodology and results of the study conducted by Serra et al. (2022), where a cardiac electrophysiology simulator based on a cellular automaton (CA) was informed with controlled biophysical simulations to predict sudden cardiac death in subjects who had suffered myocardial infarction.

From a microscopic view, the Hodgkin & Huxley (Hodgkin & Huxley, 1990) ion-gating representation of the cellular layer was used to model action potentials for ground-truth simulations. In this model, electrical dynamics are represented by a circuit of parallel components, where multiple resistors are connected to a capacitor, accounting for the flow of current across the cellular membrane through active ion channels. From a macroscopic modeling perspective, biophysical dynamics of cardiac tissue were determined by the monodomain reaction-diffusion system (see Section 3.1.3).

To model cardiac electrophysiology, an asynchronous CA with 3-dimensional elements was programmed. Each of the nodes of the automaton represented a portion of myocardial tissue. Events were processed sequentially and queued with a first-in-first-out criterion, updating a binary state variable that modeled cell depolarization (active) and repolarization (inactive). When a traveling wavefront of activation reached an inactive node or an external stimulus was applied, the node changed to an active state where it could, in turn, excite adjacent nodes before reaching a resting state and deactivate. Surrounding nodes could only be excited with a non-zero time delay depending on distance and local conduction velocity, which adjusted according to a fixed mathematical model for action potential in ventricular myocytes (Serra

et al., 2022; ten Tusscher et al., 2004).

The automaton-based method was validated using a slab model of tissue and a left ventricular geometry with scar. For the slab model, results show that the CA could generate activation time maps with small difference from ground-truth values in both isotropic and anisotropic configurations for electric potential diffusion, and could also replicate spiral waves similar to the ones observed in biophysical simulations, exhibiting the same rotation period and meandering of the spiral tip (see Figure 2.1). Moreover, by mimicking clinical pacing protocols, ventricular tachycardia could be induced with the CA model. The method reproduced spatial distribution of action potentials and electrotonic effects similar to the biophysical simulations (Serra et al., 2022), as observed in Figure 2.2.




Figure 2.1: Comparative of simulated rotor spiral waves using a CA and a biophysical model in a slab of tissue. Simulations are shown for five consecutive times, 120[ms] after the rotor stabilized. The CA simulation is color-coded by time in milliseconds (ms) from the activation (red tones) to the resting state (blue tones). The biophysical simulation shows the equivalent phenomenon, where colors represent transmembrane potential in millivolts (mV) (Serra et al., 2022). Figure adapted from Serra et al. (2022).

The CA approach captures the electrodynamics of cardiac cells and replicates the behavior of arrhythmia in macroscopic ventricular models, only requiring a desktop machine to run simulations, rendering activation time maps $300\times$ faster than a biophysical solver on a high-performance computer (Serra et al., 2022). Furthermore, the automaton implementation can be extended to efficiently mimic other alterations in the electric cycle of the heart, such as focal atrial tachycardia, flutter, and atrial fibrillation sustained by rotors (Serra et al., 2022).

Although this method could emulate AF in a considerably less amount of time than monodomain and could be used to estimate overall inducibility, there is a significant difference between the local activation of nodes from the CA model and the state of depolarized myocytes that could lead to some error in the determination of fibrillation events. This can be observed from the 3D left-ventricular simulations in Figure 2.2, where an irregular cloud of cells remain active near the base of the ventricle and seems to interact with the ectopic pulse coming from the pacing region, which is not observed during the biophysical simulation. Additionally, the CA and the monodomain model appear to display distinct activation patterns from the source of stimulation, suggesting that some of the directionality in the spread of

voltage might not be properly captured.




Figure 2.2: Simulation time instants (ms) for CA and biophysical model. A ventricular tachycardia (VT) episode resulted from applying an S1–S2 protocol at the location labeled as ‘Stm’. First row shows the endocardial layer (light gray) and the activity in the border zone (BZ) of scar region in the CA model (the core zone, ‘CZ’, and healthy tissue are transparent), while second row shows the equivalent in the biophysical model with all tissue types visible (CZ in dark gray color). Colors in the CA simulation represent time (ms), and in the biophysical simulation represent potential (mV). A conduction block can be observed a few milliseconds after the S2 (white cross). VT is subsequently sustained across the slow conduction channel (SCC). White arrows show the direction of the wavefront (Serra et al., 2022). Figure adapted from Serra et al. (2022).

2.2. ML-based Geometrical Features Approach

Here, the work presented by Bifulco et al. (2023) is summarized. Motivated by the problem of arrhythmia recurrence in persistent AF patients, arrhythmogenic properties of post-ablation substrate were assessed *in silico* by analyzing the mechanistic interactions between patterns of residual fibrosis and ablation-induced scarring through a set of spatial features designed to explain post-ablation anchored reentry.

The methodology presented consisted in clustering non-conductive regions from left atrial models and then measuring for each cluster their corresponding area, perimeter, surrounding residual fibrosis and proximity to non-conductive tissue. Then, a random forest classifier was trained with the latter attributes in a supervised setting to determine whether non-conductive tissue was arrhythmogenic after an ablation procedure. A region was tagged as pro-arrhythmic if it caused anchored reentry, and non-arrhythmogenic otherwise. The inducibility of anchored reentry was determined by manually sub-classifying episodes of arrhythmia from simulations of the monodomain system performed with the Finite Element

method (Bifulco et al., 2023).

The area of each cluster was computed as the total surface of scar, vein and valve tissue composing the cluster; perimeter, was obtained by adding up the lengths of the outer facets of each group; surrounding fibrosis was defined as the percentage of excitable fibrotic cells within a $\sim 1\text{cm}$ band set in the periphery of the non-conductive region; and the proximity index was calculated as the total element layers stacked from the boundary of the cluster before encountering another group of non-conductive tissue (Bifulco et al., 2023). Parameter extraction is illustrated in Figure 2.3 below.




Figure 2.3: Feature extraction for a potentially arrhythmogenic non-conductive tissue region, including the area, perimeter, surrounding fibrosis, and proximity to nearest non-conductive tissue (Bifulco et al., 2023). Figure adapted from Bifulco et al. (2023).

The influence of spatial features from non-conductive tissue to characterize arrhythmogenicity via classification was quantified with SHAP analysis (Lundberg & Lee, 2017), which is displayed in Figure 2.4. Perimeter was shown to be the most important attribute in the classification of pro-arrhythmic substrate after cauterization, where an increased likelihood for anchored reentry presented for perimeters in the range $\sim 15\text{-}60\text{cm}$; followed closely by area, with an optimal predictive interval $\sim 2\text{-}20\text{cm}^2$. Non-conductive regions with values outside these ranges tended to progressively weigh in favor of non-arrhythmogenicity. Furthermore, increasing surrounding fibrosis translated to a linear increase in the likelihood of identification of pro-arrhythmic substrate (Bifulco et al., 2023).

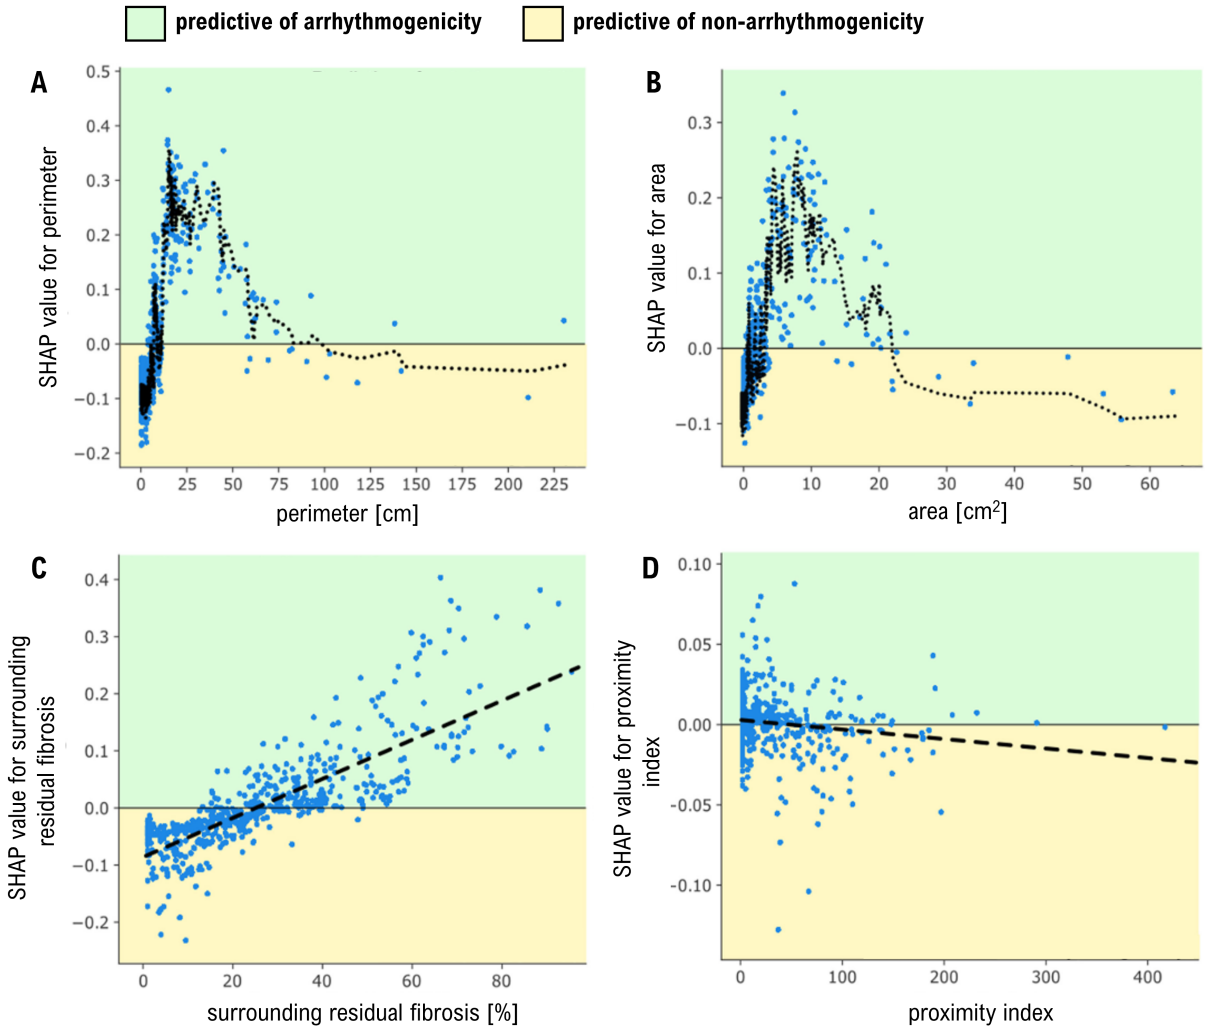


Figure 2.4: A, dependence plot for perimeter. B, dependence plot for area. C, dependence plot of surrounding residual fibrosis percentage. D, dependence plot for proximity index. SHAP indicates Shapley additive explanations (Bifulco et al., 2023). Figure adapted from Bifulco et al. (2023).

The findings from this study suggest that the juxtaposition of ablation-induced scar and residual fibrosis bear a strong structural substrate for reentry mechanisms in subjects with persistent AF, while the formation of isthmus is not a main driver for arrhythmia recurrence in post-ablation models (Bifulco et al., 2023).

Although the feature-extraction procedure elucidates the relationship between geometrical properties of scar regions and their propensity to trigger AF, it is not sufficient to establish pointwise inducibility in the cardiac domain, since the methodology proposed only considers clusters of non-conductive tissue rather than evaluating specific points from the geometry.

2.3. Multi-Fidelity Gaussian Process Classification

In this section, the method developed by Gander et al. (2022) is summarized. This work exploits the statistical correlation between low-fidelity and high-fidelity models of the elec-

trophysiology equations by informing multi-fidelity Gaussian process (GP) classifiers with both type of models to predict AF inducibility on the atrial surface from the high-fidelity source. This multi-fidelity approach offsets the monodomain computational burden to faster coarser models, with the goal to reduce the cost of estimation for a targeted accuracy and improving performance for limited computational resources (Gander et al., 2022).

This study employed zero-mean Gaussian process priors alongside Matérn-like covariance kernel functions (Gander et al., 2022) to model inducibility outcomes from the different fidelity sources. Since a non-euclidean geometry was considered, the covariance kernels were estimated by using the eigenpairs of the Laplace-Beltrami operator equipped with homogeneous Neumann boundary conditions. To construct the multi-fidelity classifier, the relationship between the latent functions of the GPs governing the high-fidelity and the low-fidelity models was postulated to be an auto-regressive prior, where the high-fidelity function corresponded to a scaled low-fidelity prior plus a nuisance function explaining the difference between models. Consequently, a multi-fidelity joint distribution over the latent functions could be obtained.

To assess the performance of classifiers in this setting, synthetic examples from a single atrial geometry under three fibrosis scenarios and three ablation cases were used. For each fibrosis/ablation case, the evaluation of the multi-fidelity GP approach was carried out by using a dataset where each element of the set was composed of a feature vector containing a point coordinate in the atrial surface and a binary tag representing the inducibility output from an AF simulation after applying a stimulation protocol from that point. Here, the data used for the inducibility tags in the training set were dichotomically comprised of two sources: an expensive and hard-to-acquire high-fidelity source based in the monodomain system on a fine mesh, and an inexpensive and faster-to-compute low-fidelity model obtained with a coarser grid. On the other hand, the set used for testing only had samples from the high-fidelity model. The multi-fidelity GP classifiers were trained using an active learning approach, where singleton samples were iteratively added by promoting the selection of locations near the decision boundary (Gander et al., 2022). Their performance was then tested against single-fidelity classifiers and a nearest neighbor algorithm, both trained with the high-fidelity data.

Results showed that single and multi-fidelity models with an active learning approach yielded better results than a fixed training design, and that both cases outperformed the nearest-neighbor estimator (Gander et al., 2022), showing a high dependence on the length-scale of the classifier (see Figure 2.5 A). Moreover, comparing the accuracy with the single-fidelity model, the proposed method rendered better predictions for AF inducibility when a reduced budget of 40 pacing sites was available (see Figure 2.5 B).

The multi-fidelity Gaussian process classifier can reduce the computational overhead of monodomain simulations while competitively maintaining or improving accuracy to the single-fidelity strategy. Moreover, the proposed method does not demand intrusive changes to existing implementations (Gander et al., 2022) and could lead to clinical translation to evaluate ablation strategies in real patient-specific anatomies. However, although allowing to reduce the number of high-resolution simulations in exchange for low-fidelity information, the multi-fidelity approach still relies on the monodomain system and depends on the cardiac

geometry explicitly, making it necessary to re-run simulations to train the classifier again after fibrosis is changed or the anatomy is modified.

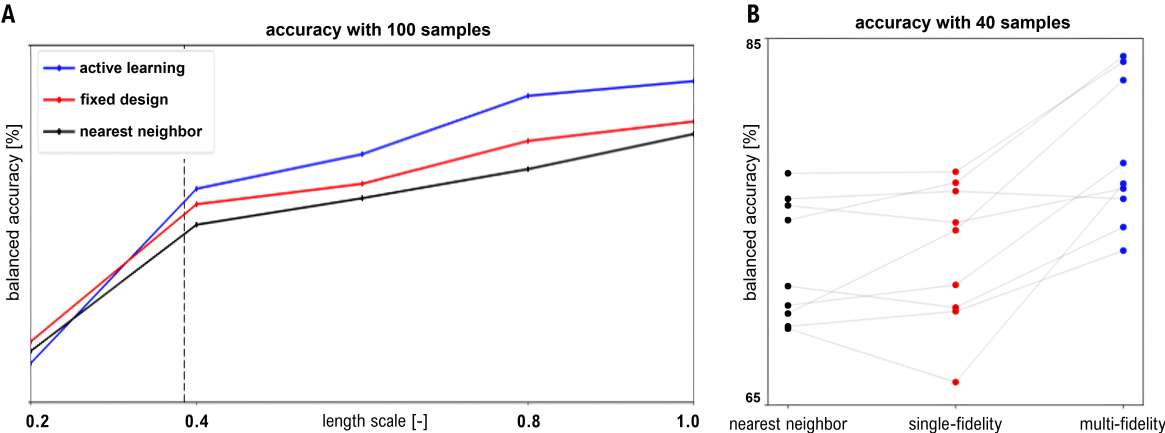


Figure 2.5: A, average balanced accuracy change for increasing length scale when the classifiers are trained with 100 samples. The dashed vertical line represents the average geodesic distance between training points of the fixed design. B, balanced accuracy comparison for the nearest neighbor, single-fidelity, and multi-fidelity classifiers, for all nine model scenarios and with a fixed budget of 40 high-fidelity simulations (Gander et al., 2022). Figure adapted from Gander et al. (2022).

Chapter 3

General Methodology

3.1. Atrial Fibrillation Modeling

In this section, three mathematical models that describe the propagation of voltage in the heart are summarized. First, the *microscopic bidomain model* is introduced, where cardiac tissue is conceptualized as a composite of two distinct conductive domains that interact by actively passing current across a conductance barrier separating the two media. Then, the *macroscopic bidomain model* is presented as a homogenized version of the previous model, where physical quantities are obtained through local averaging. Finally, the *macroscopic monodomain model* is derived by simplifying the bidomain model under the assumption of proportionality between the conductivities of intra- and extracellular spaces.

3.1.1. Microscopic Bidomain Model

Heart muscle is an ensemble of elongated excitable cells coupled longitudinally and laterally by aggregates of intracellular ionic channels (Colli Franzone et al., 2014; Goodenough & Paul, 2009). From a modeling perspective, cardiac tissue consists of an ohmic system composed by a conducting intracellular medium Ω_i overlaid by an active membrane Γ_m and immersed in a conductive matrix Ω_e . This can be thought of as a collection of roughly cylindrical conjoined cables submerged in a conductive material, separated from it by a series of capacitors, where current flow varies because of changes in the resistance of the electrical units and the variations in permittivity of the separating layer.

Mathematically, Ω_i and Ω_e are regarded as two connected open sets of \mathbb{R}^3 , and the cardiac domain Ω_H as a whole can be defined as their union alongside membrane bounds: $\Omega_H = \Omega_i \cup \Gamma_m \cup \Omega_e$ (see Figure 3.1). The microstructural effects in current flow caused by local variations in conductance because of media heterogeneity, gap junctions, collagen allocation and presence of blood vessels are accounted by the positive-definite conductivity tensors $\Sigma_i(\mathbf{x})$ and $\Sigma_e(\mathbf{x})$ (Colli Franzone et al., 2014).

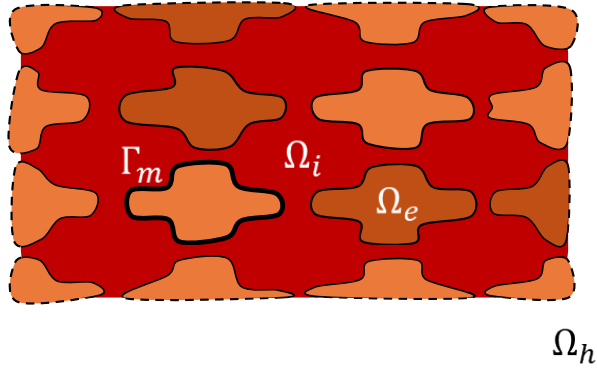


Figure 3.1: Illustration of cellular connectivity from slice of cardiac tissue. The structure presented repeats periodically to form extended layers of heart muscle.

Let V_i, V_e denote the potentials for the intra- and extracellular matrix, respectively. By Ohm's law, their corresponding current densities \mathbf{J}_i and \mathbf{J}_e satisfy the following equations:

$$\mathbf{J}_{i,e} = -\boldsymbol{\Sigma}_{i,e} \nabla V_{i,e} \quad \text{in } \Omega_{i,e} \times (0, T].$$

Considering an external septal stimuli and assuming that active sources only lie on the cellular membrane, the transmembrane current per unit area I_m can be expressed in terms of a capacitor term, an ion gating component and a stimulus current (Colli Franzone et al., 2014; Bader et al., 2021):

$$I_m = C_m \frac{\partial V_m}{\partial t} + I_{\text{ion}} + I_{\text{stim}} \quad \text{on } \Gamma_m \times (0, T];$$

where C_m is the membrane capacitance, $V_m := V_i - V_e$ is the transmembrane potential, I_{stim} is the external stimulus and I_{ion} is the ionic current. The latter is a function depending on the electrophysiological model. Namely, $I_{\text{ion}} = I_{\text{ion}}(V_m, \mathbf{w})$, where \mathbf{w} is a vector containing the ion gating and concentration variables that satisfies the following ODE system:

$$\begin{cases} \frac{\partial \mathbf{w}}{\partial t} = \mathbf{g}(V_m, \mathbf{w}) & \text{on } \Gamma_m \times (0, T]; \\ \mathbf{w}(\cdot, 0) = \mathbf{w}_0 & \text{on } \Gamma_m. \end{cases}$$

Let $\mathbf{n}_i, \mathbf{n}_e$ be the exterior boundary normals of Ω_i and Ω_e , respectively. Kirchoff's junction rule implies that normal current flux is continuous through the membrane. Under the assumption of an outward current, the following transmission boundary conditions are obtained (Colli Franzone et al., 2014):

$$\begin{cases} \mathbf{J}_i \cdot \mathbf{n}_i = \chi I_m & \text{on } \Gamma_m \times (0, T]; \\ -\mathbf{J}_e \cdot \mathbf{n}_e = \chi I_m & \text{on } \Gamma_m \times (0, T]; \end{cases}$$

with χ being the surface-to-volume ratio, computed as the fraction between membrane surface and its enclosed volume for a given length of fiber.

Moreover, under the assumption that intracellular and extracellular regions are source-free, the conservation law for electric charges provides the following equation:

$$\nabla \cdot \mathbf{J}_{i,e} = 0 \quad \text{in } \Omega_{i,e} \times (0, T].$$

With the additional supposition of an insulated cardiac domain, homogeneous Neumann conditions are imposed for the potentials in the outer boundaries $\Gamma_{i,e} := \partial\Omega_{i,e} \setminus \Gamma_m$, namely

$$\mathbf{J}_{i,e} \cdot \mathbf{n}_{i,e} = 0 \quad \text{on } \Gamma_{i,e} \times (0, T].$$

Putting together the equations above, with proper initial conditions for the potential difference V_m , the microscopic bidomain system is assembled for Ω_H :

$$(B_m) \quad \begin{cases} -\nabla \cdot (\boldsymbol{\Sigma}_{i,e} \nabla V_{i,e}) = 0 & \text{in } \Omega_{i,e} \times (0, T]; \\ \frac{\partial \mathbf{w}}{\partial t} = \mathbf{g}(V_m, \mathbf{w}) & \text{on } \Gamma_m \times (0, T]; \\ (\boldsymbol{\Sigma}_{i,e} \nabla V_{i,e}) \cdot \mathbf{n} = -\chi \left(C_m \frac{\partial V_m}{\partial t} + I_{\text{ion}} + I_{\text{stim}} \right) & \text{on } \Gamma_m \times (0, T]; \\ (\boldsymbol{\Sigma}_{i,e} \nabla V_{i,e}) \cdot \mathbf{n} = 0 & \text{on } \Gamma_{i,e} \times (0, T]; \\ \mathbf{w}(\cdot, 0) = \mathbf{w}_0; \quad V_m(\cdot, 0) = V_0 & \text{on } \Gamma_m; \end{cases}$$

where \mathbf{n} is \mathbf{n}_i on Γ_m and equal to $\mathbf{n}_{i,e}$ on $\partial\Omega_{i,e} \setminus \Gamma_m$.

3.1.2. Macroscopic Bidomain System

The system presented in the previous section corresponds to a microscopic concept of cardiac structure that does not take into consideration its global properties. To capture the macroscopic behavior of heart muscle, a model that averages myocyte conduction characteristics into a continuum representation of tissue that can hold momentarily stationary with respect to individual cell activity is used. It stems from the consideration of interpenetrating extracellular and intracellular spaces (Colli Franzone et al., 2014; Schmitt, 1969).

Let Ω_H be a heart volume such that in every point from the domain both cytoplasmic and interstitial media coexist and are separated by a continuously distributed cellular membrane. For every point $\mathbf{x} \in \Omega_H$, denote by $\mathbf{J}_{i,e}(\mathbf{x}, t)$ the macroscopic current densities computed from local averaging and I_m the transmembrane current per unit area. From the current conservation law, for any neighborhood $V \subseteq \Omega_H$ of \mathbf{x} , the average flux entering interstitial volume and exiting cytoplasmic space equals the average current across the membrane (Colli Franzone et al., 2014). More specifically,

$$\frac{1}{|V|} \int_{\partial V} \mathbf{J}_e \cdot \mathbf{n} ds = -\frac{1}{|V|} \int_{\partial V} \mathbf{J}_i \cdot \mathbf{n} ds = \frac{1}{|V|} \int_V \chi I_m dx,$$

where \mathbf{n} is the outward-pointing normal on ∂V . Taking the limit $|V| \rightarrow 0$, the equation for conservation of electric charges is recovered (Colli Franzone et al., 2014):

$$\begin{cases} \nabla \cdot \mathbf{J}_i = -\chi I_m & \text{in } \Omega_H \times (0, T]; \\ \nabla \cdot \mathbf{J}_e = \chi I_m & \text{in } \Omega_H \times (0, T]; \end{cases}$$

Moreover, regarding extra- and intracellular media as ohmic, current densities can be written in terms of the macroscopic potentials V_e and V_i as

$$\mathbf{J}_{i,e} = -\boldsymbol{\Sigma}_{i,e} \nabla V_{i,e}.$$

Applying an external stimulus on the membrane and taking into consideration the insulated heart hypothesis, the macroscopic bisyncytial model is expressed in its parabolic-parabolic formulation as follows (Colli Franzone et al., 2014; Bader et al., 2021):

$$(B_M) \quad \begin{cases} -\nabla \cdot (\boldsymbol{\Sigma}_i \nabla V_i) = -\chi \left(C_m \frac{\partial V_m}{\partial t} + I_{\text{ion}} + I_{\text{stim}} \right) & \text{in } \Omega_H \times (0, T]; \\ -\nabla \cdot (\boldsymbol{\Sigma}_e \nabla V_e) = \chi \left(C_m \frac{\partial V_m}{\partial t} + I_{\text{ion}} + I_{\text{stim}} \right) & \text{in } \Omega_H \times (0, T]; \\ \frac{\partial \mathbf{w}}{\partial t} = \mathbf{g}(V_m, \mathbf{w}) & \text{in } \Omega_H \times (0, T]; \\ (\boldsymbol{\Sigma}_{i,e} \nabla V_{i,e}) \cdot \mathbf{n} = 0 & \text{on } \partial\Omega_H \times (0, T]; \\ \mathbf{w}(\cdot, 0) = \mathbf{w}_0; \quad V_m(\cdot, 0) = V_0 & \text{in } \Omega_H. \end{cases}$$

3.1.3. Macroscopic Monodomain System

The solution from the macroscopic bidomain system presented in the previous section exhibits a wavefront propagation with a rapid and steep upstroke during the excitation phase of the heartbeat. This causes numerical methods to require meshes with high resolution for both time and space in order to accurately estimate the transmembrane potential in cardiac models (Colli Franzone et al., 2014; Dal et al., 2012). Thus, high processing speed and a significant amount of memory are needed to numerically solve the corresponding bivariate reaction-diffusion system for large scale simulations. This motivates a reduction of the model to a less demanding approximation of the electrical behavior of the heart: the monodomain system, that tackles the issue of different electrophysiological properties between cytoplasmic space and the extracellular matrix by taking the bulk medium conductivity.

Let $\mathbf{J} := \mathbf{J}_i + \mathbf{J}_e$ and $\boldsymbol{\Sigma} := \boldsymbol{\Sigma}_i + \boldsymbol{\Sigma}_e$. Furthermore, the conductivity tensors are regarded in the form of an orthonormal basis $\{\mathbf{a}_l, \mathbf{a}_t, \mathbf{a}_n\}$ describing the local direction of cardiac tissue (Colli Franzone et al., 2014; Johnston & Johnston, 2020):

$$\boldsymbol{\Sigma}_{i,e} = \sigma_{i,e}^l \mathbf{a}_l(\mathbf{x}) \otimes \mathbf{a}_l(\mathbf{x}) + \sigma_{i,e}^t \mathbf{a}_t(\mathbf{x}) \otimes \mathbf{a}_t(\mathbf{x}) + \sigma_{i,e}^n \mathbf{a}_n(\mathbf{x}) \otimes \mathbf{a}_n(\mathbf{x});$$

where \mathbf{a}_l is parallel to the fiber, $\mathbf{a}_t, \mathbf{a}_n$ are tangent and normal to the fiber's laminar sheet, and $\sigma_{i,e}^l, \sigma_{i,e}^t, \sigma_{i,e}^n$ denote the intra- and extracellular conductivities measured along the respective directions, which are assumed constant (Colli Franzone et al., 2014; Johnston & Johnston, 2020).

Since $V_m = V_i - V_e$ and the current densities are given by the potentials $V_{i,e}$, \mathbf{J} can be expressed as

$$\mathbf{J} = -\boldsymbol{\Sigma}_i \nabla V_m - \boldsymbol{\Sigma} \nabla V_e.$$

Furthermore, given that $\Sigma_{i,e}$ are positive definite, Σ is invertible, and ∇V_e can then be written as

$$\nabla V_e = -\Sigma^{-1}\Sigma_i\nabla V_m - \Sigma^{-1}\mathbf{J}.$$




Figure 3.2: Illustration of muscle composition in ventricular tissue with orthonormal vectors aligned with laminae and fiber direction. Figure based on illustration from Kocica et al. (2006).

Therefore, the second equation from (B_M) is equivalent to (Colli Franzone et al., 2014)

$$\nabla \cdot (\Sigma_e \Sigma^{-1} \Sigma_i V_m) + \nabla \cdot (\Sigma_e \Sigma^{-1} \mathbf{J}) = \chi \left(C_m \frac{\partial V_m}{\partial t} + I_{\text{ion}} + I_{\text{stim}} \right),$$

Taking into consideration a transverse isotropy regime for conductivities ($\sigma_{i,e}^t = \sigma_{i,e}^n$) and equal anisotropy ratios ($\sigma_e^l/\sigma_i^l = \sigma_e^t/\sigma_i^t = \sigma_e^n/\sigma_i^n$), the conductivity tensor in the second divergence term from the previous equation simplifies to (Colli Franzone et al., 2014; Hurtado & Henao, 2014)

$$\Sigma_e \Sigma^{-1} = \frac{\sigma_e^l}{\sigma_i^l + \sigma_e^l} \mathbf{I}.$$

This result, along with the fact that the two first bidomain equations yield $\nabla \cdot \mathbf{J} = 0$, reduces the original evolution system to a PDE describing only the dynamics of transmembrane potential:

$$\nabla \cdot (\Sigma_m \nabla V_m) = \chi \left(C_m \frac{\partial V_m}{\partial t} + I_{\text{ion}} + I_{\text{stim}} \right) \quad \text{in } \Omega_H \times (0, T];$$

with $\Sigma_m := \Sigma_e \Sigma^{-1} \Sigma_i$ being the conductivity tensor of the bulk medium, having the explicit form

$$\Sigma_m = \frac{\sigma_e^l \sigma_i^l}{\sigma_e^l + \sigma_i^l} (\mathbf{a}_l \otimes \mathbf{a}_l) + \frac{\sigma_e^t \sigma_i^t}{\sigma_e^t + \sigma_i^t} (\mathbf{I} - \mathbf{a}_l \otimes \mathbf{a}_l).$$

To derive a Neumann boundary condition for the approximate model, it suffices to use

$\Sigma_e \Sigma^{-1} = \sigma_e^l / (\sigma_i^l + \sigma_e^l) \mathbf{I}$ and recall that $\Sigma^{-1} \Sigma_i \nabla V_m = -\nabla V_e - \Sigma^{-1} \mathbf{J}$, which yields

$$(\Sigma_m \nabla V_m) \cdot \mathbf{n} = -(\Sigma_e \nabla V_e) \cdot \mathbf{n} - \frac{\sigma_e^l}{\sigma_i^l + \sigma_e^l} (\mathbf{J}_i + \mathbf{J}_e) \cdot \mathbf{n} = 0 \quad \text{on } \partial\Omega_H \times (0, T].$$

Putting together the deduced relations for V_m , along with the ionic model, the monodomain system is obtained:

$$(M) \quad \begin{cases} -\nabla \cdot (\Sigma_m \nabla V_m) = -\chi \left(C_m \frac{\partial V_m}{\partial t} + I_{\text{ion}} + I_{\text{stim}} \right) & \text{in } \Omega_H \times (0, T]; \\ \frac{\partial \mathbf{w}}{\partial t} = \mathbf{g}(V_m, \mathbf{w}) & \text{in } \Omega_H \times (0, T];, \\ (\Sigma_m \nabla V_m) \cdot \mathbf{n} = 0 & \text{on } \partial\Omega_H \times (0, T]; \\ \mathbf{w}(\cdot, 0) = \mathbf{w}_0; \quad V_m(\cdot, 0) = V_0 & \text{in } \Omega_H. \end{cases}$$

The variability in conductivity across different cardiac regions and the presence of structural abnormalities are encoded in the conductivity tensor Σ_m . Thus, the solution to the monodomain system can change depending on the availability of information regarding the anatomical parts of the patient's heart, the direction of muscle fibers and fibrosis distribution. The dependency between the monodomain solution and the variation in conductivity can be exploited to identify AF substrate by creating markers that are flexible to the data that is accessible through Σ_m .

3.2. Pacing Protocols

The efficacy of therapeutic approaches for treating AF can be tested by measuring the post-procedural vulnerability of the cardiac muscle to initiate sustained episodes of arrhythmia. In clinical practice, electrical stimulation is used to diagnose and guide treatment (Azzolin et al., 2021), making AF inducibility a landmark for targeting proarrhythmic regions and consequently determining optimal outputs for conventional therapy, such as catheter ablation. However, the progressive nature of this disorder and inherent complexity explain its prevalence even after standard-of-care treatments are performed (Azzolin et al., 2021; Chae et al., 2007; Verma et al., 2015; Kawai et al., 2019; Gharaviri et al., 2021). Moreover, the multifaceted causes of heart rhythm abnormalities and intrinsic processes that rise susceptibility for AF lead to a variety of criteria for defining inducibility, deciding pace-point collocation and modulating each stimulation intensity. This results in the creation of various pacing techniques to quantify AF vulnerability in different levels.

Patient-specific *in silico* models have shown to be effective for the design of better treatments for AF (Azzolin et al., 2021; Boyle et al., 2019; Loewe et al., 2019; McDowell et al., 2015), as they not only enable the comparison between different therapeutic procedures through simulation, but allow testing multiple pacing protocols to measure the propensity of cardiac tissue to initiate and maintain fibrillation events. In this section, three state-of-the-art protocols are reviewed: phase singularity distribution (PSD), rapid pacing (RP) and pacing at the end of the effective refractory period (PEERP). In this work, the last two techniques are used to characterize pro-arrhythmic sites in the atria.

3.2.1. Phase Singularity Distribution

A *phase singularity* (PS) is defined as the tip of a rotational front from the electrogram (EGM) phase map $\Phi(\mathbf{x}, t)$, which is a function that uniquely describes the excitation state of a cell \mathbf{x} at any time t (Iyer & Gray, 2001; Clayton et al., 2006) (see Figure 3.3). Some techniques used to construct cardiac phase maps are the Sawtooth and the Hilbert Transform methods, which are illustrated in Figure 3.4. The first method converts the signal through a linear interpolation from π to $-\pi$ between consecutive local activation times, and the second uses the argument of the analytical form of the EGM, where the real part corresponds to the signal and the imaginary part to its Hilbert transform (Lootens et al., 2024).




Figure 3.3: Phase singularity below left pulmonary veins on atrial model with counterclockwise spiral rotation from electrical potential recordings (mV).

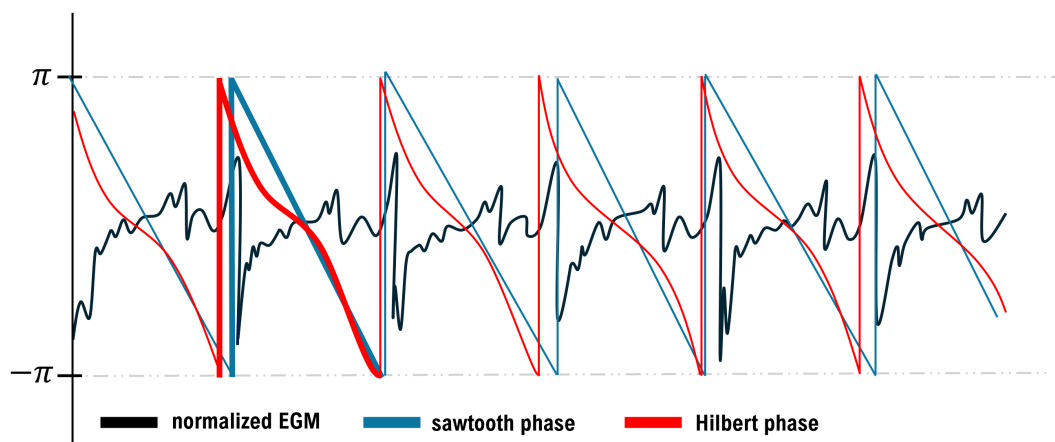


Figure 3.4: Illustration of converted EGM signal to phase maps with Sawtooth and Hilbert Transform methods. Figure adapted from Lootens et al. (2024).

The *phase singularity distribution* method (PSD) creates initial conditions for a monodomain propagation with the prior knowledge of wave dynamics and the location of reentry circuits in the heart surface (Matene & Jacquemet, 2012; Azzolin et al., 2021). It starts by positioning an arbitrary amount of phase singularities on a monolayer approximate of the atria, and then it initializes all state variables of the system from a re-entrant activation map computed through an eikonal-based shortest path algorithm that considers an anisotropic medium and assumes known conduction velocities (Matene & Jacquemet, 2012). This method requires the cycle length of the re-entrant front, the number of PSs and their direction of rotation (Azzolin et al., 2021; Matene & Jacquemet, 2012).

Although effective for simulating different types of AF and creating arbitrary large fibrillation episodes (Matene & Jacquemet, 2012), the PSD can only be applied to computational modeling, as it includes *a priori* information from the location and behavior of active sources for the maintenance of arrhythmia that are not known in practice (Azzolin et al., 2021).

3.2.2. Rapid Pacing

In the clinical setting, *rapid pacing* (RP) is a commonly used protocol to test arrhythmia inducibility in patients (Azzolin et al., 2021). This method consists in the application of a progression of stimuli with decreasing coupling intervals (CI), where a CI is defined as the time interval comprising a foregoing action potential followed by the onset of an ectopic pulse (de Vries et al., 2018), as illustrated in Figure 3.5 A.




Figure 3.5: A, electrogram of single cardiomyocyte under RP stimulation protocol. B, electrogram of cardiomyocyte with PEERP pacing strategy. The colored vertical lines correspond to the instants of stimulation. Note that the coupling interval from the PEERP protocol depends on the effective refractory period.

The RP protocol is parameterized by a collection of CI lengths $L := \{l_i\}_{i=1}^n$, the number of stimuli before each CI decrease $N := \{N_i\}_{i=1}^n$ and an indicator of successful initiation of fibrillation $I \in \{E, B\}$, that checks AF induction after every beat (B) or at the end of the protocol (E) (Azzolin et al., 2021).

For example, the most typical rapid pacing plan corresponds to a train of CIs in the decreasing range $200 - 130[ms]$ with time step $10[ms]$ and AF indicator at the end of the protocol. This is, $L = \{200, 190, \dots, 140, 130\}$, $N = \{1, 1, \dots, 1, 1\}$ and $I = E$ (Krummen et al., 2012; Zahid et al., 2016; Boyle et al., 2019; as cited in Azzolin et al., 2021):

$$RP(L, N, I) = \{200, 190, \dots, 140, 130, \mathbf{c}\}.$$

Another instance of a RP stimulation can be given by the set of CIs as in the previous example, but with duplicate repetitions for each interval and checking arrhythmia initiation at the end of every beat:

$$RP(L, \{2, \dots, 2\}, B) = \{200, \mathbf{c}, 200, \mathbf{c}, 190, \mathbf{c}, 190, \mathbf{c}, \dots, 130, \mathbf{c}, 130, \mathbf{c}\}.$$

Despite the flexibility of the RP method to render point inducibility through a multi-parameter setting, and its versatility to be utilized in both computational simulations and *in vivo* experiments, the protocol bears the difficulty of reproducibility, because of the strong dependence in the choice of variables and individual criteria for fixing a configuration to consistently quantify AF vulnerability. However, given that RP is commonly used in clinical applications, its use is justified for the first study performed for this thesis.

3.2.3. Pacing at the End of the Effective Refractory Period

Pacing at the end of the effective refractory period (PEERP) is a method that depends only on the maximum number of pulses to apply to a specific location. This protocol fixes a set of pace-points in the cardiac geometry and delivers a premature beat as soon as the underlying tissue is able to spread a new front of propagation (Azzolin et al., 2021). The instant for stimulation initiation is queued by the effective refractory period (ERP), which corresponds to the largest CI where the impulse fails to locally spread (Katritsis & Morady, 2022; Issa et al., 2009) (see Figure 3.6). The end of the ERP is estimated with a binary search algorithm that computes the minimum time at which a new depolarization wave can successfully propagate to a set of neighboring cells from the pacing location (Azzolin et al., 2021). In practice, local estimation of ERP is performed with a pacing protocol that looks for the longest stimulus-stimulus repolarizing interval (Azzolin et al., 2021; Issa et al., 2009). However, stimulus-free methods could also be used for this task (Verrier et al., 2016).

Although PEERP works under the assumption of high resolution in repolarization intervals and brings about more time complexity to put in the external pulse at the end of the ERP, this protocol was selected to assess AF inducibility for the second study related to this thesis work, as it maximizes chances of reentry when adaptively selecting the stimulus and because it only depends on a fixed current strength and maximum number of beats, thus facilitating standardization (Azzolin et al., 2021).

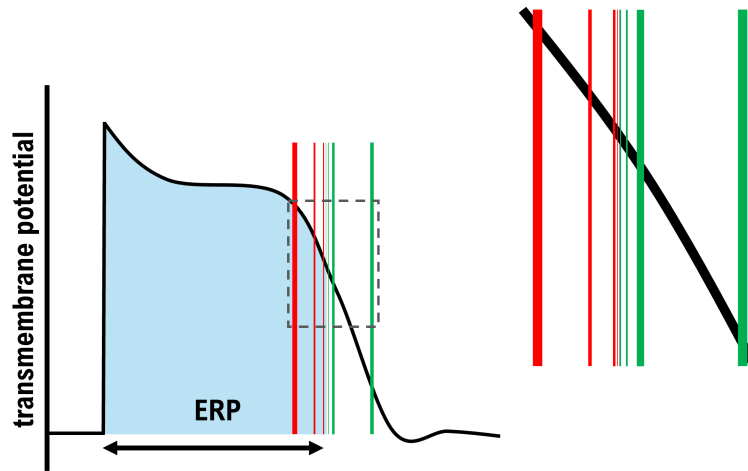


Figure 3.6: ERP representation for myocyte action potential and binary search strategy for end of ERP estimation. The colored bands indicate the pulse application and their width the state of iteration. The color green indicates that the delivered stimulus successfully spread in neighboring cells, and red otherwise.

3.3. The Fibrotic Kernel Signature

3.3.1. The Heat Kernel Signature

The heat kernel signature (hks) is a time-dependent point descriptor used in shape analysis. It embeds intrinsic geometrical information from a figure by accounting for the diffusion process through its surface after applying an infinitely large amount of heat for an infinitely small window of time on a fixed location (Sun et al., 2009). Mathematically, this signature is defined as a spatial restriction of the kernel induced by the Laplace-Beltrami operator on manifolds (Banduc et al., 2024).

Let M be a compact Riemannian manifold with Lipschitz boundary embedded in a n -dimensional space. The heat distribution u on M provided an initial condition $u_0 \in L^2(M, \mathbb{R})$ for an observation interval $[0, T]$ is governed by the following equation:

$$(H_0) \quad \begin{cases} \frac{\partial u}{\partial t} - \Delta_M u = 0 & \text{in } M \times (0, T]; \\ u = u_0 & \text{in } M \times \{t = 0\}; \end{cases}$$

where Δ_M denotes the Laplace-Beltrami operator in M . The solutions for (H_0) are built from the heat semigroup $T(t) = e^{t\Delta_M}$, where the function defined by $u(t) := T(t)u_0$ satisfies the heat equation for the initial distribution u_0 and suitable boundary conditions on ∂M . (Cazenave & Haraux, 1999; Grigor'yan, 2018; Choulli et al., 2015)

Moreover, there exists a unique function (Norris, 1997; Grigor'yan, 1997; Choi & Kim,

2013) $K : (0, \infty) \times M \times M \rightarrow \mathbb{R}$ such that for any $u_0 \in L^2(M, \mathbb{R})$

$$T(t)u_0(\mathbf{x}) = \int_M K(t, \mathbf{x}, \mathbf{y})u_0(\mathbf{y})d\mathbf{y},$$

The *heat kernel* $K(t, \mathbf{x}, \mathbf{y})$ is smooth, positive, symmetric, obeys the semigroup property $K(t + s, \mathbf{x}, \mathbf{y}) = (K(t, \mathbf{x}, \cdot), K(s, \cdot, \mathbf{y}))_{L^2}$ and satisfies the diffusion problem for an initial Dirac-delta temperature (Grigor'yan, 2018; Grieser, 2004)

$$(H_\delta) \quad \begin{cases} \frac{\partial K}{\partial t} - \Delta_M K = 0 & \text{in } M \times M \times (0, \infty); \\ \lim_{t \rightarrow 0^+} K(t, \mathbf{x}, \mathbf{y}) = \delta(\mathbf{x}, \mathbf{y}) & \text{in } M \times M; \end{cases}$$

where $\delta(\mathbf{x}, \cdot)$ denotes a Dirac distribution on $\mathbf{x} \in M$.

The heat kernel lists several useful properties to analyze shapes undergoing isometric transformations, small deformations and multi-scale comparison. In the properties stated below, manifolds are assumed compact and Riemannian, with Lipschitz boundary.

Proposition 3.1 (Intrinsic Property - Sun et al., 2009) *Let M_1, M_2 be two manifolds, and let $K_i \in C^\infty((0, \infty) \times M_i \times M_i, \mathbb{R})$ be the heat kernel induced by Δ_{M_i} . If $\mathcal{T} : M_1 \rightarrow M_2$ is an isometry, then $K_1(t, \mathbf{x}, \mathbf{y}) = K_2(t, \mathcal{T}\mathbf{x}, \mathcal{T}\mathbf{y}), \forall \mathbf{x}, \mathbf{y} \in M_1, \forall t > 0$.*

Proposition 3.2 (Informative Property - Sun et al., 2009) *Let $\mathcal{T} : M_1 \rightarrow M_2$ be a surjective map between two manifolds M_1 and M_2 , and K_i as described above. If $K_1(t, \mathbf{x}, \mathbf{y}) = K_2(t, \mathcal{T}\mathbf{x}, \mathcal{T}\mathbf{y}), \forall \mathbf{x}, \mathbf{y} \in M_1, \forall t > 0$, then \mathcal{T} is an isometry.*

Proposition 3.3 (Multi-scale Property - Dodziuk, 1983; Sun et al., 2009) *Let $M_1 \subseteq M_2$ and K_i be the Dirichlet heat kernel induced by Δ_{M_i} . Then, for all $r > 0$, $K_2(t, \mathbf{x}, \mathbf{y}) - K_1(t, \mathbf{x}, \mathbf{y}) = O(t^r)$ uniformly on $M_1 \times M_1$ as $t \rightarrow 0^+$, and $K_1(t, \mathbf{x}, \mathbf{y}) \leq K_2(t, \mathbf{x}, \mathbf{y}), \forall t > 0, \forall \mathbf{x}, \mathbf{y} \in M_1$. Moreover, let M be a manifold with Dirichlet heat kernel K and $\{M_i\}_{i=1}^\infty$ a strictly expanding sequence of open subsets of M such that $\bigcup_{i=1}^\infty M_i = M$. Then, the family of Dirichlet heat kernels $\{K_i\}_{i=1}^\infty$ corresponding to $\{M_i\}_{i=1}^\infty$ converges to K pointwise: $\lim_{i \rightarrow \infty} K_i(t, \mathbf{x}, \mathbf{y}) = K(t, \mathbf{x}, \mathbf{y}), \forall t > 0, \forall \mathbf{x}, \mathbf{y} \in M$.*

The heat kernel arises as a natural fit for a point signature. Furthermore, since it describes a diffusion process, its spatial variations manifest from its changes in time (Sun et al., 2009). In consequence, up to mild assumptions regarding the multiplicity of the eigenvalues of Δ_M , the properties previously shown are overdetermined for K , making it possible to restrict its space component without dropping valuable information from the geometry (Sun et al., 2009).

Definition 3.1 *Let M be a compact Riemannian manifold with Lipschitz boundary and K be the heat kernel derived from Δ_M . The heat kernel signature (hks) is defined as the diagonal*

argument of K :

$$\begin{aligned} hks : (0, \infty) \times M &\rightarrow \mathbb{R} \\ (t, \mathbf{x}) &\mapsto K(t, \mathbf{x}, \mathbf{x}). \end{aligned}$$

This definition can be made proper when considering the expansion of the kernel in terms of the eigenpairs $\{(\lambda_i, \varphi_i)\}_{i=0}^{\infty}$ of the operator $-\Delta_M$. This is

$$K(\mathbf{x}, \mathbf{y}) = \sum_{i=0}^{\infty} e^{-t\lambda_i} \varphi_i(\mathbf{x}) \varphi_i(\mathbf{y}).$$

The hks can be regarded as the measurement of a physical quantity in a point $\mathbf{x} \in M$ at time $t > 0$ after instantly applying a concentrated source of that quantity in that same specific location at time 0. For instance, this could be interpreted as the time course of an ink drop as it diffuses through a glass of water (Sahli Costabal et al., 2023).

This signature inherits the essential properties from the heat kernel. However, an hypothesis regarding the eigenvalues of the Laplace-Beltrami operator must be added to preserve its informative property.

Theorem 3.1 (hks Informative Property - Sun et al., 2009) *Let $\mathcal{T} : M_1 \rightarrow M_2$ be an homeomorphism between two compact Riemannian manifolds M_1 and M_2 with smooth boundary, and let hks_i be the heat kernel signature induced by Δ_{M_i} . If the eigenvalues of the Laplace-Beltrami operator on each manifold are simple and $hks_1(t, \mathbf{x}) = hks_2(t, \mathcal{T}\mathbf{x})$, $\forall \mathbf{x} \in M_1, \forall t > 0$, then \mathcal{T} is an isometry.*

Although exhibiting notable performance and efficiency among spectral methods for tasks such as segmentation, shape matching and feature localization (Boscaini et al., 2016; Wang & Solomon, 2019; Banduc et al., 2024), the hks is based in a homogeneous phenomenon of propagation. Thus, the functionality of the hks can be extended to problems involving heterogeneity and direction in the propagation of a quantified variable using second order elliptic operators in the general form $L_{\mathbf{A}}u = -\nabla_M \cdot (\mathbf{A} \nabla_M u)$, where $\mathbf{A} = \mathbf{A}(\mathbf{x})$ is a symmetric and uniformly bounded elliptic tensor.

3.3.2. Fibrosis Information Inclusion

To study AF vulnerability from intrinsic geometrical features and conduction properties in the cardiac domain with main focus on fibrosis distribution and fiber direction, a simplification of the monodomain model (M) is proposed.

Considering a Dirac impulse through a fixed point \mathbf{x}_{stim} and neglecting the reaction term I_{ion} , the monodomain system reduces to a single diffusion equation

$$\begin{cases} -\nabla \cdot (\mathbf{G}_m \nabla V_m) + \frac{\partial V_m}{\partial t} = 0 & \text{in } \Omega_H \times (0, T]; \\ (\mathbf{G}_m \nabla V_m) \cdot \mathbf{n} = 0 & \text{on } \partial\Omega_H \times (0, T]; \\ \lim_{t \rightarrow 0^+} V_m(t, \cdot) = \delta(\mathbf{x}_{\text{stim}}, \cdot) & \text{in } \Omega_H; \end{cases}$$

where $\mathbf{G}_m := (\chi C_m)^{-1} \Sigma_m$. Depending on data availability, \mathbf{G}_m can be modeled with a tensor \mathbf{F} inscribing prior knowledge of electrical conductivity through cardiac tissue. Instead of taking the Dirichlet Laplacian, the fks is obtained as a generalization of the hks given by a fibrosis-informed elliptic operator $\mathcal{L}_{\mathbf{F}}u = -\nabla \cdot (\mathbf{F}\nabla u)$ with Neumann boundary conditions in the entire volume Ω_H equipped with the usual cartesian coordinates (Banduc et al., 2024). Assuming a sufficiently regular cardiac geometry, $\mathcal{L}_{\mathbf{F}}$ induces a heat kernel in Ω_H .

Definition 3.2 (Fibrotic Kernel Signature) *Let $\Omega_H \subseteq \mathbb{R}^3$ be a cardiac domain with Lipschitz boundary and $\mathbf{F}(\mathbf{x})$ be a symmetric elliptic conductivity tensor in Ω_H with bounded coefficients encoding fibrosis distribution and fiber direction. The fibrotic kernel signature (fks) is defined as the diagonal argument of the heat kernel $K_{\mathbf{F}} : (0, \infty) \times \overline{\Omega}_H \times \overline{\Omega}_H \rightarrow \mathbb{R}$ induced by the operator $\mathcal{L}_{\mathbf{F}}u = -\nabla \cdot (\mathbf{F}\nabla u)$ with Neumann boundary conditions:*

$$\begin{aligned} fks : (0, \infty) \times \Omega_H &\rightarrow \mathbb{R} \\ (t, \mathbf{x}) &\mapsto K_{\mathbf{F}}(t, \mathbf{x}, \mathbf{x}); \end{aligned}$$

The regularity of the domain allows to decompose its corresponding kernel by the eigenpairs of the differential operator $\mathcal{L}_{\mathbf{F}}$. Since Ω_H is bounded and has a Lipschitz boundary, then $-\mathcal{L}_{\mathbf{F}}$ has compact resolvent and, by the *Spectral Theorem* (Arendt, 2006; Dodziuk, 1981), there exists a sequence $\{\phi_i\}_{i=0}^{\infty}$ of functions in $H^1(\Omega_H, \mathbb{R})$ that form a basis of $L^2(\Omega_H, \mathbb{R})$ and a collection of non-negative real numbers $\{\lambda_i\}_{i=0}^{\infty}$ tending to infinity, such that

$$(E_{\mathbf{F}}) \begin{cases} \mathcal{L}_{\mathbf{F}}\phi_i = \lambda_i\phi_i & \text{in } \Omega_H; \\ (\mathbf{F}\nabla\phi_i) \cdot \mathbf{n} = 0 & \text{on } \partial\Omega_H. \end{cases}$$

Moreover $K_{\mathbf{F}}$ yields the semigroup $\{e^{-t\mathcal{L}_{\mathbf{F}}}\}_{t \geq 0}$ of compact and non-negative operators. Thus, by *Mercer's Theorem* (Dodziuk, 1981; Sun et al., 2009), $K_{\mathbf{F}}$ allows the following representation:

$$K_{\mathbf{F}}(t, \mathbf{x}, \mathbf{y}) = \sum_{i=0}^{\infty} e^{-t\lambda_i} \phi_i(\mathbf{x}) \phi_i(\mathbf{y});$$

where the series converges uniformly and absolutely. Then, the fks admits the eigendecomposition

$$fks(t, \mathbf{x}) = \sum_{i=0}^{\infty} e^{-t\lambda_i} \phi_i(\mathbf{x})^2.$$

3.3.3. Signature Computation with Finite Elements

Numerically, the problem $(E_{\mathbf{F}})$ is solved with the \mathbb{P}_1 Finite Element method.

Let Ω_h be a polyhedral approximation of Ω_H and \mathcal{T}_h a tetrahedral mesh of $\overline{\Omega}_h$. The finite elements of order 1 associated with \mathcal{T}_h are defined by the discrete space (Allaire, 2007)

$$V_h = \left\{ v \in \mathcal{C}(\overline{\Omega}_h, \mathbb{R}) : v|_T \in \mathbb{P}_1, \forall T \in \mathcal{T}_h \right\}.$$

Since the integral formulation of $(E_{\mathbf{F}})$ is given by

$$\int_{\Omega_H} (\mathbf{F}\nabla\varphi_i) \cdot \nabla\psi = \lambda_i \int_{\Omega_H} \varphi_i\psi, \quad \forall \psi \in H^1(\Omega_H, \mathbb{R});$$

then the internal approximation of the problem turns to looking for a collection of eigenpairs $\{(\lambda_i, \varphi_i)\}_{i=0}^{\dim V_h - 1}$ in $\mathbb{R} \times V_h$ such that $0 \leq \lambda_0 \leq \lambda_1 \leq \dots \leq \lambda_{\dim V_h - 1}$ and $\{\varphi_i\}_{i=0}^{\dim V_h - 1}$ is a basis of V_h orthonormal in $L^2(\Omega_h, \mathbb{R})$ (Allaire, 2007) satisfying

$$(E_{\mathbf{F}}^h) \begin{cases} \varphi_i = \sum_{k=0}^{\dim V_h - 1} \alpha_i^k \psi_k, & \{\alpha_i^k\}_{k=0}^{\dim V_h - 1} \subseteq \mathbb{R}; \\ \mathbf{S}_h \vec{\alpha}_i = \lambda_i \mathbf{M}_h \vec{\alpha}_i, & \vec{\alpha}_i := (\alpha_i^0, \alpha_i^1, \dots, \alpha_i^{\dim V_h - 1})^T. \end{cases}$$

Here, $\{\psi_k\}_{k=0}^{\dim V_h - 1}$ is the finite element basis of V_h . \mathbf{S}_h and \mathbf{M}_h denote, respectively, the *stiffness matrix* and the *mass matrix* of the discrete problem, defined as

$$(\mathbf{S}_h)_{ij} := \int_{\Omega_h} (\mathbf{F}\nabla\psi_i) \cdot \nabla\psi_j; \quad (\mathbf{M}_h)_{ij} := \int_{\Omega_h} \psi_i\psi_j.$$

Then, the fks can be computed in an atrial mesh provided the solution $\{(\lambda_i, \varphi_i)\}_{i=0}^{\dim V_h - 1}$ of the eigenproblem $(E_{\mathbf{F}}^h)$ through the estimate

$$fks(\mathbf{x}, t) \simeq \sum_{i=0}^n e^{-\lambda_i t} \varphi_i(\mathbf{x})^2, \quad 1 \ll n \ll \dim V_h - 1.$$

For the experiments performed in this work, the discrete spectral problem $(E_{\mathbf{F}}^h)$ was solved using the **FEniCSx** software (Baratta et al., 2023; Scroggs, Dokken, et al., 2022; Scroggs, Baratta, et al., 2022; Alnaes et al., 2014) along with the **SLEPc** library (Hernandez et al., 2005).

3.3.4. Method Workflow for AF Prediction

The problem of predicting which points in the atria could potentially trigger AF events can be handled in a supervised learning setting by solving a classification task. Let $\mathbf{1}_{AF}(\mathbf{p}, \mathcal{M}) : \Omega_H \rightarrow \{0, 1\}$ be an indicator function for a given cardiac electrophysiology model \mathcal{M} , where $\mathbf{1}_{AF}(\mathbf{p}, \mathcal{M})(\mathbf{x}_{\text{stim}}) = 1$ if pacing from $\mathbf{x}_{\text{stim}} \in \Omega_H$ with a protocol \mathbf{p} can develop a sustained fibrillation episode according to the model \mathcal{M} , and $\mathbf{1}_{AF}(\mathbf{p}, \mathcal{M})(\mathbf{x}_{\text{stim}}) = 0$ otherwise. For instance, \mathbf{p} could be set as the PEERP protocol, while \mathcal{M} could correspond to the parameters of the monodomain model (M) (see Figure 3.7). Typically, pace-points are positioned on a smooth surface inside of Ω_H .




Figure 3.7: Example of inducibility criterion $\mathbf{1}_{AF}$ for PEERP protocol and monodomain model.

The function $\mathbf{1}_{AF}$ not only depends on the atrial geometry, but also on the pacing strategy, the criteria applied for categorizing arrhythmic behavior and the parameters necessary for solving the equations given by the model. Thus, a formal definition can be given for point inducibility in terms of AF testing:

Definition 3.3 (Inducible Point) *Let Ω_H be an atrial domain. For fixed \mathbf{p}, \mathcal{M} , let $\mathbf{1}_{AF}(\mathbf{p}, \mathcal{M}) : \Omega_H \rightarrow \{0, 1\}$ be as described above. A point $\mathbf{x}_{stim} \in \Omega_H$ is said to be inducible if*

$$\mathbf{1}_{AF}(\mathbf{p}, \mathcal{M})(\mathbf{x}_{stim}) = 1.$$

This brings the following definition for surface inducibility:

Definition 3.4 (Overall Inducibility) *For a pacing protocol \mathbf{p} and electrophysiology model \mathcal{M} , overall inducibility S_{AF} is defined as the fraction of tissue where AF can be induced:*

$$S_{AF} := \frac{1}{|\Omega_H|} \int_{\Omega_H} \mathbf{1}_{AF}(\mathbf{p}, \mathcal{M})(\mathbf{x}) d\mathbf{x}$$

The goal of classification is to learn the set of AF-inducible stimulation points $\Omega_H^{AF} := \mathbf{1}_{AF}(\mathbf{p}, \mathcal{M})^{-1}(\{1\})$ and, more specifically, approximating the overall inducibility of the model, because ablation strategies look for minimizing the inducible surface. Here, the objects of interest to account for are the atrial architecture, fibrosis distribution and fiber direction.

The current approach to estimate the inducible set Ω_H^{AF} is to run monodomain-based simulations at multiple stimulation points evenly distributed on the atrial domain (Banduc et al., 2024). Ideally, for a given amount of locations $P \in \mathbb{N}$, the family of pacing points $X_{stim} := \{\mathbf{x}_{stim}^i\}_{i=1}^P$ maximizes the pairwise geodesical distance d on a mid-wall smooth surface $\mathcal{S}_H \subseteq \Omega_H$:

$$\min_{\substack{i,j=1,\dots,P \\ i \neq j}} d(\mathbf{x}_{stim}^i, \mathbf{x}_{stim}^j) = \max_{\substack{X \subseteq \mathcal{S}_H \\ |X|=P}} \min_{\substack{\mathbf{x}, \mathbf{y} \in X \\ \mathbf{x} \neq \mathbf{y}}} d(\mathbf{x}, \mathbf{y}).$$

In practice, a *farthest point sampler* algorithm is employed to representatively select the pacing sites, which iteratively incorporates the farthest element from the current set until matching the desired set size. Usually $P = 100$, but given the convergence of overall inducibility for an increasing number of pacing locations, an initial speed-up can be done by

only considering a sample of 40 simulations, as suggested in Boyle et al., 2019.

The main idea of this work is to learn the inducibility function $\mathbf{1}_{AF}$ with the fks time series from a finite set of simulations for a selected pair $(\mathbf{p}, \mathcal{M})$, such that for new cardiac anatomies or geometry changes due to ablation, the inducible set can be estimated quickly, without the need to re-run simulations (Banduc et al., 2024). Here, a classifier $C(\cdot)$ is trained with the features corresponding to the fks on a discrete collection of times $T \subseteq (0, \infty)$ to predict the outputted tags $\{\mathbf{1}_{AF}(\mathbf{x}_{\text{stim}}^i)\}_{i=1}^P$ from the model:

$$C\left(\left(fks(t, \mathbf{x}_{\text{stim}}^i)\right)_{t \in T}\right) \simeq \mathbf{1}_{AF}(\mathbf{x}_{\text{stim}}^i), \quad \forall i \in \{1, \dots, P\}.$$

Although C learns inducibility targets from the fks evaluated in the subset X_{stim} , the signature as a function is independent of the pacing points chosen for simulation, thus allowing to compute $\mathbf{1}_{AF}$ in the entire atrial geometry.

After training with several cardiac cases and fibrosis scenarios, the fks is expected to accurately estimate the inducibility map $\mathbf{1}_{AF}$ for a new domain $\Omega_{H'}$:

$$C\left(\left(fks(t, \mathbf{x})\right)_{t \in T}\right) \simeq \mathbf{1}_{AF}(\mathbf{x}), \quad \forall \mathbf{x} \in \Omega_{H'}.$$

With the objective to enable comparability between the fks across different atrial geometries, it is necessary to scale the signature and define an interval that describes its stages of evolution.

Following the heuristics suggested in Sun et al. (2009), for the discrete setting Ω_h , the fks is divided by an estimate of the heat trace:

$$fks(t, x) \leftarrow \frac{\sum_{i=0}^n e^{-t\lambda_i} \varphi_i(\mathbf{x})^2}{\sum_{i=0}^n e^{-t\lambda_i}}, \quad 1 \ll n \ll \dim V_h - 1.$$

The idea of normalization comes from the exponential decrease in the difference between signatures from distinct locations, making spatial variations negligible at larger times when compared to smaller ones, thus becoming necessary to scale these variations to give uniform significance to the difference between signatures through all times.

Taking into account that $\lambda_0 = 0$ with constant function $\varphi_0 = |\Omega_H|^{-1} \mathbf{1}_{\Omega_H}$ for the eigenproblem $(E_{\mathbf{F}})$, the summations are in practice carried out starting from the smallest non-negative eigenvalue (λ_1), denoted as λ_{\min} in what follows. Then, the discrete set of times T to evaluate the signature is defined by a logarithmic progression $\{t_0, \dots, t_{N_T}\}$ in the interval $[t_{\min}, t_{\max}]$, where the endpoints of the set are determined by λ_{\min} and the largest eigenvalue (λ_{\max}) computed for the differential operator inducing the kernel:

$$t_{\min} = \alpha / \lambda_{\max}; \quad t_{\max} = \beta / \lambda_{\min};$$

with α, β positive hyper-parameters.

The notion of a domain-dependent time interval discretized with a logarithmic progression to evaluate the signature arises from the requirement of establishing consistent comparison between domains, where the signature must be evaluated at characteristic stages of diffusion. To achieve this, the pronounced changes of the fks for initial times and long-term stability are captured with an even spacing on logarithmic scale, and the relevant interval of observation is controlled through the terms $e^{-t_{\min}\lambda_{\max}}$ and $e^{-t_{\max}\lambda_{\min}}$ of the series. Here, t_{\min} marks the beginning of measurements by fixing the minimal value of each exponential component of the fks expansion at $e^{-\alpha}$, while t_{\max} indicates the endpoint for observation at the instant the signature decreases to a value in the order of $e^{-\beta}$.

Finally, the pipeline for the fks method can be summarized in two phases, which are schematically shown in Figure 3.8:

- *Offline phase:* On this phase, a set of meshed anatomies with different fibrosis scenarios $\{\Omega_h^i\}_{i=1}^{N_H}$ is used for training. For each polyhedral domain Ω_h^i , a set of pacing points $\{\mathbf{x}_j^i\}_{j=1}^{P_i}$ is extracted from the node set of its mesh and computationally tested for arrhythmia inducibility using a specified protocol and electrical model of the heart, obtaining a set of binary tags $\{y_j^i\}_{j=1}^{P_i}$, where $y_j^i = 1$ if pacing from $\mathbf{x}_j^i \in \Omega_h^i$ degrades into a fibrillation episode and $y_j^i = 0$ if not. On a separate process, a scaled fibrotic kernel signature fks_i is computed for a user-specified tensor \mathbf{F}_i for a logarithmic progression of times $\{t_k^i\}_{k=1}^{N_T}$ as previously described, where N_T is fixed across all cases. Then, the signature is sub-sampled for each collection of pacing locations and paired with the AF inducibility outcomes, creating a fks-inducibility dataset

$$\mathcal{D} := \bigcup_{i=1}^{N_H} \left\{ \left(\left(fks_i(t_k^i, \mathbf{x}_j^i) \right)_{k=1}^{N_T}, y_j^i \right) \right\}_{j=1}^{P_i} \subseteq \mathbb{R}^{N_T} \times \{0, 1\}.$$

Once \mathcal{D} is assembled, a simple machine learning classifier $C(\cdot)$ is trained with all the data available. The costliest part of the offline phase in terms of computational time and memory comes from the monodomain simulations, because its associated system must be solved individually for each pacing point used for the construction of \mathcal{D} .

- *Online phase:* On this stage, a new meshed anatomy $\Omega_h^{N_H+1}$ is introduced. The fks is computed for all of its points and for the times $\{t_k^{N_H+1}\}_{k=1}^{N_T}$ as in the previous phase, outputting the set

$$\left\{ \left(fks(t_k^{N_H+1}, \mathbf{x}) \right)_{k=1}^{N_T} \right\}_{\mathbf{x} \in \Omega_h^{N_H+1}} \subseteq \mathbb{R}^{N_T}.$$

Then, an estimate $\mathbf{1}_{AF}^C$ of the inducibility map $\mathbf{1}_{AF}$ is defined by using this data as input:

$$\mathbf{1}_{AF}^C(\mathbf{x}) := C \left(\left(fks(t_k^{N_H+1}, \mathbf{x}) \right)_{k=1}^{N_T} \right), \quad \forall \mathbf{x} \in \Omega_h^{N_H+1}.$$

With this, overall inducibility is approximated through the integral over all points in

the mesh:

$$S_{AF} \simeq \frac{1}{\text{Vol}(\Omega_h^{N_H+1})} \int_{\Omega_h^{N_H+1}} \mathbf{1}_{AF}^C(\mathbf{x}).$$

For evaluating the performance of the classifier, a similar procedure to the offline phase can be followed to create a testing set with AF inducibility outcomes for the new case.

Since evaluating the classifier is inexpensive, the major complexity in time and memory usage of the online phase comes with the fks computation. As reported in Banduc et al., 2024, the signature can be up to 40- or 100-fold faster to estimate overall inducibility than a typical monodomain-based method.

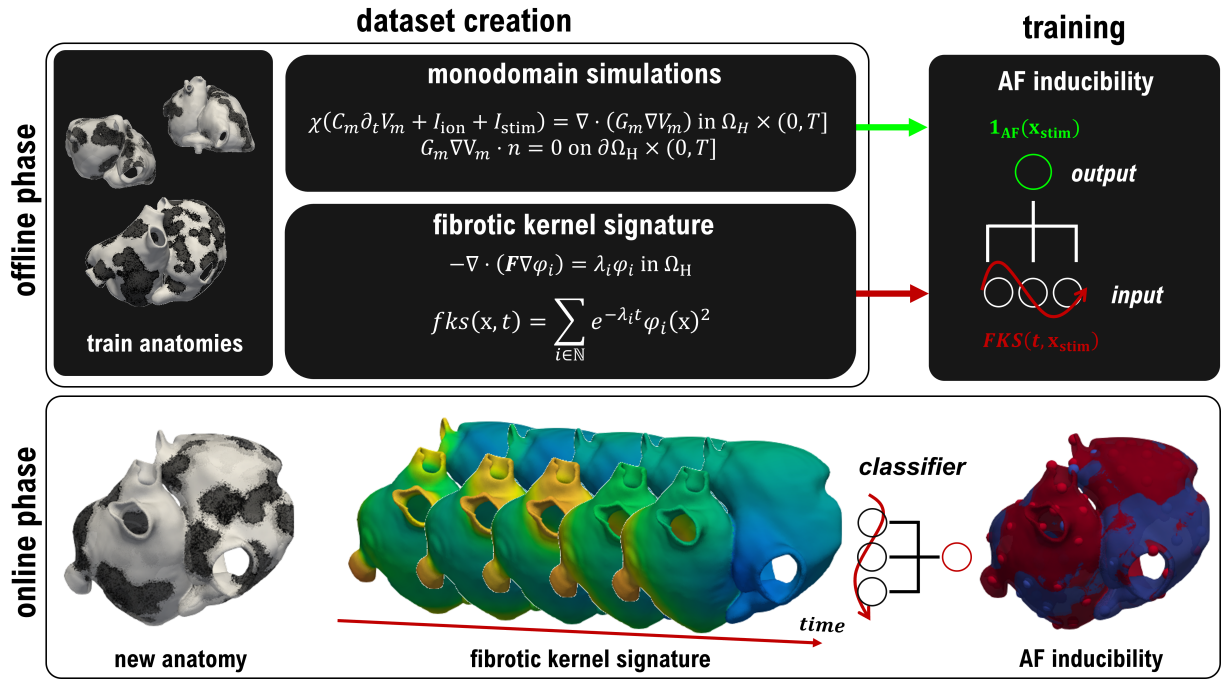


Figure 3.8: fks method pipeline. In the offline phase (top), classifiers are trained with the fks features to predict monodomain inducibility tags. Then, in the online phase (bottom), the fks is computed for a new geometry and passed through the classifier to generate an inducibility map. Figure adapted from Banduc et al. (2024).

Chapter 4

Study 1

4.1. Introduction

In this study, the method for AF inducibility in patient-specific cardiac models with the fibrotic kernel signature is proposed (Sahli Costabal et al., 2023). The fks is deduced from the diffusion component of the monodomain model, where the conductivity tensor is reduced to a function accounting for tissue heterogeneity caused by fibrosis. The signature is then introduced as an extension of the hks, encoding both the intrinsic geometry of the atrial region of the heart and the spread of electrical potential by incorporating conductivity changes in presence of fibrotic tissue into the elliptic operator inducing the heat kernel.

The relationship between the proposed signature and variability in AF vulnerability generated by ablation lines and scar tissue was learned from testing in a supervised setting 9 cases that combined different ablation strategies and fibrosis distributions for a single atrial geometry. In general, the fks does not require re-training for new anatomies, fibrosis patterns, and ablation lines, and is fast to compute when compared to standard models like the monodomain system. In the classification task, maximum balanced accuracy ranging from 75.8 to 95.8% with simple machine learning classifiers was achieved when tested on single points, and overall inducibility of each model was predicted with small error (Sahli Costabal et al., 2023).

4.2. Methods

4.2.1. Atrial Modeling

This study utilized the simulations and inducibility outcomes from Gander et al. (2022), where a previously developed human atrial fibrillation model on a volumetric domain with variable wall thickness was employed (Gharaviri et al., 2020; Potse et al., 2018). Three cases of endomysial fibrosis were considered, one moderate (m), with fibrosis covering 50% of the atria, and two severe (s_1, s_2), with fibrosis at 70%. For each fibrosis pattern, three ablation scenarios were surveyed: no ablation (NA), pulmonary vein isolation (PVI) and PVI with posterior box isolation (PVI+BOX). Since scarring may affect the structure of the atria in the cauterized region and fibrosis can induce cardiac remodeling (Dzeshka et al., 2015), each fibrosis distribution $a \in \{m, s_1, s_2\}$ and ablation case $f \in \{\text{NA}, \text{PVI}, \text{PVI+BOX}\}$ was assigned a specific geometry Ω_a^f . These models are exhibited in Figure 4.1.

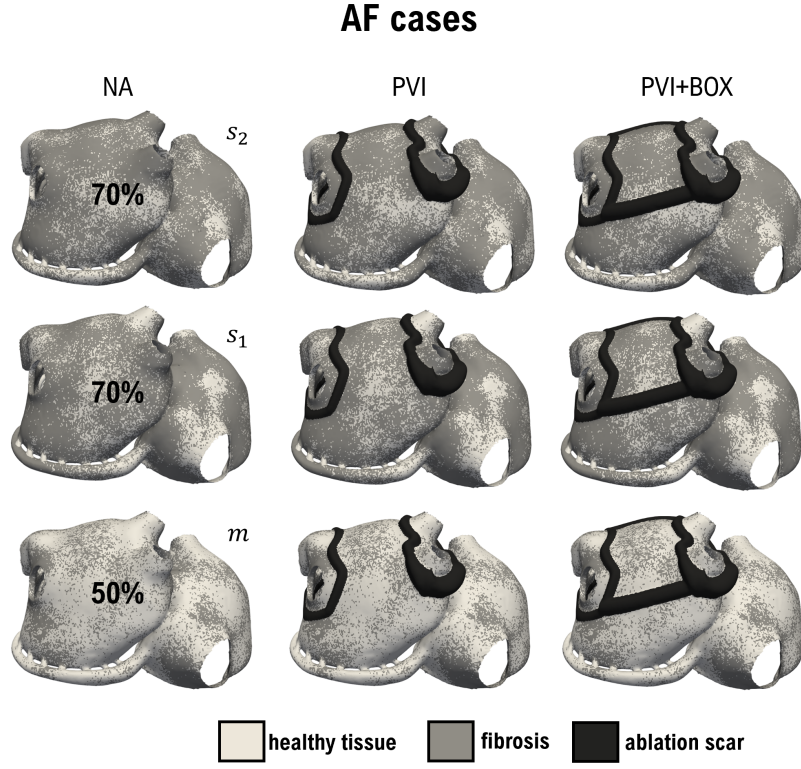


Figure 4.1: Atrial models considered in first study. Different shades of gray are used to mark healthy tissue, fibrosis and ablation scars.

For the simulation of AF episodes, the monodomain model (M) of the heart was used (Sahli Costabal et al., 2023):

$$\begin{cases} -\nabla \cdot (\Sigma_m \nabla V_m) = -\chi \left(C_m \frac{\partial V_m}{\partial t} + I_{\text{ion}}(V_m, \mathbf{w}) + I_{\text{stim}}(\mathbf{x}, t) \right) & \text{in } \Omega_a^f \times (0, T], \\ \frac{\partial \mathbf{w}}{\partial t} = \mathbf{g}(V_m, \mathbf{w}) & \text{in } \Omega_a^f \times (0, T]. \end{cases}$$

The membrane capacitance C_m was set to $1\mu\text{Fcm}^{-2}$ and the surface-to-volume ratio χ was constant and equal to 800cm^{-1} . The Courtemanche-Ramirez-Nattel model (Courtemanche et al., 1998) was considered for the nonlinear ionic current I_{ion} across the membrane and the stimulus I_{stim} corresponded to a rapid pacing protocol comprising 14 stimuli of strength $800\mu\text{Acm}^{-2}$ and duration 4ms (Gander et al., 2022), where inducibility was checked roughly 20ms after the end of the protocol. Here, for each pacing point $\mathbf{x}_{\text{stim}} \in \Omega_a^f$, $\mathbf{1}_{AF}(\mathbf{x}_{\text{stim}}) = 0$ if both the transmembrane voltage V_m and state vector \mathbf{w} approached asymptotically to their resting states V_0, \mathbf{w}_0 ; and $\mathbf{1}_{AF}(\mathbf{x}_{\text{stim}}) = 1$ otherwise. A unique set of 100 pacing locations was selected, such that all the stimulation points were common to all geometries, totaling 900 AF inducibility outputs.

The numerical solutions of the system were computed with the Propag-5 software on the Swiss National Supercomputing Centre (CSCS), using a second order difference scheme in space (see Clayton & Panfilov, 2008) and an explicit first-order Euler scheme for time stepping (Potse et al., 2006; Krause et al., 2012; as cited in Gander et al., 2022).

4.2.2. Fibrotic Kernel Signature in the Atria

Inspired by the diffusion operator of the monodomain equation and taking into consideration that fibrosis is the main structural driver of AF, an isotropic elliptic operator $\mathcal{L}_{\text{iso}}u = -\nabla \cdot (\mathbf{F}_{\text{iso}}(\mathbf{x})\nabla u)$ was taken to describe passive diffusion of electrical potential in the geometry to characterize the chance of triggering an arrhythmia (Sahli Costabal et al., 2023). The fibrosis-informed tensor \mathbf{F}_{iso} was of the form

$$\mathbf{F}_{\text{iso}}(\mathbf{x}) = \sigma_{\text{fib}}(\mathbf{x})\mathbf{I},$$

where σ_{fib} is a positive uniformly bounded scalar function in $\overline{\Omega}_a^f$ and \mathbf{I} is the 3-dimensional identity matrix. In this work, $\sigma_{\text{fib}}(\mathbf{x})$ equals 1 if \mathbf{x} belongs to healthy tissue, 0.5 if the point is in a fibrotic region and 0.01 if \mathbf{x} belongs to an ablation scar. The function was estimated with discontinuous Galerkin elements.

To obtain the fks, the geometry was normalized and the first 100 non-constant eigenfunctions were computed, sorted by eigenvalue magnitude. For each fibrosis pattern f and ablation strategy a , the eigenproblem was solved with \mathbb{P}^1 elements on the tetrahedral mesh of $\overline{\Omega}_a^f$, which had around 700,000 nodes. The logarithmic time progression of the signature was constructed as suggested in Section 3.3.4, with α, β set to $\ln(10^4)$. To accentuate the effect of ablation lesions, for each fibrosis distribution, the signature of each ablation case was scaled by the heat trace of the ablation-free model (Sahli Costabal et al., 2023).

4.2.3. Prediction of Atrial Fibrillation

Once the fks was computed across all cardiac domains, the subset of signatures associated with the pacing locations was sampled and paired with the inducibility tags. Since the dataset was conformed by a single geometry in 9 distinct scenarios and evaluated at the same locations, the performance of the method was assessed with a leave-one-out by case cross-validation technique (Sahli Costabal et al., 2023). First, classifiers were trained with 8 out of 9 configurations of fibrosis and ablation, adding up to 800 data points, and then tested with the remaining case, with 100 unseen fks-inducibility results (see Figure 4.2).

Three classifiers were considered: k-nearest neighbors (KNN), random forest (RF) and gradient boosting (GB), as implemented in `scikit-learn` (Pedregosa et al., 2011; as cited in Sahli Costabal et al., 2023). Additionally, a majority voting naïve classifier was implemented to evaluate whether classifiers were predicting based on point location or learning new information from the signature (Sahli Costabal et al., 2023), with a tie-breaking criterion in favor of the negative class. For example, if 5 out of 8 models were inducible at a given pacing location, the point was assigned $y = 1$, and if 4 out of 8 were inducible, then the given tag was $y = 0$.

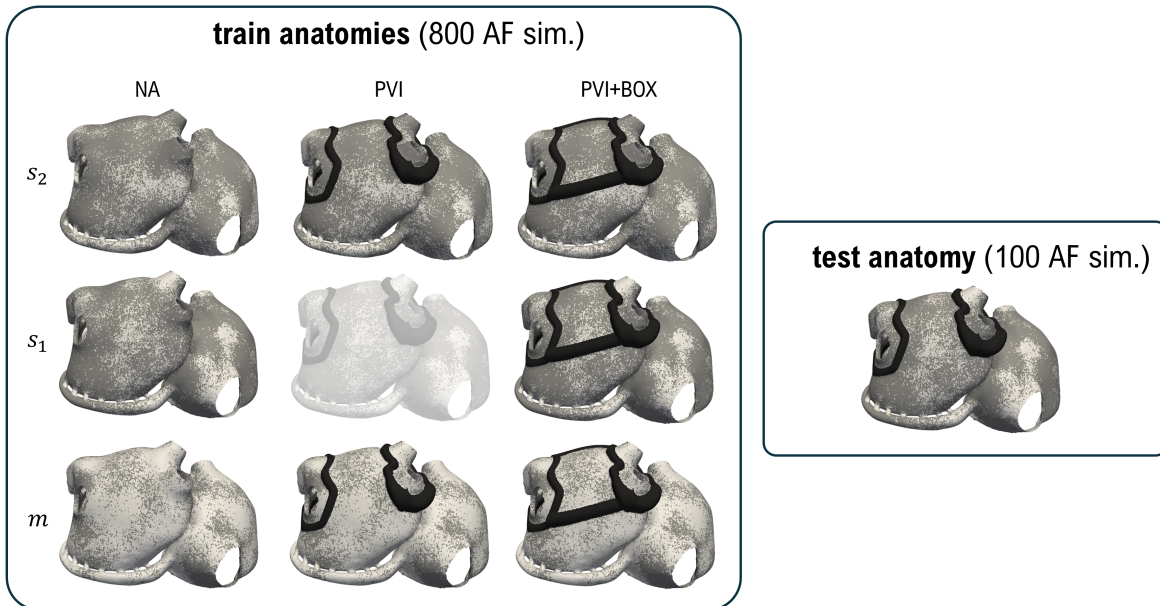


Figure 4.2: Instance of leave-one-out by case technique. A sample of size 100 is taken from each configuration (adding up to 800 samples), excluding the case with s_1 fibrosis distribution and PVI ablation, which is set aside for testing.

Two metrics were used to measure the performance of the naïve classification and fks-based methods: balanced accuracy (Mosley, 2013), which is the average of sensitivity and specificity, and overall inducibility, which was estimated by the ratio of inducible points predicted on each sample. The latter is the most important to prove effectiveness in a proposed ablation treatment (Sahli Costabal et al., 2023).

4.3. Results

The computation of the fks in the entire geometry was reported to have taken approximately 10 minutes on a modern workstation (Apple M2 Ultra, with 24-core CPU and 64GB unified memory), while running one monodomain simulation from a single pacing site took 1 hour on a GPU node at the CSCS. Disregarding the hardware differences and discounting the simulations needed for training, this represents a 600-fold speed-up for the prediction of AF inducibility (Sahli Costabal et al., 2023).

Different profiles of the fks over time are shown in Figure 4.3 for 20 randomly selected points across all cases. It can be noticed that non-inducible pacing locations generally have higher signature than inducible points, and the application of ablation lines modifies the signature more significantly at the late stages of the diffusion process. This is complemented from what is observed in Figure 4.5. The fks tends to rise in the regions enclosed by the lines of the PVI and PVI+BOX lesions, indicating electrical accumulation in the parts of the domain where originally there was a free flow of current and post-procedural non-inducibility is expected. PVI predominantly affects the signature in the area of the pulmonary veins, and including BOX ablation extends the observed electrical barrier effect towards the rest of the geometry in the roof of the left atrium (Sahli Costabal et al., 2023).




Figure 4.3: fks over time for 20 randomly selected points for 9 different cases: moderate fibrosis (m), severe fibrosis - case 1 (s_1) and severe fibrosis - case 2 (s_2). The red curves represent points where AF was induced and the blue curves are points where AF was not induced. Figure reprinted from Sahli Costabal et al. (2023).




Figure 4.4: Overall inducibility prediction for leave-one-out by case cross-validation with RF. Figure reprinted from Sahli Costabal et al. (2023).

Results regarding the predictions of AF are summarized in Figures 4.4 and 4.6. When analyzing data efficiency for balanced accuracy score in Figure 4.6, RF performs better than KNN and GB in most cases. The accuracy obtained with the fks-trained classifiers ranges from 75.8% for the severe scenario s_1 with no ablation to 95.8% for the moderate case m with PVI+BOX. The naïve classifier shows similar performance than the rest of the classifiers. However, for some instances, the classifiers trained with the signature have higher or lower accuracy than the majority voting algorithm, indicating that they are not simply memorizing the locations of the training points and matching them to the test cases.

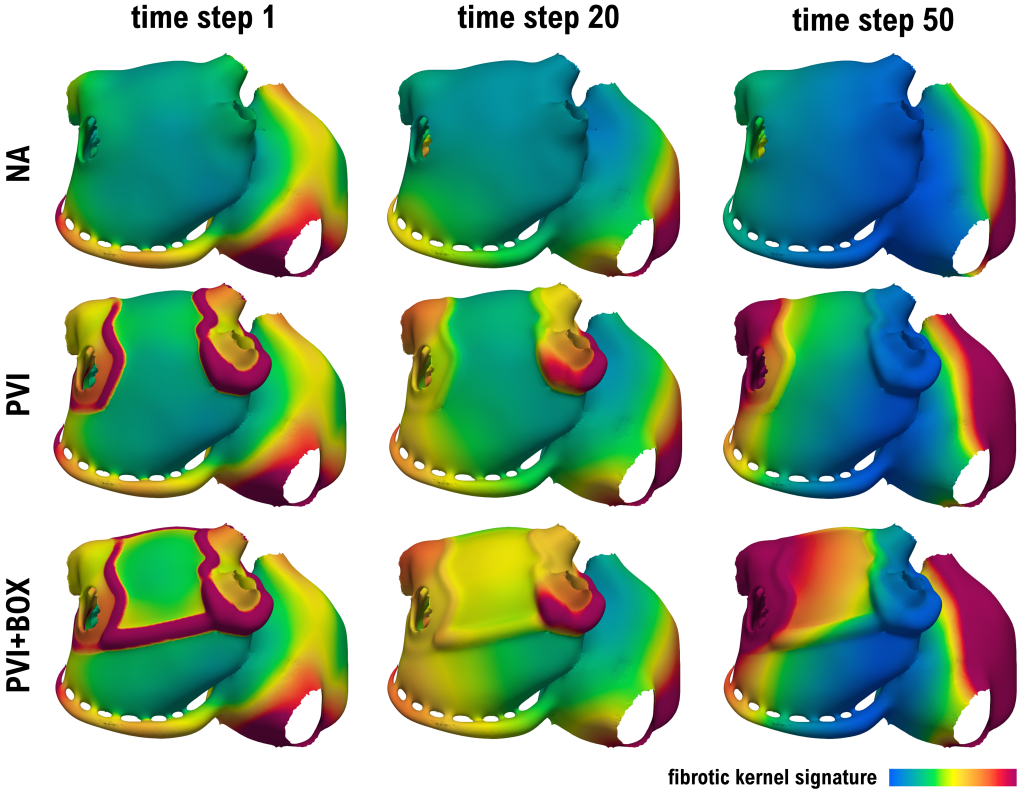


Figure 4.5: Three time steps for fks map in moderate fibrosis case under three ablation patterns: no ablation (NA), pulmonary vein isolation (PVI) and PVI with box enclosure (PVI+BOX).

In Figure 4.4 substantial differences are displayed for overall inducibility prediction when ablation lines were applied to each fibrosis model, showing a marked reduction in tissue prone to AF. PVI decreased inducibility in the range of 11 to 14 percent points, while PVI+BOX reached reductions between 16 and 19 percent points. As expected, the majority voting classifier predicted the same level of inducibility independent of the treatment used, obtaining values between 40 and 43 percent points. The RF trained with the fks on 800 samples could detect overall inducibility variation with higher fidelity, obtaining errors from 1 to 6 percent points. Also, the trend that the NA case was more inducible than PVI and that, in turn, PVI was more inducible than PVI+BOX was also correctly captured by the fks-based classifier.

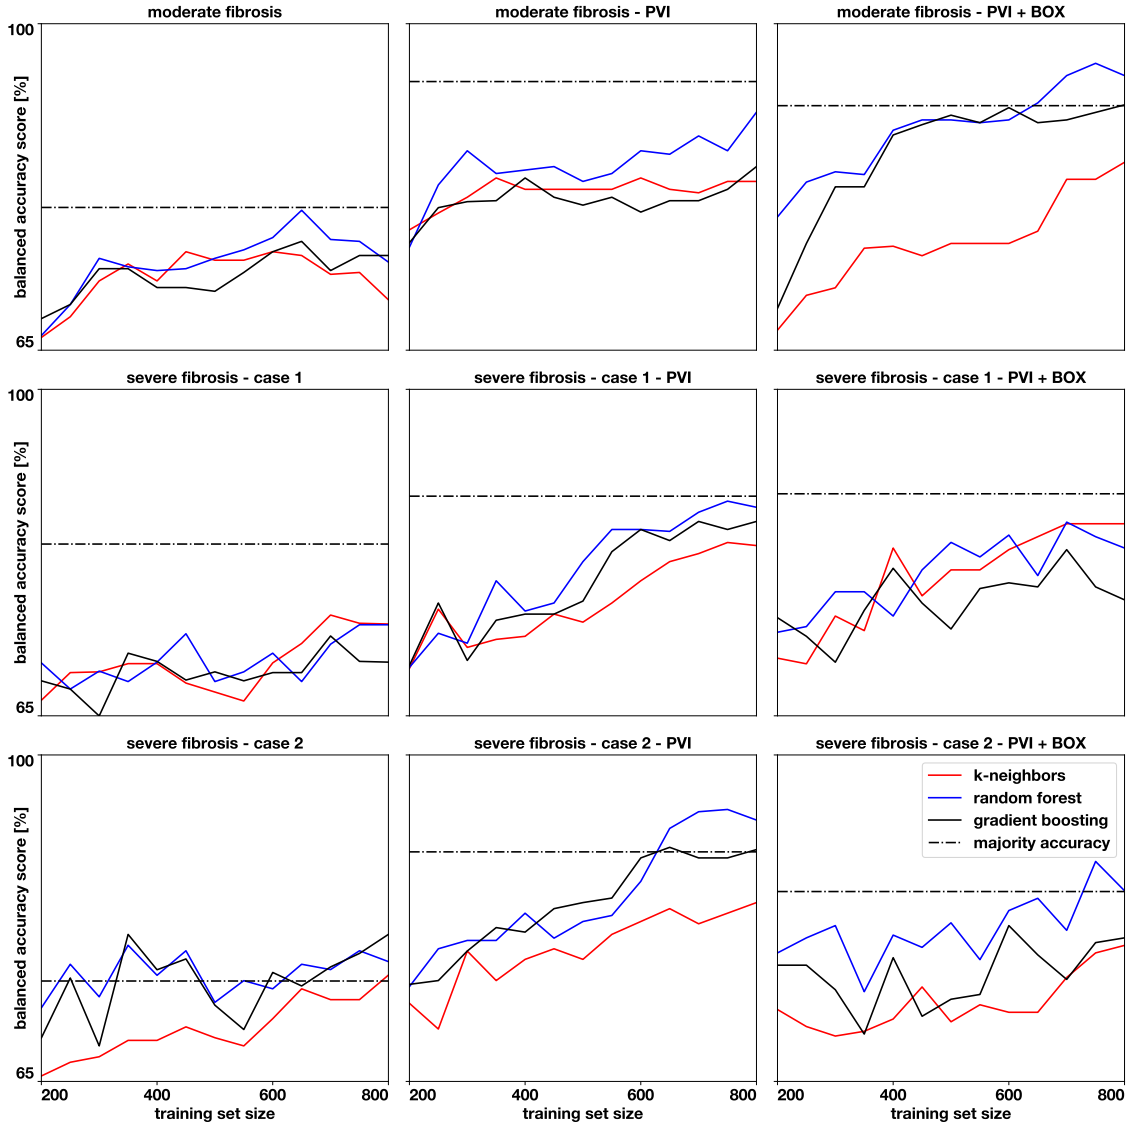


Figure 4.6: Data efficiency for balanced accuracy across all cases following a leave-on-out by case strategy. The solid lines show the performance of k-nearest neighbors, random forest and gradient boosting as the training data increases. The dashed line corresponds to the naïve approach. Figure reprinted from Sahli Costabal et al. (2023).

4.4. Discussion and Conclusions

In this study, an efficient simulation-free method to accurately predict the inducibility of AF was presented: the fibrotic kernel signature (fks). This was achieved by creating a time-dependent descriptor of fibrosis distribution and cardiac architecture, rooted in the passive diffusion process of electric potential spread across the atria. The fks can be computed for all points in a geometry at a fraction of the time spent in running the monodomain simulations typically used for testing AF vulnerability, speeding-up the calculation of inducibility maps, which is crucial to plan personalized ablation treatments within the time constraints of the clinical setting (Azzolin et al., 2021; Boyle et al., 2019; as cited in Sahli Costabal et al., 2023). The fks prediction requires only the essentials in terms of computing resources: a desktop

computer is sufficient to run the signature in a reasonably fine mesh of the cardiac structure.

Paired with simple machine learning algorithms, the fks is capable to estimate pointwise inducibility and approximate the proportion of surface prone to AF with great accuracy. As seen in Figure 4.6, the accuracy of classifiers trained with the fks tends to increase with data availability. Thus, it is expected that performance could be improved with a larger dataset for training. The fks method allows to take advantage of previously ran simulations in different anatomies in order to strengthen the predictions of the signature. Additionally, no mapping between distinct anatomies is required when evaluating the fks, since the geometrical information flows into the fks, on which the classifier depends directly. More importantly, the fks is easily extendable thus to include local fiber direction, just by redefining the fibrosis-informed elliptic operator inducing the kernel. For instance, the fiber direction and fibrosis, combined in the conductivity tensor, can be estimated from electro-anatomical mapping system data, as recently proposed (Grandits et al., 2021; Kotadia et al., 2020; Ruiz Herrera et al., 2022; as cited in Sahli Costabal et al., 2023).

Although the method introduced shows promising results in terms of computational efficiency to elaborate effective ablation treatments for atrial fibrillation, this study has some limitations. First, despite the variety of fibrosis patterns and ablation strategies changing the geometry and electrical behavior of both the monodomain and the fks models, the point descriptor used was only tested in a single anatomy, and the pacing locations were all the same, making necessary to verify whether the signature generalizes to arbitrary pacing locations and other patient cases. Numerical experiments did not consider fiber orientation nor cell type distribution into the fks construction. A preliminary study including myocardial fibers was conducted, but there were no conclusive improvements in the results. This could be explained by the regularizing effect of the diffusion model, which could attenuate the influence of local fiber direction when restricting measurements of temperature to a single point. Also, the mesh used had a resolution of 0.4mm, which possibly affected inducibility outcomes when compared to finer meshes (Gander et al., 2022; as cited in Sahli Costabal et al., 2023). Refining the mesh will affect the time to solve the eigenproblem required by the fks decomposition, but given that only modest hardware was used, larger models could be managed with more computational resources. On the other hand, the current mesh resolution was in line or already finer than those from other studies (Gander et al., 2022; as cited in Sahli Costabal et al., 2023). Finally, uncertainty in the fibrosis pattern was not accounted for, which is known to be highly affected by the threshold strategy. Ideally, uncertainty could be introduced into the fks definition as an average of the original fks with an associated covariance. This information then could be used in the classification problem.

In summary, this study proposes a novel method to predict atrial fibrillation without running simulations. The fibrotic kernel signature time-series combined with machine learning techniques could enable faster and better planning of therapeutic treatment in atrial fibrillation patients.

Chapter 5

Study 2

5.1. Introduction

This study is an extension of the previous work done in Sahli Costabal et al., 2023, and reviewed in Chapter 4. Here, the motivation and concept of fks is revisited as a simplification of the monodomain system that arises in cardiac electrophysiology to a single diffusion model, characterized by the heat kernel of an elliptic operator encoding fibrosis patterns and other cardiac structures. The further spatial restriction of the resulting kernel and observation of key stages of decay in electrical potential outputs a time-dependent point descriptor that captures atrial architecture and conductivity variations across the cardiac muscle at multiple scales, building up to a method that can lift the burden of monodomain-based simulations to test AF inducibility by reusing information from observed patient models to rapidly detect substrate for arrhythmia.

In comparison to the previous study, several enhancements are introduced to guarantee the robustness, reproducibility and fidelity of the fks method. First, the number of monodomain simulations performed was more than double the size of the original set, constituting one of the largest studies in terms of AF simulations. Three different anatomies were considered for training, as opposed to the only one from the previous study. Furthermore, pacing points were relocated for different fibrosis scenarios when the same anatomy was utilized, and the methodology was tested on three unseen geometries. Finally, the pacing protocol was changed to the PEERP strategy and simulations were ran on an open cardiac electrophysiology simulator to facilitate cross-validation and methodology adoption and refinement by others (Banduc et al., 2024).

The fks computations were still carried out by a modern workstation, but remained faster than the monodomain solver executed in the CSCS. In the classification problem, high precision and accuracy scores were obtained and overall inducibility was estimated with a mean absolute error of 2.76 percent points (Sahli Costabal et al., 2023).

5.2. Methods

5.2.1. Atrial Modeling

As in the previous study, the electrical activity of the atria was modeled with the monodomain equation (M) in volumetric cardiac domains. The values of transmembrane capacitance C_m and surface-to-volume ratio χ were set to $1\mu\text{Fcm}^{-2}$ and 800cm^{-1} , respectively. The ionic current $I_{\text{ion}}(V_m, \mathbf{w})$ was defined by a variation of the Courtemanche-Ramirez-Nattel model (Courtemanche et al., 1998) reflecting AF-induced remodeling (Loewe et al., 2014; as cited in Banduc et al., 2024), and the stimulus term $I_{\text{stim}}(\mathbf{x}, t)$ was changed to a PEERP protocol (see Section 3.2.3), where the S2 stimulus is adaptively selected at the end of the effective refractory period for a fixed current strength and stimulation region.

A transverse isotropy regime was imposed for the atrial model, as done in Section 3.1.3. Thus, the conductivity tensor Σ_m had an explicit form depending only on fiber direction \mathbf{a}_l and the fiber and cross-fiber conductivities of the bulk medium σ_l, σ_t . The specific values on each region are specified in Table 5.1 along with the observed conduction velocities.

Table 5.1: Summary of fiber (f) and transverse (t) conductivity and the resulting conduction velocities measured for planar wavefront propagation on a computational grid of average spatial discretization of $\Delta\mathbf{x} \approx 425\mu\text{m}$ for normal atrial tissue, Bachmann’s bundle, inferior/superior Vena Cava and electrically remodeled tissue (Banduc et al., 2024). Table reprinted from Banduc et al. (2024).

		σ^f	σ^e	σ	CV
		[mScm ⁻¹]	[mScm ⁻¹]	[mScm ⁻¹]	[cms ⁻¹]
normal atria	f	2.0	2.0	1.0	53.0
	t	0.25	1.0	0.2	15.5
Bachmann’s bundle	f	2.0	4.0	1.33	62.0
	t	0.25	1.0	0.2	15.5
Inferior/Superior Vena Cava	f	1.0	1.0	0.5	33.0
	t	1.0	1.0	0.5	33.0
remodeled	f	2.0	2.0	1.0	44.0
	t	0.25	1.0	0.2	13.5

Fibrosis tissue was included in the regions of the atria exhibiting most often this type of substrate in subjects with AF, which are the pulmonary veins antrum, the left lateral wall, anterior wall, posterior wall and septum in the left atrium (Bifulco et al., 2023; Benito et al., 2018; Higuchi et al., 2018; as cited in Banduc et al., 2024), as well as right atrial appendage and septum in the right atrium (Bifulco et al., 2023; Cao et al., 2010; as cited in Banduc et al., 2024). These regions were identified in the different geometries via the universal atrial coordinates (Roney, Pashaei, et al., 2019; as cited in Banduc et al., 2024). For each anatomy and fibrosis pattern, three different severity stages were implemented: low (l), covering from 24.7 to 34.0% of the atrial volume, medium (m), ranging from 30.5 to 42.6%, and high (h), affecting from 35.9 to 44% of tissue (see Table 5.2).

To account for structural remodeling and the presence of scarring, the elements at the core of fibrotic regions were removed. This approach modeled the macroscopic barrier phenomenon caused by the electrical decoupling of myocytes due to percolation of fibrosis in the tissue (Vigmond et al., 2016; as cited in Banduc et al., 2024). Moving away from the core, fibrosis was gradually mixed with remodeled active tissue, where several ionic conductances were re-scaled to consider effects of cytokine-related remodeling (Roney, Whitaker, et al., 2019; as cited in Banduc et al., 2024) (-50% g_{K1} , -40% g_{Na} and -50% g_{CaL}).

AF inducibility was tested with the **OpenCARP** software (Augustin et al., 2021; Plank et al., 2021), which was used to numerically solve the monodomain equations applying finite elements in space (Plank et al., 2021; openCARP consortium et al., 2023; Colli Franzone et al., 2014) and a semi-implicit Crank-Nicolson scheme for time stepping (Allaire, 2007) complemented with an operator splitting technique (Plank et al., 2021; openCARP consortium et al., 2023). The Rush-Larsen scheme (Rush & Larsen, 1978) was employed to update gating variables and ionic concentrations (Plank et al., 2008; as cited in Banduc et al., 2024). The computational domain corresponded to a tetrahedral mesh with an average resolution of 425 μ m and the time step was fixed at 0.05ms.

For inducing AF, this study made use of the PEERP protocol available in **carputils** (Augustin et al., 2021). This pacing strategy was selected among other methods because it maximizes chances of reentry, is open-source and depends in fewer human-defined parameters. The number of ectopic beats was set to 3 and the stimulus was defined by a current of strength 70 μ Acm⁻² applied at a fixed pacing point during a 4ms interval. APD₉₄ was used for the refractory period calculations (Banduc et al., 2024). The determination of initial states characterizing the tissue limit cycle, as well as model parameter calibration and conduction velocity measurements were performed using the **ForCEPSS** framework (Gsell et al., 2024; as cited in Banduc et al., 2024) included in **carputils** (Plank et al., 2021; openCARP consortium et al., 2023; as cited in Banduc et al., 2024).

For each pacing location \mathbf{x}_{stim} , $\mathbf{1}_{AF}(\mathbf{x}_{stim}) = 0$ if no activity of the tissue was registered during a 20ms time window, and $\mathbf{1}_{AF}(\mathbf{x}_{stim}) = 1$ if AF or flutter was observed. A reentry apparatus was classified as flutter if the maximum dominant frequency was at most of 4.7Hz or if at least 50% of the atrial domain reached the maximum dominant frequency. The remainder cases were regarded as AF (Azzolin, Eichenlaub, Nagel, Nairn, Sanchez, et al., 2023; as cited in Banduc et al., 2024).

5.2.2. Data

A total of 24 cases were included in this work. For each, 101 stimulation points were attempted to run for inducibility. The pacing locations were equally distributed in the sense of geodesical distance, using the farthest point sampling criterion described in Section 3.3.4. For the 3 anatomies considered for training (shown in Figure 5.1), one was studied under 3 different fibrosis distributions and the remaining were evaluated for 2 fibrosis patterns each, all of which had three levels of increasing severity (see Table 5.2), summing up to 2,121 attempted simulations through a total of 21 cases. The test set was comprised of other 3 atrial meshes (shown in Figure 5.2), each with a distinct fibrosis pattern of different level, totaling 303 attempted monodomain runs across 3 cases (Banduc et al., 2024).

Table 5.2: Summary statistics for all the simulations performed. Active points refers to the pacing points that were not passive tissue. Re-entries correspond to the sum of AF and flutter cases (Banduc et al., 2024). Table adapted from Banduc et al. (2024).

anatomy	fibrosis case	fibrosis level	fibrosis ratio	simulations	active points	AF	flutter	re-entries	inducibility	average walltime
training										
1	1	l	28.3%	101	101	74	0	74	73.3 %	02:58:36
		m	37.6%	101	101	68	0	68	67.3 %	03:12:24
		h	43.8%	100	100	59	1	60	60.0 %	02:58:33
	2	l	24.7%	101	101	71	0	71	70.3 %	02:57:24
		m	30.5%	100	100	66	0	66	66.0 %	03:10:20
		h	36.1%	101	101	70	0	70	69.3 %	03:00:47
	3	l	34.0%	100	100	74	0	74	74.0 %	03:07:53
		m	39.5%	101	99	63	1	64	63.4 %	02:55:33
		h	42.7%	99	98	68	0	68	68.7 %	03:12:36
2	1	l	31.0%	101	100	77	0	77	76.2 %	02:15:37
		m	36.9%	100	100	64	0	64	64.0 %	02:08:55
		h	42.6%	101	101	62	0	62	61.4 %	02:26:13
	2	l	26.7%	99	98	71	0	71	71.7 %	02:23:48
		m	32.5%	100	100	71	0	71	71.0 %	02:45:12
		h	38.6%	100	99	72	0	72	72.0 %	02:24:09
3	1	l	28.3%	101	101	75	0	75	74.3 %	02:33:53
		m	36.1%	101	100	72	16	88	87.1 %	02:17:30
		h	44.0%	101	101	62	0	62	61.4 %	02:43:45
	2	l	26.2%	101	100	78	0	78	77.2 %	02:40:49
		m	33.9%	101	101	71	0	71	70.3 %	02:47:17
		h	35.9%	101	100	76	0	76	75.2 %	02:28:47
testing										
4	2	l	26.4%	81	79	60	0	60	74.1 %	02:16:38
5	3	m	42.6%	91	89	48	0	48	52.7 %	01:51:47
6	1	h	40.2%	88	87	49	16	65	73.9 %	02:17:58

Pre-procedural CT scans from patients included in the local AF ablation registry of the Medical University of Graz were used. This registry was approved by the ethics committee of the Medical University of Graz (reference number 26-217 ex 13/14) and all patients gave written informed consent. Six subjects were selected (4 male, 2 female) undergoing their first catheter ablation for persistent AF. CT data were automatically segmented (Payer et al., 2017; as cited in Banduc et al., 2024), obtaining labels of all cardiac and vascular blood pools, and the left ventricular myocardium. Volumetric biatrial walls were generated by extruding right and left atrial blood pools, and by meshing of the extruded walls using **Meshtool** (Neic et al., 2020; as cited in Banduc et al., 2024). A fully automated Laplace–Dirichlet–Rule–Based method was employed to annotate the different atrial regions and for generating fiber orientation, based on the algorithm proposed in Azzolin, Eichenlaub, Nagel, Nairn, Sánchez, et al., 2023; as cited in Banduc et al. (2024).

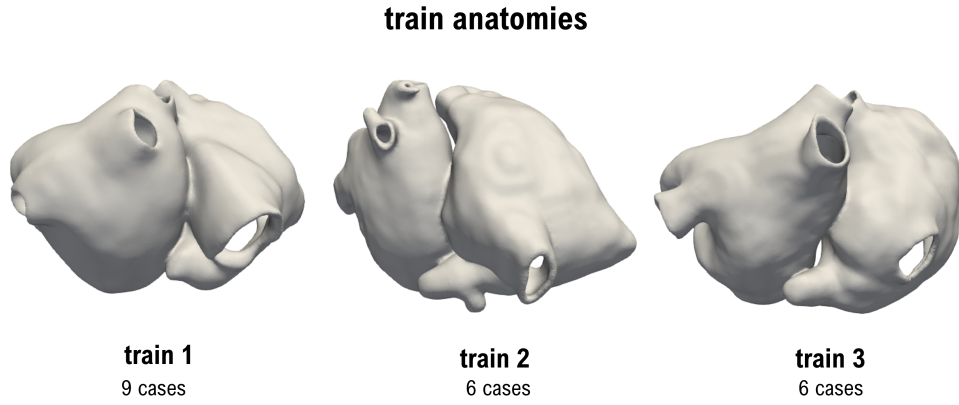


Figure 5.1: Train geometries used in second study, the first was comprised of 3 fibrosis configurations with 3 severity patterns each, while the second and third had 2 fibrosis patterns with 3 severity scenarios each (Banduc et al., 2024).




Figure 5.2: Test geometries used in second study. Each case corresponded to a different severity scenario (Banduc et al., 2024).

The average mesh size was of around 900,000 nodes (Table 5.3). All simulations were performed on the Swiss National Supercomputing Centre (CSCS) using one node with 12 CPU cores (Banduc et al., 2024).

Table 5.3: Mesh information. Table reprinted from Banduc et al. (2024).

anatomy	num. nodes	num. cells
1	1,135,921	6,054,425
2	835,509	4,323,575
3	974,811	5,047,035
4	942,965	4,812,877
5	671,827	3,434,286
6	937,972	4,821,153

5.2.3. Fibrotic Kernel Signature in the Atria

Local fiber orientation takes a significant role in directionality of electrical wave-front propagation in the atria, participating along with fibrosis in the prevalence and progression of AF due to enhanced anisotropy (Kharche et al., 2014). Thus, subject to data availability, the fks isotropic operator encompassing fibrosis distribution can be complemented with a positive-definite uniformly bounded tensor $\mathbf{F}_{\text{aniso}}$ encoding fiber direction in Ω_h .

Although fibrosis and fiber direction jointly affect the behavior of electrical conductance, their impact could be weighted differently for the generation of arrhythmia. For this reason, the two cases \mathbf{F}_{iso} and $\mathbf{F}_{\text{aniso}}$ were taken separately, resulting in the computation of two fks options that were later concatenated, as illustrated in Figure 5.3. For the experiments, the corresponding tensors were of the form

$$\mathbf{F}_{\text{iso}}(\mathbf{x}) = \sigma_{\text{fb}}(\mathbf{x})\mathbf{I}; \quad \mathbf{F}_{\text{aniso}}(\mathbf{x}) = \Sigma_m(\mathbf{x});$$

where σ_{fb} equals 1 in healthy tissue and 0.01 in fibrotic regions. For each case, as in the previous study, the fks was computed using the same methodology for eigenvalue extraction, time progression and normalization of the cardiac domain, taking on average 17 minutes (see Table 5.4) on a local workstation (Apple M2 Ultra, with 24-core CPU and 64GB unified memory) on the same meshes used for the monodomain simulations. Once the feature vectors were obtained for both iso- and anisotropic options, their concatenation was done along the same axis, obtaining an array of size $2N_T$, with $N_T = 50$ being the number of timesteps for the signature.

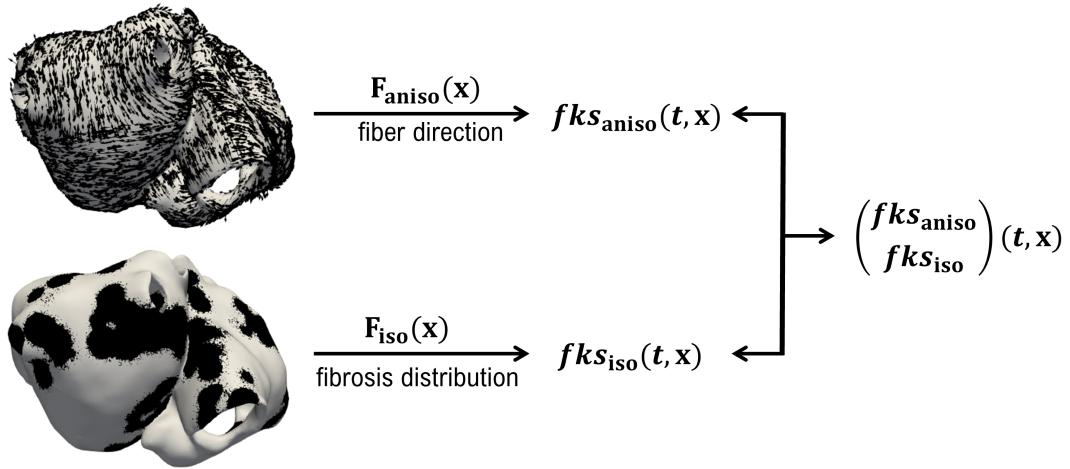


Figure 5.3: fks extraction from aniso- (top) and isotropic (bottom) options. The first case considers the monodomain conductivity tensor to account for fiber direction, and the second an heterogeneous conductivity function modeling conduction variations due to fibrosis (Banduc et al., 2024).

Table 5.4: Computation time for fks eigenproblem on each model and conductivity option (Banduc et al., 2024).

anatomy	fibrosis case	fibrosis level	time aniso.	time iso.
training				
1	1	l	00:21:10	00:22:20
		m	00:20:42	00:19:58
		h	00:20:41	00:20:27
	2	l	00:21:11	00:21:54
		m	00:21:09	00:19:54
		h	00:19:49	00:20:22
	3	l	00:20:56	00:20:58
		m	00:22:25	00:20:45
		h	00:19:49	00:19:58
2	1	l	00:14:33	00:13:50
		m	00:14:23	00:13:54
		h	00:14:27	00:13:36
	2	l	00:14:22	00:13:53
		m	00:14:19	00:13:37
		h	00:14:20	00:13:37
3	1	l	00:13:26	00:13:47
		m	00:13:26	00:14:23
		h	00:14:25	00:13:31
	2	l	00:13:41	00:12:50
		m	00:13:33	00:14:27
		h	00:13:42	00:13:29
testing				
4	2	l	00:14:30	00:14:32
5	3	m	00:08:52	00:08:55
6	1	h	00:17:00	00:17:12

5.2.4. Prediction of Atrial Fibrillation

For each case presented, after the monodomain outputs for inducibility and fks features were computed, each signature associated with a specific pacing location was paired with the corresponding tag, as long as the AF simulation associated to that point was terminated successfully.

For the AF inducibility classification task, a random forest (RF) algorithm with default parameters and a multi-layer perceptron (MLP) with three hidden layers were tested using the `scikit-learn` library (Pedregosa et al., 2011). Other classifiers were also considered, such as gradient boosting and nearest neighbors, but did not perform as good as the ones just described. All these classifiers took in the order of seconds to train and fractions of a second

to evaluate. Thus, the most time-consuming task came from the fks computation.

The performance of each classifier was measured with the F1 score, which is defined as the harmonic mean of precision and recall, and overall inducibility, which was estimated on each case as the ratio of inducible points among the pacing set. The F1 metric is suitable for imbalanced classes, such as the case of the dataset considered, while overall inducibility acts as an endpoint metric, since it is used in practice for quantifying the propensity of cardiac tissue to AF. Other standard metrics were also used, such as accuracy, sensitivity and specificity.

Given that fibrosis patches were expected to predominantly have non-inducible points, a naïve classifier C_{naive} based on scar distribution was evaluated against the fks method to set a baseline result and show that classifiers trained with the signature were not only learning a fibrosis indicator function. For $\mathbf{x} \in \bar{\Omega}_h$, if \mathbf{x} belonged to healthy tissue, then $C_{\text{naive}} = 1$, and if \mathbf{x} was sited inside a fibrosis region, then $C_{\text{naive}} = 0$. That is, pacing from healthy sections degrades into arrhythmia and stimulating scar tissue does not (Banduc et al., 2024). Besides the naïve approach, the inducibility ratio of 40 simulations was used as a second estimator for overall inducibility, as done in Boyle et al., 2019.

5.3. Results

5.3.1. Atrial Fibrillation Simulations

In Table 5.2 a summary of the dataset that was constructed is presented. From the total of 2,424 attempted simulations, 2,371 of them were successfully executed, with an average wall time for each case between 2 and 3 hours, summing up to 264 days of computation. Inducibility ratios ranged from 52.7 to 87.1% and were not associated with the fibrosis level. Most of the cases induced AF and only a few cases of flutter were observed (1.4% of all simulations) (Banduc et al., 2024).

Overall inducibility values stabilize as more simulations are used to compute them, as seen in Figure 5.4. It can be noticed from the three leftmost panels that the rate of convergence is similar, regardless the variability of inducible points and the level of fibrosis. In the rightmost panel, the error evolves when more pacing points are used when compared to the output obtained with all available points. As expected, this error is reduced as more points are included albeit there are some oscillations. When computing overall inducibility with 40 simulations, the median absolute error is of 4.17 percent points and the maximum is of 13.76 percent points (Banduc et al., 2024).




Figure 5.4: Convergence of inducibility. Changes in computed inducibility are shown as more simulations are considered. Cases are grouped by fibrosis severity. Each of the three leftmost panels show 8 different cases, while the rightmost sub-plot shows the evolution of percent error when considering the inducibility computed with all simulations as ground truth (Banduc et al., 2024). Figure reprinted from Banduc et al. (2024).

5.3.2. Prediction of Inducibility

In Figures 5.5 and 5.6 the fks for the test anatomies are shown at initial, intermediate and advanced stages of diffusion. The isotropic version of the signature presents high variability between cases, increasing in areas infiltrated by fibrosis, while the fks computed with the monodomain conductivity tensor tends to be consistent across the different geometries, with some structures displaying distinct values, such as the pulmonary veins. Then, in Figure 5.7, some examples from the fks are shown for the test anatomies and both signature options, showing a characteristic behavior of non-inducible points, which generally reach greater values of the signature, and inducible sites, which in most cases turn out to be more stable in time (Banduc et al., 2024).

When trained with the fks, machine learning classifiers can accurately predict whether pacing points could trigger arrhythmia. RF obtained F1 scores ranging from 84.4 to 86.2% on the test set with 10 random initializations. The MLP classifier obtained F1 scores between 83.6 and 85.7%. Figure 5.8 shows how the F1 score varies as more data is available for the 3 different test cases. First, it is observed although RF and MLP reach similar results, for the second test anatomy RF notably outperforms the MLP. In all scenarios, the classifiers get better outputs than the naïve estimator, but this difference is of the order of roughly 5%. For the amount of data used for training, it can be noticed that there is not a marked dependency on the number of sample size used for training, except when very few points are used. A similar behavior can be observed with other standard metrics (see Figure .A.1) (Banduc et al., 2024).

Overall inducibility predicted for the test cases are presented in Figure 5.9. The ground truth values were computed from all the available simulations on each case (see Table 5.2). Here, the fks method had the lowest mean absolute error, with 2.76 percent points for the RF classifier, which is nearly doubled by the fibrosis-based naïve approach, with 5.26 percent points, and the 40-simulations estimator, with 5.86 percent points of error (Banduc et al., 2024).




Figure 5.5: Isotropic fks map evolution across the three test cases for the 1st, the 20th and the 50th time step computed. Figure reprinted from Banduc et al. (2024).




Figure 5.6: Anisotropic fks map evolution across the three test cases for the 1st, the 20th and the 50th time step computed. Figure reprinted from Banduc et al. (2024).




Figure 5.7: fks samples from test anatomies for isotropic (top) and anisotropic (bottom) options of the signature. For each case, 20 points are randomly selected from the pacing sites and their corresponding fks evolution in time is shown. A curve is colored red if the pacing point induced AF, and blue if not. Figure reprinted from Banduc et al. (2024).

Finally, in Figure 5.10, inducibility maps for the test anatomies are displayed. Here, a considerable correlation between the location of fibrosis and the inducible points can be seen. That is, when applying the stimulation protocol in fibrotic tissue, there is a high chance of not inducing AF, while there is a large possibility of triggering arrhythmia when pacing from healthy tissue. This supports the fact that the naïve classifier obtains good accuracy metrics and overall inducibility estimates. Interestingly, the fks-based RF can identify non-inducibility in both fibrosis sections and non-fibrotic regions. Notably, near the inferior vena cava, test cases 1 and 3 are not inducible, while test case 2 is. The classifier trained with the signature can correctly predict this variation, even though the pacing location is always positioned on healthy tissue (Banduc et al., 2024).




Figure 5.8: Data efficiency for F1 score across test cases for a random forest (blue) and a multi-layer perceptron classifier (mlp, red). The solid line represents the median F1 score, while the shaded area represents the range between minimum and maximum F1 score over 10 runs with different random seeds (Banduc et al., 2024). Figure reprinted from Banduc et al. (2024).

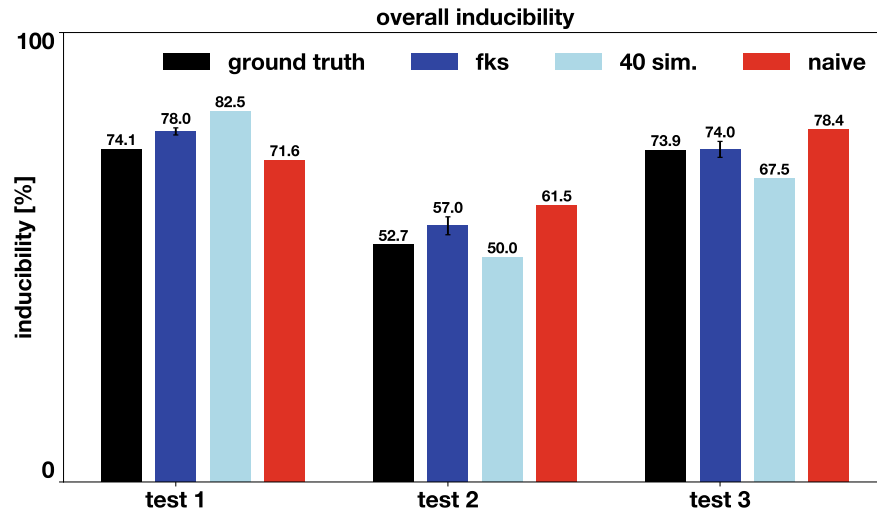


Figure 5.9: Overall inducibility predicted by the fks-trained RF and the naïve classifier. The error bars for the fks classifier represent the standard deviation over 10 runs with different random seeds. For reference, the inducibility computed using only 40 simulations is also included. The ground truth inducibility was computed using all available simulations (Banduc et al., 2024). Figure reprinted from Banduc et al. (2024).




Figure 5.10: Inducibility maps and pacing locations for test cases. In the top row, fibrotic patterns are shown in gray/black, inducible sites are colored in red and non-inducible points are colored in blue. In the bottom row, predicted inducibility maps from the RF classifier are shown and colored under the same criterion used for inducible points (Banduc et al., 2024). Figure reprinted from Banduc et al. (2024).

5.4. Discussion and Conclusions

In this study, a simulation-free method to predict atrial fibrillation based in electrical passive diffusion was successfully implemented, where a set of features that intrinsically incorporate patient-specific fibrosis patterns and anisotropy in conduction was engineered to identify through learning algorithms pro-arrhythmic substrate in the atria. This tool can greatly accelerate the design of different treatments, disposing the need of running an extensive number of simulations to assess AF vulnerability.

To train classifiers with the fks features, an extensive dataset consisting of 2,371 monodomain runs for AF inducibility was built, which constitutes one of the largest in terms of number of simulations (Roney, Whitaker, et al., 2019; as cited in Banduc et al., 2024). With a total computational runtime of 264 days, the database created provides unique insights into the calculations of inducibility. For instance, it was shown that considering only 40 simulations, which is a number of simulations that had already been used in previous studies to predict overall inducibility, can lead to large errors. It can be noticed that it is possible to trigger AF in many of the cases when pacing from healthy tissue. This is to be expected from the pacing protocol selected, since PEERP is designed to maximize chances of initiating fibrillation (Azzolin et al., 2021; as cited in Banduc et al., 2024). Nonetheless, a diversity of inducibility levels were observed, showing that the fibrosis and cardiac geometry play a major role in this quantity.

After training simple machine learning classifiers it is possible to accurately predict whether pacing locations are prone to AF and estimate overall inducibility with small error in unseen anatomies with specific fibrosis patterns at a fraction of the cost of running monodomain simulations. The cost of computing together the two options of the fks is less than the cost of running a single simulation. Even if these 2 costs were equated, the speed up when attempting to compute overall inducibility would still be from 40- to 100-fold depending on the number of points needed.

The naïve approach, which evaluates whether a pacing point is located on healthy tissue or scar, does a good job in predicting AF. However, this classifier will probably not generalize well to other pacing locations, fibrotic models or ablation patterns, because it is only based on local knowledge. For instance, the idea of this classifier was tested with endomyocardial fibrosis (Gharaviri et al., 2020; Pezzuto et al., 2019; as cited in Banduc et al., 2024), and it was not a good predictor (Sahli Costabal et al., 2023; as cited in Banduc et al., 2024). Furthermore, it is known that inducibility is a non-local phenomenon, which is why ablation lines can prevent AF. The naïve estimator will predict the same output, regardless if there is an ablation lesion nearby. Thus, the accuracy results of the naïve approach are circumstantial to this study and cannot be expected to be replicated in different scenarios.

Even though this study shows great promise to reduce the computational time of AF prediction, it has some limitations. First, the method is constrained in accuracy to perfectly estimate inducibility. The fks should fully characterize the changes in fibrosis and geometry, but some cases cannot be correctly predicted. This could be explained by the pre-pacing step, which is common for all the points and might prompt some directionality in how AF forms, which is not currently included in the proposed method. The reduction from reaction-

diffusion of the monodomain system to passive diffusion may also explain the limitations of the model. Another option is that structures that are responsible for the generation of arrhythmia are far away from the pacing location and cannot be sufficiently characterized by the signature. Alternative signatures based on other kernels might perform better (Wang & Solomon, 2019; as cited in Banduc et al., 2024). In this work, a homogeneous ionic model for healthy tissue was considered, even though there is some spatial heterogeneity to be taken into account (Krueger et al., 2013; as cited in Banduc et al., 2024). Finally, a relatively coarse mesh of element size $425\mu\text{m}$ was used, but this is in line with what other clinical studies have employed (Boyle et al., 2019; Azzolin et al., 2021; as cited in Banduc et al., 2024).

In summary, the fks shows to be a fast and robust method for the prediction of AF, that could be used to improve therapy planning. It is hoped that the adoption of this technique will enable the design of optimal ablation strategies within the demands of clinical practice.

Chapter 6

General Discussion

For this thesis work, two studies were conducted to introduce and evaluate the capacity of the fibrotic kernel signature to efficiently capture the electrical properties and multi-scale architecture of human atrial domains, to assess AF vulnerability through the prediction of monodomain outputs with specific protocols of stimulation. The first study encompassed a single volumetric anatomy, where distinct ablation lines were applied under three different scenarios of endomyocardial fibrosis, and the second study served as an extension of this work to multiple atrial geometries subject to high variability in patchy fibrosis.

The computation of the fks was performed in a modern workstation, taking 10 minutes in a mesh with around 700,000 nodes and 17 minutes for meshes with an average of 900,000 nodes. On the other hand, running on the CSCS, the monodomain solvers took for each of the pacing sites 1 hour on the first case, and between 2 to 3 hours on the second case. From what is observed in Figure 5.4, the inducibility ratio stabilizes past 40 simulations, which is consistent with the naïve speed-up approach of considering pacing sets of size 40 instead of a collection of 100 points.

It is observed that the isotropic option of the signature tends to have increased values in regions that are not prone to AF (see Figures 4.3 and 4.5 from Study 1, and Figures 5.5 and 5.7 from Study 2). This behavior is also replicated in sections of tissue enclosed by ablation, where non-inducibility is expected. The anisotropic option exhibits distinct decay for different structures within the same anatomy and stays invariant to changes in geometry (Figure 5.6 from Study 2). Although the presence of muscle fiber effectively affects the spread of electrical potential across the atria, its influence may be considerably attenuated by the passive diffusion model from which the fks is rooted, because of the smoothing effect of the heat kernel.

Different machine learning algorithms were employed for the AF classification task: random forest, k-nearest neighbors, gradient boosting and multi-layer perceptron. All of the fks-trained estimators could accurately predict inducibility, with random forest outperforming the rest of classifiers across the diversity of fibrosis scenarios and ablation strategies, as observed in Figures 4.6 and 5.8. Classifiers presenting a consistent difference in accuracy in comparison to the majority voting approach from the first study shows that the signature captures additional information from the phenomenon described in the monodomain model and not simply memorizes outputs based on pacing location. This is reaffirmed by what is

observed from the second study, where high-accuracy results were still obtained even though the pacing set changed for each anatomy, fibrosis pattern and severity case. Furthermore, utilizing a naïve scar-based approach allows to separate classifiers trained with the fks from a trivial fibrosis indicator. Their average agreement in prediction is reported to be in the range of 64.0 to 76.5%. Although the naïve estimator achieved competitive accuracy with the method proposed, these results were incidental and are not guaranteed to extend to other configurations, pacing techniques or electrophysiological models. The well-behavior of this approach is explained by the inducibility maximization capacities of the PEERP protocol and the patchy distribution of fibrosis.

From the overall inducibility outcomes on each study (see Figures 4.4 and 5.9), it is observed that the signature method can estimate this value with small absolute error. The random forest classifier trained with the fks features properly predicts inducibility ratios in all the tested geometries when no electric isolation procedures are performed and presents the expected post-ablation reduction of inducible surface. Although the naïve approach from the second study returns close estimates to the ground truth values, these results will not reproduce when ablation lines are added to the model.

Even though the method was tested to predict AF inducibility, the fks could be used to predict other quantities that depend on the fibrotic pattern and the geometry, such as the density of reentrant drivers (Boyle et al., 2019), ectopic activity, or even mechanically derived quantities. Preliminary results show that the signature can be used to approximate via regression the phase singularity density in some atrial models, explaining a fair amount of variance in this quantity, with decreased mean absolute error and performing better than an averaging dummy regressor.

Aside from the limitations discussed on each study, it comes as an inherent restriction of the fks model to disregard the cellular ionic structure to characterize a phenomenon that is governed by a reaction component. The exclusive consideration of passive diffusion sets a bound for the capability of the signature to predict fibrillation. This can be seen, for instance, between signatures of different points within the same anatomy presenting very similar values but opposite inducibility outcomes, where the difference may be drawn from ionic reaction of myocytes after stimulation.

Given that the fks comes from the extension of the hks model via anisotropy derived from the monodomain cardiac model, the intrinsic and informative properties from the original signature should be inherited up to certain extent by the fks. However, changing the boundary conditions and adding an elliptic tensor arises the requirement of more detailed properties involving the encoding of extrinsic quantities and the multi-scale behavior of the signature.

Firstly, including a conductivity tensor into the diffusion operator raises to the question if a diagonal restriction of the associated heat kernel retains the changes in conduction residing in the original matrix. From what has been seen empirically, there is some correlation between the fks values and presence of fibrosis. Thus, from an inverse problem perspective, full information regarding the signature within the domain should at least allow to reconstruct some estimate for the fibrosis-informed isotropic tensor. Finally, the homogeneous Neumann boundary condition from the fks leads to the problem of domain-monotonicity from the heat

kernel. For instance, as opposed to the Dirichlet case, the Neumann heat kernel reaches progressively lower values with domain expansion when approaching the stationary state, since there is convergence to the square of the first eigenfunction of the associated elliptic operator over the set, which is a constant equals the inverse of the volume of the domain. However, there are counterexamples that also demonstrate that the Neumann case is not in general monotonically decreasing with domain expansion (Bass & Burdzy, 1993; Hsu, 1994).

Chapter 7

Conclusion

In this thesis work, a novel method to efficiently estimate atrial fibrillation inducibility has been developed: the fibrotic kernel signature (fks). Inspired by the heat kernel signature, the fks is a time-dependent point descriptor derived from the spatial restriction of the kernel induced by an intrinsic fibrosis-informed elliptic operator in the human atria. This signature encodes cardiac architectural properties and conduction variability due to the presence of scar tissue and muscle fibers, which allows it to characterize pro-arrhythmic substrate for patient-specific models. The fks dispenses from the computational burden of typical AF simulations by reducing the monodomain system to a problem of diffusion of electric potential across the geometry after instantaneously applying a stimulus of high magnitude to a fixed stimulation point. The computation of the signature can be performed on a modern machine with modest hardware, taking a considerably less amount of time to execute than a monodomain-based solver.

When paired with simple machine learning classifiers, the fks can accurately predict if whether pacing from a specific location for a given stimulation protocol may degrade into a sustained event of fibrillation. Furthermore, the proposed method allows to reliably estimate overall inducibility, which is a quantity that serves as an endpoint metric to evaluate the effectiveness of therapeutic treatments for AF.

It has been proven that the fks can adapt to the inducibility outcomes from different fibrosis events, being shown to obtain accurate results in at least two scarring models: reactive interstitial fibrosis, which is the one presented in the first study, and reparative patchy fibrosis, which was used in the second study. Similarly, via training of machine learning classifiers, the fks can predict inducibility tags from different pacing protocols, thus capturing pacing-independent mechanisms of AF residing in the heart structure. This can be seen from the results in accuracy obtained for both rapid pacing and PEERP protocols.

The signature method reflects the action of isolating cauterization in the periphery of ablated regions, exhibiting a characteristic barrier effect after the generation of the lesion, expressing the reduction in overall inducibility after the procedure. The ablation patterns tested were pulmonary veins isolation and roof lines, since they are the standard-of-care enclosures in clinical practice, but other ablation lines could be added to reduce arrhythmia vulnerability in other locations.

In addition to the speed-up for inducibility map generation and the benefits exhibited above, the fks has several other advantages that are worth mentioning in the context of patient-specific modeling. First, it can be adapted to data availability of the cardiac domain, where fibrosis, cardiac fibers and other structures and processes affecting the spread of electric potential across the atria can be included as required, as long as they can be represented through the geometry or incorporated via the conductivity tensor. Second, it can be used in several configurations to construct a set of signatures that allow to weight the influence of different anatomical components of the heart into the AF model, where multiple classifiers can be trained with different signature options and then combined into an ensemble to obtain an averaged output. Third, the employment of features from a straightforward model and the utilization of well-known classification algorithms allows reproducibility of the fks method and interpretability of the results, which is convenient when planning an optimal ablation strategy for the patient. Even in algorithms that rely on attributes of the reentrant patterns to design the ablation strategy, the fks predictions can be used as prior knowledge, informing where to pace to successfully induce AF (Banduc et al., 2024).

Finally, the associated process of passive diffusion rooted in the monodomain reaction-diffusion system leads to a trade-off between model complexity to better explain the AF phenomenon and the computational resources needed to implement these models within the restrictions of the clinical context, which translates into the fks method not reaching perfect accuracy even if trained with large amounts of data, albeit greatly reducing the computational time for measuring AF inducibility.

In the future, it is planned to enhance the fks method by incorporating the activation time from sinus rhythm as one of the features of the classifier, as this could be quickly computed with an eikonal solver (Pezzuto et al., 2017; as cited in Banduc et al., 2024), and include some information from cellular action potential, such that spatial heterogeneity can be accounted for. Another proposed task to complement the reviewed studies is to evaluate the signature to compute phase singularity density estimates through regression, where preliminary experiments have already been carried out and show promising results.

In summary, the fks provides a powerful methodology for encoding the electrical properties of cardiac muscle and describing the underlying pro-arrhythmic substrate in the atria. Given the time restraints and computational resources required to predict inducibility, the proposed method is a step forward to improve non-invasive therapy planning and is expected to be used as a proficient tool to optimally design personalized treatment for AF patients undergoing ablation surgery within the demands of clinical practice.

Bibliography

- Allaire, G. (2007). *Numerical analysis and optimization: An introduction to mathematical modelling and numerical simulation*. OUP Oxford.
- Allessie, M., Ausma, J., & Schotten, U. (2002). Electrical, contractile and structural remodeling during atrial fibrillation. *Cardiovascular research*, *54*(2), 230–246.
- Alnaes, M. S., Logg, A., Ølgaard, K. B., Rognes, M. E., & Wells, G. N. (2014). Unified form language: A domain-specific language for weak formulations of partial differential equations. *ACM Transactions on Mathematical Software*, *40*. doi: 10.1145/2566630
- Arendt, W. (1999/2000). *Evolution equations governed by elliptic operators*. University of Ulm.
- Arendt, W. (2002). Chapter 1 semigroups and evolution equations: Functional calculus, regularity and kernel estimates. In C. Dafermos & E. Feireisl (Eds.), (Vol. 1, p. 1-85). North-Holland. doi: [https://doi.org/10.1016/S1874-5717\(04\)80003-3](https://doi.org/10.1016/S1874-5717(04)80003-3)
- Arendt, W. (2006). *Heat kernels*. (Retrieved from <https://www.math.tecnico.ulisboa.pt/>)
- Arendt, W., & ter Elst, A. (1997). Gaussian estimates for second order elliptic operators with boundary conditions. *Journal of Operator Theory*, *38*(1), 87–130.
- Augustin, C., Bayer, J., Bishop, M., Caforio, F., Campos, F., Costa, C. M., ... Wülfers, E. M. (2021). openCARP (v5.0). *RADAR4KIT, Karlsruhe*. doi: 10.35097/389
- Azzolin, L., Eichenlaub, M., Nagel, C., Nairn, D., Sanchez, J., Unger, L., ... Loewe, A. (2023). Personalized ablation vs. conventional ablation strategies to terminate atrial fibrillation and prevent recurrence. *Europace*, *25*(1), 211–222.
- Azzolin, L., Eichenlaub, M., Nagel, C., Nairn, D., Sánchez, J., Unger, L., ... Loewe, A. (2023). Augmenta: Patient-specific augmented atrial model generation tool. *Computerized Medical Imaging and Graphics*, *108*, 102265. doi: <https://doi.org/10.1016/j.compmedimag.2023.102265>
- Azzolin, L., Schuler, S., Dössel, O., & Loewe, A. (2021). A reproducible protocol to assess arrhythmia vulnerability in silico: Pacing at the end of the effective refractory period. *Frontiers in physiology*, *12*, 420.
- Bader, F., Bendahmane, M., Saad, M., & Talhouk, R. (2021, October). Derivation of a new macroscopic bidomain model including three scales for the electrical activity of cardiac tissue. *Journal of Engineering Mathematics*, *131*(1). doi: 10.1007/s10665-021-10174-8
- Banduc, T., Azzolin, L., Manninger, M., Scherr, D., Plank, G., Pezzuto, S., & Sahli Costabal, F. (2024). Simulation-free prediction of atrial fibrillation inducibility with the fibrotic kernel signature. *Medical Image Analysis*.

- Baratta, I. A., Dean, J. P., Dokken, J. S., Habera, M., Hale, J. S., Richardson, C. N., ... Wells, G. N. (2023). *DOLFINx: the next generation FEniCS problem solving environment*. preprint. doi: 10.5281/zenodo.10447666
- Bass, R., & Burdzy, K. (1993). On domain monotonicity of the neumann heat kernel. *Journal of Functional Analysis*, *116*(1), 215-224. doi: <https://doi.org/10.1006/jfan.1993.1111>
- Benito, E. M., Cabanelas, N., Nuñez-Garcia, M., Alarcón, F., Figueras I Ventura, R. M., Soto-Iglesias, D., ... Mont, L. (2018, 12). Preferential regional distribution of atrial fibrosis in posterior wall around left inferior pulmonary vein as identified by late gadolinium enhancement cardiac magnetic resonance in patients with atrial fibrillation. *Europace*, *20*(12), 1959–1965.
- Bifulco, S. F., Macheret, F., Scott, G. D., Akoum, N., & Boyle, P. M. (2023). Explainable machine learning to predict anchored reentry substrate created by persistent atrial fibrillation ablation in computational models. *Journal of the American Heart Association*, *12*(16), e030500.
- Bishop, M., & Plank, G. (2024). Stochastic aspects of virtual heart modelling - stochastic behaviour of arrhythmia induction in virtual heart models suggests caution in offering mechanistic insights. *NCVR*.
- Blomström-Lundqvist, C., Gizurarson, S., Schwieler, J., Jensen, S. M., Bergfeldt, L., Kennebäck, G., ... Mörtzell, D. (2019). Effect of catheter ablation vs antiarrhythmic medication on quality of life in patients with atrial fibrillation: The captaf randomized clinical trial. *JAMA*, *321*(11), 1059–1068. doi: 10.1001/jama.2019.0335
- Boscaini, D., Masci, J., Rodolà, E., Bronstein, M., & Cremers, D. (2016, 05). Anisotropic diffusion descriptors. *Computer Graphics Forum*, *35*. doi: 10.1111/cgf.12844
- Boyle, P. M., Zghaib, T., Zahid, S., Ali, R. L., Deng, D., Franceschi, W. H., ... others (2019). Computationally guided personalized targeted ablation of persistent atrial fibrillation. *Nature biomedical engineering*, *3*(11), 870–879.
- Cao, H., Xue, L., Wu, Y., Ma, H., Chen, L., Wang, X., ... Chen, Y. (2010, 8). Natriuretic peptides and right atrial fibrosis in patients with paroxysmal versus persistent atrial fibrillation. *Peptides*, *31*(8), 1531–9.
- Cazenave, T., & Haraux, A. (1999). *An introduction to semilinear evolution equations* (Revised Edition ed.). Clarendon Press.
- Chae, S., Oral, H., Good, E., Dey, S., Wimmer, A., Crawford, T., ... Chugh, A. (2007). Atrial tachycardia after circumferential pulmonary vein ablation of atrial fibrillation: Mechanistic insights, results of catheter ablation, and risk factors for recurrence. *Journal of the American College of Cardiology*, *50*(18), 1781-1787. doi: <https://doi.org/10.1016/j.jacc.2007.07.044>
- Chen, S.-A., Hsieh, M.-H., Tai, C.-T., Tsai, C.-F., Prakash, V. S., Yu, W.-C., ... Chang, M.-S. (1999). Initiation of atrial fibrillation by ectopic beats originating from the pulmonary veins. *Circulation*, *100*(18), 1879-1886. doi: 10.1161/01.CIR.100.18.1879
- Choi, J., & Kim, S. (2013, 01). Green's function for second order parabolic systems with neumann boundary condition. *Journal of Differential Equations*, *254*. doi: 10.1016/j.jde.2013.01.003
- Choulli, M., Kayser, L., & Ouhabaz, E. (2015, 02). Observations on gaussian upper bounds for

neumann heat kernels. *Bulletin of the Australian Mathematical Society*, 92. doi:10.1017/S0004972715000611

- Clayton, R., & Panfilov, A. (2008). A guide to modelling cardiac electrical activity in anatomically detailed ventricles. *Progress in Biophysics and Molecular Biology*, 96(1), 19-43. (Cardiovascular Physiome) doi: <https://doi.org/10.1016/j.pbiomolbio.2007.07.004>
- Clayton, R., Zhuchkova, E., & Panfilov, A. (2006). Phase singularities and filaments: Simplifying complexity in computational models of ventricular fibrillation. *Progress in Biophysics and Molecular Biology*, 90(1), 378-398. (From Funny Current to Current Physiome) doi: <https://doi.org/10.1016/j.pbiomolbio.2005.06.011>
- Colli Franzone, P., Pavarino, L. F., & Scacchi, S. (2014). *Mathematical cardiac electrophysiology* (Vol. 13). Springer.
- Courtemanche, M., Ramirez, R. J., & Nattel, S. (1998). Ionic mechanisms underlying human atrial action potential properties: insights from a mathematical model. *American Journal of Physiology-Heart and Circulatory Physiology*, 275, H301-H321.
- Dal, H., Göktepe, S., Kaliske, M., & Kuhl, E. (2012). A fully implicit finite element method for bidomain models of cardiac electrophysiology. *Computer Methods in Biomechanics and Biomedical Engineering*, 15(6), 645-656.
- Davies, E. (1987). Ultracontractivity and the heat kernel for schrijdinger operators and dirichlet laplacians..
- Davies, E. B. (1989). *Heat kernels and spectral theory*. Cambridge University Press.
- de Vries, L. J., Martirosyan, M., van Domburg, R. T., Wijchers, S. A., Géczy, T., & Szili-Torok, T. (2018, Jan). Coupling interval variability of premature ventricular contractions in patients with different underlying pathology: an insight into the arrhythmia mechanism. *Journal of interventional cardiac electrophysiology : an international journal of arrhythmias and pacing*, 51(1), 25-33. doi: 10.1007/s10840-017-0309-8
- Dodziuk, J. (1981). Eigenvalues of the laplacian and the heat equation. *The American Mathematical Monthly*, 88(9), 686-695. doi: 10.2307/2320674
- Dodziuk, J. (1983, 01). Maximum principle for parabolic inequalities and the heat flow on open manifolds. *Indiana Univ. Math. J.*, 32, 703-716. doi: 10.1512/iumj.1983.32.32046
- Dzeshka, M. S., Lip, G. Y., Snezhitskiy, V., & Shantsila, E. (2015). Cardiac fibrosis in patients with atrial fibrillation: Mechanisms and clinical implications. *Journal of the American College of Cardiology*, 66(8), 943-959. doi: <https://doi.org/10.1016/j.jacc.2015.06.1313>
- Gander, L., Pezzuto, S., Gharaviri, A., Krause, R., Perdikaris, P., & Sahli Costabal, F. (2022). Fast characterization of inducible regions of atrial fibrillation models with multi-fidelity gaussian process classification. *Frontiers in Physiology*, 260.
- Gharaviri, A., Bidar, E., Potse, M., Zeemering, S., Verheule, S., Pezzuto, S., ... Schotten, U. (2020). Epicardial fibrosis explains increased endo-epicardial dissociation and epicardial breakthroughs in human atrial fibrillation. *Frontiers in Physiology*, 11(68).
- Gharaviri, A., Pezzuto, S., Potse, M., Verheule, S., Conte, G., Krause, R., ... Auricchio, A. (2021). Left atrial appendage electrical isolation reduces atrial fibrillation recurrences:

- simulation study. *Circulation: Arrhythmia and Electrophysiology*, 14(1), e009230. doi: 10.1161/CIRCEP.120.009230
- Glück, J., & Mui, J. (2024). *Non-positivity of the heat equation with non-local robin boundary conditions*.
- Goodenough, D. A., & Paul, D. L. (2009, July). Gap junctions. *Cold Spring Harbor perspectives in biology*, 1(1), a002576. (Review) doi: 10.1101/cshperspect.a002576
- Grandits, T., Pezzuto, S., Costabal, F. S., Perdikaris, P., Pock, T., Plank, G., & Krause, R. (2021). Learning atrial fiber orientations and conductivity tensors from intracardiac maps using physics-informed neural networks. In *Functional imaging and modeling of the heart: 11th international conference, fimh 2021, stanford, ca, usa, june 21-25, 2021, proceedings* (pp. 650–658).
- Griepentrog, J., Kaiser, H., & Rehberg, J. (2001, 01). Resolvent and heat kernel properties for second order elliptic differential operators with general boundary conditions in lp. *Advances in Mathematical Sciences and Applications*, 11.
- Grieser, D. (2004). *Notes on heat kernel asymptotics*.
- Grigor'yan, A. (1997, 01). Gaussian upper bounds for the heat kernel on arbitrary manifolds. *Journal of Differential Geometry - J DIFFEREN GEOM*, 45. doi: 10.4310/jdg/1214459753
- Grigor'yan, A. (2006, 01). Heat kernels on weighted manifolds and applications. *Contemporary Mathematics*, 398. doi: 10.1090/conm/398/07486
- Grigor'yan, A. (2018, March). *Lectures on heat kernels on riemannian manifolds*. Course Notes.
- Gsell, M., Neic, A., Bishop, M., Gillette, K., Prassl, A., Augustin, C., ... Plank, G. (2024). Forceps – a framework for cardiac electrophysiology simulations standardization. *Computer Methods and Programs in Biomedicine*.
- Haissaguerre, M., Jaïs, P., Shah, D. C., Takahashi, A., Hocini, M., Quiniou, G., ... Clémenty, J. (1998). Spontaneous initiation of atrial fibrillation by ectopic beats originating in the pulmonary veins. *New England Journal of Medicine*, 339(10), 659–666.
- Hernandez, V., Roman, J. E., & Vidal, V. (2005). SLEPc: A scalable and flexible toolkit for the solution of eigenvalue problems. *ACM Trans. Math. Software*, 31(3), 351–362.
- Higuchi, K., Cates, J., Gardner, G., Morris, A., Burgon, N. S., Akoum, N., & Marrouche, N. F. (2018, 1). The spatial distribution of late gadolinium enhancement of left atrial magnetic resonance imaging in patients with atrial fibrillation. *JACC: Clinical Electrophysiology*, 4(1), 49–58.
- Hinderer, S., & Schenke-Layland, K. (2019). Cardiac fibrosis – a short review of causes and therapeutic strategies. *Advanced Drug Delivery Reviews*, 146, 77–82. (Wound healing and fibrosis – State of play) doi: <https://doi.org/10.1016/j.addr.2019.05.011>
- Hodgkin, A., & Huxley, A. (1990). A quantitative description of membrane current and its application to conduction and excitation in nerve. *Bulletin of Mathematical Biology*, 52(1), 25–71. doi: [https://doi.org/10.1016/S0092-8240\(05\)80004-7](https://doi.org/10.1016/S0092-8240(05)80004-7)
- Hsu, E. (1994, 03). A domain monotonicity property of the neumann heat kernel. *OSAKA JOURNAL OF MATHEMATICS*, 31.

- Hurtado, D. E., & Henao, D. (2014). Gradient flows and variational principles for cardiac electrophysiology: Toward efficient and robust numerical simulations of the electrical activity of the heart. *Computer Methods in Applied Mechanics and Engineering*, 273, 238–254. doi: 10.1016/j.cma.2014.02.002
- Hytönen, T., van Neerven, J., Veraar, M., & Weis, L. (2016). Bochner spaces. In *Analysis in banach spaces : Volume i: Martingales and littlewood-paley theory* (pp. 1–66). Cham: Springer International Publishing. doi: 10.1007/978-3-319-48520-1_1
- Issa, Z. F., Miller, J. M., & Zipes, D. P. (2009). Chapter 2 - electrophysiological testing. In Z. F. Issa, J. M. Miller, & D. P. Zipes (Eds.), *Clinical arrhythmology and electrophysiology* (p. 27-56). Philadelphia: W.B. Saunders. doi: <https://doi.org/10.1016/B978-1-4160-5998-1.00005-7>
- Iyer, A. N., & Gray, R. A. (2001). An experimentalist’s approach to accurate localization of phase singularities during reentry. *Annals of Biomedical Engineering*, 29(1), 47–59. doi: 10.1114/1.1335538
- Joglar, J. A., Chung, M. K., Armbruster, A. L., Benjamin, E. J., Chyou, J. Y., Cronin, E. M., ... others (2024). 2023 acc/aha/accp/hrs guideline for the diagnosis and management of atrial fibrillation: a report of the american college of cardiology/american heart association joint committee on clinical practice guidelines. *Circulation*, 149(1), e1–e156.
- Johnston, B. M., & Johnston, P. R. (2020). Approaches for determining cardiac bidomain conductivity values: progress and challenges. *Medical & Biological Engineering & Computing*, 58(12), 2919–2935. doi: 10.1007/s11517-020-02272-z
- Kaboudian, A., Cherry, E. M., & Fenton, F. H. (2019). 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.
- Katritsis, D. G., & Morady, F. (2022). 6 - basic intervals and atrial and ventricular conduction curves. In D. G. Katritsis & F. Morady (Eds.), *Clinical cardiac electrophysiology* (p. 54-71.e1). Elsevier. doi: <https://doi.org/10.1016/B978-0-323-79338-4.00015-7>
- Kawai, S., Mukai, Y., Inoue, S., Yakabe, D., Nagaok, K., Sakamoto, K., ... Tsutsui, H. (2019). Non-pulmonary vein triggers of atrial fibrillation are likely to arise from low-voltage areas in the left atrium. *Sci Rep*, 9, 12271. doi: 10.1038/s41598-019-48669-1
- Kharche, S., Castro, S., Thomas, B., Colman, M., Jarvis, J., Smaill, B., ... Zhao, J. (2014). Role of fiber orientation in atrial arrhythmogenesis. In *Computing in cardiology 2014* (p. 1041-1044).
- Kocica, M., Corno, A., Carreras, F., Ballester, M., Moghbel, M., Cueva, C., ... Torrent-Guasp, F. (2006, 05). The helical ventricular myocardial band: global, three-dimensional, functional architecture of the ventricular myocardium. *European journal of cardio-thoracic surgery : official journal of the European Association for Cardio-thoracic Surgery*, 29 Suppl 1, S21-40. doi: 10.1016/j.ejcts.2006.03.011
- Kornej, J., Börschel, C. S., Benjamin, E. J., & Schnabel, R. B. (2020). Epidemiology of atrial fibrillation in the 21st century: novel methods and new insights. *Circulation research*, 127(1), 4–20.
- Kotadia, I., Whitaker, J., Roney, C., Niederer, S., O’Neill, M., Bishop, M., & Wright, M.

- (2020). Anisotropic cardiac conduction. *Arrhythmia & Electrophysiology Review*, 9(4), 202.
- Krause, D., Potse, M., Dickopf, T., Krause, R., Auricchio, A., & Prinzen, F. (2012). Hybrid parallelization of a large-scale heart model. In *Facing the multicore-challenge ii* (pp. 120–132). Springer.
- Krueger, M. W., Dorn, A., Keller, D. U., Holmqvist, F., Carlson, J., Platonov, P. G., ... Dössel, O. (2013). In-silico modeling of atrial repolarization in normal and atrial fibrillation remodeled state. *Medical & biological engineering & computing*, 51, 1105–1119.
- Krummen, D. E., Bayer, J. D., Ho, J., Ho, G., Smetak, M. R., Clopton, P., ... Narayan, S. M. (2012). Mechanisms of human atrial fibrillation initiation clinical and computational studies of repolarization restitution and activation latency. *Circulation: Arrhythmia and Electrophysiology*, 5(6), 1149–1159. doi: 10.1161/CIRCEP.111.969022
- Loewe, A., Poremba, E., Oesterlein, T., Luik, A., Schmitt, C., Seemann, G., & Dössel, O. (2019). Patient-specific identification of atrial flutter vulnerability—a computational approach to reveal latent reentry pathways. *Frontiers in Physiology*, 9, 1910.
- Loewe, A., Wilhelms, M., Dössel, O., & Seemann, G. (2014). Influence of chronic atrial fibrillation induced remodeling in a computational electrophysiological model. *Biomed. Eng.*, 59(S1), S929–S932. doi: 10.1515/bmt-2014-5012
- Lootens, S., Janssens, I., Van Den Abeele, R., Wülfers, E. M., Bezerra, A. S., Verstraeten, B., ... Vandersickel, N. (2024). Directed graph mapping exceeds phase mapping for the detection of simulated 2d meandering rotors in fibrotic tissue with added noise. *Computers in Biology and Medicine*, 171, 108138. doi: <https://doi.org/10.1016/j.compbiomed.2024.108138>
- Lundberg, S., & Lee, S.-I. (2017, 12). A unified approach to interpreting model predictions..
- Marini, C., Santis, F. D., Sacco, S., Russo, T., Olivieri, L., Totaro, R., & Carolei, A. (2005). Contribution of atrial fibrillation to incidence and outcome of ischemic stroke. *Stroke*, 36(6), 1115–1119. doi: 10.1161/01.STR.0000166053.83476.4a
- Matene, E., & Jacquemet, V. (2012, Nov). Fully automated initiation of simulated episodes of atrial arrhythmias. *Europace : European pacing, arrhythmias, and cardiac electrophysiology : journal of the working groups on cardiac pacing, arrhythmias, and cardiac cellular electrophysiology of the European Society of Cardiology*, 14 Suppl 5, v17–v24. doi: 10.1093/europace/eus271
- McDowell, K. S., Zahid, S., Vadakkumpadan, F., Blauer, J., MacLeod, R. S., & Trayanova, N. a. (2015). Virtual Electrophysiological Study of Atrial Fibrillation in Fibrotic Remodeling. *Plos One*, 10(2), e0117110. doi: 10.1371/journal.pone.0117110
- Mosley, L. (2013). A balanced approach to the multi-class imbalance problem..
- Muffoletto, M., Qureshi, A., Zeidan, A., Muizniece, L., Fu, X., Zhao, J., ... Aslanidi, O. (2021). Toward patient-specific prediction of ablation strategies for atrial fibrillation using deep learning. *Frontiers in Physiology*, 12, 674106.
- Neic, A., Gsell, A. F. M., Karabelas, E., Prassl, A. J., & Plank, G. (2020). Automating image-based mesh generation and manipulation tasks in cardiac modeling workflows using meshtool. *SoftwareX*. doi: 10.1016/j.softx.2020.100454

- Neic, A., Liebmann, M., Hoetzel, E., Mitchell, L., Vigmond, E. J., Haase, G., & Plank, G. (2012, August). Accelerating cardiac bidomain simulations using graphics processing units. *IEEE transactions on bio-medical engineering*, 59(8), 2281–90. doi: 10.1109/TBME.2012.2202661
- Niederer, S., Kerfoot, E., Benson, A., Bernabeu, M., Bernus, O., Bradley, C., ... Smith, N. (2011, nov). Verification of cardiac tissue electrophysiology simulators using an N-version benchmark. *Philosophical Transactions. Series A, Mathematical, Physical, and Engineering Sciences*, 369(1954), 4331–51. doi: 10.1098/rsta.2011.0139
- Niederer, S., Mitchell, L., Smith, N., & Plank, G. (2011). Simulating human cardiac electrophysiology on clinical time-scales. *Frontiers in physiology*, 2, 14.
- Norris, J. R. (1997). Heat kernel asymptotics and the distance function in lipschitz riemannian manifolds. *Acta Mathematica*, 179(1), 79–103. doi: 10.1007/BF02392720
- openCARP consortium, Augustin, C., Boyle, P. M., Loechner, V., Colin, R., Gsell, M., ... Linder, M. (2023). *openCARP*. doi: 10.35097/1845
- Payer, C., Štern, D., Bischof, H., & Urschler, M. (2017). Multi-label whole heart segmentation using cnns and anatomical label configurations. In *International workshop on statistical atlases and computational models of the heart* (pp. 190–198).
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., ... Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12, 2825–2830.
- Pezzuto, S., Gharaviri, A., Schotten, U., Potse, M., Conte, G., Caputo, M. L., ... Auricchio, A. (2018). Beat-to-beat p-wave morphological variability in patients with paroxysmal atrial fibrillation: an in silico study. *EP Europace*, 20(suppl_3), iii26–iii35. doi: 10.1093/europace/euy227
- Pezzuto, S., Hake, J., & Sundnes, J. (2016). Space-discretization error analysis and stabilization schemes for conduction velocity in cardiac electrophysiology. *International Journal for Numerical Methods in Biomedical Engineering*, 32(10), e02762.
- Pezzuto, S., Kal’avský, P., Potse, M., Prinzen, F. W., Auricchio, A., & Krause, R. (2017). Evaluation of a rapid anisotropic model for ECG simulation. *Frontiers in Physiology*, 8(MAY). doi: 10.3389/fphys.2017.00265
- Pezzuto, S., Quaglino, A., & Potse, M. (2019). On sampling spatially-correlated random fields for complex geometries. In *International conference on functional imaging and modeling of the heart* (pp. 103–111).
- Plank, G., Loewe, A., Neic, A., Augustin, C., Huang, Y.-L. C., Gsell, M., ... Vigmond*, E. (2021). The openCARP simulation environment for cardiac electrophysiology. *Computer Methods and Programs in Biomedicine*, 208, 106223. doi: 10.1016/j.cmpb.2021.106223
- Plank, G., Zhou, L., Greenstein, J. L., Cortassa, S., Winslow, R. L., O’Rourke, B., & Trayanova, N. A. (2008, September). From mitochondrial ion channels to arrhythmias in the heart: computational techniques to bridge the spatio-temporal scales. *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences*, 366(1879), 3381–409. (ISBN: 1364-503X (Print)\r1364-503X (Linking)) doi: 10.1098/rsta.2008.0112
- Potse, M., Dubé, B., Richer, J., Vinet, A., & Gulrajani, R. M. (2006). 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.
- Potse, M., Gharaviri, A., Pezzuto, S., Auricchio, A., Krause, R., Verheule, S., & Schotten, U. (2018). Anatomically-induced fibrillation in a 3d model of the human atria. In *2018 computing in cardiology conference (cinc)* (Vol. 45). IEEE. doi: 10.22489/CinC.2018.366
- Roney, C. H., Pashaei, A., Meo, M., Dubois, R., Boyle, P. M., Trayanova, N. A., . . . Vigmond, E. J. (2019). Universal atrial coordinates applied to visualisation, registration and construction of patient specific meshes. *Medical Image Analysis*, 55, 65–75. doi: 10.1016/j.media.2019.04.004
- Roney, C. H., Whitaker, J., Sim, I., O’Neill, L., Mukherjee, R. K., Razeghi, O., . . . Niederer, S. A. (2019). A technique for measuring anisotropy in atrial conduction to estimate conduction velocity and atrial fibre direction. *Computers in Biology and Medicine*, 104(July 2018), 278–290. doi: 10.1016/j.compbimed.2018.10.019
- Rosillo de Souza, G., Krause, R., & Pezzuto, S. (2024). Boundary integral formulation of the cell-by-cell model of cardiac electrophysiology. *Engineering Analysis with Boundary Elements*, 158, 239-251. doi: 10.1016/j.enganabound.2023.10.021
- Ruiz Herrera, C., Grandits, T., Plank, G., Perdikaris, P., Sahli Costabal, F., & Pezzuto, S. (2022). Physics-informed neural networks to learn cardiac fiber orientation from multiple electroanatomical maps. *Engineering with Computers*, 38(5), 3957–3973.
- Rush, S., & Larsen, H. (1978). A practical algorithm for solving dynamic membrane equations. *IEEE Transactions on Biomedical Engineering*, 25(4), 389–392.
- Saad-Omer, S. M., Ryad, R., Limbana, T., Zahid, T., & Jahan, N. (2020). Catheter ablation vs. medical treatment in patients with atrial fibrillation. *Cureus*, 12(8), e9700. doi: 10.7759/cureus.9700
- Sahli Costabal, F., Banduc, T., Gander, L., & Pezzuto, S. (2023). The fibrotic kernel signature: Simulation-free prediction of atrial fibrillation. In *International conference on functional imaging and modeling of the heart* (pp. 87–96).
- Schmitt, O. H. (1969). Biological information processing using the concept of interpenetrating domains. In K. N. Leibovic (Ed.), *Information processing in the nervous system: Proceedings of a symposium held at the state university of new york at buffalo 21st–24th october, 1968* (pp. 325–331). Berlin, Heidelberg: Springer Berlin Heidelberg. doi: 10.1007/978-3-642-87086-6_18
- Schotten, U., Verheule, S., Kirchhof, P., & Goette, A. (2011). Pathophysiological mechanisms of atrial fibrillation: a translational appraisal. *Physiological Reviews*, 91(1), 265–325. doi: 10.1152/physrev.00031.2009
- Scroggs, M. W., Baratta, I. A., Richardson, C. N., & Wells, G. N. (2022). Basix: a runtime finite element basis evaluation library. *Journal of Open Source Software*, 7(73), 3982. doi: 10.21105/joss.03982
- Scroggs, M. W., Dokken, J. S., Richardson, C. N., & Wells, G. N. (2022). Construction of arbitrary order finite element degree-of-freedom maps on polygonal and polyhedral cell meshes. *ACM Transactions on Mathematical Software*, 48(2), 18:1–18:23. doi: 10.1145/3524456
- Serra, D., Romero, P., Garcia-Fernandez, I., Lozano, M., Liberos, A., Rodrigo, M., . . . Se-

- bastian, R. (2022). An automata-based cardiac electrophysiology simulator to assess arrhythmia inducibility. *Mathematics*, *10*(8), 1293.
- Sun, J., Ovsjanikov, M., & Guibas, L. (2009). A concise and provably informative multi-scale signature based on heat diffusion. *Computer Graphics Forum*, *28*(5), 1383–1392.
- ten Tusscher, K., Noble, D., Noble, P., & Panfilov, A. (2004). A model for human ventricular tissue. *American Journal of Physiol – Heart and Circulation Physiology*, *286*(4), H1573–89. doi: 10.1152/ajpheart.00794.2003
- Trohman, R. G., Huang, H. D., & Sharma, P. S. (2023). Atrial fibrillation: primary prevention, secondary prevention, and prevention of thromboembolic complications: part 1. *Frontiers in cardiovascular medicine*, *10*, 1060030. doi: 10.3389/fcvm.2023.1060030
- Tsao, C. W., Aday, A. W., Almarzooq, Z. I., Anderson, C. A., Arora, P., Avery, C. L., ... others (2023). Heart disease and stroke statistics—2023 update: a report from the american heart association. *Circulation*, *147*(8), e93–e621.
- Verma, A., Jiang, C.-y., Betts, T. R., Chen, J., Deisenhofer, I., Mantovan, R., ... others (2015). Approaches to catheter ablation for persistent atrial fibrillation. *New England Journal of Medicine*, *372*, 1812–1822.
- Verrier, R. L., Fuller, H., Justo, F., Nearing, B. D., Rajamani, S., & Belardinelli, L. (2016, 4). Unmasking atrial repolarization to assess alternans, spatiotemporal heterogeneity, and susceptibility to atrial fibrillation. *Heart Rhythm*, *13*(4), 953–61.
- Vigmond, E., Pashaei, A., Cochet, H., Amroui, S., & Haissaguerre, M. (2016). Percolation as a Mechanism to Explain Atrial Fractionated Electrograms and Reentry in a Fibrosis Model Based on Imaging Data. *Heart Rhythm*, 1–8. doi: 10.1016/j.hrthm.2016.03.019
- Wang, Y., & Solomon, J. (2019). Intrinsic and extrinsic operators for shape analysis. In *Handbook of numerical analysis* (Vol. 20, pp. 41–115). Elsevier.
- Yi, F., Hou, W., Zhou, C., Yin, Y., Lu, S., Duan, C., ... Zhang, H. (2019). Radiofrequency ablation versus antiarrhythmic drug therapy for atrial fibrillation: Meta-analysis of safety and efficacy. *Journal of cardiovascular pharmacology*, *73*(4), 241–247. doi: 10.1097/FJC.0000000000000654
- Zahid, S., Cochet, H., Boyle, P. M., Schwarz, E. L., Whyte, K. N., Vigmond, E. J., ... Trayanova, N. A. (2016). Patient-derived models link re-entrant driver localization in atrial fibrillation to fibrosis spatial pattern. *Cardiovascular Research*, *110*(3), 443–454. doi: 10.1093/cvr/cvw073

Annex

Annex .A. Performance with Standard Metrics

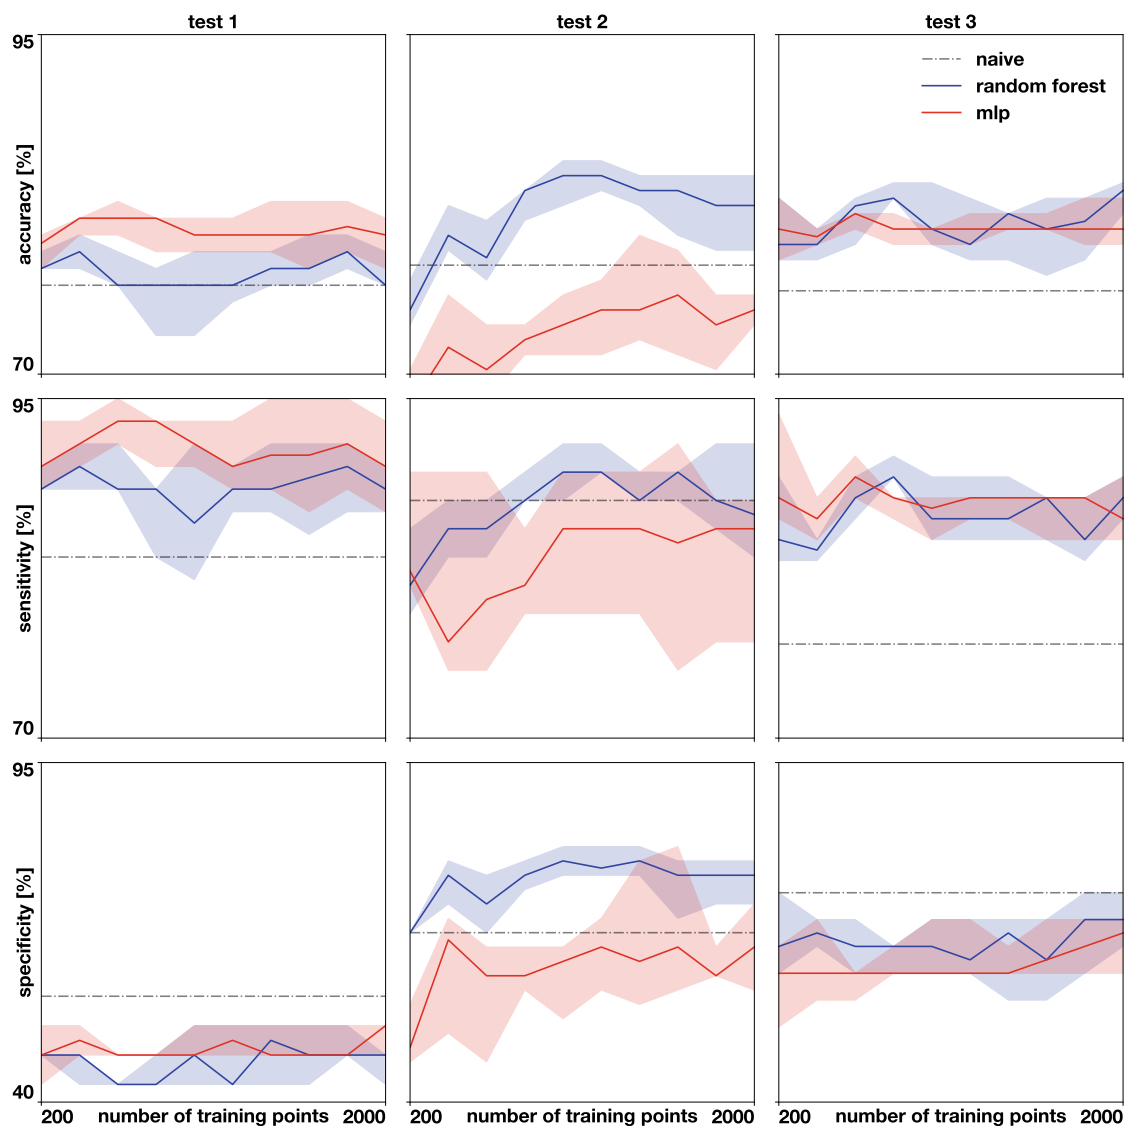


Figure .A.1: Data efficiency for accuracy, sensitivity and specificity across test cases for a random forest (blue) and a multi-layer perceptron classifier (mlp, red). The solid lines represent the median values, while the shaded area represents the range between minimum and maximum scores over 10 runs with different random seeds (Banduc et al., 2024).

In addition to the F1 score used for measuring the performance of the proposed method in the second study, other standard metrics for evaluating the models were considered and are exhibited in Figure .A.1: Accuracy, which is the ratio of correct predictions; sensitivity, which is the fraction of true positives and the size of the positive class; and specificity, which is the fraction of true negatives and the size of the negative class.

From the results presented, it can be observed that the fks-based estimators generally reach higher levels of accuracy and outperform the naïve classifier, where the random forest algorithm yields better results. However, the predictive power of the negative class tends to be significantly less effective for the first test case, indicating that the classifier could be overestimating inducibility on this setting.