The fibrotic kernel signature: simulation-free prediction of atrial fibrillation

Francisco Sahli Costabal[1][0000−0002−2612−463X], Tomás Banduc[1], Lia Gander[2], and Simone Pezzuto[3][0000−0002−7432−0424]

[1] Department of Mechanical and Metallurgical Engineering, School of Engineering and Institute for Biological and Medical Engineering, Schools of Engineering, Medicine and Biological Sciences, Pontificia Universidad Católica de Chile, Santiago, Chile fsc@ing.puc.cl

[2] Center for Computational Medicine in Cardiology, Euler Institute, Università della Svizzera italiana, Lugano, Switzerland lia.gander@usi.ch

[3] Department of Mathematics, Università di Trento, Via Sommarive 14, 38123 Povo, Italy simone.pezzuto@unitn.it

Abstract. We propose a fast classifier that is able to predict atrial fibrillation inducibility in patient-specific cardiac models. Our classifier is general and it does not require re-training for new anatomies, fibrosis patterns, and ablation lines. This is achieved by training the classifier on a variant of the Heat Kernel Signature (HKS). Here, we introduce the "fibrotic kernel signature" (FKS), which extends the HKS by incorporating fibrosis information. The FKS is fast to compute, when compared to standard cardiac models like the monodomain equation. We tested the classifier on 9 combinations of ablation lines and fibrosis patterns. We achieved maximum balanced accuracies with the classifiers ranging from 75.8% to 95.8%, when tested on single points. The classifier is also able to predict very well the overall inducibility of the model. We think that our classifier can speed up the calculation of inducibility maps in a way that is crucial to create better personalized ablation treatments within the time constraints of the clinical setting.

Keywords: Heat kernel signature · Fibrotic Kernel Signature · Atrial Fibrillation · Fibrosis · Patient-Specific Modeling

1 Introduction

Fibrosis is one of the main drivers of Atrial Fibrillation (AF), the most common cardiac arrhythmia [16]. Fibrosis significantly increases tissue heterogeneity and anisotropy in conduction, which in turn enhance AF inducibility and complexity. Its distribution in the atria is patient-specific and progresses with AF in a vicious loop: the more fibrosis is present in the tissue, the more AF events are likely to occur, which trigger more fibrosis deposition [13].

Ideally, therapeutic approaches to AF as catheter ablation should return the best outcome when tailored to patient fibrosis distribution, at least as claimed in recent retrospective and prospective studies [2, 8, 9]. These studies show that,

thanks to patient-specific *in silico* models, it is possible to estimate AF inducibility for various ablation scenarios, and then select the best treatment for the patient. However, testing AF inducibility is costly, both in terms of time and required computational resources.

The assessment of AF inducibility can be understood as a classification problem. Given a patient-specific anatomy, fibrosis pattern, and ablation lines, the objective is determining whether a pacing protocol may lead to a stable AF event or not [4]. Testing a fixed number of well-distributed, pacing location is a standard protocol to estimate inducibility [2]. However, the protocol needs to be repeated when ablation lines are added or a new anatomy is to be tested. Since the total computational cost can be very high, some authors proposed an adaptive pacing protocol [1] or surrogate models of AF [14].

In this work, we propose a classifier for AF inducibility that does not require retraining when ablation lines and fibrosis change. The classifier is based on the Heat Kernel Signature (HKS) [15], a time series that effectively encodes local geometrical and topological information of a domain. Along with its variants, HKS is popular in shape analysis. Mathematically, the HKS is based on the heat (or diffusion) equation. It may be interpreted as the concentration time course of a ink drop as it diffuses throughout the domain. Here, we extend the HKS by incorporating the fibrosis pattern into the diffusion operator. In this way, once trained, the classifier only requires the HKS, which is cheap to compute compared to a standard monodomain simulation.

The manuscript is structured as follows: we review the AF modeling framework in Sec. 2.1, which has been used to generate the dataset. In Sec. 2.2 we present the fibrotic kernel signature, and apply it to the definition of the classifier, in Sec. 2.3. We conclude with results (Sec. 3) and discussion (Sec. 4).

2 Methods

2.1 Cardiac atrial modeling

The monodomain equation is the most common model in simulating atrial fibrillation (AF). It reads as follows:

$$\chi\Big(C_{\mathrm{m}}\frac{\partial V}{\partial t} + I_{\mathrm{ion}}(V,\boldsymbol{z}) - I_{\mathrm{stim}}(\boldsymbol{x},t)\Big) = \mathrm{div}(\boldsymbol{G}\nabla V), \tag{1}$$

where $V(\boldsymbol{x},t)$ is the transmembrane voltage as a function of the spatial position $x \in \Omega$, $\bar{\Omega}$ being the active tissue of the atrial domain, and the time $t \geq 0$. The other parameters are: χ, the surface-to-volume ratio; C_{m}, the membrane capacitance; $I_{\mathrm{stim}}(\boldsymbol{x},t)$ is the stimulation current; and $\boldsymbol{G}$, the monodomain electric conductivity. The nonlinear term $f(V,\boldsymbol{z})$ encompasses all ion currents flowing through the cellular membrane, which are numerous for physiological models. Here, we consider the Courtemanche-Ramirez-Nattel model [3]. Fibrosis is modeled by reducing the intra-cellular conductivity in the cross-fiber direction, which in turn affects the tensor $\boldsymbol{G}$. Specifically, we encode the presence of fibrosis by

reducing the cross-fiber conductivity in $\boldsymbol{G}$. See [5] for the parameter values and for the numerical method to solve the monodomain equation (1).

AF is a self-sustained, chaotic activation of the atria. There are several ways for triggering it in the atrial model. A clinically feasible approach consists in a train of stimuli delivered at some specific location, with a decreasing interval between each stimulus. We define the *inducibility function* $\mathcal{I} : \Omega \to \{0, 1\}$ as follows: $\mathcal{I}(\boldsymbol{y}) = 1$ if the stimulation protocol delivered at $\boldsymbol{y}$ successfully induced AF, and zero otherwise. We check whether AF is induced or not in the model by checking whether the integral over Ω and a window of time of the currents (diffusion, stimulus, and ionic) is non-zero.

We consider 9 different models of atrial fibrillation on a fixed geometry with: 3 different fibrotic patterns (moderate - 50%, severe case 1 - 70% and severe case 2 - 70%), and 3 ablations (no ablation, PVI, and PVI + BOX ablation). For each case, we run 100 pacing locations [4]. In total, the dataset has 900 simulations of inducibility that we will use to assess the accuracy of the proposed methodology.

2.2 Fibrotic kernel signature on the atria

Simulating AF at human scale is computationally expensive. Therefore, we are interested in *learning* the classifier $\mathcal{I}$ from a sparse set of simulations. In this work, we use the heat kernel signature (HKS) [15], which is a technique to characterize points in geometries, and can be used for segmentation and shape matching among other applications. This point descriptor is based on the heat diffusion process on a given shape that captures concisely the intrinsic information from a geometry, up to isometry in an efficient, stable and multi-scaled way. The HKS is also invariant to rotations and translations.

The HKS can be computed efficiently using the heat kernel $k_t(\boldsymbol{x}, \boldsymbol{y})$, which represents evolution of the temperature over time t at point $\boldsymbol{y}$ when the initial temperature is a Dirac delta $\delta_{\boldsymbol{x}}(\boldsymbol{y})$ applied at point $\boldsymbol{x}$. Then, the HKS is a vector defined as $\text{hks}(\boldsymbol{x})_i = k_{t_i}(\boldsymbol{x}, \boldsymbol{x})$, for a finite number of time steps $t_i > 0$. These time steps are computed in a logarithmic progression and the signature is later normalized. Intuitively, the HKS represents how the temperature evolves over time in a point after applying an impulse at $t = 0$ at that location.

Instead of solving the diffusion equation, the HKS can be effectively calculated as a sum involving the eigenfunctions $\phi_i(\boldsymbol{x})$ and eigenvalues λ_i of the Laplace operator on the shape on interest:

$$k_t(\boldsymbol{x}, \boldsymbol{y}) = \sum_{i=0}^{+\infty} e^{-\lambda_i t} \phi_i(\boldsymbol{x}) \phi_i(\boldsymbol{y}) \tag{2}$$

Here, we introduce the *fibrotic kernel signature* (FKS) by incorporating information regarding the fibrotic pattern in the signature. Inspired by the diffusion operator of the monodomain equation (1), we first consider the elliptic operator $\mathcal{L}u \coloneqq -\operatorname{div}(\sigma_f(\boldsymbol{x})\nabla u)$ with homogeneous Neumann boundary condition, for $\sigma_f \in L^\infty(\Omega)$ and $u \in H^1(\Omega)$, that similarly to the Laplace operator has a countable spectrum of eigenvalues and eigenfunctions [12]. We remark that now

the spectrum depends on the fibrosis pattern. Then, we use Eq. (2) to compute the FKS. The function $\sigma_f(\boldsymbol{x})$ is 1 where there is healthy tissue and 0.5 where there is fibrotic tissue.

To compute the FKS, we use the first 100 eigenvalues, ordered by magnitude. Since the operator $\mathcal{L}$ is symmetric, eigenvalues are real. It is also possible to show that the first eigenvalue is zero, and the others are positive. To solve the eigenproblem we use the Finite Element Method with linear Lagrange elements on a hexahedral mesh with $\approx$ 700 000 nodes. We implemented the solver in `DOLFINx` interface and `SLEPc` library on Python. We set the `SLEPc` solver for a generalized non-Hermitian case. We use the same normalization of the signature per fibrotic pattern. We compute the normalization and time steps for the base case without ablation and also apply it to the cases with the same fibrotic pattern with ablation. We represent the ablation by setting $\sigma_f(\boldsymbol{x}) = 0.001$ in the ablation lines, effectively creating a barrier for the heat.

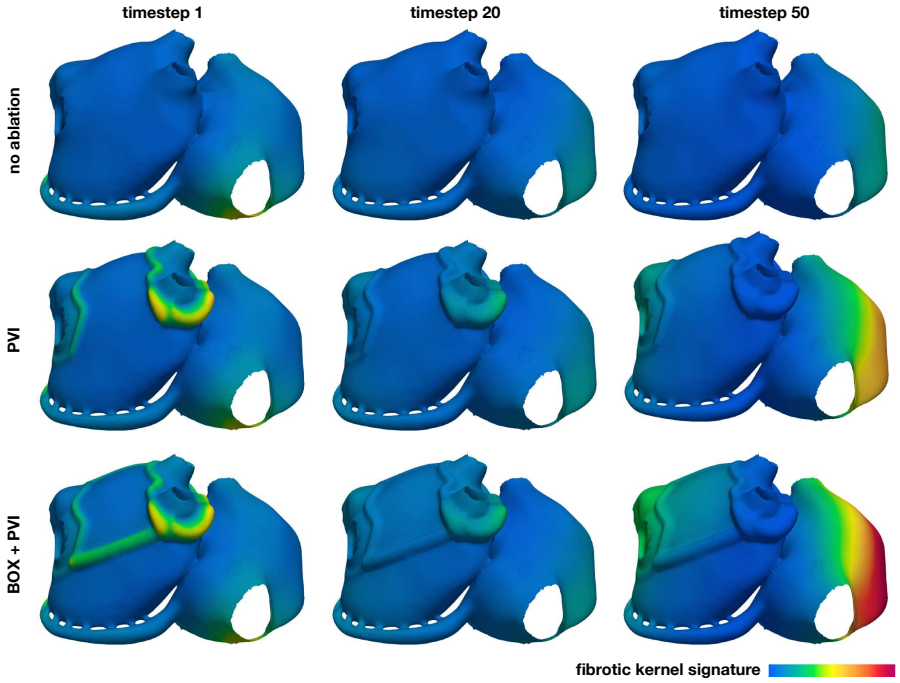

Fig. 1. Fibrotic kernel signature for moderate fibrosis case and three ablation patterns: no ablation, pulmonary vein isolation and PVI + BOX ablation. The no ablation case is taken as reference model for time-scaling and signature normalisation.

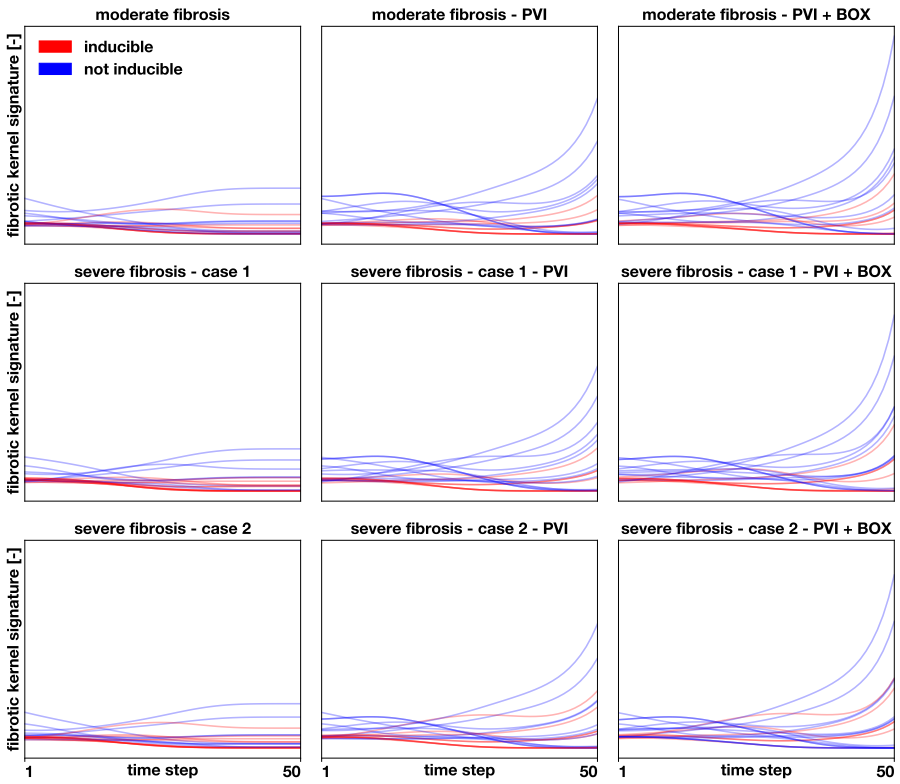

Fig. 2. Fibrotic kernel signature over time for 20 randomly selected points for 9 different cases. The red curves represent points where AF was induced and the blue curves are points where AF was not inducible. The non-ablation case is taken as reference model for time-scaling and signature normalisation. The abscissa represents the time vector, that is the input of the classifier.

2.3 Prediction of atrial fibrillation

Once we have computed the FKS for all the cases, we can use our dataset of 900 simulations and associate the signature at those locations to its inducibility. Then, we will train machine learning classifiers to learn to distinguish between signatures associated with inducible cases and signatures where AF was not induced. Concretely, we try to approximate the inducibility function $\mathcal{I}(\boldsymbol{y}) \approx f\big(\mathrm{fks}(\boldsymbol{y})\big)$ depending on the fibrotic kernel signature at a given location $\boldsymbol{y}$. Here $f(\cdot)$ represents some machine learning classifier. Once the classfier has been trained, we can predict the inducibility function without running simulations on a particular case. To assess the performance of this method, we perform leave-one-out cross-validation by case. We train with the simulations of 8 out of 9 cases, totalling up to 800 data points and test with 100 simulations of the unseen

case. We consider three classifiers: the k-neighbours classifier, random forest, and gradient boosting classifier as implemented in `scikit-learn` [10]. To assess how much the classifier is actually learning, we propose a naïve alternative, which we call "majority voting". Since the dataset is evaluated at the same locations for all 9 cases, we take a majority vote of the labels at a given location for the 8 cases to predict the label of the remaining case. For example, at a given location, if 5 out 8 models were inducible, we predict this point as inducible for the excluded case. This naïve classifier allows to determine whether the our method is just predicting based on the location of the point or is learning additional information. We use 2 metrics to evaluate the performance of these methods: balanced accuracy and the overall inducibility, which is computed as the fraction of points are predicted as inducible. This last metric is the most important to determine whether a proposed ablation treatment is effective or not.

3 Results

The computation of the FKS for each case took approximately 10 minutes on a modern workstation. Running one simulation for testing inducibility took an hour on a single GPU node at Swiss National Supercomputing Centre (CSCS). Disregarding the hardware differences, this represents a speed-up of 600 if we predict the inducibility of one case based on 100 simulations, discounting the simulations needed for training.

In Figure 1 we show the FKS for one fibrotic pattern at 3 different time steps and for the base case, the PVI and PVI + BOX ablation cases. We see that the PVI ablation affects conductivity in the periphery of the pulmonary veins, acting as barrier for heat propagation towards the rest of the atria, allowing heat accumulation on regions where base case presents lower temperatures. When adding the BOX ablation extends this accumulation of temperature in the roof of the left atrium.

Different examples of the FKS over time are shown in Figure 2 for the 20 randomly selected locations which correspond across the 9 cases considered. We note that locations that are not inducible tend to have a higher signature than the inducible points. We also observe that the application of the ablation patterns tend to modify the signatures, especially towards the end of time.

The results regarding the prediction of AF are summarized in Figures 3 and 4. When we analyze the balanced accuracy of our method for different levels of training data in Figure 3, we see that in general random forest tends to perform better in all cases. We achieve maximum balanced accuracies with the classifiers ranging from 75.8% for the severe fibrosis - case 1, to 95.8% in the case moderate fibrosis + PVI + BOX. The majority voting classifier tends to show a similar performance than the classifiers based on the FKS. However, for some cases, the FKS classifiers have higher or lower accuracy than the majority voting classifier, indicating that the FKS classifiers are not simply memorizing the location of the training points and matching it the test cases.

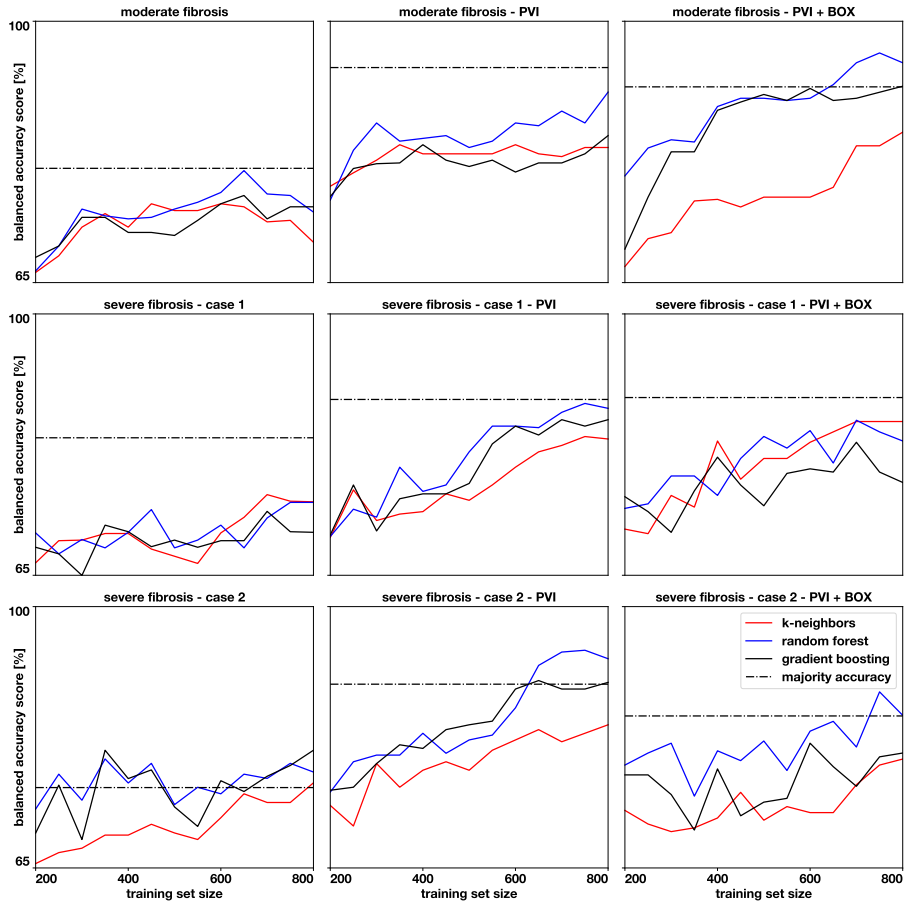

Fig. 3. Balanced accuracies for 9 different cases. We show in the solid lines the performance of 3 different machine learning classifiers as the training data increases. The dashed line represents a baseline naive classifier based on the inducibility of the other 8 cases.

Regarding the predictions of overall inducibility, we see bigger differences in Figure 4. First, we note that the ablation lines applied to the models have a marked effect on the inducibility. Applying PVI reduces the inducibility between 11 to 14 percentage points and applying PVI + BOX reduces the inducibility between 16 to 19 percentage points. The majority voting classifier, as expected, tends to predict a similar level of inducibility independent of the treatment applied to the case. For all cases, it predicts between 40 and 43% inducibility. The random forest classifier based on the FKS trained with 800 points can detect the changes in inducibility much better. The errors range between 1 to 6 percentage points inducibility. Also, the trend that the base case is more inducible than

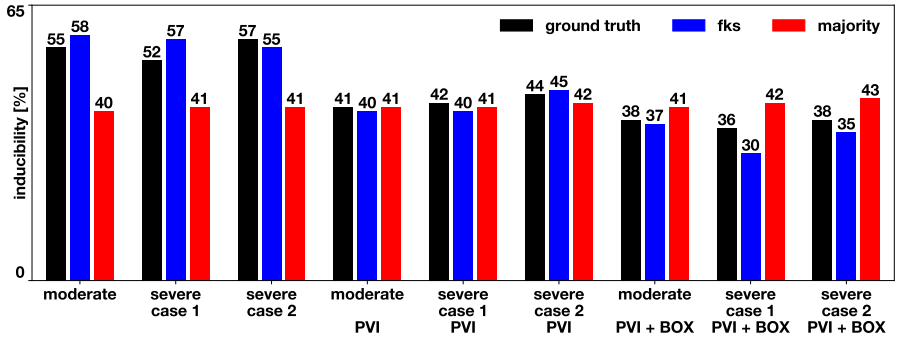

Fig. 4. Inducibility prediction for 9 cases from the FKS trained with 800 points from other cases, compared to the majority naive classifier. The FKS classifier shows better performance for most cases.

PVI and even more inducible than PVI + BOX is also correctly predicted by the FKS-based classifier.

4 Discussion

In this work, we present the fibrotic kernel signature, an efficient way to predict the inducibility of AF without running a computer simulation. We achieve this by creating a descriptor of the fibrotic substrate and the anatomy, here named "fibrotic kernel signature" (FKS). The FKS can be computed for all the points in the model at a fraction of the cost than running the simulations required for computing the inducibility for a case. We think that speeding up the calculation of inducibility maps is crucial to create better personalized ablation treatments within the time constraints of the clinical setting [2, 1]. FKS prediction does not require high-performance computing facilities: a desktop computer is sufficient.

When combined the FKS with simple machine learning algorithms, we see that we can predict the point-wise inducibility with good accuracy and the global inducibility with excellent accuracy. In Figure 3 we observe that the accuracies tend to increase as more data is available. We expect that the performance would be improved as we train the classifier with more cases. Our method allows to take advantage of all the simulations that we could run for different patients in order the improve the predictions with the FKS. We remark that no mapping between different anatomies is required when evaluating the classifier since the geometrical information flows into the FKS, on which the classifier depends. Importantly, the FKS is easily extendable thus to include the local fiber direction, just by redefining the elliptic operator. For instance, the fiber direction and fibrosis, combined in the conductivity tensor, can be estimated from electroanatomical mapping system data, as recently proposed [7, 6, 11].

Our study has some limitations. First, we only tested our method in a single geometry. Although the different patterns of fibrosis and ablation effectively change the geometry for the monodomain equation and the FKS, we still need to verify whether our method generalizes to other patient anatomies. Another limitation is that the FKS does not consider the fiber orientation and the cell type distribution. We conducted preliminary studies including the fibers, but there was no improvement in the results. Finally, we use a mesh resolution of 0.4 mm, which we know affects the inducibility [4], when compared to finer meshes. This will affect to time to solve the eigenproblem required by the FKS, but given that we are currently using only modest hardware, we could manage the larger models. On the other hand, our current mesh resolution is in line or already finer than those from other studies [2, 1]. We also did not account for uncertainty in the fibrosis pattern, which is known to be highly affected by the threshold strategy. Ideally, uncertainty could be introduced into our FKS definition to compute a mean FKS with associated covariance; this information then could be used in the classification problem.

In summary, we propose a novel method to predict atrial fibrillation without running simulations. We believe that the fibrotic kernel signature combined with machine learning techniques will enable faster and better planning of ablation treatments in atrial fibrillation patients.

Acknowledgments

FSC and TB acknowledge the support of the project FONDECYT-Iniciación 11220816. FSC also acknowledges the support of the project ERAPERMED-134 from ANID. This work was also financially supported by the Theo Rossi di Montelera Foundation, the Metis Foundation Sergio Mantegazza, the Fidinam Foundation, and the Horten Foundation to the Center for CCMC. SP also acknowledges the CSCS-Swiss National Supercomputing Centre (No. s1074).

References

1. Azzolin, L., Schuler, S., Dössel, O., Loewe, A.: A reproducible protocol to assess arrhythmia vulnerability in silico: Pacing at the end of the effective refractory period. Frontiers in physiology **12**, 420 (2021)
2. Boyle, P.M., Zghaib, T., Zahid, S., Ali, R.L., Deng, D., Franceschi, W.H., Hakim, J.B., Murphy, M.J., Prakosa, A., Zimmerman, S.L., et al.: Computationally guided personalized targeted ablation of persistent atrial fibrillation. Nature biomedical engineering **3**(11), 870–879 (2019)
3. Courtemanche, M., Ramirez, R.J., Nattel, S.: Ionic mechanisms underlying human atrial action potential properties: insights from a mathematical model. American Journal of Physiology-Heart and Circulatory Physiology **275**, H301–H321 (1998)
4. Gander, L., Pezzuto, S., Gharaviri, A., Krause, R., Perdikaris, P., Sahli Costabal, F.: Fast characterization of inducible regions of atrial fibrillation models with multi-fidelity gaussian process classification. Frontiers in Physiology p. 260 (2022)

5. Gharaviri, A., Bidar, E., Potse, M., Zeemering, S., Verheule, S., Pezzuto, S., Krause, R., Maessen, J.G., Auricchio, A., Schotten, U.: Epicardial fibrosis explains increased endo-epicardial dissociation and epicardial breakthroughs in human atrial fibrillation. Frontiers in Physiology **11**(68) (2020)
6. Grandits, T., Pezzuto, S., Costabal, F.S., Perdikaris, P., Pock, T., Plank, G., Krause, R.: 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. Springer (2021)
7. Kotadia, I., Whitaker, J., Roney, C., Niederer, S., O'Neill, M., Bishop, M., Wright, M.: Anisotropic cardiac conduction. Arrhythmia & Electrophysiology Review **9**(4), 202 (2020)
8. Loewe, A., Poremba, E., Oesterlein, T., Luik, A., Schmitt, C., Seemann, G., Dössel, O.: Patient-specific identification of atrial flutter vulnerability–a computational approach to reveal latent reentry pathways. Frontiers in Physiology **9**, 1910 (2019)
9. McDowell, K.S., Zahid, S., Vadakkumpadan, F., Blauer, J., MacLeod, R.S., Trayanova, N.a.: Virtual Electrophysiological Study of Atrial Fibrillation in Fibrotic Remodeling. Plos One **10**(2), e0117110 (2015). https://doi.org/10.1371/journal.pone.0117110
10. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. Journal of Machine Learning Research **12**, 2825–2830 (2011)
11. Ruiz Herrera, C., Grandits, T., Plank, G., Perdikaris, P., Sahli Costabal, F., Pezzuto, S.: Physics-informed neural networks to learn cardiac fiber orientation from multiple electroanatomical maps. Engineering with Computers **38**(5), 3957–3973 (2022)
12. Salsa, S.: Partial differential equations in action: from modelling to theory, vol. 99. Springer (2016)
13. Schotten, U., Verheule, S., Kirchhof, P., Goette, A.: Pathophysiological mechanisms of atrial fibrillation: a translational appraisal. Physiological Reviews **91**(1), 265–325 (2011). https://doi.org/10.1152/physrev.00031.2009
14. Serra, D., Romero, P., Garcia-Fernandez, I., Lozano, M., Liberos, A., Rodrigo, M., Bueno-Orovio, A., Berruezo, A., Sebastian, R.: An automata-based cardiac electrophysiology simulator to assess arrhythmia inducibility. Mathematics **10**(8), 1293 (2022)
15. Sun, J., Ovsjanikov, M., Guibas, L.: A concise and provably informative multi-scale signature based on heat diffusion. Computer Graphics Forum **28**(5), 1383–1392 (2009)
16. Tsao, C.W., Aday, A.W., Almarzooq, Z.I., Alonso, A., Beaton, A.Z., Bittencourt, M.S., Boehme, A.K., Buxton, A.E., Carson, A.P., Commodore-Mensah, Y., et al.: Heart disease and stroke statistics—2022 update: a report from the american heart association. Circulation **145**(8), e153–e639 (2022)