

# Physics-informed neural networks to learn cardiac fiber orientation from multiple electroanatomical maps

Carlos Ruiz Herrera<sup>†,1,4</sup>, Thomas Grandits<sup>†,2,6</sup>, Gernot Plank<sup>6</sup>, Paris Perdikaris<sup>3</sup>, Francisco Sahli Costabal<sup>1,4\*</sup> and Simone Pezzuto<sup>5</sup>

<sup>1</sup>Department of Mechanical and Metallurgical Engineering, School of Engineering, Pontificia Universidad Católica de Chile, Santiago, Chile.

<sup>2</sup>Institute of Mathematics and Scientific Computing, University of Graz, Graz, Austria.

<sup>3</sup>Department of Mechanical Engineering and Applied Mechanics, University of Pennsylvania, Philadelphia, USA.

<sup>4</sup>Institute for Biological and Medical Engineering, Schools of Engineering, Medicine and Biological Sciences, Pontificia Universidad Católica de Chile, Santiago, Chile.

<sup>5</sup>Center for Computational Medicine in Cardiology, Euler Institute, Università della Svizzera italiana, via Buffi 13, Lugano, 6900, Switzerland.

<sup>6</sup>Gottfried Schatz Research Center - Division of Biophysics, Medical University of Graz, Graz, Austria.

\*Corresponding author(s). E-mail(s): [fsc@ing.puc.cl](mailto:fsc@ing.puc.cl);

Contributing authors: [caruiz2@uc.cl](mailto:caruiz2@uc.cl);

[thomas.grandits@uni-graz.at](mailto:thomas.grandits@uni-graz.at); [gernot.plank@medunigraz.at](mailto:gernot.plank@medunigraz.at);

[gpg@seas.upenn.edu](mailto:gpg@seas.upenn.edu); [simone.pezzuto@usi.ch](mailto:simone.pezzuto@usi.ch);

## Abstract

We propose FiberNet, a method to estimate *in-vivo* the cardiac fiber architecture of the human atria from multiple catheter recordings of the

---

<sup>†</sup>These authors contributed equally to this work

electrical activation. Cardiac fibers play a central role in the electro-mechanical function of the heart, yet they are difficult to determine *in-vivo*, and hence rarely truly patient-specific in existing cardiac models. FiberNet learns the fiber arrangement by solving an inverse problem with physics-informed neural networks. The inverse problem amounts to identifying the conduction velocity tensor of a cardiac propagation model from a set of sparse activation maps. The use of multiple maps enables the simultaneous identification of all the components of the conduction velocity tensor, including the local fiber angle. We extensively test FiberNet on synthetic 2-D and 3-D examples, diffusion tensor fibers, and a patient-specific case. We show that 3 maps are sufficient to accurately capture the fibers, also in the presence of noise. With fewer maps, the role of regularization becomes prominent. Moreover, we show that the fitted model can robustly reproduce unseen activation maps. We envision that FiberNet will help the creation of patient-specific models for personalized medicine. The full code is available at <http://github.com/fsahli/FiberNet>.

**Keywords:** Cardiac Fibers, Physics-Informed Neural Networks, Cardiac Electrophysiology, Anisotropic conduction velocity, Eikonal Equation, Deep learning

## 1 Introduction

In recent years, the future vision of precision cardiology through digital twinning has gained traction [1, 2]. In such scenarios, a digital replica of a patient’s heart is generated from a variety of measurements in the clinic. Such replica can be used in a multitude of ways, from predicting intervention outcomes to building cohorts for augmented drug trials. However, building such cardiac digital twins faces several difficulties, such as dealing with noisy measurements, or the problem of having multiple possible parameterizations explaining the encountered measurements, mathematically known as ill-posedness [3].

The fiber distribution in the heart is a key determinant of cardiac function and, as such, it has a prominent role in digital twinning [4]. The electrical stimulus activating the myocardium travels at different speeds, depending on whether the propagation occurs along or across the fibers [5]. More precisely, the electrical conduction within the myocardium is *anisotropic*, because the electrical propagation is faster in the direction of the fibers. There are multiple determinants of anisotropic conduction, and some may uncover a pathological condition [6]. Therefore, an imprecise knowledge of the anisotropic conduction in digital twinning may yield wrong predictions or even non-physiological behaviors.

In patient-specific models, cardiac fibers are generally arranged following some rules, dictated by prior histological knowledge on their distribution in the heart [4, 7]. These algorithms are well-established for the ventricles [8]. More recently, some rule-based methods and atlas-based for the atria have

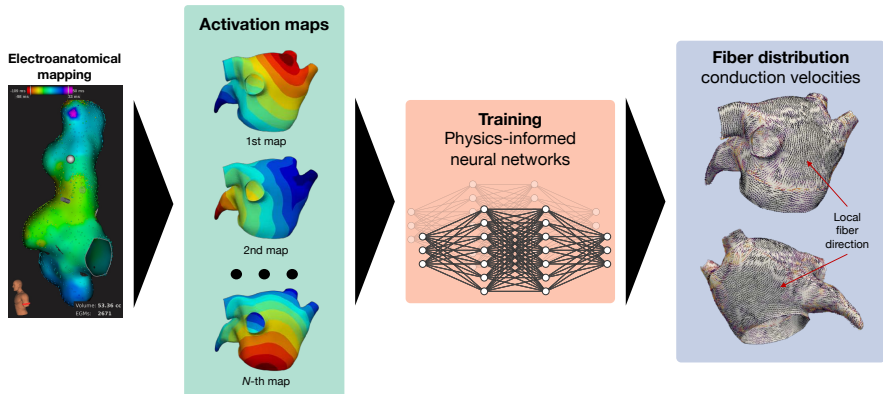
also been proposed [9–13]. In all cases, the fiber field is anatomically-tailored, but technically speaking not yet patient-specific. Diffusion-Tensor Magnetic Resonance (DT-MR) imaging of the heart, the gold-standard tool for sub-millimeter imaging of cardiac fibers [14–16], is unfortunately not feasible *in-vivo* for the atrial wall.

A clinically-viable way to infer the local fiber direction is based on the conductivity properties of the tissue. Clinically, the conduction velocity (CV) can be indirectly determined with an electroanatomical mapping system [17, 18], a minimally-invasive, catheter-based tool to record the local activation of the endocardial wall. Electroanatomical maps are common in clinical electrophysiology, because they are routinely acquired before ablation therapy, and are relatively easy to export in text format for further inspection. By comparing the local CV obtained from multiple maps, it is therefore possible to derive the direction of fastest conduction, that is the fiber direction [19].

In this work, we similarly aim at determining the fibers from multiple maps. Rather than working locally, we solve an inverse problem where the conductivity tensor can be deduced from the activation map through a propagation model, namely the eikonal equation [20, 21]. This formulation has a number of advantages. First, as the fiber field is recovered, we simultaneously fit a predictive model that, in principle, can faithfully reproduce the observed activations. Second, the model can extrapolate, in a physiological manner, the activation map in regions where data is scarce or absent [22]. Third, the inverse problem can be easily informed by prior histological knowledge on the fiber distribution through a regularization term or, more generally, a Bayesian approach, for instance by using rule-based fiber fields [13]. Therefore, the fiber field could be extrapolated as well in a consistent and physiological manner.

Physics Informed Neural Networks (PINNs) are a recently developed variant of machine learning-based methods to efficiently solve inverse problems that are governed by partial differential equations [23]. They have been shown to accurately model complex physical problems with a small number of known data points. Unlike regular neural networks which typically require vast amounts of labeled data to make accurate predictions, PINNs can learn much quicker to simulate this kind of system due to the incorporation of physical laws – represented as systems of partial differential equations – into their loss functions. Moreover, PINNs are a genuinely mesh-free method, overcoming the issue of generating meshes of complex domains like the heart. In fact, the model and data are penalized at collocation and data points, respectively, while neural networks representing the quantities of interests can be evaluated at any spatial location. Thus, no topological information from the geometrical model is required.

In this work, we propose FiberNet, a PINN-based method to solve the inverse problem of identifying fibers in the heart from electroanatomical maps (Fig. 1). This is achieved by simultaneously fitting multiple neural networks to multiple electroanatomical maps while using a common network that predicts the conduction velocity tensor at different locations. In this sense, FiberNet



**Fig. 1** FiberNet translates a set of electroanatomical maps into a continuous estimate of the fiber field and conduction velocity. These estimates can be used for simulating cardiac activation in a predictive model based on the eikonal equation. Internally, FiberNet uses physics-informed neural networks to constrain the parameter space during the training phase.

extends our previous work on the same problem [21, 22, 24], which only used a single activation map and was based on a different representation of the conductivity tensor. We extensively validate FiberNet with several numerical experiments and real data. The first set of experiments consists of a completely synthetic example with a 2-D planar geometry and a 3-D atrial geometry, where we compare the single- and multi-map approaches. Secondly, we apply FiberNet to an atlas of diffusion tensor image fiber fields of human atria [25]. Here, the ground-truth activation times are generated with the Fast Iterative Method (FIM) simulation for triangular surfaces [26, 27]. Moreover, we validate the predictive capabilities of FiberNet by first fitting the fibers on training maps, and then generating an unseen map and attempting to predict it using the learned fibers. Finally, we apply FiberNet to clinically-obtained EAMs of the left atrium, thus proving the feasibility of the approach in the clinical setting.

This paper is structured as follows. In Section 2 we review some of the existing methods to estimate the conduction velocity from sparse activation recordings, a highly relevant topic for this work. FiberNet is introduced in Section 3, along with the PINN framework. The identifiability of the fibers from single and multiple maps and the role of the regularization are discussed in Section 3.3. The last part of our work is devoted to numerical validation and applications. In Section 4 we verify the robustness to noise of FiberNet in a synthetic 2-D example and on a 3-D atrial geometry. An extensive benchmark with a publicly available atrial fiber atlas [12, 25] is given in Section 5. Finally, in Section 6 we apply FiberNet to a clinical data set. A discussion and outlook are provided in Section 7.

## 2 Related works

The problem of identifying cardiac fibers from electrical data of atrial activation has been considered only very recently, see e.g. [19, 21, 22, 24]. More broadly, however, the topic is related to the identification of the local conduction velocity (CV) of the tissue, for which several methods have been proposed thus far. In this section, we briefly review the most common approaches. We group them into two categories: data-driven and physics-driven methods. In the former, the CV is directly estimated from the data through either geometrical or statistical arguments; in the latter, the conduction parameters of a cardiac model are fitted to the data, hence the CV is a byproduct of the forward model.

### 2.1 Data-driven methods

Trivially, the CV is the ratio between the distance traveled by an activation front in a given amount of time. The activation time can be detected from local electrograms obtained with contact mapping systems. The location of the electrodes on the catheter is also tracked with good accuracy. Therefore, with sufficient coverage of the atrial surface, it is possible to estimate the local CV in a purely data-driven manner, with no prior assumption on the underpinning physiology.

A first class of methods estimates the local CV  $\theta(\mathbf{x})$  at some given location  $\mathbf{x}$  on the basis of the temporal differences in activation between  $\mathbf{x}$  and its neighbors [17, 28–31]. These methods are easy to implement and very fast to execute. Here, the main challenge is dealing with noise and inconsistency in the data. Both the location and activation time are subject to noise, hence their incremental ratio, the CV, is greatly susceptible to large variations due to uncertainty. In this scenario, least-squares approaches [29, 30] may perform better than purely geometric methods [17, 28]. Nonetheless, a careful pre-processing of the input data is always necessary, in order to avoid unrealistically fast or low CVs [17, 32]. In some cases, very high CVs may actually be due to physiological phenomena, such as breakthroughs and front collisions where, however, the CV is not well-defined.

Rather than estimating the CV from point-wise measurements, other authors suggested to first interpolate or recover a smooth activation map from the data, say  $\tilde{\phi}(\mathbf{x})$ , and then compute the conduction velocity from its definition, that is

$$\theta(\mathbf{x}) = \|\nabla\tilde{\phi}\|^{-1}. \quad (1)$$

This has been recently done with Gaussian Process Regression (GPR) on manifolds [18]. In GPR, the properties of the kernel (smoothness, correlation length) can be optimized to capture the physics of the problem and reduce generalization error. Furthermore, GPR is probabilistic by nature and enables uncertainty quantification in CV estimates and active learning [33].

Since cardiac conduction is anisotropic [6], the value of  $\theta(\mathbf{x})$  also depends on the propagation direction  $\mathbf{p}$ . In all previous methods, therefore, the estimated

CV depends on the activation map, since  $\mathbf{p} = \nabla\tilde{\phi}(\mathbf{x})/\|\nabla\tilde{\phi}(\mathbf{x})\|$ . With the use of multiple activation maps, or by accounting for prior knowledge on the fiber distribution, it is however possible to simultaneously estimate the longitudinal and transverse CV, as recently proposed in [19]. In this case, it is tacitly assumed that the CV in the myocardium obeys a Riemannian metric, that is  $\theta(\mathbf{x}, \mathbf{p}) = \sqrt{\mathbf{D}(\mathbf{x})\mathbf{p} \cdot \mathbf{p}}$ , for some symmetric, positive definite tensor field  $\mathbf{D}(\mathbf{x})$ . In the 2-dimensional case, as for the atrial surface, the tensor  $\mathbf{D}(\mathbf{x})$  is determined by three independent parameters, e.g., the fiber angle and the CV along and across the fibers. Therefore, at least 3 independent activation maps are required to uniquely identify  $\mathbf{D}$ .

Interestingly, from Eq. (1) with  $\theta(\mathbf{x}, \mathbf{p})$  as above, we recover the *anisotropic eikonal equation* [20], see also Eq. (3) in Sec. 3.2. Therefore,  $\mathbf{D}(\mathbf{x})$  may be recovered by imposing Eq. (1) for multiple activation maps  $\tilde{\phi}_i(\mathbf{x})$ ,  $i = 1, \dots, M$  at each point  $\mathbf{x}$  of the surface [19]. The corresponding algebraic system for  $M = 3$ , however, may have no unique solution. For instance, this may happen when  $\nabla\tilde{\phi}_i(\mathbf{x})$  and  $\nabla\tilde{\phi}_j(\mathbf{x})$ , for  $i \neq j$ , are parallel at some  $\mathbf{x}$ . Similarly, it may happen that the fitted  $\mathbf{D}(\mathbf{x})$  is not positive-definite, hence there is no physiological solution. An alternative approach, valid for an arbitrary number of maps and always ensuring at least one solution, is based on minimizing the residual

$$\sum_{i=1}^M \left( \sqrt{\mathbf{D}(\mathbf{x})\nabla\tilde{\phi}_i(\mathbf{x}) \cdot \nabla\tilde{\phi}_i(\mathbf{x})} - 1 \right)^2, \quad (2)$$

on the space of parameters defining  $\mathbf{D}$ . It is worth noting that finding  $\mathbf{D}$  can be understood as a fitting-an-ellipse-to-points problem, for which more robust algorithms are available [34].

Finally, it is worth mentioning that there exist methods for estimating the CV directly from the temporal dynamics of the electrograms, without the need for the activation map [35, 36].

## 2.2 Physics-driven methods

Data-driven methods may suffer when data is scarce or unevenly distributed, as it often happens with electroanatomical maps. In absence of data, prior physiological knowledge may, however, be enforced to still recover a plausible CV. Physics-driven (or physiology-driven) approaches follow this path. For cardiac electrophysiology, several propagation models can be used to constrain the CV to the activation map. These models can be based on the eikonal equation [21, 22, 24, 33], reaction-diffusion systems [37–39], or a multi-fidelity combination thereof [40]. The typical workflow aims at minimizing the observational residual between simulated and recorded activation, by optimizing the (distributed) conductivity parameters of the model. Once the optimal parameters of the model have been obtained, the local CV is trivially computed.

Physics-driven approaches are generally more robust than purely data-driven methods, as they allow one to weigh data fidelity against model fidelity,

through a regularization term. More importantly, physics-driven methods potentially provide a predictive model of cardiac electrophysiology, thus they can be employed in personalized therapeutic approaches [41]. However, they have a significantly higher computational footprint, both in terms of memory and time. Moreover, some other parameters, such as the early activation sites, may potentially influence the CV reconstruction [21]. In this respect, a good trade-off between purely data-driven and physics-driven methods consists in accounting for the physics only weakly [23, 24, 33], for instance through a penalization term in the loss function such as Eq. (2), rather than enforcing the model point-wise. This observation further motivates the method presented below.

## 3 Methods

### 3.1 Propagation model

Let  $\mathcal{S} \subset \mathbb{R}^3$  be a smooth orientable surface representing, for instance, the left atrial endocardium. We model cardiac activation with the anisotropic eikonal equation [20]. We do not consider diffusion or curvature correction terms. The eikonal equation models the arrival times  $\phi(\mathbf{x})$  resulting from the spread of an electric activation wavefront within the myocardium, propagating with direction-dependent CV. The equation reads as follows

$$\sqrt{\mathbf{D}(\mathbf{x})\nabla\phi(\mathbf{x}) \cdot \nabla\phi(\mathbf{x})} = 1, \quad (3)$$

where  $\mathbf{D}(\mathbf{x}) \in \mathbb{R}^{3 \times 3}$  is a symmetric, positive-definite tensor field representing the conductivity. Specifically,

$$\theta(\mathbf{x}, \mathbf{p}) = \sqrt{\mathbf{D}(\mathbf{x})\mathbf{p} \cdot \mathbf{p}} \quad (4)$$

is the conduction velocity in direction  $\mathbf{p}$ .

We define the local fiber direction as the direction of fastest propagation and we denote it by  $\mathbf{l}(\mathbf{x}) \in \mathbb{R}^3$ . Since the propagation is constrained on the atrial surface,  $\mathbf{l}(\mathbf{x})$  is orthogonal to  $\mathbf{n}(\mathbf{x})$ , the normal direction of  $\mathcal{S}$ . A generic way to represent  $\mathbf{D}$  is therefore as follows:

$$\mathbf{D}(\mathbf{x}) = v_l^2(\mathbf{x})\mathbf{l}(\mathbf{x}) \otimes \mathbf{l}(\mathbf{x}) + v_t^2(\mathbf{x})\mathbf{t}(\mathbf{x}) \otimes \mathbf{t}(\mathbf{x}), \quad (5)$$

where  $\mathbf{t}(\mathbf{x}) \in \mathbb{R}^3$  is the transverse vector field, orthogonal to both  $\mathbf{l}$  and  $\mathbf{n}$ , while  $v_l$  and  $v_t$  are the conduction velocities respectively along  $\mathbf{l}$  and  $\mathbf{t}$ . We note that the velocity in the normal direction is zero.

Given an orthonormal basis of the tangent space of  $\mathcal{S}$  at a location  $\mathbf{x}$ , denoted by  $\{\mathbf{v}_1(\mathbf{x}), \mathbf{v}_2(\mathbf{x})\}$ , we formulate the fiber and transverse directions as

$$\mathbf{l}(\mathbf{x}) = a\mathbf{v}_1(\mathbf{x}) \otimes \mathbf{v}_1(\mathbf{x}) + \sqrt{1-a^2}\mathbf{v}_2(\mathbf{x}) \otimes \mathbf{v}_2(\mathbf{x}), \quad (6)$$

$$\mathbf{t}(\mathbf{x}) = -\sqrt{1 - a^2} \mathbf{v}_1(\mathbf{x}) \otimes \mathbf{v}_1(\mathbf{x}) + a \mathbf{v}_2(\mathbf{x}) \otimes \mathbf{v}_2(\mathbf{x}), \quad (7)$$

where  $a \in [-1, 1]$  is the cosine of the angle between the fiber direction and  $\mathbf{v}_1(\mathbf{x})$ . While any orthonormal basis on the tangent space would be feasible (even piecewise-constant), later regularization techniques greatly benefit from the smoothness of the basis. We therefore employed the vector heat method [42], implementing parallel transport of vectors on manifolds, to propagate a first initial vector across the entire atrial surface, thus to obtain  $\mathbf{v}_1(\mathbf{x})$ . The second vector field  $\mathbf{v}_2(\mathbf{x})$  is obtained by orthogonalization. See also Appendix A and Figure A1.

In summary, the conductivity tensor is defined by 3 scalar fields,

- $a(\mathbf{x}) \in [-1, 1]$ , the cosine of the fiber angle in the  $\{\mathbf{v}_1, \mathbf{v}_2\}$  basis,
- $e_1(\mathbf{x}) := v_l^2(\mathbf{x}) \geq 0$ , the square of the longitudinal velocity, and
- $e_2(\mathbf{x}) := v_t^2(\mathbf{x}) \geq 0$ , the square of the transversal velocity.

### 3.2 Learning fibers from multiple maps

We consider the problem of simultaneously identifying  $a(\mathbf{x})$ ,  $e_1(\mathbf{x})$  and  $e_2(\mathbf{x})$  from a set of  $M \geq 1$  electroanatomical maps. For convenience, we define

$$\mathbf{d}(\mathbf{x}) = [a(\mathbf{x}), e_1(\mathbf{x}), e_2(\mathbf{x})],$$

so that the conductivity tensor  $\mathbf{D}(\mathbf{x})$ , through Eq. (5)-(7), is a function of  $\mathbf{d}(\mathbf{x})$ , that is

$$\mathbf{D}(\mathbf{x}) = \tilde{\mathbf{D}}(\mathbf{d}(\mathbf{x})).$$

We solve the identification problem with the PINN framework, extending our previous work [24]. Here, we approximate *each* activation map  $\phi_m(\mathbf{x})$ ,  $m = 1, \dots, M$ , and the conductivity parameters vector  $\mathbf{d}(\mathbf{x})$  with a total number of  $M + 1$  artificial neural networks,

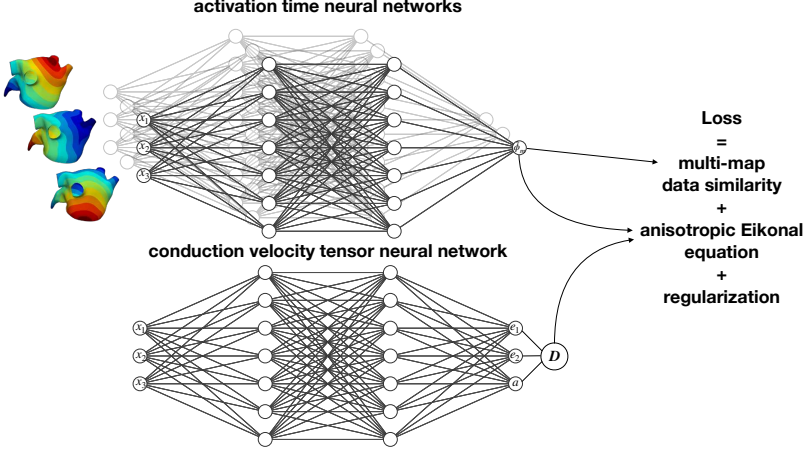
$$\begin{aligned} \phi_m(\mathbf{x}) &\approx \hat{\phi}_m(\mathbf{x}) := \text{NN}(\mathbf{x}; \boldsymbol{\theta}_{\phi_m}), \quad m = 1, \dots, M \\ \mathbf{d}(\mathbf{x}) &\approx \hat{\mathbf{d}}(\mathbf{x}) := \text{NN}(\mathbf{x}; \boldsymbol{\theta}_D), \end{aligned}$$

where  $\boldsymbol{\theta}_{\phi_m}$  and  $\boldsymbol{\theta}_D$  denote the trainable parameters of the networks. The approximated conductivity tensor trivially follows,  $\hat{\mathbf{D}}(\mathbf{x}) = \tilde{\mathbf{D}}(\hat{\mathbf{d}}(\mathbf{x}))$ . The output layer of  $\hat{\phi}_m(\mathbf{x})$  is the sigmoid function, hence it ranges in  $[0, 1]$ , whereas for  $\hat{\mathbf{d}}(\mathbf{x})$  we consider a sigmoid scaled by a factor  $C > 0$  for  $e_1$  and  $e_2$ , and a hyperbolic tangent activation function for the angle  $a$ .

To train our model, we take into consideration the data from the maps, the residual from the eikonal equation, and regularization terms. In Figure 2 we present a schematic of our approach. We build a loss function that includes all these components, as follows:

$$\mathcal{L}(\boldsymbol{\theta}_\phi, \boldsymbol{\theta}_D) := \mathcal{L}_{\text{data}} + \alpha_m \mathcal{L}_{\text{eiko}} + \alpha_e \mathcal{L}_{\text{cv}} + \alpha_a \mathcal{L}_{\text{ang}}, \quad (8)$$





**Fig. 2** Schematic representation of the physics-informed neural networks employed in this work. We use multiple neural networks to represent each of the activation times received for training. We use one neural network to represent the conduction velocity tensor. From here, we can compute the fiber orientation. We train all the networks simultaneously to satisfy the data from the activation time maps and the Eikonal equation, which links activation times to conduction velocities.

where  $\alpha_m$ ,  $\alpha_e$  and  $\alpha_a$  are hyper-parameters. The hyper-parameter  $\alpha_m$  controls the relevance of eikonal equation, while the hyper-parameters  $\alpha_e$  and  $\alpha_a$  control the amount of regularization applied to the conduction velocities and angles, respectively. We also collect all the neural network weights  $\{\theta_{\phi_m}\}_{m=1}^N$  in a single vector  $\theta_\phi$ .

In detail, we consider a set of data points, denoted by  $\mathbf{x}_{m,i} \in \mathcal{S}$ , where  $m = 1, \dots, M$  indicates the map number, and  $i = 1, \dots, N_m$  the point number. For each point, we consider a recorded activation time, denoted by  $\phi_m(\mathbf{x}_{m,i})$ . Each activation map is also divided by a factor  $T_{\max}$ , so that the data ranges in  $[0, 1]$ . Furthermore, we have a set of collocation points  $\mathbf{y}_j \in \mathcal{S}$ ,  $j = 1, \dots, N_C$ , where the eikonal model is enforced. The terms in the loss function are as follows,

$$\begin{aligned} \mathcal{L}_{\text{data}} &= \frac{1}{M} \sum_{m=1}^M \frac{1}{N_m} \sum_{i=1}^{N_m} \left( \hat{\phi}(\mathbf{x}_{m,i}) - \phi(\mathbf{x}_{m,i}) \right)^2, \\ \mathcal{L}_{\text{eiko}} &= \frac{1}{MN_C} \sum_{m=1}^M \sum_{j=1}^{N_C} \left( T_{\max} \sqrt{\hat{\mathbf{D}}(\mathbf{y}_j) \nabla \hat{\phi}_m(\mathbf{y}_j) \cdot \nabla \hat{\phi}_m(\mathbf{y}_j)} - 1 \right)^2, \\ \mathcal{L}_{\text{cv}} &= \frac{1}{N_C} \sum_{j=1}^{N_C} \left( H_{\delta_e}(\nabla e_1(\mathbf{y}_j)) + H_{\delta_e}(\nabla e_2(\mathbf{y}_j)) \right), \\ \mathcal{L}_{\text{ang}} &= \frac{1}{N_C} \sum_{j=1}^{N_C} H_{\delta_a}(\nabla a(\mathbf{y}_j)). \end{aligned}$$

We note that the eikonal model is not enforced for all  $\mathbf{x} \in \mathcal{S}$ , as generally done in PDE-constrained optimization, but rather on a set of collocation points  $\{\mathbf{y}_j\}_{j=1}^{N_C}$ . In this respect, the method is *mesh-free*, since there is no need for a triangulation of  $\mathcal{S}$  to represent the quantities of interest. Moreover, the eikonal equation is never actually solved, an important aspect since we do not have a precise knowledge of the earliest activation sites.

For the regularization terms, we consider the Huber Total Variation function

$$H_\delta(\mathbf{q}) = \begin{cases} \frac{1}{2\delta} \|\mathbf{q}\|^2, & \text{if } \|\mathbf{p}\| \leq \delta, \\ \|\mathbf{q}\| - \frac{1}{2}\delta, & \text{otherwise.} \end{cases} \quad (9)$$

The Huber total variation penalization is used in this problem as an additional restriction for the minimization because it has been shown to favor piecewise constant solutions [43]. We use a different penalization for the fiber angle and the conduction velocities to be able to have more control over their individual regularity, as shown in Section 4.

In summary, we train the neural networks with the loss function (8) to find the weights  $\boldsymbol{\theta}_\phi$  and  $\boldsymbol{\theta}_D$ , with hyper-parameters  $\alpha_m$ ,  $\alpha_e$ ,  $\alpha_a$ ,  $\delta_e$ ,  $\delta_a$ .

### 3.3 Identifiability

In view of the discussion in Section 2.1, multiple activations are required to reconstruct the fiber field. More precisely, at least 3 maps, along with some conditions, are needed to uniquely determine  $\mathbf{d}(\mathbf{x})$ . We show here that thanks to the proposed regularization mechanisms this is not the case for FiberNet. In fact, we can always obtain an estimate of the fiber field, even from a single map, as done in [24]. However, multiple maps will significantly improve the estimation of the fiber field and the CVs, as we extensively show in Section 4.

As an example, let us consider the problem in a 2-D square domain, and suppose that the true conductivity tensor is  $\mathbf{D} \in \mathbb{R}^{2 \times 2}$ , a symmetric, positive-definite, constant tensor. Here, the fiber direction is the principal eigenvector of  $\mathbf{D}$ . If we start an activation at the origin, the true activation map is as follows:

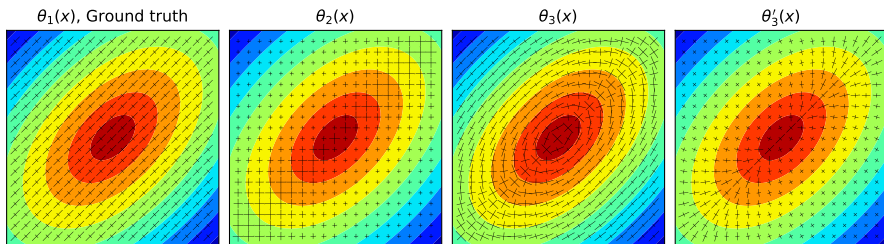
$$\phi(\mathbf{x}) = \sqrt{\mathbf{D}^{-1} \mathbf{x} \cdot \mathbf{x}}. \quad (10)$$

We construct now multiple conductivity tensors yielding the same activation map. We can rewrite the eikonal model as

$$\theta(\mathbf{x}, \mathbf{p}) \|\nabla \phi(\mathbf{x})\| = 1, \quad \mathbf{p} = \frac{\nabla \phi(\mathbf{x})}{\|\nabla \phi(\mathbf{x})\|},$$

where  $\theta(\mathbf{x}, \mathbf{p})$  is the conduction velocity in the propagation direction  $\mathbf{p}$ . Then, as anticipated above, the most trivial choice is:

$$\theta_1(\mathbf{x}, \mathbf{p}) = \sqrt{\mathbf{D} \mathbf{p} \cdot \mathbf{p}}.$$



**Fig. 3** Multiple conductivity tensors reproducing the same activation map. The tensor is represented with its eigenvectors rescaled by the eigenvalues. The solution with lowest Huber norm is  $\theta_1(\mathbf{x}, \mathbf{p})$ , which also corresponds to the ground truth.

We can also reproduce Eq. (10) with an *isotropic* model but spatially-varying conduction velocity:

$$\theta_2(\mathbf{x}, \mathbf{p}) = \sqrt{\frac{\mathbf{D}^{-1} \mathbf{x} \cdot \mathbf{x}}{\mathbf{D}^{-1} \mathbf{x} \cdot \mathbf{D}^{-1} \mathbf{x}}}.$$

Even by enforcing a space-varying anisotropic model,

$$\theta_3(\mathbf{x}, \mathbf{p}) = \sqrt{\tilde{\mathbf{D}}(\mathbf{x}) \mathbf{p} \cdot \mathbf{p}},$$

we can still find multiple choices, besides  $\tilde{\mathbf{D}}(\mathbf{x}) = \mathbf{D}$ . Consider the following transversely isotropic tensor:

$$\tilde{\mathbf{D}}(\mathbf{x}) = e_2(\mathbf{x}) \mathbf{I} + (e_1(\mathbf{x}) - e_2(\mathbf{x})) \mathbf{l}(\mathbf{x}) \otimes \mathbf{l}(\mathbf{x}),$$

where  $\mathbf{l}(\mathbf{x})$  is the fiber direction. After substituting it in the previous expression, we have

$$\theta_3(\mathbf{x}, \mathbf{p}) = \sqrt{e_2(\mathbf{x}) + (e_1(\mathbf{x}) - e_2(\mathbf{x})) (\mathbf{l}(\mathbf{x}) \cdot \mathbf{p})^2},$$

thus, by choosing  $\mathbf{l}$  such that  $\mathbf{l}(\mathbf{x}) \cdot \mathbf{p} = 0$  for all  $\mathbf{x}$  and  $e_2(\mathbf{x}) = \theta_2^2(\mathbf{x})$ , we obtain a solution for any choice of  $e_1(\mathbf{x}) > e_2(\mathbf{x})$ . Alternatively, by taking  $\mathbf{l}(\mathbf{x}) = \mathbf{p}$  and  $e_1(\mathbf{x}) = \theta_2^2(\mathbf{x})$ , we again have an infinite number of solutions by varying  $e_2(\mathbf{x}) < e_1(\mathbf{x})$ . Therefore, as expected, a single map is not sufficient to fully qualify the tensor  $\mathbf{D}$  from an algebraic point of view. See Figure 3 for a visual example.

Which fiber field is the learning algorithm going to prefer? Suppose we have a map sampled from Eq. (10) in absence of noise. Then, all the above choices for  $\theta$  cancel the fidelity term *and* the eikonal penalization term in the loss function, because they faithfully reproduce the true activation map. All these cases are global minima of the loss function (8), at least when no regularization is present. The neural networks would simply interpolate the data, in a least-squares sense. In the presence of regularization, however, the landscape of the

loss function differs significantly. In fact,  $\theta_1$  would likely be favored because, being constant in space, it has the smallest Huber norm amongst all possible choices. However, this choice is mostly dictated by the *prior* assumption encoded by the regularization itself, rather than the data, thus the reconstructed fibers may sensibly differ from the ground truth, especially in the presence of heterogeneity. Therefore, a reliable method for reconstructing the fibers should rely on multiple maps, as proposed in our approach.

## 4 Numerical assessment

We implemented FiberNet using Tensorflow [44]. For all experiments, we use 4 CPUs of an AMD EPYC 7702 64-Core Processor in parallel for a fixed number of iterations of ADAM [45] with mini-batches and the default hyper-parameters. We ran 4 different sets of experiments: a synthetic 2-dimensional case, a synthetic 3-dimensional case on a patient-derived geometrical model, 7 different cases where the fiber orientations were measured with DT-MR imaging, and one case where 3 different maps were obtained for a patient. For each of these experiments, we report both activation time and fiber orientation error. For the activation times, we report the root mean squared error. We quantify the fiber error as:  $\frac{1}{N_S} \sum_S \arccos(\mathbf{f} \cdot \hat{\mathbf{f}})$ , where  $\mathbf{f}$  is the ground truth direction of the fibers used to generate data and  $\hat{\mathbf{f}}$  is the predicted fiber direction, which is obtained as the eigenvector associated with largest eigenvalue of  $\hat{\mathbf{D}}$ . The data fidelity error is measured as the root mean squared error (RMSE) of  $(\phi - \hat{\phi})$  of all the surface points. For the case of having multiple maps the same RMSE measurement is used but the error is of  $(\phi_m - \hat{\phi}_m), \forall m$ . We fix the following hyper-parameters for all these experiments:  $\{\alpha_a : 10^{-9}, \delta_e : 10^{-3}, \delta_a : 10^{-3}\}$ . The remaining hyper-parameters used in each experiment are specified in each section.

### 4.1 Synthetic 2D example

For the first set of numerical experiments, we use a flat domain defined as  $\Omega := [-1, 1] \times [-1, 1]$  discretized with a regularly spaced grid of  $35 \times 35$  points. We generate a triangular mesh with these points. We create synthetic maps by solving the eikonal equation with a fast iterative method [27]. We select 5 earliest activation sites using a random Latin hypercube design, to avoid selecting points that are close and effectively generate the same map. We set the conduction velocity tensors in the domain with a piece-wise constant function:

$$\mathbf{D}(x, y) = \begin{cases} \begin{bmatrix} 1 & 0 \\ 0 & \frac{1}{2} \end{bmatrix} & x + y < 0, \\ \begin{bmatrix} \frac{1}{2} & 0 \\ 0 & 1 \end{bmatrix} & \text{otherwise.} \end{cases} \quad (11)$$

From the solution, we select 245 points, using a Latin hypercube design, as data to train the physics-informed neural networks. We feed the neural

networks with either 1, 2, 3, or 5 of the generated maps. We split data points between maps, such that the total amount of points remains constant. The synthetic fiber orientations and the 5 maps used to train the model are shown in Figure 4. We set the hyperparameters to  $[\alpha_m : 10^{-2}, \alpha_a : 10^{-9}, \delta_e : 10^{-3}, \delta_a : 10^{-3}]$ , while for  $\alpha_e$ , the regularization of the conduction velocities, we run a sensitivity study between the values  $10^{-9}$  and  $10^{-3}$ . We use 5 hidden layers of 10 neurons for each of the networks that predict  $\hat{\phi}$  and 5 hidden layers of 5 neurons each to approximate  $\hat{D}$ . Each network is then trained for 3000 Adam mini-batch iterations with batch-size 32. Each of these tests is repeated 5 times using the same set of activation sites and sample points to quantify the variability of the training process.

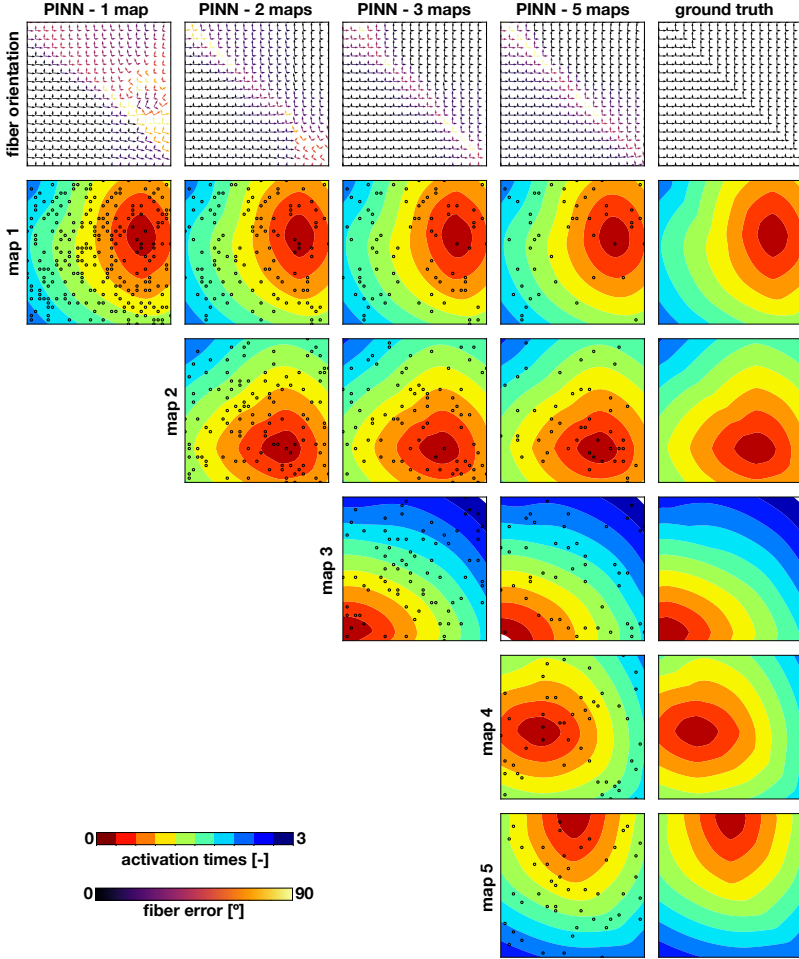
In Figure 4, we show the results for a representative case with median error for  $\alpha_e = 10^{-5}$ . First, we observe that the quality of the approximation of the activation maps decreases as the number of maps increases since there are fewer data points per map. However, the accuracy in the reconstruction of the maps does not translate into a more precise reconstruction of the fiber field. For instance, when only one map is used, the activation map is near perfectly reconstructed, but we can observe large fiber errors, some of them up to nearly  $90^\circ$ . With 2 maps, the errors are greatly reduced, however, there is a region of large fiber error near the lower right corner. For 3 and 5 maps, the errors are concentrated near the transition in fiber orientation, which is to be expected, as there are not enough data points to clearly define that boundary.

The fiber orientation errors for varying degrees of  $\alpha_e$  are presented in Figure 5. As shown in the right panel, we observe that for a larger amount of maps, the fiber error is decreased. Nonetheless, the gains from going from 1 to 3 maps are considerably bigger than going from 3 to 5 maps. We also see that the results tend to be more robust to the level of regularization when we feed the model with 3 or 5 maps. The results for 1 and 2 maps show great variability, which reflects the ill-posedness of this inverse problem when less than 3 maps are available.

## 4.2 Atrial geometry with rule-based fibers

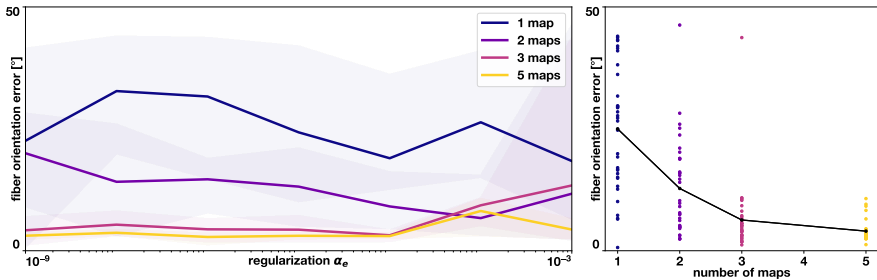
In this numerical experiment we consider a patient-specific geometry of the left atrium. Fibers have been semi-automatically assigned from histological in a previous work [46], shown in Figure 6. The reference longitudinal  $v_l$  and transverse velocity  $v_t$  are respectively set equal to 0.6 m/s and 0.4 m/s. The smooth basis necessary for the cosine of the fiber angle  $a(\mathbf{x})$  is obtained from the atrial basis provided in Appendix A (see also Figure A1). With the reference conductivity tensor, we generate 5 activation maps using an eikonal solver [27], by pacing at 5 different locations well apart from each other. The first pacing site is randomly placed, whereas the subsequent ones are obtained through the farthest point sampling approach.

We train physics-informed neural networks that are fed either with 1, 3, or 5 maps. We use 870 measurement points that are split between the different cases, such that the total number of points remains constant. The density



**Fig. 4** Activation times simulated by the runs with median error for hyper-parameters:  $\{\alpha_m : 10^{-2}, \alpha_e : 10^{-5}, \alpha_a : 10^{-9}, \delta_e : 10^{-3}, \delta_a : 10^{-3}\}$  for 1, 3 and 5 activation maps and the corresponding estimated conduction velocities (first row). The sampled points used for each run are represented as black circles and the colors are scaled to the 0.0-3.0 range (with white meaning out of range) for the activation time maps and to the range 0-90° for the fiber orientation error in the conduction velocity maps.

ranges 9.9 samples/cm<sup>2</sup> when we use 1 map to 1.9 samples/cm<sup>2</sup>. We set the hyperparameter to  $\alpha_m : 10^{-4}$ . We use neural networks of 7 hidden layers with 20 neurons to approximate  $\hat{\phi}$  and one neural network with 5 hidden layers with 20 neurons each to generate  $\hat{\mathbf{D}}$ . Each network is then trained for 30 000 Adam mini-batch iterations with batch-size 32. We also perform a sensitivity test for the regularization of conduction velocities and vary the parameter  $\alpha_e$  between  $10^{-9}$  and  $10^{-3}$ .



**Fig. 5** Left panel: fiber orientation error for  $\alpha_m = 10^{-2}$  and  $\alpha_a = 10^{-9}$ . Error values of all runs (shaded areas) and averages for 5 runs (solid lines) in 2D case of PINN with 1, 3 & 5 maps as inputs using different values of the  $\alpha_e$  conduction velocity total variation penalization weight. Right panel: fiber orientation errors for all regularization values for different number of maps. The solid line represents the mean.

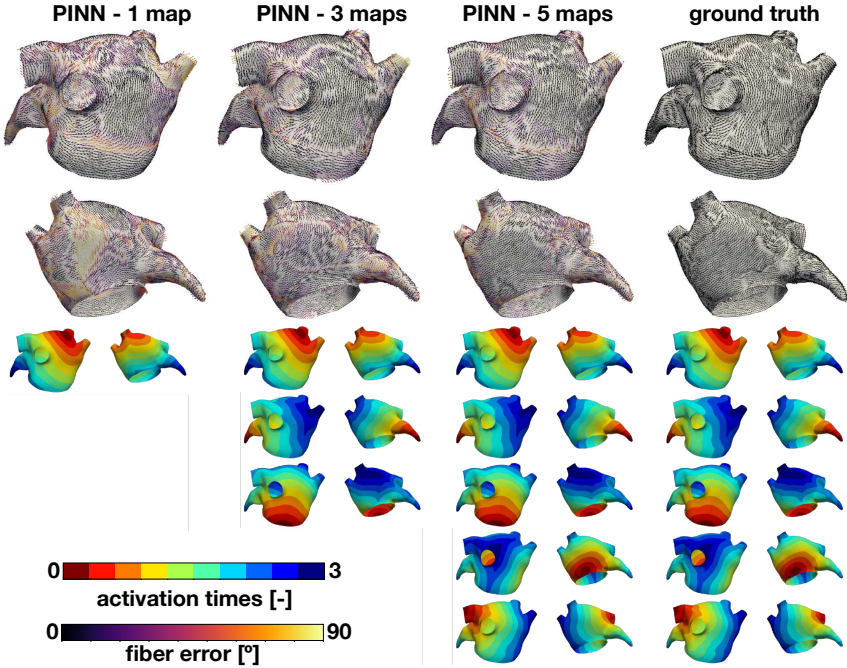
The results of the effect of conduction velocity regularization on the fiber error in this case are presented in Figure 7. First, we observe that the fiber error is considerably higher when one map is used, compared to the cases with 3 and 5 maps. We see that increasing the regularization tends to decrease the error for all cases. Nonetheless, the case with one map is more sensitive to the amount of regularization applied than the cases of 3 and 5 maps. As in the 2D case, we see that increasing from 1 to 3 maps has a much bigger effect on fiber error than going from 3 to 5. We also note that the variability in the results decreases for all cases compared to the 2D case.

We show an example of these results with the case of median error and  $\alpha_e : 10^{-5}$  in Figure 6. In the right column, we can see the ground truth fiber distribution and the five maps fed to the different models. In general, we observe that all the maps are well reconstructed. However, we see regions of large fiber errors when training with one map. We see that the cases trained with 3 and 5 maps tend to concentrate the errors around the zones where there are sharp transitions in fiber directions, such as near the mitral annulus. Other regions of high error for all cases are near the pulmonary veins, which have a circumferential fiber orientation, which is hard to capture, most likely due to a lack of data.

## 5 Atrial geometries with diffusion tensor fibers

For the next numerical experiment, we obtain fiber distributions for 7 different left atrium geometries from diffusion tensor magnetic resonance imaging [12, 16]. The images were taken *ex-vivo* with an isotropic resolution of 0.4 mm and mapped to a biatrial mesh (Figure 8, bottom row). Here, we attempt to learn the endocardial fibers, generating synthetic activation maps by solving the eikonal equation as described in the previous sections with these fibers. We randomly select 1020 points on each of the surfaces and split them between 1, 3, and 5 maps to obtain the activation time data for training. Since the different geometries have different surface areas, this numbers of measurements



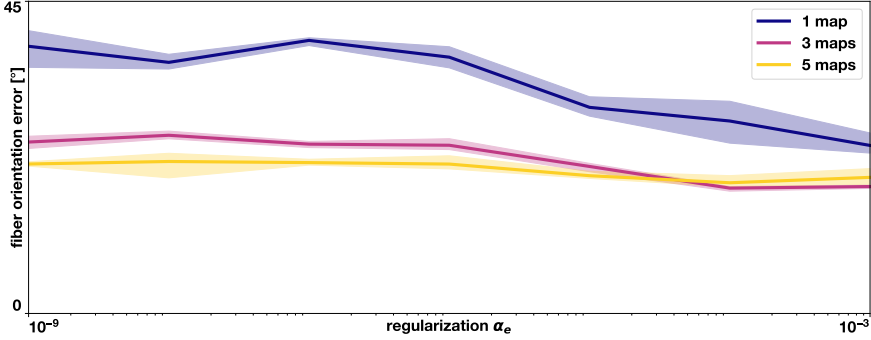


**Fig. 6** Predicted fiber orientations and activation times learned with 1, 3 and 5 maps. We show the results for the run with median error and hyper-parameters  $\{\alpha_m : 10^{-4}, \alpha_e : 10^{-5}\}$ .

corresponds to densities that range from 9.06 to 18.68 samples/cm<sup>2</sup> for one map, and 1.81 to 3.73 samples/cm<sup>2</sup> when we used 5 maps. In this experiment, we also test the robustness to noise of our method. We perturb the generated activation maps with Gaussian noise with zero mean and standard deviation of 0.1, 1 and 2 ms. We also include an additional validation step to test the overall performance of the method. We take the predicted and ground truth fibers and generate an additional map that is initiated from an activation site that is different from the maps used for training. Then, we compare these two maps and compute the error. In this way, we quantify how predictive are the learned fibers to model different scenarios, not included in the original dataset. For all cases, we set the hyper-parameters to  $\{\alpha_m : 10^{-4}, \alpha_e : 10^{-5}\}$ , which are the same used in the previous section.

The results of these experiments are presented in Figure 8 and 9, and Table B1. In Figure 8, we show the fiber predictions and errors for the 7 different cases. We first note that the fibers obtained from DT-MR imaging are not smooth, with recurrent abrupt changes in direction. We also note that the approximation of the fibers improves as the number of maps is increased, which can be noted in the Figure 8, as the high error regions (in yellow) are less frequent for the cases with more maps. This qualitative result is confirmed in Table B1, where the median fiber errors when training with one map range

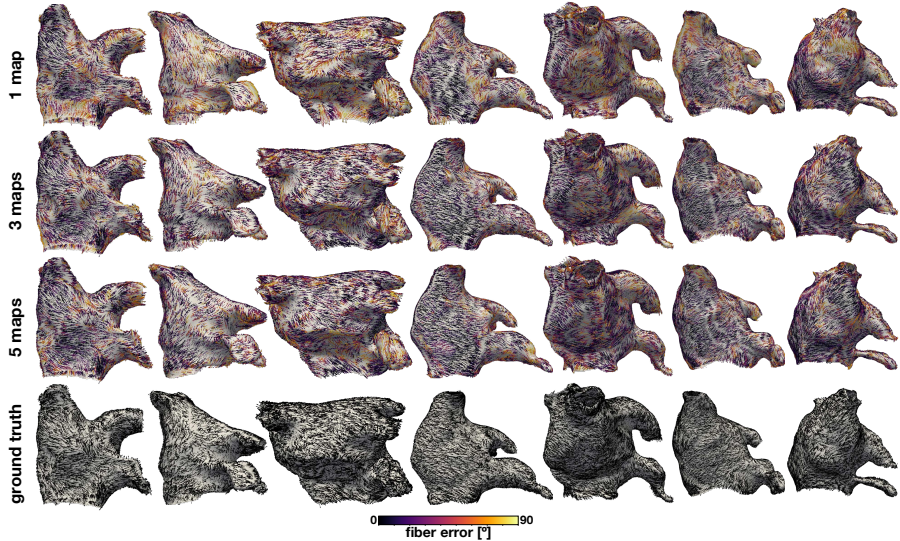




**Fig. 7** Fiber orientation error for  $\alpha_m = 10^{-4}$  and  $\alpha_a = 10^{-9}$ . Error values of all runs (shaded area) and averages for 5 runs (solid lines) in 3D atrial geometry case with 1, 3 & 5 maps as inputs using different values of the  $\alpha_e$  conduction velocity total variation penalization weight.

between  $24.8^\circ$  and  $30.2^\circ$ , and they are decreased to the range of  $18.3^\circ$  and  $23.2^\circ$  when training with 3 maps, and the range of  $16.2^\circ$  and  $23.3^\circ$  when training with 5 maps. When we add noise to the activation measurements, we observe in the case of 0.1 ms that there is no clear trend, as some of the cases tend to improve and some cases tend to worsen their accuracy. Nonetheless, the variations in fiber errors are small, less than  $5^\circ$  for all cases. When we add 1 ms of noise, the performance of the method decreases for all cases and the number of maps. The fiber error increases, on average,  $13.2^\circ$  when training with one map,  $12.4^\circ$  when training with 3 maps, and  $7.5^\circ$  when training with 5 maps. When we add 2 ms of noise, the error increases further, on average, to  $61.0^\circ$  with 1 map,  $19.1^\circ$  with 3 maps and  $13.1^\circ$  with 5 maps. We note here that the hyper-parameters used for the noise study were calibrated for a noiseless case with synthetic geometry and fibers. Thus, it is expected that tuning these parameters, especially  $\alpha_m$ , which controls the relevance of the eikonal equation, might improve the results. Overall, we see from the trend that for a fixed amount of data points it is always better to distribute them in different maps, especially in the presence of noise. The decrease in error is much more pronounced when increasing from 1 to 3 maps than when increasing from 3 to maps.

For the validation with an additional map, we observe in Figure 9 that in general, using the learned fibers can lead to an accurate prediction of an unseen activation pattern. In the left panel, we see an example using 3 maps for the noiseless case. The map created with the predicted fibers is a smoother version of the one created with ground truth fibers, as the original data presented significant spatial variations. For the cases with 0 and 0.1 ms of noise, we see that number of maps used for training does not influence the predictive capabilities of the methods. For the noiseless case, the root mean squared errors in activation time, on average, range from 2.7 ms when training with 5 maps to 4.4 ms when training with one map. However, when we inject 1 ms



**Fig. 8 Fiber predictions on diffusion tensor fibers.** We show the results for 7 different cases when we trained FiberNet with 1, 3 and 5 maps and no noise in the activation time data. Errors in the cases with noise can be found in Table B1. All the views are anteroposterior.

of noise, these differences are more pronounced, with an average root mean squared errors ranging from 5.4 ms when training with 5 maps to 17.7 ms when training with one map.

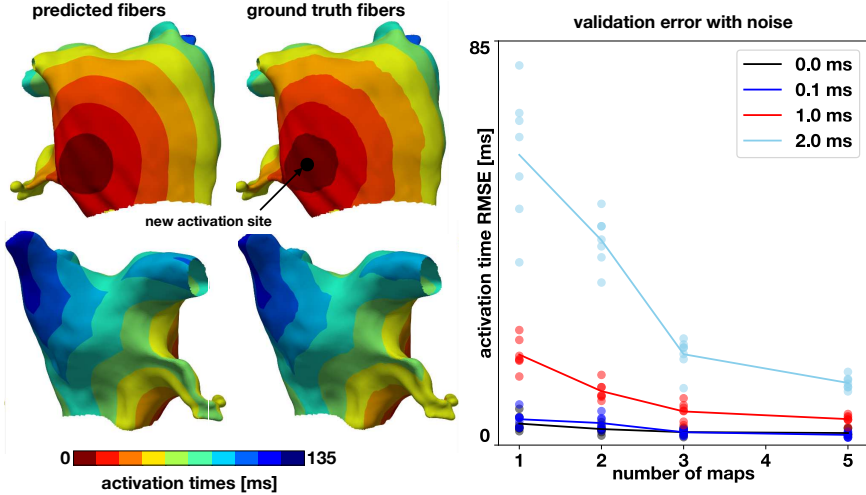
## 6 Patient-specific geometry with multiple maps

In this final Section, we report the application of the proposed methodology to a patient-specific case.

### 6.1 Data acquisition and pre-processing

Data from a single patient who underwent a pulmonary veins isolation (PVI) procedure has been collected at the Institute of Cardiocentro Ticino with oral and written informed consent for the investigation. The study has been performed in compliance with the Declaration of Helsinki. Data consisted of an electrophysiological study of the left atrium (LA) right before and after PVI performed with an electroanatomical mapping system (RHYTHMIA HDx, Boston Scientific, USA) equipped with an ultra-density mapping catheter (INTELLIMAP ORION, Boston Scientific, USA). A total of 3 electroanatomical maps have been acquired before ablation. Each map contained several electrical recordings of the extracellular voltage of the endocardial atrial wall, as summarized in Table 1.

Each recording consisted of the 3-D location of the recording electrode and 700 ms of unfiltered electric signal (sampling resolution is 953 Hz). Up to 64 recordings could be collected simultaneously, with only a portion of them in



**Fig. 9** Validation of conduction velocity tensor predictions with an additional map. Left panel, an example of one of the cases tested, where we create an activation from a point unseen during training. The top row shows posteroanterior views and the bottom row shows anteroposterior views. Right panel, we show a summary of the results of this validation for different number of maps and different levels of noise.

Map Number	Rhythm	Pacing Location	Acquired EGMs	Accepted EGMs
1	Paced	CS Distal	12 405	3 663 (29.5%)
2	Paced	CS Proximal	2 713	442 (16.3%)
3	SR	—	1 563	159 (10.2%)
Total			16 681	4264 (25.6%)

**Table 1** Summary of the electroanatomical mapping data of the patient-specific case. SR: Sinus Rhythm, CS: Coronary Sinus, EGMs: Electrograms.

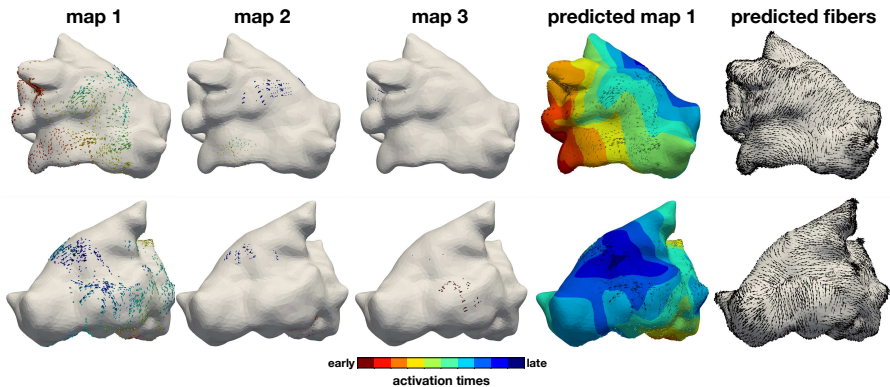
contact with the wall. Unipolar and bipolar electrograms have been automatically aligned in time to the R peak of the surface 12-lead ECG, simultaneously recorded. Pre-processing of the maps consisted in excluding electrograms with (i) low unipolar amplitude ( $< 0.05$  mV); (ii) poor contact as indicated by the system; (iii) inconsistent surface P wave; (iv) discrepancy of  $> 20$  ms between unipolar and bipolar activation time. The activation time was computed from the unipolar signal as the steepest negative deflection in the signal after the application of zero-phase forward and reverse 4-th order Butterworth with a cutoff frequency of 120 Hz. As reported in Table 1, the pre-processing excluded roughly 75% of the points, most of them because of low amplitude. The analysis has been performed with MATLAB version R2021a.

The left atrial anatomy has been obtained from the mapping system. The original triangular mesh has been re-meshed with fTetWild [47] using the default parameters, with a final median edge length of 2.1 mm (5 112 points). Finally, the 3-D location of the electrodes was projected onto the atrial surface.

It is worth mentioning that the mesh is only required for sampling collocation points and visualization, as by itself the proposed method is *mesh free*.

## 6.2 Results

The pre-processed maps are shown in Figure 10, the first 3 columns. We apply our method to these 3 maps and predict the fibers. We use the same hyper-parameters as in the previous sections. Although we do not have access to the real fiber orientations and we cannot compute the fiber error, we can check how well we have approximated the activation maps provided. Overall, we obtain the root mean squared error of 2.09 ms on the 3 activation maps. We show the fit for map 1, which is the most complete, in Figure 10, fourth column, demonstrating an excellent agreement with the measurements, which are shown as discrete points. Finally, in the predicted fibers we observe some of the expected features of the left atrium fibers in Figure 10, last column. We see that fibers go from the anterior to the posterior region through the atrial roof, and we also observe some regions where the fibers are aligned to the mitral annulus.



**Fig. 10 Results of patient-specific case.** On the first 3 columns, we show the data points acquired colored by activation time, on the fourth column we show the learned map with the data points on top and on the last column we show the predicted fiber orientations.

## 7 Concluding remarks

In this work, we present FiberNet, a novel method to estimate the complete conduction velocity tensor from a set of electroanatomical maps. The decomposition of the tensor simultaneously provides a patient-specific estimate of the fiber directions and conduction velocities. We tested our methodology in 2- and 3-dimensional synthetic examples, *ex-vivo* fibers obtained from DT-MR imaging, and clinical data. Furthermore, we validated our approach by creating an additional map, unseen during training, and we are able to accurately

predict the activation pattern. This step shows that the learned conductivity tensors can be used reliably for personalized modeling and digital twinning. From a methodological perspective, we showed that is not possible to estimate the conduction velocity tensors from a single map. In theory, at least 3 different maps are necessary, although having the correct regularization might reduce the error. Each map should provide different propagation directions in the tissue, albeit not necessarily orthogonal. In our numerical experiments, we see that for the same number of measurements, distributing them across different maps reduces the fiber error. Unsurprisingly, the error of estimating each of the maps increases, but this does not affect the accuracy of the fiber estimates.

FiberNet is currently limited to work on surfaces, which is a reasonable approximation for the atria under several circumstances [11]. We plan to extend FiberNet to work with solid bodies and apply it to the ventricles. Transmurally, fibers rotate from an endocardial to an epicardial direction, and these directions may sensibly differ [7]. In these cases, besides histological knowledge, we would require additional data to learn the activation times transmurally, which might come from body surface potentials and prior histological knowledge [7, 8, 13]. Along this line, we are currently using minimally invasive data, which might limit the applicability of our approach. We plan to extend it to work with the surface electrocardiogram, by converting the activation times into extracellular potentials through the forward bidomain model [48–50]. Our method takes on the order of one hour to train. Even though this is an order of magnitude larger than for some of the local methods [17, 19], this is a competitive time to solve an inverse problem [21]. Furthermore, to our knowledge, this is the first global method to estimate conduction velocity tensor from multiple maps. We plan to accelerate the training process by transfer learning [51] and also by incorporating prior knowledge of the fiber distribution in the basis of vectors that we use to locally parametrize the surface. In this way, the functions that are approximated by the neural network are simpler and can be learned in fewer training steps due to spectral bias [52]. FiberNet cannot be applied during AF in its current formulation, because it assumes that all data points for a single map are from the same beat or, at least from the same activation pattern. However, it could be applied to a patch of tissue covered by the catheter, and it should perform comparably to a local method in estimating the CV. It is worth remarking that during AF, the CV varies beat to beat, due to the restitution properties of the tissue. Finally, FiberNet still needs further validation, especially in its capability to extrapolate activation and fibers, as we only tested it using either known fibers or known activation maps. We plan to study and optimize the performance under noisy conditions to allow the application of this method in the clinical setting. We plan to compare our approach to existing rule-based methods for generating fibers, as recently proposed [10, 13]. In the clinical case, it is unclear whether FiberNet estimated the endocardial fiber orientation, or rather a transmurally homogenized fiber orientation. The results likely depend on the degree of endo-epicardial dissociation [46]. Ideally, we will acquire data from an animal model

where we could obtain electroanatomical maps and *ex-vivo* fiber orientations either from histological measurements or diffusion tensor imaging [53].

FiberNet is a new step in the road of personalized medicine. We hope this method will aid the creation of patient-specific models for accurate diagnosis and procedure planning.

**Acknowledgments.** This work was funded by an Open Seed Fund CORFO 14ENi2-26862 to CRH and FSC, the ANID – Millennium Science Initiative Program – NCN19-161 to FSC. We also acknowledge the School of Engineering computing cluster at Pontificia Universidad Católica de Chile for providing the computational resources for this study. SP and FSC acknowledge the financial support of the Leading House for Latin American Region (Grant Agreement No. RPG 2117). This work was also financially supported by the Theo Rossi di Montelera Foundation, the Metis Foundation Sergio Mantegazza, the Fidinam Foundation, the Horten Foundation to the Center for Computational Medicine in Cardiology. SP also acknowledges the CSCS-Swiss National Supercomputing Centre (Production Grant No. s1074) and the Swiss Heart Foundation (No. FF20042). Finally, we are very grateful to Prof. Angelo Auricchio and Dr. Giulio Conte from Instituto Cardiocentro Ticino, EOC (Lugano, Switzerland) for providing the patient-specific data.

**Compliance with Ethical Standards.** The authors have no relevant financial or non-financial interests to disclose. The clinical data used in this study was obtained with oral and written informed consent for the investigation. The study has been performed in compliance with the Declaration of Helsinki.

## References

- [1] Corral-Acero, J., Margara, F., Marciniak, M., Rodero, C., Loncaric, F., Feng, Y., Gilbert, A., Fernandes, J.F., Bukhari, H.A., Wajdan, A., Martinez, M.V., Santos, M.S., Shamohammdi, M., Luo, H., Westphal, P., Leeson, P., DiAchille, P., Gurev, V., Mayr, M., Geris, L., Pathmanathan, P., Morrison, T., Cornelussen, R., Prinzen, F., Delhaas, T., Doltra, A., Sitges, M., Vigmond, E.J., Zacur, E., Grau, V., Rodriguez, B., Remme, E.W., Niederer, S., Mortier, P., McLeod, K., Potse, M., Pueyo, E., Bueno-Orovio, A., Lamata, P.: The ‘Digital Twin’ to enable the vision of precision cardiology. *European Heart Journal* **41**(48), 4556–4564 (2020). <https://doi.org/10.1093/eurheartj/ehaa159>. Accessed 2021-04-18
- [2] Peirlinck, M., Costabal, F.S., Yao, J., Guccione, J., Tripathy, S., Wang, Y., Ozturk, D., Segars, P., Morrison, T., Levine, S., et al.: Precision medicine in human heart modeling. *Biomechanics and modeling in mechanobiology*, 1–29 (2021)
- [3] O’Sullivan, F.: A statistical perspective on ill-posed inverse problems. *Statistical science*, 502–518 (1986)



- [4] Ho, S., Sanchez-Quintana, D.: The importance of atrial structure and fibers. *Clinical Anatomy: The Official Journal of the American Association of Clinical Anatomists and the British Association of Clinical Anatomists* **22**(1), 52–63 (2009)
- [5] Clerc, L.: Directional differences of impulse spread in trabecular muscle from mammalian heart. *The Journal of physiology* **255**(2), 335–346 (1976)
- [6] 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)
- [7] Streeter Jr, D.D., Spotnitz, H.M., Patel, D.P., Ross Jr, J., Sonnenblick, E.H.: Fiber orientation in the canine left ventricle during diastole and systole. *Circulation research* **24**(3), 339–347 (1969)
- [8] Bayer, J.D., Blake, R.C., Plank, G., Trayanova, N.a.: A novel rule-based algorithm for assigning myocardial fiber orientation to computational heart models. *Annals of biomedical engineering* **40**(10), 2243–54 (2012). <https://doi.org/10.1007/s10439-012-0593-5>
- [9] Gonzales, M.J., Sturgeon, G., Krishnamurthy, A., Hake, J., Jonas, R., Stark, P., Rappel, W.-J., Narayan, S.M., Zhang, Y., Segars, W.P., *et al.*: A three-dimensional finite element model of human atrial anatomy: new methods for cubic hermite meshes with extraordinary vertices. *Medical image analysis* **17**(5), 525–537 (2013)
- [10] Wachter, A., Loewe, A., Krueger, M.W., Dössel, O., Seemann, G.: Mesh structure-independent modeling of patient-specific atrial fiber orientation. *Current Directions in Biomedical Engineering* **1**(1), 409–412 (2015). <https://doi.org/10.1515/cdbme-2015-0099>
- [11] Roney, C.H., Pashaei, A., Meo, M., Dubois, R., Boyle, P.M., Trayanova, N.A., Cochet, H., Niederer, S.A., Vigmond, E.J.: Universal atrial coordinates applied to visualisation, registration and construction of patient specific meshes. *Medical Image Analysis* **55**, 65–75 (2019)
- [12] Roney, C.H., Bendikas, R., Pashakhanloo, F., Corrado, C., Vigmond, E.J., McVeigh, E.R., Trayanova, N.A., Niederer, S.A.: Constructing a human atrial fibre atlas. *Annals of biomedical engineering* **49**, 233–250 (2021)
- [13] Piersanti, R., Africa, P.C., Fedele, M., Vergara, C., Dedè, L., Corno, A.F., Quarteroni, A.: Modeling cardiac muscle fibers in ventricular and atrial electrophysiology simulations. *Computer Methods in Applied Mechanics and Engineering* **373**, 113468 (2021)

- [14] Lombaert, H., Peyrat, J.-M., Croisille, P., Rapacchi, S., Fanton, L., Cheriet, F., Clarysse, P., Magnin, I., Delingette, H., Ayache, N.: Human atlas of the cardiac fiber architecture: study on a healthy population. *IEEE transactions on medical imaging* **31**(7), 1436–1447 (2012)
- [15] Teh, I., McClymont, D., Burton, R.A., Maguire, M.L., Whittington, H.J., Lygate, C.A., Kohl, P., Schneider, J.E.: Resolving fine cardiac structures in rats with high-resolution diffusion tensor imaging. *Scientific reports* **6**(1), 1–14 (2016)
- [16] Pashakhanloo, F., Herzka, D.A., Ashikaga, H., Mori, S., Gai, N., Bluemke, D.A., Trayanova, N.A., McVeigh, E.R.: Myofiber architecture of the human atria as revealed by submillimeter diffusion tensor imaging. *Circulation: arrhythmia and electrophysiology* **9**(4), 004133 (2016)
- [17] Cantwell, C.D., Roney, C.H., Ng, F.S., Siggers, J.H., Sherwin, S.J., Peters, N.S.: Techniques for automated local activation time annotation and conduction velocity estimation in cardiac mapping. *Computers in Biology and Medicine* **65**, 229–242 (2015). <https://doi.org/10.1016/j.combiomed.2015.04.027>
- [18] Coveney, S., Corrado, C., Roney, C.H., O’Hare, D., Williams, S.E., O’Neill, M.D., Niederer, S.A., Clayton, R.H., Oakley, J.E., Wilkinson, R.D.: Gaussian process manifold interpolation for probabilistic atrial activation maps and uncertain conduction velocity. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **378**(2173), 20190345 (2020)
- [19] Roney, C.H., Whitaker, J., Sim, I., O’Neill, L., Mukherjee, R.K., Razeghi, O., Vigmond, E.J., Wright, M., O’Neill, M.D., Williams, S.E., *et al.*: A technique for measuring anisotropy in atrial conduction to estimate conduction velocity and atrial fibre direction. *Computers in biology and medicine* **104**, 278–290 (2019)
- [20] Colli Franzone, P., Guerri, L., Rovida, S.: Wavefront propagation in an activation model of the anisotropic cardiac tissue: asymptotic analysis and numerical simulations. *Journal of mathematical biology* **28**(2), 121–176 (1990)
- [21] Grandits, T., Pezzuto, S., Lubrecht, J.M., Pock, T., Plank, G., Krause, R.: PIEMAP: Personalized inverse eikonal model from cardiac electro-anatomical maps. In: Puyol Anton, E., Pop, M., Sermesant, M., Campello, V., Lalande, A., Lekadir, K., Suinesiaputra, A., Camara, O., Young, A. (eds.) STACOM. M&Ms and EMIDEC Challenges. *Lecture Notes in Computer Science*, vol. 12592, pp. 76–86. Springer, Cham (2021)
- [22] Lubrecht, J.M., Grandits, T., Gharaviri, A., Schotten, U., Pock, T., Plank,



- G., Krause, R., Auricchio, A., Pezzuto, S.: Automatic reconstruction of the left atrium activation from sparse intracardiac contact recordings by inverse estimate of fiber structure and anisotropic conduction in a patient-specific model. *EP Europace* **23**, 63–70 (2021)
- [23] Raissi, M., Perdikaris, P., Karniadakis, G.E.: Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comp. Phys.* **378**, 686–707 (2019)
- [24] Grandits, T., Pezzuto, S., Sahli Costabal, F., Perdikaris, P., Pock, T., Plank, G., Krause, R.: Learning atrial fiber orientations and conductivity tensors from intracardiac maps using physics-informed neural network. In: Ennis, D.B., Perotti, L.E., Wang, V.Y. (eds.) *Functional Imaging and Modeling of the Heart. FIMH 2021. Lecture Notes in Computer Science*, vol. 12738, pp. 650–658. Springer, Cham (2021). [https://doi.org/978-3-030-78710-3\\_62](https://doi.org/978-3-030-78710-3_62)
- [25] Roney, C., Bendikas, R., Pashakhanloo, F., Corrado, C., Vigmond, E., McVeigh, E., Trayanova, N., Niederer, S.: Constructing a Human Atrial Fibre Atlas. Zenodo (2020). <https://doi.org/10.5281/zenodo.3764917>
- [26] Fu, Z., Jeong, W., Pan, Y., Kirby, R., Whitaker, R.: A Fast Iterative Method for Solving the Eikonal Equation on Triangulated Surfaces. *SIAM Journal on Scientific Computing* **33**(5), 2468–2488 (2011). <https://doi.org/10.1137/100788951>. Accessed 2018-07-26
- [27] Grandits, T.: A Fast Iterative Method Python package. *Journal of Open Source Software* **6**(66), 3641 (2021). <https://doi.org/10.21105/joss.03641>. Accessed 2021-10-23
- [28] Verma, B., Oesterlein, T., Loewe, A., Luik, A., Schmitt, C., Dössel, O.: Regional conduction velocity calculation from clinical multichannel electrograms in human atria. *Computers in Biology and Medicine* **92**, 188–196 (2018). <https://doi.org/10.1016/j.combiomed.2017.11.017>
- [29] van Schie, M.S., Heida, A., Taverne, Y.J.H.J., Bogers, A.J.J.C., de Groot, N.M.S.: Identification of local atrial conduction heterogeneities using high-density conduction velocity estimation. *EP Europace* **23**(11), 1815–1825 (2021)
- [30] Pagani, S., Dede', L., Frontera, A., Salvador, M., Limite, L.R., Manzoni, A., Lipartiti, F., Tsitsinakakis, G., Hadjis, A., Della Bella, P., Quarteroni, A.: A computational study of the electrophysiological substrate in patients suffering from atrial fibrillation. *Frontiers in Physiology* **12**, 927 (2021)
- [31] Good, W.W., Gillette, K.K., Zenger, B., Bergquist, J.A., Rupp, L.C.,

- Tate, J., Anderson, D., Gsell, M.A.F., Plank, G., MacLeod, R.S.: Estimation and validation of cardiac conduction velocity and wavefront reconstruction using epicardial and volumetric data. *IEEE Transactions on Biomedical Engineering* **68**(11), 3290–3300 (2021). <https://doi.org/10.1109/TBME.2021.3069792>
- [32] Nothstein, M., Luik, A., Jadidi, A., Sánchez, J., Unger, L.A., Wülfers, E.M., Dössel, O., Seemann, G., Schmitt, C., Loewe, A.: Cvar-seg: An automated signal segmentation pipeline for conduction velocity and amplitude restitution. *Frontiers in Physiology* **12**, 746 (2021). <https://doi.org/10.3389/fphys.2021.673047>
- [33] Sahli Costabal, F., Yang, Y., Perdikaris, P., Hurtado, D.E., Kuhl, E.: Physics-informed neural networks for cardiac activation mapping. *Frontiers in Physics* **8**, 42 (2020)
- [34] Gander, W., Golub, G.H., Strebler, R.: Least-squares fitting of circles and ellipses. *BIT Numerical Mathematics* **34**(4), 558–578 (1994)
- [35] Massé, S., Magtibay, K., Jackson, N., Asta, J., Kusha, M., Zhang, B., Balachandran, R., Radisic, M., Deno, D.C., Nanthakumar, K.: Resolving myocardial activation with novel omnipolar electrograms. *Circulation: Arrhythmia and Electrophysiology* **9**(7), 004107 (2016)
- [36] Gaeta, S., Bahnson, T.D., Henriquez, C.: High-resolution measurement of local activation time differences from bipolar electrogram amplitude. *Frontiers in Physiology* **12**, 536 (2021)
- [37] Yang, H., Veneziani, A.: Estimation of cardiac conductivities in ventricular tissue by a variational approach. *Inverse Problems* **31**(11), 115001 (2015)
- [38] Barone, A., Gizzi, A., Fenton, F., Filippi, S., Veneziani, A.: Experimental validation of a variational data assimilation procedure for estimating space-dependent cardiac conductivities. *Computer Methods in Applied Mechanics and Engineering* **358**, 112615 (2020)
- [39] Irakoze, É., Jacquemet, V.: Multiparameter optimization of nonuniform passive diffusion properties for creating coarse-grained equivalent models of cardiac propagation. *Computers in Biology and Medicine* **138**, 104863 (2021). <https://doi.org/10.1016/j.combiomed.2021.104863>
- [40] Chegini, F., Kopaničáková, A., Krause, R., Weiser, M.: Efficient identification of scars using heterogeneous model hierarchies. *EP Europace* **23**(Supplement\_1), 113–122 (2021)

- [41] Arevalo, H.J., Vadakkumpadan, F., Guallar, E., Jebb, A., Malamas, P., Wu, K.C., Trayanova, N.A.: Arrhythmia risk stratification of patients after myocardial infarction using personalized heart models. *Nature Communications* **7**(May), 11437 (2016). <https://doi.org/10.1038/ncomms11437>
- [42] Sharp, N., Soliman, Y., Crane, K.: The vector heat method. *ACM Transactions on Graphics (TOG)* **38**(3), 1–19 (2019)
- [43] Chambolle, A., Pock, T.: An introduction to continuous optimization for imaging. *Acta Numerica* **25**, 161–319 (2016). <https://doi.org/10.1017/S096249291600009X>. Accessed 2018-10-05
- [44] Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S., Murray, D.G., Steiner, B., Tucker, P., Vasudevan, V., Warden, P., Wicke, M., Yu, Y., Zheng, X.: Tensorflow: A system for large-scale machine learning. In: 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16), pp. 265–283 (2016)
- [45] Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization. arXiv:1412.6980 [cs] (2017). arXiv: 1412.6980. Accessed 2020-07-29
- [46] 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)
- [47] Hu, Y., Schneider, T., Wang, B., Zorin, D., Panozzo, D.: Fast tetrahedral meshing in the wild. *ACM Trans. Graph.* **39**(4) (2020). <https://doi.org/10.1145/3386569.3392385>
- [48] Pezzuto, S., Kal'avskÿ, P., Potse, M., Prinzen, F.W., Auricchio, A., Krause, R.: Evaluation of a rapid anisotropic model for ECG simulation. *Frontiers in physiology* **8**, 265 (2017)
- [49] Potse, M.: Scalable and accurate ecg simulation for reaction-diffusion models of the human heart. *Frontiers in physiology* **9**, 370 (2018)
- [50] Gillette, K., Gsell, M.A., Prassl, A.J., Karabelas, E., Reiter, U., Reiter, G., Grandits, T., Payer, C., Štern, D., Urschler, M., *et al.*: A framework for the generation of digital twins of cardiac electrophysiology from clinical 12-leads ecgs. *Medical Image Analysis* **71**, 102080 (2021)
- [51] Lejeune, E., Zhao, B.: Exploring the potential of transfer learning for metamodels of heterogeneous material deformation. *Journal of the*

Mechanical Behavior of Biomedical Materials **117**, 104276 (2021)

- [52] Wang, S., Wang, H., Perdikaris, P.: On the eigenvector bias of Fourier feature networks: From regression to solving multi-scale PDEs with physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering* **384**, 113938 (2021)
- [53] Roney, C.H., Bendikas, R., Pashakhanloo, F., Corrado, C., Vigmond, E.J., McVeigh, E.R., Trayanova, N.A., Niederer, S.A.: Constructing a human atrial fibre atlas. *Annals of Biomedical Engineering* **49**(1), 233–250 (2021)

## Appendix A Smooth basis

When defining a conduction velocity tensor through its eigencomponents, as was done in Section 3.1, any two non-coinciding (preferably orthonormal) vectors of the tangent space  $\mathcal{S}$  may be a feasible choice. Computationally, this 2D basis could be chosen locally, e.g. on a per-triangle basis. However, as the chosen Huber-norm regularization in (8) penalizes variation of the parameter vector  $\mathbf{d}$ , it is greatly beneficial to define a smooth basis.

To this end, we computed the first basis  $\mathbf{v}_1$  of the tangent space using the vector heat method [42]. The vector heat method quickly computes parallel transport of a chosen vector on the whole manifold by shortly diffusing a given initial vector field using the heat equation with the connection Laplacian  $\Delta^\nabla = -\nabla_C^* \nabla_C$  for  $\nabla_C$  being the Levi-Civita connection of the manifold. The rescaled diffused vectors then closely approximate the parallel transport of the initial vector field (for numerical and qualitative comparisons, we refer to the original paper [42]).

The initial vector to be transported across the manifold is manually selected for each manifold, such that the discontinuities of the vector field are minimal. Discontinuities might arise at the cut-locus, i.e. regions equidistant from the given vector, as there are multiple possible parallel transports at such locations. Figure A1 shows such a computed parallel transport from the anterior wall of the atria. The resulting vector field is smooth in most parts, exhibiting a very low overall variation. Such a basis was chosen manually for all meshes.

## Appendix B DT-MR imaging errors



**Fig. A1** Result of transporting a single vector on the anterior wall of the atria across the entire manifold. The resulting vector field is smooth except for parts of high curvature such as the left atrial appendage or pulmonary veins.

noise [ms]	# of maps	case 1	case 2	case 3	case 4	case 5	case 6	case 7
0.0	1	31.4 (14.2 - 54.5)	30.4 (13.4 - 55.4)	26.2 (10.7 - 52.4)	28.4 (11.9 - 55.3)	26.9 (11.3 - 51.8)	23.0 (9.7 - 46.6)	21.5 (9.1 - 45.7)
0.0	3	22.5 (9.7 - 44.8)	22.1 (9.6 - 44.6)	19.7 (8.2 - 40.8)	21.5 (9.2 - 42.2)	21.6 (9.3 - 42.8)	18.6 (7.7 - 38.5)	17.4 (7.7 - 36.0)
0.0	5	21.6 (9.4 - 44.5)	20.1 (8.6 - 41.9)	17.5 (7.3 - 39.0)	18.5 (8.0 - 38.4)	19.8 (8.3 - 40.5)	17.9 (7.7 - 38.1)	17.0 (7.6 - 33.9)
0.1	1	34.5 (15.6 - 60.4)	26.5 (11.5 - 51.6)	24.6 (11.1 - 48.9)	30.7 (13.3 - 56.4)	27.2 (11.2 - 52.6)	28.0 (12.1 - 54.0)	25.7 (11.3 - 49.8)
0.1	3	22.9 (10.1 - 45.5)	24.1 (11.0 - 47.1)	19.4 (8.2 - 40.2)	20.9 (9.0 - 42.3)	21.9 (9.6 - 43.8)	19.7 (8.5 - 40.4)	18.0 (7.9 - 36.4)
0.1	5	20.9 (9.1 - 43.7)	22.1 (9.5 - 43.6)	17.5 (7.3 - 39.0)	19.7 (8.8 - 39.9)	22.4 (10.1 - 45.0)	19.8 (8.3 - 41.3)	14.9 (6.5 - 31.2)
1.0	1	45.6 (22.4 - 67.9)	44.5 (22.0 - 67.0)	37.2 (18.4 - 59.9)	41.4 (19.4 - 65.6)	43.6 (20.9 - 66.5)	42.6 (20.8 - 65.1)	42.5 (20.8 - 66.2)
1.0	3	38.6 (18.0 - 62.8)	29.5 (12.9 - 54.3)	30.7 (14.2 - 54.9)	31.5 (14.4 - 55.2)	30.5 (13.7 - 55.5)	31.8 (14.3 - 55.2)	30.6 (13.7 - 53.4)
1.0	5	33.0 (14.8 - 58.6)	26.6 (11.9 - 50.5)	24.6 (11.2 - 48.0)	26.7 (11.2 - 50.4)	25.2 (10.7 - 49.1)	25.5 (11.5 - 47.8)	26.0 (11.9 - 48.1)
2.0	1	46.7 (25.3 - 68.0)	43.6 (20.6 - 67.7)	36.4 (18.1 - 59.8)	43.2 (21.5 - 66.1)	43.1 (19.3 - 65.9)	42.8 (20.7 - 66.2)	47.8 (24.6 - 69.7)
2.0	3	44.1 (22.0 - 67.2)	41.8 (20.3 - 66.3)	35.5 (18.0 - 59.7)	40.5 (19.8 - 63.5)	37.8 (18.3 - 61.0)	37.4 (16.9 - 61.6)	35.7 (16.8 - 58.9)
2.0	5	42.9 (21.9 - 64.8)	41.1 (20.4 - 65.2)	31.3 (13.6 - 55.8)	35.8 (16.6 - 59.7)	34.7 (15.8 - 59.5)	33.7 (15.4 - 57.6)	34.8 (16.1 - 59.7)

**Table B1** Fiber errors for the 7 fiber distributions obtained from DT-MR imaging. Results are presented for different levels of noise standard deviation and the number of maps used for training. Results are presented as median (25% - 75% percentile).