Digital twinning of cardiac electrophysiology models from the surface ECG: A geodesic backpropagation approach

Thomas Grandits, Jan Verhülsdonk, Gundolf Haase, Alexander Effland, and Simone Pezzuto

Abstract— The eikonal equation has become an indispensable tool for modeling cardiac electrical activation accurately and efficiently. In principle, by matching clinically recorded and eikonal-based electrocardiograms (ECGs), it is possible to build patient-specific models of cardiac electrophysiology in a purely non-invasive manner. Nonetheless, the fitting procedure remains a challenging task.

The present study introduces a novel method, Geodesic-BP, to solve the inverse eikonal problem. Geodesic-BP is well-suited for GPU-accelerated machine learning frameworks, allowing us to optimize the parameters of the eikonal equation to reproduce a given ECG.

We show that Geodesic-BP can reconstruct a simulated cardiac activation with high accuracy in a synthetic test case, even in the presence of modeling inaccuracies. Furthermore, we apply our algorithm to a publicly available dataset of a biventricular rabbit model, with promising results.

Given the future shift towards personalized medicine, Geodesic-BP has the potential to help in future functionalizations of cardiac models meeting clinical time constraints while maintaining the physiological accuracy of state-of-the-art cardiac models.

Index Terms— Cardiac Digital Twin, Eikonal Model, ECG, Electrophysiology, Backpropagation, Machine Learning

I. INTRODUCTION

Cardiac digital twinning is a potential pillar of future's precision cardiology [1]. A digital twin is a computational replica of the patient's heart, from its anatomy and structure up to its functional response. Commonly, it is obtained by fitting the parameters of a generic computational model to the clinical data of the patient. Here, we focus on cardiac electrophysiology, where we aim at personalizing a cardiac

T. Grandits and G. Haase are with the Department of Mathematics and Scientific Computing, University of Graz, Austria (e-mail: {thomas.grandits,gundolf.haase}@uni-graz.at), T. Grandits is additionally affiliated with NAWI Graz.

J. Verhülsdonk and A. Effland are with the Institute for Applied Mathematics, University of Bonn, Germany (e-mail: {verhuelsdonk,effland}@iam.uni-bonn.de).

S. Pezzuto is with the Laboratory of Mathematics for Biology and Medicine, Department of Mathematics, University of Trento, Italy, and with the Center for Computational Medicine in Cardiology, Euler Institute, Università della Svizzera italiana, Switzerland (e-mail: simone.pezzuto@unitn.it).

conduction model based on the eikonal equation, from non-invasive electrocardiographic (ECG) recordings [2], [3].

Eikonal-based propagation models are an excellent basis for cardiac electrophysiology digital twins. They provide a good balance between computational efficiency and physiological accuracy for cardiac activation maps [4], [5], [6]. When supplemented with the lead field approach, they can also provide accurate ECGs [7], [3]. Hence, the parameters of eikonal-based models could be fitted from non-invasive clinical data—surface potentials and cardiac imaging—in an efficient manner. This goal can be achieved using optimization methods based on stochastic sampling strategies [8], [3], Bayesian optimization [9], gradient-descent [10], or deep learning approaches [11].

The numerical efficiency of the optimization procedure relies on two factors: the computational cost of the forward problem, and the number of samples required to achieve convergence within a given tolerance. Gradient-based methods can address the latter, if gradient computation is efficient and robust. For the former, many efficient eikonal solvers exist, but not all of them are suitable for automatic differentiation. A major criticality is that they internally solve a local optimization problem [12], [13], one for each element of the grid and for possibly many iterations. This may represent a bottleneck in the computation of the gradient. A possible solution is to compute the gradient for the continuous problem, which can be shown is equivalent to the computation of geodesic paths [10], [14], and then discretize. The *optimize-then-discretize* strategy is however not ideal, because it may introduce numerical noise in the gradient and degrade the convergence rate.

In this paper, we introduce a novel inverse eikonal solver—named *Geodesic-BP*—and use it to efficiently solve the ECG-based digital twinning problem for ventricular models. To this end, we start by introducing a highly-parallelizable anisotropic eikonal solver for triangulations in arbitrary dimension. The key contribution is a novel approach for the solution of the local optimization problem and based on a linear convex non-smooth optimization method. Locally, the linear geodesic path within each simplicial element is approximated with a guaranteed bound, similar to many presented eikonal solvers [12]. This solver can be implemented using current machine learning (ML) libraries to efficiently calculate the solution of the eikonal equation. The forward solver is suitable for the computation of the gradient via the backpropagation algo-
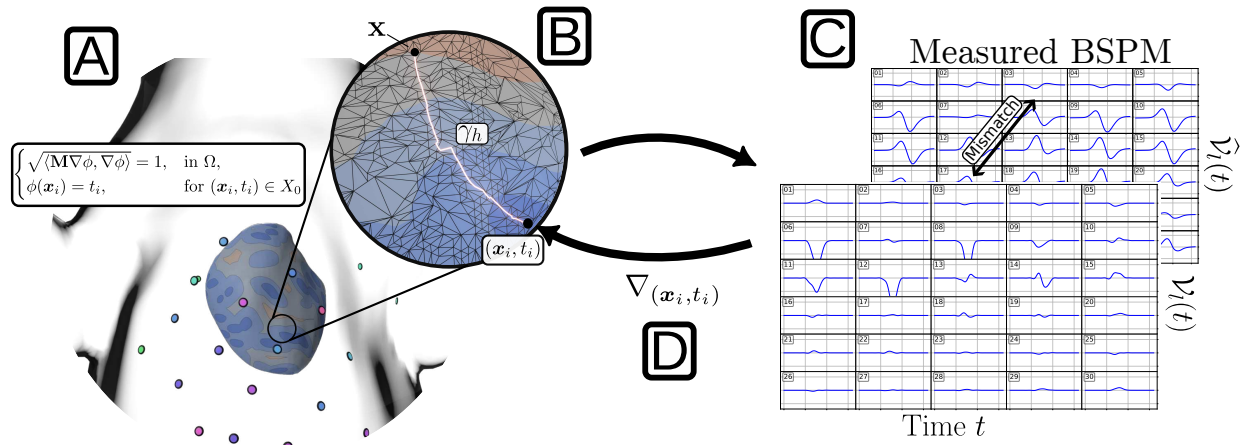
Fig. 1. The anisotropic eikonal equation is an efficient tool to compute the electrical activation of a heart (A), here shown on a public rabbit torso model. Many solvers utilize the relation of the eikonal equation to geodesic paths, which in the case of $\mathcal{P}_1$-meshes become piecewise-linear paths γ_h between the initial conditions $(\boldsymbol{x}_i, t_i)$ and points on the surfaces $\boldsymbol{x}$ (B). By applying a non-linear transformation explained in Section II-D, we can compute the measurable electrical activation on the torso surface and compute the mismatch against the measured potentials at sparse points (C). Implementing the approach in a differentiable manner using a backpropagation framework such as pytorch, we can efficiently calculate the gradient w.r.t. the initial conditions and use it to optimally choose $(\boldsymbol{x}_i, t_i)$ to match the measured potentials (D).

rithm, effectively enabling a *discretize-then-optimize* strategy (Geodesic-BP: Geodesic Back-Propagation). We will show that the backpropagation through such a solver is equivalent to piecewise-linear geodesic backtracking [15], [16], [10]. Finally, we will demonstrate the potential of these methods for the problem of estimating initial eikonal conditions to match a given body surface potential map (BPSM) in a synthetic 2D study, and real 3D torso model.

In summary, the major contributions are as follows:

1) A novel, highly-parallelizable anisotropic eikonal solver, well-suited for GPGPU architectures and ML libraries.
2) A backpropagation-based gradient computation, internally related to the computation of geodesics.
3) An efficient implementation of the method.
4) We demonstrate how to use Geodesic-BP for the identification of cardiac activation sequence from the non-invasive ECG recordings.
5) We show that the inverse procedure is robust to model perturbations.

An outline of Geodesic-BP is provided in Figure 1.

II. MATERIAL AND METHODS

In this section, the modeling assumptions for the electrical propagation inside the heart are introduced along with an efficient method for computing this activation in a parallelizable way. Then, we elaborate on the computation of the electrical activation given measurements on the torso leads, which ultimately leads to the formulation of the inverse ECG problem used throughout all experiments.

A. Cardiac activation

We model cardiac activation with the anisotropic eikonal equation (see [17], [7], [6]), which reads as follows

$$\begin{cases} \sqrt{\langle \mathbf{M}(\boldsymbol{x})\nabla\phi(\boldsymbol{x}), \nabla\phi(\boldsymbol{x})\rangle} = 1, & \text{in } \Omega, \\ \phi(\boldsymbol{x}_i) = t_i, & \text{for } (\boldsymbol{x}_i, t_i) \in X_0, \end{cases} \quad (1)$$

where $\Omega \subset \mathbb{R}^d$, $d \geq 2$, is a bounded domain, $\phi : \Omega \to \mathbb{R}$ is the activation map, $\langle \cdot, \cdot \rangle$ is the inner product in $\mathbb{R}^d$, and $\mathbf{M} : \Omega \to \mathcal{S}_d$ is a symmetric positive definite (s.p.d.) tensor field, where $\mathcal{S}_d \subset \mathbb{R}^{d \times d}$ denotes the space of d-times s.p.d. tensors, used to account for the orthotropic conduction velocity inside the heart [17]. The onset of activation is determined by the set of initial conditions: $X_0 = \{(\boldsymbol{x}_i, t_i)\}_{i=1}^K$, where $\boldsymbol{x}_i \in \Omega$ and t_i are respectively the location and the onset timing of the i-th activation site. For convenience, we additionally define the norm in the metric $\mathbf{D}$, i.e. $\|\boldsymbol{x}\|_\mathbf{D} = \sqrt{\langle \mathbf{D}\boldsymbol{x}, \boldsymbol{x}\rangle}$ for $\mathbf{D} \in \mathcal{S}_d$.

Next, we discretize the domain Ω (assumed polytopal) by a simplicial triangulation. We denote by $V = \{\boldsymbol{v}_i\}_{i=1}^{n_v}$ and $\mathcal{E} = \{E_j\}_{j=1}^{n_e}$ respectively the set of vertices and elements of the triangulation. Each element $E_j \in \mathcal{E}$ is a d-simplex (triangle in 2-D, tetrahedron in 3-D) with $d+1$ vertices. We also define $E_{j \setminus i} \subset \partial E_j$ as the $(d-1)$-dimensional face of E_j opposite to vertex $\boldsymbol{v}_i$. Note that $E_{j \setminus i}$ is a $(d-1)$-simplex (triangle in 3-D, segment in 2-D). For a general unstructured grid, we can approximate the activation time with a piecewise linear finite element function $\phi : \Omega \to \mathbb{R}$, see [18]. (For the sake of simplicity, we use the same symbol as for the continuous case.) For the nodal values we use the shorthand notation $\phi_i = \phi(\boldsymbol{v}_i)$. Following [12], [19], [20], we approximate $\mathbf{D}(\boldsymbol{x}) = \mathbf{M}^{-1}(\boldsymbol{x})$ as a piecewise-constant function and set $\mathbf{D}_j = \mathbf{D}(\boldsymbol{x})|_{E_j}$.

We numerically solve the anisotropic eikonal equation (1) by locally enforcing the following optimality condition for the nodal values:

$$\phi_i = \min_{E_j \in \omega_i} \phi_{j \setminus i}^*, \quad (2)$$

where $\omega_i = \{E \in \mathcal{E} : \boldsymbol{v}_i \text{ is a vertex of } E\}$ is the set of elements containing $\boldsymbol{v}_i$ and $\phi_{j \setminus i}^*$ is the minimum of the following optimization problem:

$$\phi_{j \setminus i}^* = \min_{\boldsymbol{x} \in E_{j \setminus i}} \phi(\boldsymbol{x}) + \|\boldsymbol{v}_i - \boldsymbol{x}\|_{\mathbf{D}_j}. \quad (3)$$

The optimality condition (2) is known as Hopf-Lax formula. It is possible to show that, under regularity assumptions on Ω, $\mathbf{D}$,
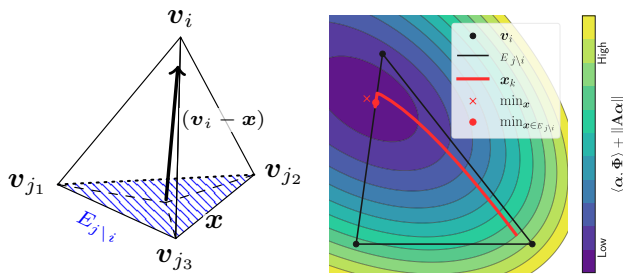
Fig. 2. Conceptual visualization of a linear upwind discretization (5) on the example of a single tetrahedron (left). Finding the minimum upwind origin $\boldsymbol{x} = \sum_n \alpha_n \boldsymbol{v}_{j_n}$ can be cast into a non-smooth linear optimization problem, solvable using the FISTA algorithm (right).

and the triangulation, the numerical solution of (2) correctly convergences to the unique viscosity solution of (1) when the triangulation is refined [21]. As explained later, this principle is at the basis of Geodesic-BP.

B. Solution of the local problem

In what follows, we show how to efficiently and in a parallelizable way solve (3). We observe that eqs. (2) and (3) consists of a series of constrained minimization problems, one for each element of the patch. Clearly, each local problem in (3) is convex, because the objective function is convex in $\boldsymbol{x}$ (note that ϕ is linear on $E_{j \setminus i}$) and the constraint is a simplex, thus convex as well. Instead of relying on exact formulas, e.g., available in 2-D and 3-D from [20], we propose here a new dimension-independent algorithm for the solution of (3).

With no loss of generality, we suppose that in (3) the vertices of E_j are $\{\boldsymbol{v}_{j_1}, \ldots, \boldsymbol{v}_{j_{d+1}}\}$ and ordered such that the last vertex $\boldsymbol{v}_{j_{d+1}}$ always corresponds to $\boldsymbol{v}_i$. Next, we consider the set of barycentric coordinates for the face $E_{j \setminus i}$:

$$C_d = \left\{ \boldsymbol{\alpha} \in \mathbb{R}^d : \langle \boldsymbol{\alpha}, \mathbf{1} \rangle = 1 \wedge \boldsymbol{\alpha} \geq 0 \right\}, \quad (4)$$

such that each point $\boldsymbol{x}$ in a face $E_{j \setminus i}$ can be uniquely written as $\boldsymbol{x} = \sum_{n=1}^{d} \alpha_n \boldsymbol{v}_{j_n}$. Similarly, the $\mathcal{P}_1$ function $\phi(\boldsymbol{x})$ on $E_{j \setminus i}$ is given by the linear combination of the vertex values: $\phi(\boldsymbol{x}) = \sum_{n=1}^{d} \alpha_n \phi_{j_n}$. Thus, Eq. (3) becomes

$$
\begin{aligned}
\phi_{j \setminus i}^* &= \min_{\boldsymbol{\alpha} \in C_d} \sum_{n=1}^{d} \alpha_n \phi_{j_n} + \left\| \sum_{n=1}^{d} \boldsymbol{v}_{j_n} \alpha_n - \boldsymbol{v}_i \right\|_{\mathbf{D}_j} \\
&= \min_{\boldsymbol{\alpha} \in C_d} \langle \boldsymbol{\alpha}, \boldsymbol{\Phi}_{j \setminus i} \rangle + \left\| \mathbf{A}_{j \setminus i} \boldsymbol{\alpha} \right\|_2,
\end{aligned} \quad (5)
$$

where $\boldsymbol{\Phi}_{j \setminus i} = (\phi_{j_1}, \ldots, \phi_{j_d})^\top$ is the vector of nodal values of $\phi|_{E_j}$, except for ϕ_i, and

$$\mathbf{A}_{j \setminus i} = \mathbf{D}_j^{1/2} \left(\begin{array}{ccc} | & & | \\ \boldsymbol{v}_{j_1} - \boldsymbol{v}_i & \ldots & \boldsymbol{v}_{j_d} - \boldsymbol{v}_i \\ | & & | \end{array} \right).$$

A graphical representation of the local optimization problem is given in Figure 2.

Eq. (5) now takes a common form of forward/backward-splitting algorithms, where the function to minimize is a linear combination of two convex functions with one potentially being non-smooth. Such problems can be minimized using the

FISTA algorithm [22], which alternatingly optimizes using a proximal step into the non-smooth function and a gradient step in the direction of the smooth function. The smooth part of (5) was thus chosen as $h(\boldsymbol{\alpha}) = \|\mathbf{A}\boldsymbol{\alpha}\|_2$, and the non-smooth part as $g(\boldsymbol{\alpha}) = \langle \boldsymbol{\alpha}, \boldsymbol{\Phi} \rangle + \delta_{C_d}(\boldsymbol{\alpha})$, where δ_C is the indicator function on the set C such that

$$\delta_C(\boldsymbol{x}) = \begin{cases} 0, & \text{if } \boldsymbol{x} \in C, \\ \infty, & \text{otherwise.} \end{cases}$$

The gradient of the smooth part $h(\boldsymbol{\alpha})$ is given by

$$\nabla h(\boldsymbol{\alpha}) = \frac{\mathbf{A}^\top \mathbf{A} \boldsymbol{\alpha}}{h(\boldsymbol{\alpha})}, \quad (6)$$

which is well-defined since $h(\boldsymbol{\alpha}) > 0$ for $\boldsymbol{\alpha} \in C_d$ and non-collapsed elements. The proximal step in the non-smooth part $g(\boldsymbol{\alpha})$ is then

$$
\begin{aligned}
\text{prox}_{\frac{1}{L} g}(\tilde{\boldsymbol{\alpha}}) &= \arg\min_{\boldsymbol{\alpha}} g(\boldsymbol{\alpha}) + \frac{L}{2} \|\boldsymbol{\alpha} - \tilde{\boldsymbol{\alpha}}\|^2 \\
&= \text{proj}_{C_d}\left(\tilde{\boldsymbol{\alpha}} - L^{-1} \boldsymbol{\Phi} \right),
\end{aligned}
$$

where L is a constant. The projection proj_{C_d} onto simplices is available in closed form [23, Figure 1]. The ISTA algorithm consists of the following update step:

$$\boldsymbol{\alpha}_{k+1} = \text{proj}_{C_d}\left(\left(\boldsymbol{\alpha}_k - L^{-1} \nabla h(\boldsymbol{\alpha}_k) \right) - L^{-1} \boldsymbol{\Phi} \right). \quad (7)$$

An optimal choice of L in the algorithm is the Lipschitz constant of ∇h, which we can bound as

$$L \leq \frac{\|\mathbf{A}^\top \mathbf{A}\|}{\underline{h}} \left(1 + \frac{\|\mathbf{A}\|}{\underline{h}} \right), \quad (8)$$

where $\|\cdot\|$ in (8) refers to the spectral norm of the operator and $\underline{h} = \frac{1}{d\sqrt{d}} \sigma_{\min}(\mathbf{A}) \leq \|\mathbf{A}\boldsymbol{x}\|_2 = h(\boldsymbol{x})$ is a lower bound of $h(\boldsymbol{x})$, for $\sigma_{\min}(\mathbf{A})$ being the smallest singular value of $\mathbf{A}$. Note that $\mathbf{A}$ is constant in each element, so the local Lipschitz constants can be efficiently pre-computed.

Algorithm 1: Solution of the local problem in Eq. (5)

Data: $\mathbf{A}$, L, $\boldsymbol{\Phi} = (\phi_1, \ldots, \phi_d)^\top$
Result: $\boldsymbol{\alpha}_{n_f}$
1   $\hat{\boldsymbol{\alpha}}_0 = \frac{1}{d}\mathbf{1}$
2   **for** $k = 1$ **to** n_f **do** `#FISTA main loop`
3     $h = \|\mathbf{A}\hat{\boldsymbol{\alpha}}_{k-1}\|$
4     $\nabla h = h^{-1}\mathbf{A}^\top\mathbf{A}\hat{\boldsymbol{\alpha}}_{k-1}$ `#Eq. (6)`
5     $\tilde{\boldsymbol{\alpha}}_k = (\hat{\boldsymbol{\alpha}}_{k-1} - L^{-1}\nabla h) - L^{-1}\boldsymbol{\Phi}$
6     $\boldsymbol{\alpha}_k = \text{proj}_{C_d}(\tilde{\boldsymbol{\alpha}}_k)$ `#See [23, Figure 1]`
7     $\hat{\boldsymbol{\alpha}}_k = \boldsymbol{\alpha}_k + \beta_k(\boldsymbol{\alpha}_k - \boldsymbol{\alpha}_{k-1})$ `#Acceleration`

We summarize the solution of the local problem in Algorithm 1, where we also include the acceleration on Line 7 with $\beta_k = \frac{k-1}{k+1}$, as per the FISTA method, see [22]. Note that Algorithm 1 works for arbitrary spatial dimension $d \in \mathbb{N}$. Since the local problem is convex, it has a unique minimum $\boldsymbol{\alpha}^*$. The rate of convergence is (see [22, Theorem 4.4])

$$\phi_{j \setminus i}^k - \phi_{j \setminus i}^* \leq \frac{2L \|\boldsymbol{\alpha}_0 - \boldsymbol{\alpha}^*\|^2}{(k+1)^2}. \quad (9)$$

Finally, we compute the global eikonal solution by iteratively applying Algorithm 1 to each element $E_j \in E$ and for all vertices $\boldsymbol{v}_i \in E_j$, and then take minimum: see Algorithm 2.

Algorithm 2: Geodesic-BP

Data: $\mathcal{P}_1$ mesh $M = (V, \mathcal{E})$, piecewise-constant metric $\mathbf{D}$, initial conditions $X_0 = \left\{(\boldsymbol{x}_i, t_i)\right\}_{i=1}^{K}$

Result: Geodesic distances ϕ_K

1   $\forall \boldsymbol{v} \in V : \phi_{-1}(\boldsymbol{v}) = \phi_0(\boldsymbol{v}) = \infty$

2   $\forall (\boldsymbol{x}_i, t_i) \in X_0 : \phi_{-1}(\boldsymbol{x}_i) = \phi_0(\boldsymbol{x}_i) = t_i$

     `#Compute all simplex upwind directions towards`
     `each vertex i inside element j`

3   $\tilde{E} = \bigcup_{i=1}^{d+1} \bigcup_j E_{j \setminus i}$ Compute for each element $\mathbf{A}_{j \setminus i}, L_{j \setminus i}$ `#see Eqs. (5) and (8)`

4   **for** $k = 1$ **to** K **do**

5      $\phi_k := \phi_{k-1}$

        `#Compute all updates in parallel`

6   **for parallel** $\tilde{E}_{j \setminus i} \in \tilde{E}$ **do**

           `#Collect the simplex vertices opposite`
           `of the vertex i`

7          $(\boldsymbol{v}_{j_1}, \ldots \boldsymbol{v}_{j_d}) = \tilde{E}_{j \setminus i}$

           `#Optional: Only compute updates when`
           `necessary`

8   **if** $\exists \boldsymbol{v}_{j_n} \in \tilde{E}_{j \setminus i} : (\phi_{k-2}(\boldsymbol{v}_{j_n}) - \phi_{k-1}(\boldsymbol{v}_{j_n})) \geq \varepsilon$ **then**

9             $\boldsymbol{\Phi} = (\phi_k(\boldsymbol{v}_{j_1}), \ldots, \phi_k(\boldsymbol{v}_{j_d}))^\top$

10            $\boldsymbol{\alpha}_{n_f} = $ Algorithm 1$(\mathbf{A}_{j \setminus i}, L_{j \setminus i}, \boldsymbol{\Phi})$

11   **With** *enabled gradient computations*

              `#Eq. (5)`

12              $\phi_{j \setminus i} = \left\langle \boldsymbol{\alpha}_{n_f}, \boldsymbol{\Phi} \right\rangle + \left\| \mathbf{A}_{j \setminus i} \boldsymbol{\alpha}_{n_f} \right\|$

13   **With** *enabled gradient computations*

14   **for parallel** $v_i \in V$ **do**

           `#see Eq. (2)`

15            $\phi_k(\boldsymbol{v}_i) = \min_{E_j \in \omega_i} \phi_{j \setminus i}$

        `#ε-convergence criterion`

16   **if** $\forall \boldsymbol{v} \in V : (\phi_{k-1}(\boldsymbol{v}) - \phi_k(\boldsymbol{v})) < \varepsilon$ **then**

17         $\phi_K = \phi_k$

18   **break**

C. Geodesic-BP

To utilize the computed eikonal solution in an inverse approach, we want to compute the gradients of the eikonal solution w.r.t. some quantity of interest, e.g., the initial conditions X_0. More precisely, we are interested in the gradient $\nabla_{(\boldsymbol{x}_i, t_i)} \phi(\boldsymbol{x})$ and its efficient computation. This can be done analytically for (1), or algorithmically from Algorithm 2. The former (see e.g. [10]) relies on a *optimize-then-discretize* strategy. whereas the latter employs a *discretize-then-optimize* approach. Following the latter, we compute $\nabla_{(\boldsymbol{x}_i, t_i)} \phi(\boldsymbol{v})$ for all vertices $\boldsymbol{v} \in V$ and observe that

$$\nabla_{(\boldsymbol{x}_i, t_i)} (\phi(\boldsymbol{x})) = \left\langle \boldsymbol{\alpha}, \nabla_{(\boldsymbol{x}_i, t_i)} (\phi_{j_1}, \ldots, \phi_{j_{d+1}})^\top \right\rangle, \quad (10)$$

for $\boldsymbol{x} \in E_j$. Note that we neglect the derivative of $\boldsymbol{\alpha}$ w.r.t. X_0 as a consequence of the first variation of the geodesic distance [24, Chapter 10]. The nodal values ϕ_{j_n} in succession are then derived through (3), until we terminate at the simplex containing the initial condition $\boldsymbol{x}_i \in E_k$. By defining $\boldsymbol{x}_i$ in a continuous fashion (as was done in [10]), such that for all vertices $\boldsymbol{v}_{j_n}$ of the element E_j containing $\boldsymbol{x}_i$ it holds that

$$\phi_{j_n} = t_i + \left\| \boldsymbol{v}_{j_n} - \boldsymbol{x}_i \right\|_{\mathbf{D}_j}, \quad (11)$$

the derivative readily follows:

$$\nabla_{\boldsymbol{x}_i} \phi_{j_n} = -\frac{\mathbf{D}_j \left(\boldsymbol{v}_{j_n} - \boldsymbol{x}_i \right)}{\left\| \boldsymbol{v}_{j_n} - \boldsymbol{x}_i \right\|_{\mathbf{D}_j}}, \quad \nabla_{t_i} \phi_{j_n} = 1,$$

which can be effectively calculated using automatic differentiation of (11), or subsequent nodes (10) which are connected by a geodesic path to $(\boldsymbol{x}_i, t_i)$. In particular, we use the back-propagation algorithm implemented in many ML libraries. More technical aspects are provided in Section II-E.

D. ECG computation

We describe the torso with a domain $\Omega_0 \subset \mathbb{R}^d$, whereas the heart is denoted by Ω, that is assumed to be well contained in Ω_0. The full heart-torso domain is $\Omega_T = \bar{\Omega}_0 \cup \bar{\Omega}_T$, and the torso surface is $\Sigma = \partial \Omega_T$. The electrodes are placed on the torso at fixed locations $\tilde{X}_e = \{\boldsymbol{x}_{e_i}\}_{i=1}^{N} \subset \Sigma$. Each lead $\mathcal{V}_{l_e}(t)$ of the ECG is defined as the zero-sum of the electric potential at multiple locations. Thanks to the lead-field formula, we have the following representation:

$$\mathcal{V}_{l_e}(t) = \int_\Omega \left\langle \mathbf{G}_i(\boldsymbol{x}) \nabla Z_{l_e}(\boldsymbol{x}), \nabla V_m(\boldsymbol{x}, t) \right\rangle \mathrm{d}\boldsymbol{x}, \quad (12)$$

where $\mathbf{G}_i : \Omega \to \mathcal{S}_d$ is the intracellular conductivity tensor, $V_m(\boldsymbol{x}, t) : \Omega \times [0, T] \to \mathbb{R}$ is the transmembrane potential, and $Z_{l_e} : \Omega_T \to \mathbb{R}$ is the l_e-th lead field, solution of the problem:

$$\begin{cases} - \nabla \cdot \left((\mathbf{G}_i + \mathbf{G}_e) \nabla Z_{l_e} \right) = 0, & \text{in } \Omega_T, \\ - \nabla Z_{l_e} \cdot \mathbf{n} = \delta_{\boldsymbol{x}_{l_e}} - \dfrac{1}{|X_W|} \sum_{\boldsymbol{y} \in X_W} \delta_{\boldsymbol{y}}, & \text{on } \Sigma. \end{cases} \quad (13)$$

The set X_W contains the electrodes at the left/right arm and left foot, used to form the Wilson's Central Terminal (WCT). Eq. (13) can be solved once, since the problem is time-independent. The bulk conductivity of the torso, $\mathbf{G}_i + \mathbf{G}_e$, also depends on the extracellular conductivity tensor $\mathbf{G}_e : \Omega \to \mathcal{S}_d$. The derivation of this problem and exact formulation of (13) can be found in further detail in [5], [25], [10].

Next, we assume that the transmembrane potential is travelling wave with activation times dictated by $\phi(\boldsymbol{x})$:

$$V_m(\boldsymbol{x}, t) = U(t - \phi(\boldsymbol{x})) = U(\xi). \quad (14)$$

Since we are only interested in the depolarization sequence, we use the following template action potential, see [10], [7]:

$$U(\xi) = K_0 + \frac{K_1 - K_0}{2} \tanh(2\xi / \tau), \quad (15)$$

where K_0 and K_1 are respectively the minimum and maximum transmembrane potential.

Additionally, we precompute the discretized integral ECG

This article has been accepted for publication in IEEE Transactions on Biomedical Engineering. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TBME.2023.3331876

5

operator (12) by means of a 2nd order simplicial Gaussian quadrature scheme. We used the library scikit-fem [26] to assemble this operator for our $\mathcal{P}_1$ mesh, automatically interpolating and integrating V_m at the Gaussian quadrature points. This way, we can compute the leads as $\mathcal{V}_{l_e}(t) = (BU(t - \phi))_{l_e}$ where $B : \mathbb{R}^{n_v \times N}$ is the mentioned discretized, precomputed lead field integration operator of (12), $n_v = |V|$ are the number of vertices and N is the number of leads.

For optimizing the inverse ECG problem, we want to fit $\mathcal{V}_{l_e}$ of (12) against a measured $\hat{\mathcal{V}}_{l_e}$. We consider for that purpose the minimization problem

$$\min_{(\boldsymbol{x}_i, t_i) \in X_0} \frac{1}{N |T|} \sum_{l_e} \int_T \left(\mathcal{V}_{l_e}(\phi_{X_0}, t) - \hat{\mathcal{V}}_{l_e}(t) \right)^2 \mathrm{d}t, \quad (16)$$

where ϕ_{X_0} is the eikonal solution for given initial conditions and N defines the number of measured/computed leads. Note that with B computed from (12), we can easily conclude for $\xi = t - \phi$ that $\nabla_{(\boldsymbol{x}_i, t_i)} \mathcal{V}_{l_e}(t) = -B^\top \frac{\partial U(\xi)}{\partial \xi} \nabla_{(\boldsymbol{x}_i, t_i)} \phi$, where $\nabla_{(\boldsymbol{x}_i, t_i)} \phi$ can be computed using Algorithm 2 as discussed in Section II-C. Note that backpropagation through $\nabla_{(\boldsymbol{x}_i, t_i)} \mathcal{V}_{l_e}$ will yield the equivalent result. The resulting gradients were then used for a gradient-based optimization utilizing ADAM [27].

E. Implementation aspects

There are a few practical remarks when implementing Algorithms 1 and 2 that are important to consider: When computing the local update in Algorithm 1, the number of iterations n_f is the dominating factor for performance. Since the local problem is strictly convex with a precomputed Lipschitz constant L_E, the exact error bounds (see (9)) towards the global minimum in (5) are known and could be used to compute n_f on an element basis. Note that L_j only depends on $\mathbf{A}_{j \setminus i}$, which is strongly related to the mesh geometry and may increase as element quality decreases.

Furthermore, The upper bound presented in (8) is not tight and might lead to slow convergence of the local FISTA updates in Algorithm 1. In practice, we thus recomputed $\underline{h}$ by minimizing $\boldsymbol{\alpha}^*$ through Algorithm 1 with $\boldsymbol{\Phi} = \mathbf{0}$ and using the estimate $\underline{h} = h(\boldsymbol{\alpha}^*)/1.1$.

When considering Algorithm 2, note that the parallelization denoted in line 6 is more descriptive for CPU-parallel implementations. In ML frameworks, in contrast, the parallelization comes from the vectorization of single implemented operations, which might require re-ordering of some operations. Still, all operations inside the loop starting from line 6 can be parallelized, where only a special parallelization is required for Line 14 (scatter operation). This makes it well-suited for vectorized implementations on a GPU in frameworks such as PyTorch [28].

The optional performance optimization denoted in line 8 already significantly cuts down on the required memory for backpropagation, only computing updates when the solutions of simplex elements have changed. Checking for an ε change in activations before the update can also avoid unwanted gradient eliminations in cases where the old and new values of ϕ take the same value.

Note that the local problem in Algorithm 1 could also be solved using the exact formulation derived for low-dimensional cases ($d \leq 4$) and fixed simplex dimensions as in previous works [12], [20]. In practice, however, the parallelization of this explicit solution is limited by the required condition to satisfy the constraint in (4). Such branching leads to unwanted performance losses in SIMD architectures when parallelizing over multiple elements.

Note that the number of activation sites is never modified by Algorithm 2. Some activation sites may however be overshadowed by the surrounding activation, thus effectively providing no contribution to the overall activation. When this happens, the gradient w.r.t. the overshadowed site is simply zero, thus the point is ignored by the algorithm. If an initial condition $(\boldsymbol{x}, t_i)$ is overshadowed as described at the final iteration of the optimization, we often refer to it as inactive.

III. RESULTS

In this section, details of the dataset and results of applying our method for matching an ECG for both a synthetic 2-D and a real 3-D case are provided.

A. Simulation study

1) Setup: To test and analyze our algorithm, we first start with a simple 2-D simulation study. We designed a simple setup consisting of an idealized torso with 8 electrodes distributed around the surface, with left and right arm electrodes composing Wilson's central terminal (WCT) (see Figure 3). Seven lead fields were then computed each w.r.t. WCT.
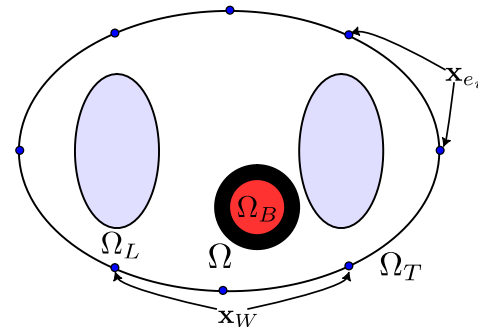

Fig. 3. The 2D simulation setup. We consider a simple symmetrical left ventricular slice $\boldsymbol{\Omega}$ to be located in a heterogeneous torso $\boldsymbol{\Omega_T}$, containing the blood pool $\boldsymbol{\Omega_B}$ and lungs $\boldsymbol{\Omega_L}$, with different conductivities. 8 electrodes $\boldsymbol{x_{e_i}}$ (blue circles) are circularly distributed around the torso, two of which are chosen as the WCT $\boldsymbol{x_W} \in \boldsymbol{X_W}$.

The LV, meshed at a resolution of $0.9 \,\mathrm{mm}$ (≈ 2700 DoFs), contained 4 initial condition tuples $(\boldsymbol{x}_i, t_i)$, located throughout half of the LV myocardium with timings between $2.5 \,\mathrm{ms}$ and $10 \,\mathrm{ms}$. For computing the ground truth, we use another model with a $0.45 \,\mathrm{mm}$ space resolution and perturbed conductivities in all sub-domains of the torso. As an initial guess, we consider 8 activation sites, all with $0 \,\mathrm{ms}$ timing and evenly distributed at the transmural center. The optimization was performed using ADAM [27] for 400 epochs using a learning rate of 0.5, taking $\approx 5 \,\mathrm{min}$ on a AMD Ryzen 7 3800X 8-core processor.
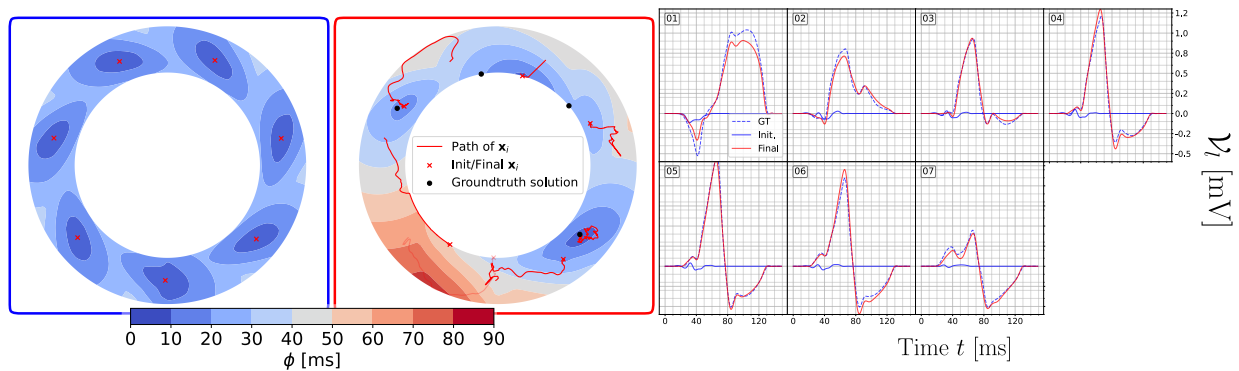
Fig. 4. Result of the ECG optimization on the idealized torso experiment using **7** leads. The initial model (left) uses a simple simultaneous activation, evenly spread throughout the myocardium. The optimization reduces the number of initial conditions and is able to match both ϕ and $\mathcal{V}_l$ closely (middle and right). The transparent red lines in the middle plot show the optimization path of initial conditions $(\boldsymbol{x_i}, \boldsymbol{t_i})$, which are inactive in the final solution.

2) Results: The results of the optimization against the initial and ground-truth solution and parameters are shown in Figure 4. We can nearly perfectly fit the ECG and approximately recover the initial conditions. Overall, the total root mean squared error (RMSE) w.r.t. the eikonal solution ϕ is only 5.8 ms. We tested the algorithm also for a varying number of initial conditions and found that it is a hyperparameter of importance, but need not be chosen exactly: A cross-validation on the 2-D torso revealed that the worst match is achieved, by underestimating the number of initial conditions. In such a case, the achievable activations are not complex enough and the model is unable to faithfully replicate both the ECG and unobservable eikonal solution. Increasing the number of initial conditions beyond the ground truth's number is neither able to better match the ECG by a significant margin, but will also not overfit and significantly increase the error on the eikonal solution. The most common cause of high errors on the eikonal solution is related to getting trapped in local minima when optimizing for the ECG, circumventable by batched optimization (not utilized in any of the shown experiments).

B. Body Potential Surface Map Reconstruction

1) Setup: The second anatomy is from a public rabbit ECGi model [29]. The anesthetized rabbit was equipped with a 32-lead ECG vest during normal sinus rhythm and the signals were recorded with a frequency of 2048 Hz. The ECG was recorded for 10 seconds, filtered, and averaged over all beats to produce a single beat. From this beat, we extracted the QRS duration to match the biventricular activation. The biventricular mesh consisted of 508 975 DoFs and 2 522 501 tetrahedra, roughly equivalent to a mean edge length of 0.82 mm in a human heart (assuming a size factor of 3 between rabbit and human hearts). Note that we also tested the optimization with a lower resolution of 143 564 DoFs and 591 975 tetrahedra, roughly equivalent to a mean edge length of 1.25 mm.

We use the electrophysiological parameters from [30], including cardiac fiber orientation, lead fields Z_l, and conduction velocity. In the conducted experiment, we chose $K_0 = -85$ mV, $K_1 = 30$ mV, $\tau_1 = 1$ ms in (15). The extracted QRS sequence was chosen during the time interval $t \in$

$[100, 165]$ ms. We optimized for the initial conditions $(\boldsymbol{x_i}, t_i)$ of the model, i.e. X_0 in Algorithm 2. For this purpose, we randomly distributed 100 points on the biventricular surface $(\boldsymbol{x_i} \in \partial\Omega)$ and with all timings initialized to the approximate Q onset $t_i = 20$ ms.

2) Results: The experiment was run for 400 epochs using a learning rate of 0.5 on an Nvidia A100 GPU (40GB), taking approximately 1.8 h (or 0.7 h on the reduced resolution). The computed results can be found on Zenodo [31].

The results for optimizing (16) with 400 epochs of ADAM using Algorithm 2 can be seen in Figure 5. The intial and optimized activation are shown in the left upper and lower side of the image. is shown together with the comparison of all ECGs (right). The initially randomly sampled points evenly cover the domain Ω, to allow for a fast excitation of the whole domain, though this leads to overexpressed deflections in many leads initially. The optimization significantly changes the activation of the heart and is able to achieve a promising match of the ECG in many leads. The loss in (16) converges after ≈ 250 epochs, with only minor variations in the parameters thereafter. Of the initial 100 points, only 91 are still active at the end of the optimization, meaning that the corresponding t_i is contained in the final solution and not earlier activated by closer initial conditions.

IV. DISCUSSION

We presented Geodesic-BP, an algorithm to efficiently solve the inverse ECG problem while being very efficient on a single GPU. We are able to fit the parameters of a patient-specific model only from the surface ECG and the segmented anatomy. The fitted model faithfully reproduced the recorded ECG in an animal experiment. In a synthetic case (Section III-A), Geodesic-BP is able to recover the ground-truth solution with high accuracy, even in the presence of perturbation in the forward model and in the lead field. The resolution of the considered 3-D heart model consists of $> 10^5$ DoFs, which is sufficiently accurate for a high-fidelity eikonal solution in a human heart, as it would correspond to roughly 0.8 mm edge length. At full resolution, all computations only take 16 s.

The considered cardiac parameter estimation problem is of high-relevance [32], [33], [1], [34], yet only a few works are
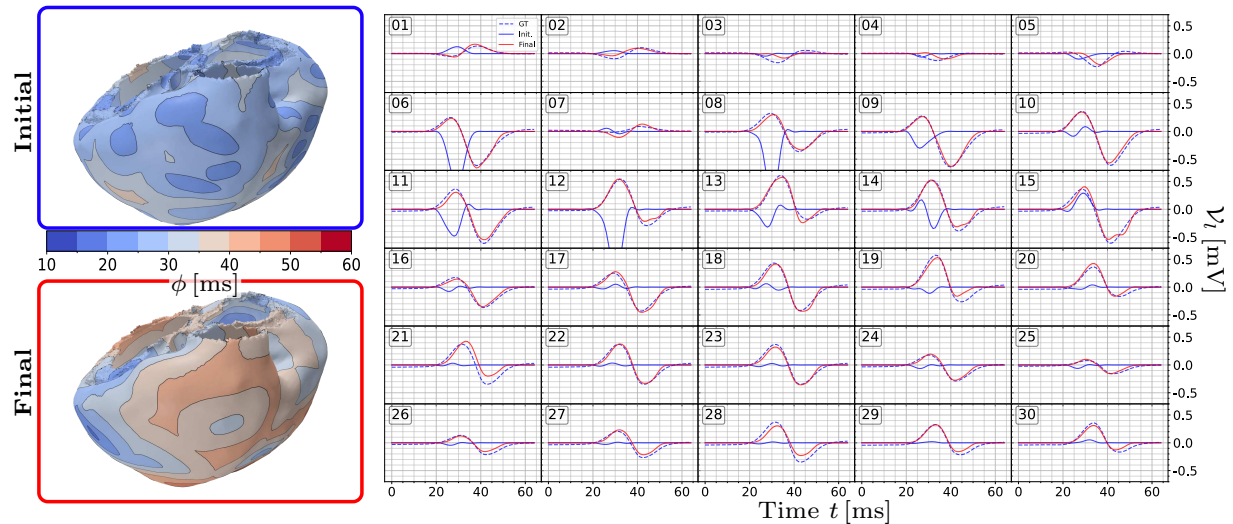
Fig. 5. Result of the ECG optimization on the rabbit heart model. We compare the initial and final model (blue/red) and show the computed activations $\phi(x)$ (left) and BPSM $\mathcal{V}_l(t)$ for each model (right). The computed ECGs (solid lines) are shown on top of the measured ECGs (dashed blue line). We show the first 30 of 31 available leads to save an extra row.

able to efficiently solve this problem using solely non-invasive ECG measurements. The considered heart model and ECG originates from a real-world rabbit rather than a human model since it is (to the knowledge of the authors) the only publicly available full torso and heart model, captured with this high-level detail and including preprocessed electrophysiological measurements. Such animal models share many properties of human cardiac electrophysiology [35]. Furthermore, the algorithm and equations involved translate 1:1 to human cardiac electrophysiology. Likewise, the methods involved are not restricted to biventricular models, and can be extended to atrial, or whole heart models under sinus rhythm. We also foresee the future possibility of functionalizing publicly available human heart and torso models, currently missing parameterization [36].

The runtime in the full rabbit heart model of 1.8 h for the computed 400 iterations is still expensive but it reduces to 0.7 h for a coarser resolution (see Section III-B.2). Two main advantages of our novel method are the high number of parameters that can be simultaneously optimized, allowing for many initial sites in contrast to other works [3], [2], and having no spatial restrictions on the onset location [9]. The first advantage is a consequence of utilizing backpropagation to compute the gradient w.r.t. our parameters, resulting in the runtime not being dependent on the number of parameters and the convergence being dictated by properties of the underlying function to optimize [37]. Thus, compared to existing works based on derivative-free optimization with the same objective, our algorithm will scale much better in terms of spatially-varying parameters such as conduction velocities [38]. The latter advantage of arbitrary onset locations follows from the use of the volumetric description of the eikonal equation and lead fields, from which we can compute parameters and the local activation times not only on the heart surface, but also transmurally. In conjunction, both of these advantages allow us to model potentially more complex activations than previously

were able by other works.

From an implementation point of view, Algorithm 2 can be readily applied to optimizing the conduction velocity tensor $\mathbf{D}$ with little modification. Based on our previous research [32], [16], however, optimizing conduction velocities either requires strong prior information in the form of regularization, or multiple activation sequences to create meaningful results. The question how well such information can be reconstructed from the BSPM, and the space of possible solutions w.r.t. initial conditions both remain topics for future research.

One major limitation of Geodesic-BP is that it cannot handle re-entry phenomena such as tachycardia and fibrillation, because it relies on the eikonal equation with focal boundary conditions. However, the eikonal model can be reformulated to allow re-entry, see [39] and it could be used in an optimization problem where the initial condition is to be identified from the ECG. Finally, we also foresee many practical improvements for reducing the computational time of the local solver as outlined Section II-E.

V. Conclusions

This paper presents Geodesic-BP, an efficient method to compute piecewise linear geodesics, and demonstrated its performance for solving the inverse ECG problem in a limited parameter setting. By creating the link between backpropagation through a FEM-solver with geodesic backtracking, we were able to show that our previous FEM-based inverse eikonal models [15], [16] are linked to our recent work [10] working with geodesic backtracking. Such inverse procedures are current research topics of high-relevance and efficient inverse eikonal models such as the present one may help in advancing patient-specific cardiac electrophysiology to practical real-world scenarios in the near future.

References

[1] J. Corral-Acero, F. Margara, M. Marciniak, C. Rodero, F. Loncaric, Y. Feng, A. Gilbert, J. F. Fernandes, H. A. Bukhari, A. Wajdan,

This article has been accepted for publication in IEEE Transactions on Biomedical Engineering. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TBME.2023.3331876

8

M. V. Martinez, M. S. Santos, M. Shamohammdi, H. Luo, P. Westphal, P. Leeson, P. DiAchille, V. Gurev, M. Mayr, L. Geris, P. Pathmanathan, T. Morrison, R. Cornelussen, F. Prinzen, T. Delhaas, A. Doltra, M. Sitges, E. J. Vigmond, E. Zacur, V. Grau, B. Rodriguez, E. W. Remme, S. Niederer, P. Mortier, K. McLeod, M. Potse, E. Pueyo, A. Bueno-Orovio, and P. Lamata, "The 'Digital Twin' to enable the vision of precision cardiology," *European Heart Journal*, vol. 41, pp. 4556–4564, Dec. 2020.

[2] S. Pezzuto, F. W. Prinzen, M. Potse, F. Maffessanti, F. Regoli, M. L. Caputo, G. Conte, R. Krause, and A. Auricchio, "Reconstruction of three-dimensional biventricular activation based on the 12-lead electrocardiogram via patient-specific modelling," *EP Europace*, vol. 23, no. 4, pp. 640–647, 2021.

[3] K. Gillette, M. A. F. Gsell, A. J. Prassl, E. Karabelas, U. Reiter, G. Reiter, T. Grandits, C. Payer, D. Štern, M. Urschler, J. D. Bayer, C. M. Augustin, A. Neic, T. Pock, E. J. Vigmond, and G. Plank, "A Framework for the generation of digital twins of cardiac electrophysiology from clinical 12-leads ECGs," *Medical Image Analysis*, vol. 71, p. 102080, July 2021.

[4] J. P. Keener, "An eikonal-curvature equation for action potential propagation in myocardium," *Journal of mathematical biology*, vol. 29, no. 7, pp. 629–651, 1991. Publisher: Springer.

[5] P. C. Franzone, L. F. Pavarino, and S. Scacchi, *Mathematical cardiac electrophysiology*, vol. 13. Springer, 2014.

[6] A. Neic, F. O. Campos, A. J. Prassl, S. A. Niederer, M. J. Bishop, E. J. Vigmond, and G. Plank, "Efficient computation of electrograms and ECGs in human whole heart simulations using a reaction-eikonal model," *Journal of Computational Physics*, vol. 346, pp. 191–211, Oct. 2017.

[7] S. Pezzuto, P. Kal'avský, M. Potse, F. W. Prinzen, A. Auricchio, and R. Krause, "Evaluation of a Rapid Anisotropic Model for ECG Simulation," *Frontiers in Physiology*, vol. 8, p. 265, 2017.

[8] J. Camps, B. Lawson, C. Drovandi, A. Minchole, Z. J. Wang, V. Grau, K. Burrage, and B. Rodriguez, "Inference of ventricular activation properties from non-invasive electrocardiography," *Medical Image Analysis*, vol. 73, p. 102143, Oct. 2021.

[9] S. Pezzuto, P. Perdikaris, and F. Sahli Costabal, "Learning cardiac activation maps from 12-lead ECG with multi-fidelity Bayesian optimization on manifolds," *IFAC-PapersOnLine*, vol. 55, pp. 175–180, 2022.

[10] T. Grandits, A. Effland, T. Pock, R. Krause, G. Plank, and S. Pezzuto, "GEASI: Geodesic-based earliest activation sites identification in cardiac models," *International Journal for Numerical Methods in Biomedical Engineering*, vol. 37, no. 8, p. e3505, 2021. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/cnm.3505.

[11] L. Li, J. Camps, A. Banerjee, M. Beetz, B. Rodriguez, and V. Grau, "Deep Computational Model for the Inference of Ventricular Activation Properties," in *Statistical Atlases and Computational Models of the Heart. Regular and CMRxMotion Challenge Papers* (O. Camara, E. Puyol-Antón, C. Qin, M. Sermesant, A. Suinesiaputra, S. Wang, and A. Young, eds.), Lecture Notes in Computer Science, (Cham), pp. 369–380, Springer Nature Switzerland, 2022.

[12] W. Jeong and R. Whitaker, "A Fast Iterative Method for Eikonal Equations," *SIAM Journal on Scientific Computing*, vol. 30, pp. 2512–2534, Jan. 2008.

[13] T. Grandits, "A Fast Iterative Method Python package," *Journal of Open Source Software*, vol. 6, p. 3641, Oct. 2021.

[14] J. Fehrenbach and L. Weynans, "Source and metric estimation in the eikonal equation using optimization on a manifold," *Inverse Problems and Imaging*, vol. 17, no. 2, pp. 419–440, 2023.

[15] T. Grandits, K. Gillette, A. Neic, J. Bayer, E. Vigmond, T. Pock, and G. Plank, "An inverse Eikonal method for identifying ventricular activation sequences from epicardial activation maps," *Journal of Computational Physics*, vol. 419, p. 109700, Oct. 2020.

[16] T. Grandits, S. Pezzuto, J. M. Lubrecht, T. Pock, G. Plank, and R. Krause, "PIEMAP: Personalized Inverse Eikonal Model from Cardiac Electro-Anatomical Maps," in *Statistical Atlases and Computational Models of the Heart. M&Ms and EMIDEC Challenges* (E. Puyol Anton, M. Pop, M. Sermesant, V. Campello, A. Lalande, K. Lekadir, A. Suinesiaputra, O. Camara, and A. Young, eds.), Lecture Notes in Computer Science, (Cham), pp. 76–86, Springer International Publishing, 2021.

[17] P. C. Franzone and L. Guerri, "Spreading of excitation in 3-d models of the anisotropic cardiac tissue. I. validation of the eikonal model," *Mathematical Biosciences*, vol. 113, pp. 145–209, Feb. 1993.

[18] A. Quarteroni and A. Valli, *Numerical approximation of partial differential equations*, vol. 23. Springer Science & Business Media, 2008.

[19] Z. Fu, W. Jeong, Y. Pan, R. Kirby, and R. Whitaker, "A Fast Iterative Method for Solving the Eikonal Equation on Triangulated Surfaces,"

[20] *SIAM Journal on Scientific Computing*, vol. 33, pp. 2468–2488, Jan. 2011.

[20] Z. Fu, R. Kirby, and R. Whitaker, "A Fast Iterative Method for Solving the Eikonal Equation on Tetrahedral Domains," *SIAM Journal on Scientific Computing*, vol. 35, pp. C473–C494, Jan. 2013.

[21] F. Bornemann and C. Rasch, "Finite-element discretization of static Hamilton-Jacobi equations based on a local variational principle," *Comput. Vis. Sci.*, vol. 9, no. 2, pp. 57–69, 2006.

[22] A. Beck and M. Teboulle, "A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems," *SIAM Journal on Imaging Sciences*, vol. 2, pp. 183–202, Jan. 2009.

[23] J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra, "Efficient projections onto the l1-ball for learning in high dimensions," in *Proceedings of the 25th international conference on Machine learning*, ICML '08, (New York, NY, USA), pp. 272–279, Association for Computing Machinery, July 2008.

[24] B. O'neill, *Semi-Riemannian geometry with applications to relativity*. Academic press, 1983.

[25] J. P. Keener and J. Sneyd, *Mathematical physiology*, vol. 1. Springer, 1998.

[26] T. Gustafsson and G. D. McBain, "scikit-fem: A Python package for finite element assembly," *Journal of Open Source Software*, vol. 5, no. 52, p. 2369, 2020.

[27] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," *arXiv:1412.6980 [cs]*, Jan. 2017. arXiv: 1412.6980.

[28] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "PyTorch: An Imperative Style, High-Performance Deep Learning Library," in *Advances in Neural Information Processing Systems*, vol. 32, Curran Associates, Inc., 2019.

[29] R. Moss, E. M. Wülfers, R. Lewetag, T. Hornyik, S. Perez-Feliz, T. Strohbach, M. Menza, A. Krafft, K. E. Odening, and G. Seemann, "A volumetric model of rabbit heart and torso including ECG data and ventricular activation sequence," Mar. 2022. DOI: 10.5281/zenodo.6340066.

[30] R. Moss, E. M. Wülfers, R. Lewetag, T. Hornyik, S. Perez-Feliz, T. Strohbach, M. Menza, A. Krafft, K. E. Odening, and G. Seemann, "A computational model of rabbit geometry and ECG: Optimizing ventricular activation sequence and APD distribution," *PLOS ONE*, vol. 17, p. e0270559, June 2022. Publisher: Public Library of Science.

[31] T. Grandits, J. Verhülsdonk, G. Haase, A. Effland, and S. Pezzuto, "Geodesic-BP Dataset," Nov. 2023. DOI: 10.5281/zenodo.8027712.

[32] C. Ruiz Herrera*, T. Grandits*, G. Plank, P. Perdikaris, F. Sahli Costabal, and S. Pezzuto, "Physics-informed neural networks to learn cardiac fiber orientation from multiple electroanatomical maps," *Engineering with Computers*, July 2022.

[33] S. Coveney, C. Corrado, C. H. Roney, R. D. Wilkinson, J. E. Oakley, F. Lindgren, S. E. Williams, M. D. O'Neill, S. A. Niederer, and R. H. Clayton, "Probabilistic Interpolation of Uncertain Local Activation Times on Human Atrial Manifolds," *IEEE Transactions on Biomedical Engineering*, vol. 67, pp. 99–109, Jan. 2020. Conference Name: IEEE Transactions on Biomedical Engineering.

[34] C. H. Roney, J. Whitaker, I. Sim, L. O'Neill, R. K. Mukherjee, O. Razeghi, E. J. Vigmond, M. Wright, M. D. O'Neill, S. E. Williams, and S. A. Niederer, "A technique for measuring anisotropy in atrial conduction to estimate conduction velocity and atrial fibre direction," *Computers in Biology and Medicine*, vol. 104, pp. 278–290, Jan. 2019.

[35] K. E. Odening, I. Baczko, and M. Brunner, "Animals in cardiovascular research: important role of rabbit models in cardiac electrophysiology," *European Heart Journal*, vol. 41, p. 2036, June 2020.

[36] K. Aras, W. Good, J. Tate, B. Burton, D. Brooks, J. Coll-Font, O. Doessel, W. Schulze, D. Potyagaylo, L. Wang, P. van Dam, and R. MacLeod, "Experimental Data and Geometric Analysis Repository—EDGAR," *Journal of Electrocardiology*, vol. 48, pp. 975–981, Nov. 2015.

[37] A. Beck, *First-order methods in optimization*. SIAM, 2017.

[38] A. R. Conn, K. Scheinberg, and L. N. Vicente, *Introduction to derivative-free optimization*. SIAM, 2009.

[39] E. Pernod, M. Sermesant, E. Konukoglu, J. Relan, H. Delingette, and N. Ayache, "A multi-front eikonal model of cardiac electrophysiology for interactive simulation of radio-frequency ablation," *Computers & Graphics*, vol. 35, no. 2, pp. 431–440, 2011.