

Smoothness and continuity of cost functionals for ECG mismatch computation^{*}

Thomas Grandits^{*,****,**} Simone Pezzuto^{***}
Gernot Plank^{**,****}

^{*} Institute of Mathematics and Scientific Computing, University of Graz, Graz, Austria (e-mail: thomas.grandits@uni-graz.at)

^{**} Gottfried Schatz Research Center - Division of Biophysics, Medical University of Graz, Graz, Austria (e-mail: gernot.plank@medunigraz.at)

^{***} Center for Computational Medicine in Cardiology, Euler Institute, Università della Svizzera italiana, Lugano, Switzerland, (e-mail: simone.pezzuto@usi.ch)

^{****} BioTechMed-Graz, Graz, Austria

Abstract: The field of cardiac electrophysiology tries to abstract, describe and finally model the electrical characteristics of a heartbeat. With recent advances in cardiac electrophysiology, models have become more powerful and descriptive as ever. However, to advance to the field of inverse electrophysiological modeling, i.e. creating models from electrical measurements such as the ECG, the less investigated field of smoothness of the simulated ECGs w.r.t. model parameters need to be further explored. The present paper discusses smoothness in terms of the whole pipeline which describes how from physiological parameters, we arrive at the simulated ECG. Employing such a pipeline, we create a test-bench of a simplified idealized left ventricle model and demonstrate the most important factors for efficient inverse modeling through smooth cost functionals. Such knowledge will be important for designing and creating inverse models in future optimization and machine learning methods.

Copyright © 2022 The Authors. This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Keywords: ECG, Cardiac modeling, Physiology, Parameter estimation, Inverse problems, Smoothness

1. INTRODUCTION

Modern health care is increasingly pushing towards personalized approaches for an improved therapeutic outcome over standard interventions, see Corral-Acero et al. (2020). Precision cardiology aims at individualizing computational models of the heart from patient-specific data—possibly with non-invasive data, such as cardiac imaging and the standard 12-lead electrocardiogram (ECG)—thus to provide to cardiologists an advanced tool to improve diagnosis and to optimize the therapeutic approach.

The keystone of model personalization is the identification of the parameters, which mathematically translates to an optimization problem. Specifically, we focus in this contribution on the case where the loss functional is solely based on the ECG. The electrophysiological inverse problem w.r.t. the ECG has been already considered in previous works with either derivative-free optimization (sometimes combinatorial), see e.g. Gillette et al. (2021); Camps et al. (2021); Pezzuto et al. (2021), or gradient-based approaches, see e.g. Grandits et al. (2021). In either

case, the parameter space is typically low-dimensional, with less than 20 parameters. To further advance the field of inverse cardiac electrophysiology, it will be indispensable to also explore higher dimensional spaces and thus allow for more varied models. However, increasing the dimensionality of the explored parameter space comes with the obvious drawback of aggravated optimization since the exploration of the parameter space will become exponentially more difficult, sometimes emphasized as the *curse of dimensionality*, see Bellman (2015). A better optimization approach relies on the computation of the local gradient of a cost function and perform best for smooth and convex functions. It is however currently still an open question, if complex electrophysiological simulations and matching them to ECGs is a challenging problem and how smooth the resulting loss function is.

While single studies have shown the effect of different parameters on the cost functionals of such inverse problems, see Grandits et al. (2020); Yang et al. (2017), they focused on a small subset of the parameters. Additionally, to the knowledge of the authors, no extensive study of the effect of cardiac electrophysiological models and ECG matching techniques on the continuity and smoothness of the resulting loss function have been conducted thus far. The problem of matching a recorded, patient-specific

^{*} This work was financially supported by the Theo Rossi di Montelera Foundation, the Metis Foundation Sergio Mantegazza, the Fidinam Foundation, the Horten Foundation and the CSCS-Swiss National Supercomputing Centre production grant s1074 to SP.

ECG to a model-based ECG is particularly subtle, because ECGs are typically not very smooth and affected by noise. Moreover, ECG amplitude and morphology may be subject to epistemic uncertainty, due poor electrode contact or other physiological aspects (respiration, epicardial fat, torso inhomogeneities) generally excluded by the modeling framework. Finally, ECG alignment and detection of the onset of activation is not always clear. All these aspects affect the *metric* to evaluate the mismatch.

In this paper, we will use several time series distance functions, applied between the optimal and current generated ECGs and visually analyze the difficulty of the resulting optimization problem in terms of visual smoothness. Evaluating cost functionals of parameter spaces with more than 3 to 4 parameters is no simple feat. Traditionally, a dimensionality reduction algorithm is first applied, such as principal component analysis, or more recently, machine learning inspired techniques such as auto-encoders or T-SNE Hinton and Roweis (2002). In this paper we will rather explore and visualize the loss function of an 8-dimensional space of important and intertwined parameters using techniques promoted for the visualization of smoothness in residual networks (ResNets) Li et al. (2018).

From a modeling perspective, we built a state-of-the-art pipeline for the preparation of simulations, from data to ECG. The reconstruction and generation of anatomical models or features, purely from imaging data (such as MRI) is a non-trivial task that many previous studies investigated, see Chen et al. (2020); Arevalo et al. (2016). Here, we will study the effect of mesh resolution of the heart on the optimization problem, ranging from 2.5 mm to 0.5 mm mean edge length. Moreover, given the importance of anisotropic conduction in the heart, being faster in fiber longitudinal direction, we will study its effect on the ECG w.r.t. the fiber angles.

Concerning the electrophysiology part, there is a vast literature on various models (bidomain, monodomain, eikonal) and their physiological accuracy, see Keener and Sneyd (1998); Sundnes et al. (2007); Jalife et al. (2011). The considered gold-standard model, the bidomain system, is however still very expensive to solve Vigmond et al. (2008). Our study requires us to repeatedly evaluate the electrical activation several thousand times, which necessitates evaluations in the timeframe of no more than a few seconds. We will therefore focus on the anisotropic eikonal model, which have been shown to offer a faithful approximation to the activation of the heart, while being fast to solve, see Franzone and Guerri (1993). When using such an eikonal model and prescribing a fixed membrane potential at the computed activation time, we can efficiently compute the ECGs measured at the body surface with only minor estimation errors, see Pezzuto et al. (2017).

The paper is organized as follows: In Sec. 2 we present the ECG model considered in this study, and the various metrics to compare ECGs. Several numerical experiments are summarized in Sec. 3 and later discussed in Sec. 4.

2. METHODS

2.1 The ECG model

We present here the full pipeline to generate the ECG model from anatomical data and parameters.

Geometric pipeline The reference anatomy was an idealized left ventricle (LV). The LV was built as a volume bounded by two truncated prolate spheroidal surfaces (see Fig. 1). The inner surface (endocardium) and outer surface (epicardium) have the same center and major axis, set to 37.5 mm, whereas the minor axis are set to 41.5 mm and 45 mm, respectively. The spheroids were truncated by a plane orthogonal to the major axis and at a distance of 35 mm from the center. The corresponding tetrahedral mesh was obtained at different resolutions using `fTetWild`, see Hu et al. (2020).

To ease the navigation in the LV, we make use of universal ventricular coordinates (UVCs) Bayer et al. (2018), as depicted in Fig. 1. The UVCs are the three functions

$$\theta(x) \in [0, 2\pi], \quad z(x) \in [0, 1], \quad \rho(x) \in [0, 1]$$

that measure, respectively, the angular position, the normalized longitudinal distance (from apex to base), the angular position and the normalized transmural distance (from endocardium to epicardium). Importantly, the map $\mathcal{U} : \Omega \ni x \mapsto (\theta, z, \rho)$ is invertible almost everywhere, so to make navigation in the LV easier. This system is built by solving a sequence of Laplace equations with given Dirichlet boundary conditions, followed by a normalization (see Bayer et al. (2018)).

The fiber and sheet structure in the LV was assigned with a rule-based approach, following Bayer et al. (2012). Thanks to the UVCs, the local basis can be expressed as

$$\mathbf{Q} = \begin{pmatrix} | & | & | \\ \hat{\mathbf{e}}_\theta & \hat{\mathbf{e}}_z & \hat{\mathbf{e}}_\rho \\ | & | & | \end{pmatrix},$$

for $\hat{\mathbf{e}}_\theta, \hat{\mathbf{e}}_z, \hat{\mathbf{e}}_\rho$ being the orthonormalized basis formed by applying Gram-Schmidt to the set $\{\nabla\theta, \nabla z, \nabla\rho\}$. In order to define the fiber \mathbf{f} , sheet \mathbf{s} and normal \mathbf{n} direction, the local basis in \mathbf{Q} was eventually rotated, first with a rotation of angle $\alpha(\rho)$ around the direction $\hat{\mathbf{e}}_\theta$, and then with a rotation of angle $\beta(\rho)$ around $\hat{\mathbf{e}}_\rho$:

$$\begin{pmatrix} | & | & | \\ \mathbf{f} & \mathbf{s} & \mathbf{n} \\ | & | & | \end{pmatrix} = \mathbf{Q} \begin{pmatrix} \cos \alpha & -\sin \alpha & 0 \\ \sin \alpha & \cos \alpha & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \beta & \sin \beta \\ 0 & -\sin \beta & \cos \beta \end{pmatrix}. \quad (1)$$

The angles $\alpha(\rho)$ and $\beta(\rho)$ were linearly interpolated, as follows:

$$\begin{aligned} \alpha(\rho) &= \alpha_1 + (\alpha_2 - \alpha_1)\rho, \\ \beta(\rho) &= \beta_1 + (\beta_2 - \beta_1)\rho. \end{aligned}$$

Electrophysiological pipeline We model the electrical activation throughout the heart, denoted by $\phi(x)$, $x \in \Omega$ with the anisotropic eikonal equation

$$\begin{cases} \sqrt{\mathbf{D}\nabla\phi \cdot \nabla\phi} = 1, & x \in \Omega, \\ \phi(x_k) = \phi_k, & k = 1, \dots, N. \end{cases} \quad (2)$$

The point-wise initial conditions were set through the UVC system, that is x_k is the projection of the closest mesh vertex of the point $\mathcal{U}^{-1}(\theta_k, z_k, \rho_k)$. The symmetric

positive-definite conduction velocity tensor \mathbf{D} , was defined as follows:

$$\mathbf{D} = K \cdot \mathbf{G}_i (\mathbf{G}_i + \mathbf{G}_e)^{-1} \mathbf{G}_e, \quad (3)$$

that is, we assumed a monodomain approximation of the electric conductivity. For the sake of simplicity, we assumed $K = 1$ in our experiments. The tensors \mathbf{G}_i and \mathbf{G}_e , respectively denoting the intra- and extra-cellular electric conductivity, were set as follows:

$$\mathbf{G}_{i,e} = \begin{pmatrix} | & | & | \\ \mathbf{f} & \mathbf{s} & \mathbf{n} \\ | & | & | \end{pmatrix} \begin{pmatrix} g_f^{i,e} & & \\ & g_s^{i,e} & \\ & & g_n^{i,e} \end{pmatrix} \begin{pmatrix} - & \mathbf{f} & - \\ - & \mathbf{s} & - \\ - & \mathbf{n} & - \end{pmatrix}, \quad (4)$$

with $g_f^e = 0.12$ S/m, $g_s^e = g_n^e = g_n^i = 0.08$ S/m. The values of g_f^i and g_s^i are variable.

For solving the eikonal equation (2), we used the `fim-python` library, see Grandits (2021), which implements the fast iterative method solver from Fu et al. (2013).

The depolarization from of the membrane potential was approximated by shifting a characteristic waveform v_m as follows:

$$v_m(x, t) = k_0 + \frac{k_1 - k_0}{2} \left[\tanh \left(2 \frac{t - \phi(x)}{\tau_1} \right) + 1 \right], \quad (5)$$

for $k_0 = -85$ mV, $k_1 = 30$ mV and $\tau_1 = 1$ ms. Finally, the electrical measurement for an electrode combination (referred to as leads) can then be computed as (see Keener and Sneyd (1998)):

$$V_j(t) = \int_{\Omega} \langle \nabla Z_j(x), \mathbf{G}_i(x) \nabla v_m(x, t) \rangle dx, \quad (6)$$

where $j = 1, 2, 3$ and Z_j refers to the lead field. Here, for the sake of simplicity, we considered ∇z_j equal to the canonical basis vectors \mathbf{e}_j , yielding a pseudo axis-aligned Frank ECG (Keener and Sneyd, 1998, p. 529).

2.2 Loss functionals

We consider the following problem: given a target ECG, denoted by $\hat{V}_j(t)$ for $j = 1, 2, 3$, we aim at identifying the free parameters in the ECG model depicted above, such that the mismatch between the target ECG and the simulated one is minimal. We discretize the time with a fixed time step $\Delta t = 1$ ms, that is $t_n = n\Delta t$ for $n = 0, \dots, M$ and $M\Delta t = T$. The simulated and target ECGs are sampled at $t = t_n$, and the resulting vector for each lead, denoted by \mathbf{V}_j and $\hat{\mathbf{V}}_j$, have respectively components $[V_j]_n = V_j(t_n)$ and $[\hat{V}_j]_n = \hat{V}_j(t_n)$.

Parameter space The free parameters in the model and their range of definition are:

- $\alpha_1 \in [-130, 10]$ and $\alpha_2 \in [-10, 130]$ (in degrees).
- $g_f^i \in [0.08, 0.6]$ and $g_s^i \in [0.04, 0.12]$ (in S/m).
- The position $x_1 \in \Omega$ and timing $\phi_0 \in [-25, 75]$ ms of the first activation site. In particular, we set x_1 from the UVCs $\theta_1 \in [-\pi, \pi]$, $\rho_1 \in [0, 1]$, and $z \in [0, 1]$.

Therefore, the dimension of the parameter space is 8.

The target ECG was obtained by taking the free parameters equal to mid-value in their range, and with the highest mesh resolution.

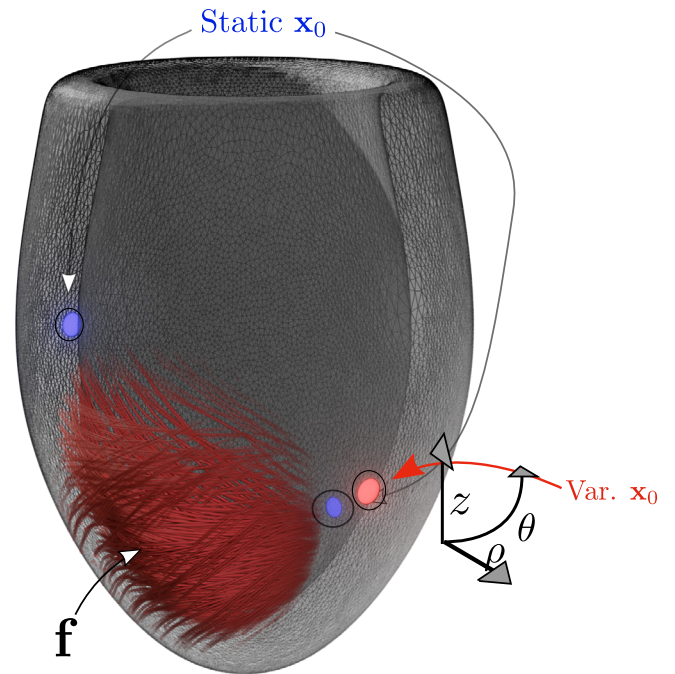


Fig. 1. LV model setup with the fiber field \mathbf{f} as streamlines. The excitation is initiated at three different locations \mathbf{x}_0 at different timings: two are fixed (blue) and the last one (red) is varying. Spatial exploration is performed in the UVC space (θ, ρ, z) that only allows feasible points inside the domain.

Loss metrics Besides standard error metrics, here we also consider dynamic time warping (DTW), see Sakoe and Chiba (1978), as recently proposed by Camps et al. (2021). In DTW, a cost matrix C is generated using a distance measure between all possible samples of two timeseries. The DTW algorithm will then find an optimal warp path through the cost matrix that optimally aligns the two signals, minimizing the cost of the warp path.

In summary, we compared the following metric distances $\mathcal{D}(\mathbf{V}_j, \hat{\mathbf{V}}_j)$:

- (1) Least-squares error: $\mathcal{D}_2(\mathbf{f}, \mathbf{g}) = \sum_n (f_n - g_n)^2$.
- (2) Dynamic Time Warping (DTW)
 - $\mathcal{D}_{\text{DTW},1}(\mathbf{f}, \mathbf{g})$ with $C_{n,m} = |f_n - g_m|$.
 - $\mathcal{D}_{\text{DTW},2}(\mathbf{f}, \mathbf{g})$ with $C_{n,m} = (f_n - g_m)^2$.
- (3) Cosine distance: $\mathcal{D}_{\text{cos}}(\mathbf{f}, \mathbf{g}) = 1 - \frac{\langle \mathbf{f}, \mathbf{g} \rangle}{\|\mathbf{f}\| \|\mathbf{g}\|}$.

Note that in all cases, the losses of all leads are simply summed up and in the case of DTW, the three cost matrices of the three leads $C_{X/Y/Z}$ are similarly added and DTW is performed on the resulting cost matrix $C = C_X + C_Y + C_Z$.

Visualization of the loss To showcase the complexity of the model w.r.t. the parameter space, it is often beneficial to visualize the loss landscape. For this, we follow Yang et al. (2017), where the authors consider two randomly sampled (possibly orthogonal) vectors from the parameter space and sample across a predefined range. This can be thought of as conceptually slicing the d -dimensional

hypercube with a hyperplane and computing the losses on this hyperplane. The main advantage of such an approach is that no dimensionality reduction is required, while the interdependence of parameters in the parameter space is not lost.

The mentioned landscape can then be computed by sampling the loss functional $\mathcal{D} : \mathbb{R}^d \rightarrow \mathbb{R}$ as

$$\mathcal{L}(a, b) = \mathcal{D}(\mathbf{x} + a\eta + b\delta), \quad (7)$$

where \mathbf{x} is the parameter vector of the considered ground-truth solution and \mathcal{D} is the chosen distance measure.

3. RESULTS

To give an overview of the loss landscape, we sampled the hyperplane with 50×50 samples, at 7 different resolutions for 5 different random direction pairs (hyperplanes), resulting in a total of 87,500 computed ECGs.

We based our analysis on one of the random hyperplanes, but the overall conclusion drawn from all 3 hyperplanes remains the same. We start by showing the computed ECG of the optimal configuration in Fig. 2 at different resolutions. Note that a spatial resolution of ≈ 1 mm is considered sufficient for most electrophysiological applications with the eikonal model, see Franzoni and Guerri (1993). However, we still have some numerical errors on the signal. This is most likely associated to the inaccuracy of capturing the very narrow wavefront of the integral (6). Such noise may be considered minor from a visual perspective, but is very likely responsible for the discontinuous loss landscapes as seen in Fig. 3.

The loss landscapes themselves also exhibit discretization artefacts especially on lower resolutions, most likely associated with the discretization of the variable initiation site, which is required to coincide with a vertex in our setup. Perhaps surprisingly, nearly all loss functions perform similarly in terms of smoothness for the tested parameter ranges and no significant advantage can be seen by applying the more sophisticated DTW instead of the straight-forward l_2 error. However, we note that in all parameter scenarios, the main peaks of the signals at least partially overlap. When trying to match non-overlapping or only slightly overlapping signals, the advantage of DTW may be more pronounced.

It is noteworthy that we also calculated the loss after applying mean filters to the ECG that visually removed most of the noise, but did only change the smoothness of the loss landscape by an insignificant amount.

4. DISCUSSION

This paper offers some insights on the impact and significance of parameters on different loss functions w.r.t. the ECG.

The complete pipeline in Sec. 2 involves a several steps, some of which involve solutions of PDEs. From the viewpoint of optimization, it is important to understand the smoothness of the loss function w.r.t. its argument, that is in our case the parameters of the forward model. When the loss function itself is non-smooth, obviously we cannot expect parametric smoothness. This is the case of the

l^1 -error or the DTW. For the other error metrics, the smoothness of the loss function will in fact rely on the parametric smoothness of the ECG $V_j(t)$. In our experiments, the most restrictive parameters in terms of regularity are the conduction velocity tensor and the initiation site. Together, they define the eikonal model for computing the activation map. It is possible to show that the activation map is Lipschitz-continuous w.r.t. the initiation site (either location or onset timing). Therefore, the gradient of the activation, parallel to the propagation direction and used in the computation of the ECG, is defined almost everywhere. As a matter of fact, in virtue of (6), the singularity set has little effect on the ECG. A similar argument applies to the conduction velocity tensor, which depends on the fiber angles and the conductivity values.

More prominent is the role of the numerical discretization. In particular, the location of the initiation sites are Dirichlet boundary conditions to the eikonal equation (2), therefore applied at discrete locations (the mesh vertices). In such a case, a projection to the closest vertex in the mesh needs to be considered, resulting in a discontinuous loss function.

Another problem is that the temporal evaluation of the ECG should follow, to some extent, the spatial grid size. First, please note that in (5) the time is a parameter. For some fixed t , v_m is generally interpolated on the computational grid (or, equivalently, evaluated at some quadrature nodes) for the computation of the ECG in (6). Therefore, when the time step is too small, the narrow function ∇v_m might be poorly approximated when barely touching any degrees-of-freedom. This translates into (numerical) oscillations on the surface ECG, as reported in Fig. 2. The threshold value for the appearance of oscillations is related to the conduction velocity, because when the wave moves fast the *saltatory* propagation disappears. In fact, the ratio between the mesh size and the time step should be comparable to the conduction velocity. This is unfortunately difficult sometimes, because the time step is a global parameter whereas mesh size and conduction velocity is local to the mesh.

The loss function itself does not appear to be very smooth in the parameter space, as deduced from the level sets (see Fig. 3). Possible sources of this noise are time discretization and spatial discretization (element size), where we suspect that (6) would especially benefit from higher order smoothness of the solution, or higher resolutions. Additionally, the Lipschitz-continuity of ϕ translates to v_m , which could be remedied or at least mitigated by an eikonal-diffusion formulation, which has also been shown to be a closer approximation to the underlying physiological processes Keener and Sneyd (1998).

Some simplifying assumptions have been made in the course of this study to make the many computations more tractable: For one, the lead field is usually not axis aligned, but rather a result of another elliptic Poisson PDE. As the leads are usually far away from the computational domain (the heart), the impact however may not be very large. Secondly, the eikonal solution is an approximation of the bidomain model, but is preferable due to its low computational demand, while still maintaining model fidelity.




Fig. 2. Plot of the computed ECGs for the optimal parameter configuration ($a = b = 0$ in (7)) at the tested different resolutions. The ECGs get successively smoother, as we reduce the element size, but at a very slow rate




Fig. 3. Loss landscapes over the whole tested parameter range at different resolutions and for multiple loss functions. Colormap ranges are individual for each contour plot.

Still, the present study provides an initial step in better understanding the complexity of fitting models to given ECGs. Such insights will be useful in future optimization and machine learning methods that rely and benefit from a smooth loss function to solve the electrophysiological inverse problem to ultimately bring us closer to cardiac digital twinning using a wide range of parameters.

ACKNOWLEDGEMENTS

We would like to thank Matthias Gsell for sharing insights on practical considerations for biventricular models, especially UVCs.

REFERENCES

Arevalo, H.J., Vadakkumpadan, F., Guallar, E., Jebb, A., Malamas, P., Wu, K.C., and Trayanova, N.A. (2016). Arrhythmia risk stratification of patients after myocardial infarction using personalized heart models. *Nature Communications*, 7(1), 11437. doi:10.1038/ncomms11437.

Bayer, J.D., Blake, R.C., Plank, G., and Trayanova, N.A. (2012). A novel rule-based algorithm for assigning myocardial fiber orientation to computational heart models. *Annals of Biomedical Engineering*, 40(10), 2243–2254. doi:10.1007/s10439-012-0593-5.

Bayer, J., Prassl, A.J., Pashaie, A., Gomez, J.F., Frontera, A., Neic, A., Plank, G., and Vigmond, E.J. (2018).

- Universal ventricular coordinates: A generic framework for describing position within the heart and transferring data. *Medical Image Analysis*, 45, 83–93. doi:10.1016/j.media.2018.01.005.
- Bellman, R.E. (2015). *Adaptive control processes: a guided tour*. Princeton university press.
- Camps, J., Lawson, B., Drovandi, C., Mincholé, A., Wang, Z.J., Grau, V., Burrage, K., and Rodriguez, B. (2021). Inference of ventricular activation properties from non-invasive electrocardiography. *Medical Image Analysis*, 73, 102143. doi:10.1016/j.media.2021.102143.
- Chen, C., Qin, C., Qiu, H., Tarroni, G., Duan, J., Bai, W., and Rueckert, D. (2020). Deep Learning for Cardiac Image Segmentation: A Review. *Frontiers in Cardiovascular Medicine*, 7, 25. doi:10.3389/fcvm.2020.00025.
- 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., Shamohamdi, 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., and Lamata, P. (2020). The “Digital Twin” to enable the vision of precision cardiology. *European Heart Journal*, 41(48), 4556–4564. doi:10.1093/eurheartj/ehaa159.
- Franzone, P.C. and Guerri, L. (1993). Spreading of excitation in 3-d models of the anisotropic cardiac tissue. I. validation of the eikonal model. *Mathematical Biosciences*, 113(2), 145–209. doi:10.1016/0025-5564(93)90001-Q.
- Fu, Z., Kirby, R., and Whitaker, R. (2013). A Fast Iterative Method for Solving the Eikonal Equation on Tetrahedral Domains. *SIAM Journal on Scientific Computing*, 35(5), C473–C494. doi:10.1137/120881956.
- Gillette, K., Gsell, M.A.F., Prassl, A.J., Karabelas, E., Reiter, U., Reiter, G., Grandits, T., Payer, C., Štern, D., Urschler, M., Bayer, J.D., Augustin, C.M., Neic, A., Pock, T., Vigmond, E.J., and Plank, G. (2021). A Framework for the generation of digital twins of cardiac electrophysiology from clinical 12-leads ECGs. *Medical Image Analysis*, 71, 102080. doi:10.1016/j.media.2021.102080.
- Grandits, T. (2021). A Fast Iterative Method Python package. *Journal of Open Source Software*, 6(66), 3641. doi:10.21105/joss.03641.
- Grandits, T., Effland, A., Pock, T., Krause, R., Plank, G., and Pezzuto, S. (2021). GEASI: Geodesic-based Earliest Activation Sites Identification in cardiac models. *International Journal for Numerical Methods in Biomedical Engineering*, 37. doi:10.1002/cnm.3505.
- Grandits, T., Gillette, K., Neic, A., Bayer, J., Vigmond, E., Pock, T., and Plank, G. (2020). An inverse Eikonal method for identifying ventricular activation sequences from epicardial activation maps. *Journal of Computational Physics*, 419, 109700. doi:10.1016/j.jcp.2020.109700.
- Hinton, G. and Roweis, S.T. (2002). Stochastic neighbor embedding. In *NIPS*, volume 15, 833–840. Citeseer.
- Hu, Y., Schneider, T., Wang, B., Zorin, D., and Panozzo, D. (2020). Fast Tetrahedral Meshing in the Wild. *ACM Trans. Graph.*, 39(4). doi:10.1145/3386569.3392385.
- Jalife, J., Delmar, M., Anumonwo, J., Berenfeld, O., and Kalifa, J. (2011). *Basic cardiac electrophysiology for the clinician*. John Wiley & Sons.
- Keener, J.P. and Sneyd, J. (1998). *Mathematical physiology*, volume 1. Springer.
- Li, H., Xu, Z., Taylor, G., Studer, C., and Goldstein, T. (2018). Visualizing the loss landscape of neural nets. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS’18*, 6391–6401. Curran Associates Inc., Red Hook, NY, USA.
- Pezzuto, S., Kal’avský, P., Potse, M., Prinzen, F.W., Auricchio, A., and Krause, R. (2017). Evaluation of a Rapid Anisotropic Model for ECG Simulation. *Frontiers in Physiology*, 8, 265. doi:10.3389/fphys.2017.00265.
- Pezzuto, S., Prinzen, F.W., Potse, M., Maffessanti, F., Regoli, F., Caputo, M.L., Conte, G., Krause, R., and Auricchio, A. (2021). Reconstruction of three-dimensional biventricular activation based on the 12-lead electrocardiogram via patient-specific modelling. *EP Europace*, 23(4), 640–647. doi:10.1093/europace/euaa330.
- Sakoe, H. and Chiba, S. (1978). Dynamic programming algorithm optimization for spoken word recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 26(1), 43–49. doi:10.1109/TASSP.1978.1163055. Conference Name: IEEE Transactions on Acoustics, Speech, and Signal Processing.
- Sundnes, J., Lines, G.T., Cai, X., Nielsen, B.F., Mardal, K.A., and Tveito, A. (2007). *Computing the electrical activity in the heart*, volume 1. Springer Science & Business Media.
- Vigmond, E.J., Weber dos Santos, R., Prassl, A.J., Deo, M., and Plank, G. (2008). Solvers for the cardiac bidomain equations. *Progress in Biophysics and Molecular Biology*, 96(1), 3–18. doi:10.1016/j.pbiomolbio.2007.07.012.
- Yang, Y., Engquist, B., Sun, J., and Hamfeldt, B.F. (2017). Application of optimal transport and the quadratic Wasserstein metric to full-waveform inversion. *GEO-PHYSICS*, 83(1), R43–R62. doi:10.1190/geo2016-0663.1. Publisher: Society of Exploration Geophysicists.