

Learning cardiac activation maps from 12-lead ECG with multi-fidelity Bayesian optimization on manifolds[★]

Simone Pezzuto^{*} Paris Perdikaris^{**}
Francisco Sahli Costabal^{***,****,†}

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

^{**} Department of Mechanical Engineering and Applied Mechanics, University of Pennsylvania, Philadelphia, Pennsylvania, USA (e-mail: pgg@seas.upenn.edu)

^{***} Department of Mechanical and Metallurgical Engineering, School of Engineering, Pontificia Universidad Católica de Chile, Santiago, Chile (e-mail: fsc@ing.puc.cl)

^{****} Institute for Biological and Medical Engineering, Schools of Engineering, Medicine and Biological Sciences, Pontificia Universidad Católica de Chile, Santiago, Chile

[†] Millennium Nucleus for Applied Control and Inverse Problems

Abstract: We propose a method for identifying an ectopic activation in the heart non-invasively. Ectopic activity in the heart can trigger deadly arrhythmias. The localization of the ectopic foci or earliest activation sites (EASs) is therefore a critical information for cardiologists in deciding the optimal treatment. In this work, we formulate the identification problem as a global optimization problem, by minimizing the mismatch between the ECG predicted by a cardiac model, when paced at a given EAS, and the observed ECG during the ectopic activity. Our cardiac model amounts at solving an anisotropic eikonal equation for cardiac activation and the forward bidomain model in the torso with the lead field approach for computing the ECG. We build a Gaussian process surrogate model of the loss function on the heart surface to perform Bayesian optimization. In this procedure, we iteratively evaluate the loss function following the lower confidence bound criterion, which combines exploring the surface with exploitation of the minimum region. We also extend this framework to incorporate multiple levels of fidelity of the model. We show that our procedure converges to the minimum only after 11.7 ± 10.4 iterations (20 independent runs) for the single-fidelity case and 3.5 ± 1.7 iterations for the multi-fidelity case. We envision that this tool could be applied in real time in a clinical setting to identify potentially dangerous EASs.

Keywords: Manifold Gaussian Process; Bayesian optimization; Eikonal model; ECG Inverse Problem; Earliest Activation Sites Identification.

1. INTRODUCTION

Modern cardiology is increasingly shifting towards personalized approaches inspired by precision medicine, see Peirlinck et al. (2021). The underpinning idea is to create a so-called digital twin of the patient’s heart, only by using non-invasively acquired data such as those from cardiac imaging and standard electrocardiography. In this way, the cardiologist of the future will have the opportunity to

investigate and optimize therapeutic interventions before their actual implementation on the patient.

The keystone problem in the creation of the digital twin is the personalization of existing cardiac models from patient-specific data. Such models are generally expensive from a computational point of view and, in order to capture the inter-patient variability, several parameters need to be optimized, see Grandits et al. (2021b). A promising approach to reduce the computational burden is to train a fast surrogate model of the high fidelity model for a selected set of parameters, e.g., as proposed by Pagani and Manzoni (2021) in the context of uncertainty quantification. Alternatively, a multi-fidelity approach is a viable and possibly preferable choice when physiology-based, low fidelity models are available. This is actually the case of cardiac electrophysiology, see e.g. Quaglino et al. (2019);

^{*} This work was funded by the Leading House for Latin American Region RPG2117 awarded to SP and FSC and the ANID – Millennium Science Initiative Program – NCN19-161 to FSC. This work was also financially supported by the Theo Rossi di Montelera Foundation, the Metis Foundation Sergio Mantegazza, the Fidinam Foundation, the Horten Foundation. We also thank the CSCS-Swiss National Supercomputing Centre for the production grant s1074 awarded to SP.

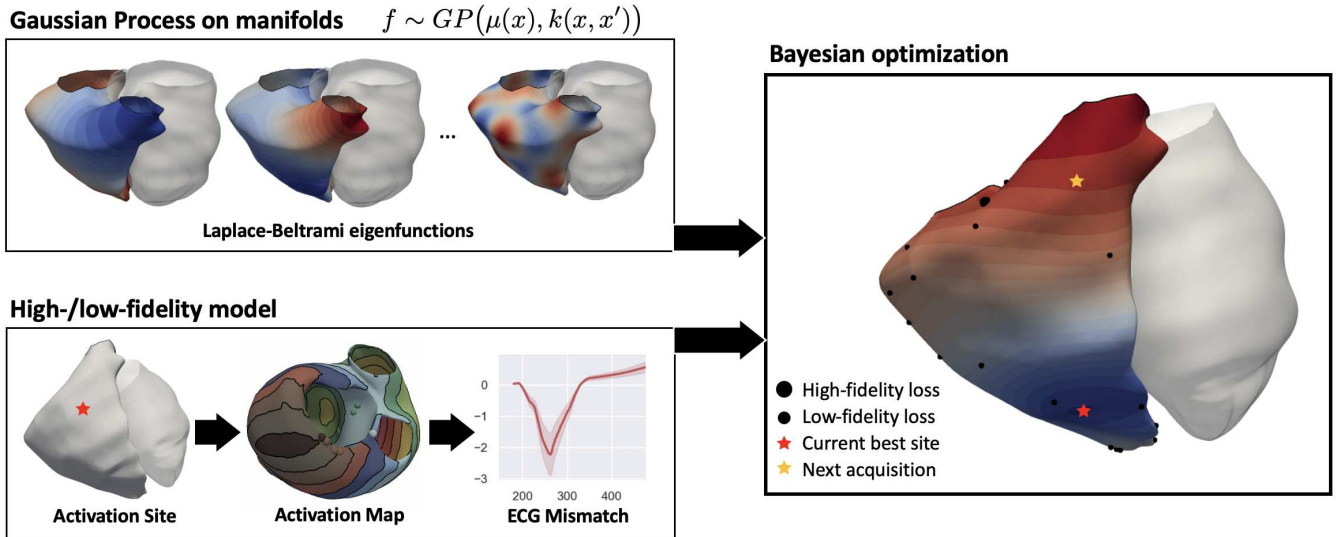


Fig. 1. Summary of the proposed method to identify an EAS from the ECG. For a tentative EAS (red star, bottom-left panel), we compute the error in ECG by simulating the activation map and the ECG with an eikonal model. The loss function is obtained from the training set through GP regression on manifolds, which is based on Laplace-Beltrami eigenfunctions (top-left panel). The loss function is updated using active learning as in Bayesian optimization (right panel). The best approximation of the best EAS is the global minimum of the loss function.

Sahli Costabal et al. (2019), since the eikonal model can effectively approximate the monodomain model in several cases and at a fraction of the computational cost of the high fidelity model, as shown by Neic et al. (2017); Pezzuto et al. (2017).

The second problem associated with digital twinning is the complexity of the parameter space. In this work we focus on the parameters determining the onset of cardiac activation, the so-called earliest activation sites (EASs). In healthy conditions, the Purkinje network is the main driver of the ventricular activation, see e.g. Sahli Costabal et al. (2016), but under pathological conditions like bundle branch block, ventricular tachycardia and ectopic activation, the EASs are much sparser and more localized, see Palamara et al. (2014).

The problem of localizing the EASs from the 12-lead electrocardiogram (ECG) has received little attention only until very recently. A common thread of all the recent literature is the use of the eikonal model at the base of a fast ECG simulator. In this context, the eikonal model is a convenient choice because EASs are simply Dirichlet boundary conditions to the forward problem, hence easy to implement. Methodologically, Camps et al. (2021) employ a stochastic method, namely the Approximate Bayesian Computation (ABC), to determine conduction velocity and EASs. Gillette et al. (2021) propose instead a combinatorial approach combined with a parametrization of the heart to “flatten” the parameter space into a hypercube. Similarly, Pezzuto et al. (2021) parametrized the endocardial surfaces in order to apply a derivative-free trust region algorithm to the problem. Finally, gradient-based methods, respectively based on geodesics computation and shape derivative, have been proposed by Grandits et al. (2021a) and Kunisch et al. (2019).

In this work, we take a Bayesian perspective to the problem of identifying a single EAS from the ECG, as for the case of ectopic activity. Bayesian optimization is a derivative-free global optimization method that only requires the model evaluation function to work. Being based on Gaussian Processes (GPs) for the approximation of the loss function, it can be adapted to work on manifolds, as proposed by Lindgren et al. (2011), and take advantage of multi-fidelity techniques, see e.g. Kennedy and O’Hagan (2000). Moreover, uncertainty quantification is naturally embedded in the method.

The proposed method, summarized in Figure 1, is explained in details in Section 2 and assessed with numerical experiments on realistic geometries, reported in Section 3. The outcome of the experiments is discussed in Section 4.

2. METHODS

2.1 Forward ECG model

In this work we employ an eikonal-based ECG model. Given an EAS located at $\mathbf{x}_0 \in \mathcal{S}$, where $\mathcal{S} \in \mathbb{R}^3$ is cardiac surface, the eikonal equation provides the activation map, denoted by $\tau(\mathbf{x}; \mathbf{x}_0)$, in the whole myocardium $\Omega \subset \mathbb{R}^3$:

$$\begin{cases} \sqrt{\mathbf{D}(\mathbf{x}) \nabla \tau(\mathbf{x}) \cdot \nabla \tau(\mathbf{x})} = 1, & \mathbf{x} \in \Omega, \\ \tau(\mathbf{x}_0) = 0, \end{cases}$$

where $\mathbf{D}(\mathbf{x})$ is the conduction velocity tensor. We assume faster conduction in the direction of myocardial fibers, that is $\mathbf{D}(\mathbf{x}) = v_t^2 \mathbf{I} + (v_\ell^2 - v_t^2) \mathbf{l} \otimes \mathbf{l}$, where v_ℓ and v_t are respectively the velocity in the local fiber direction $\mathbf{l}(\mathbf{x})$ and in the orthogonal direction. Next, the transmembrane potential is modeled as traveling wave as follows:

$$V_m(\mathbf{x}, t; \mathbf{x}_0) = U(t - \tau(\mathbf{x}; \mathbf{x}_0)).$$

The function $U(\xi)$ represents the action potential, assumed homogeneous in the domain. Since we are interested

in the activation sequence only, a reasonable approximation smooth step function as follows:

$$U(\xi) = V_0 + \frac{V_1 - V_0}{2} (\tanh(\xi) + 1),$$

where $V_0 = -80\text{mV}$ and $V_1 = 20\text{mV}$.

Finally, the ECG is a vector-valued function with components $V_k(t; \mathbf{x}_0)$, $k = 1, \dots, 12$, that reads as follows:

$$V_k(t; \mathbf{x}_0) = \int_{\Omega} \mathbf{G}_i(\mathbf{x}) \nabla V_m(\mathbf{x}, t; \mathbf{x}_0) \cdot \nabla Z_k(\mathbf{x}) d\mathbf{x},$$

where $\mathbf{G}_i(\mathbf{x}) = \sigma_{i,t} \mathbf{I} + (\sigma_{i,\ell} - \sigma_{i,t}) \mathbf{l} \otimes \mathbf{l}$ is the intra-cellular electric conductivity tensor, and $Z_k(\mathbf{x})$ the k -th lead field function. See Pezzuto et al. (2017) for further details on the implementation.

2.2 Loss function

The loss function $\mathcal{F}(\mathbf{x}_0)$ measures the mismatch between the simulated ECG $\mathbf{V}(t; \mathbf{x}_0)$ and the reference ECG $\hat{\mathbf{V}}(t)$ in the least-squares sense. It reads as follows:

$$\mathcal{F}(\mathbf{x}_0) = \sum_{k=1}^L \int_0^T (V_k(t; \mathbf{x}_0) - \hat{V}_k(t))^2 dt, \quad \mathbf{x}_0 \in \mathcal{S}.$$

The standard ECG has 12 leads, thus $L = 12$.

2.3 Manifold Gaussian Process

We construct our surrogate model of the objective function using GP regression, see Rasmussen and Williams (2006), a flexible tool for nonlinear Bayesian regression. We assume that we have a data-set of input/output pairs $\mathcal{D} = \{(\mathbf{x}_i, y_i)_{i=1}^N\} = \{\mathbf{X}, \mathbf{y}\}$. The input \mathbf{X} corresponds to N points in the manifold \mathcal{S} that represents the cardiac surface. The output $\mathbf{y} \in \mathbb{R}^N$ contains the corresponding N evaluations of the loss function \mathcal{F} at those locations. Our goal is to infer a latent function f , such that

$$\mathbf{y} = f(\mathbf{X}) + \epsilon, \quad (1)$$

where ϵ is a noise process that may corrupt our observations of \mathbf{y} . For simplicity, we assume that ϵ is Gaussian and uncorrelated, $\epsilon \sim \mathcal{N}(\mathbf{0}, \sigma_n^2 \mathbf{I})$, where σ_n^2 is an unknown variance parameter that will be learned from the data. We assume a zero-mean Gaussian process prior

$$f(\mathbf{x}) \sim \mathcal{GP}(\mathbf{0}, k(\mathbf{x}, \mathbf{x}'; \theta)). \quad (2)$$

We first identify the optimal set of kernel hyper-parameters and model parameters, $\Theta = \{\theta, \sigma_n^2\}$, and then use the optimized model to perform predictions at a set of new unobserved locations \mathbf{x}^* . The covariance kernel function $k(\mathbf{x}, \mathbf{x}'; \theta)$ that depends on a set of hyper-parameters θ and encodes any prior belief or domain expertise we may have about the underlying function f plays a central role in this process. We train our model by minimizing the negative log-marginal likelihood of the Gaussian process model, see Rasmussen and Williams (2006). Since our likelihood is Gaussian, we can evaluate it analytically in closed form,

$$\mathcal{L}(\Theta) := \frac{1}{2} \log |\mathbf{K} + \sigma_n^2 \mathbf{I}| + \frac{1}{2} \mathbf{y}^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{y} + \frac{N}{2} \log(2\pi), \quad (3)$$

where the covariance matrix $\mathbf{K} \in \mathbb{R}^{N \times N}$ follows from evaluating the kernel function $k(\cdot, \cdot; \theta)$ at the locations of the input training data \mathbf{X} . For the minimization, we

adopt a quasi-Newton optimizer L-BFGS with random restarts. Once we have trained our model on the available data, we compute the posterior predictive distribution $p(\mathbf{y}^* | \mathbf{x}^*, \mathcal{D}) \sim \mathcal{N}(\mu(\mathbf{x}^*), \Sigma(\mathbf{x}^*))$ at a new location \mathbf{x}^* by conditioning on the observed data,

$$\begin{aligned} \mu(\mathbf{x}^*) &= k(\mathbf{x}^*, \mathbf{X}) (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{y} \\ \Sigma(\mathbf{x}^*) &= k(\mathbf{x}^*, \mathbf{x}^*) - k(\mathbf{x}^*, \mathbf{X}) (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} k(\mathbf{X}, \mathbf{x}^*), \end{aligned} \quad (4)$$

where $\mu(\mathbf{x}^*)$ and $\Sigma(\mathbf{x}^*)$ denote the posterior mean and variance.

A key ingredient of the GP regression is the selection of the kernel function. A common choice is the Matérn kernel, which explicitly allows one to encode smoothness assumptions for the latent functions $f(\mathbf{x})$, see Rasmussen and Williams (2006). In a Euclidean space setting, the kernel function can be explicitly written in terms of Euclidean distance.

The form based on the Euclidean distance is not suitable to be used on manifolds, as in our case. A naïve approach is to replace the Euclidean distance between points with the geodesic distance on the manifold surface. Even though this approach may work for some cases, there is no guarantee that the resulting covariance will be positive semi-definite, see Borovitskiy et al. (2020), a key requirement for a kernel function. Here, we follow an alternative approach, implicitly based on the solution of the following stochastic partial differential equation (SPDE), see Whittle (1963); Lindgren et al. (2011):

$$\begin{cases} (\kappa^2 \mathbf{I} - \Delta)^{\alpha/2} u = \mathcal{W}, & \mathbf{x} \in \Omega, \\ \mathbf{n} \cdot \nabla (\kappa^2 \mathbf{I} - \Delta)^j u = 0, & \mathbf{x} \in \partial\Omega, j = 0, \dots, \lfloor \frac{\alpha-1}{2} \rfloor, \end{cases} \quad (5)$$

where $\alpha = \nu + d/2$ and $-\Delta$ is the Laplace-Beltrami operator on the d -dimensional manifold, and \mathcal{W} is the spatial Gaussian white noise on Ω . When $\Omega = \mathbb{R}^d$, the solution of the fractional SPDE is a Matérn random field evaluated with the Euclidean metric, see Lindgren et al. (2011). Therefore, since the SPDE in Eq. (5) trivially generalizes to manifolds with no loss of positive definiteness of the correlation kernel, thanks to the properties of the pseudo-differential operator, see Borovitskiy et al. (2020). The correlation function can be explicitly written as follows. Let $\{(\lambda_i, \psi_i)\}_{i=0}^{\infty}$ be the eigenvalue/eigenfunction pairs of the Laplace-Beltrami operator with pure Neumann boundary conditions, that is

$$\begin{cases} -\Delta \psi_i = \lambda_i \psi_i & \mathbf{x} \in \Omega, \\ -\mathbf{n} \cdot \nabla \psi_i = 0, & \mathbf{x} \in \partial\Omega, \end{cases} \quad (6)$$

for all $i \in \mathbb{N}$. Then, we can represent Matérn-like kernels on manifolds as

$$k(\mathbf{x}, \mathbf{x}'; \theta) = \frac{\eta^2}{C} \sum_{i=0}^{\infty} \left(\frac{1}{\ell^2} + \lambda_i \right)^{-\alpha} \psi_i(\mathbf{x}) \psi_i(\mathbf{x}') \quad (7)$$

where C is a normalizing constant. In practice, the eigen-decomposition is truncated to a number N_{eig} of pairs.

In this work, we discretize the manifold \mathcal{S} using a triangulated mesh and solve Eq. (6) using finite element shape functions. As such, we can obtain the stiffness matrix \mathbf{A} and mass matrix \mathbf{M} :

$$\mathbf{A}_{ij} = \mathbf{A} \int_{\mathcal{B}}^{\text{nel}} \nabla N_i \cdot \nabla N_j d\mathcal{B} \quad (8)$$

$$\mathbf{M}_{ij} = \mathbf{A} \int_{\mathcal{B}}^{\text{nel}} N_i N_j d\mathcal{B}, \quad (9)$$

where \mathbf{A} represents the assembly of the local element matrices, and N are the finite element shape functions. Then, we solve the eigenvalue problem:

$$\mathbf{A}\mathbf{v} = \lambda\mathbf{M}\mathbf{v} \quad (10)$$

In practice, to compute the kernel in Eq. (7) we use a portion of all the resulting eigenpairs, starting from the smallest eigenvalues. We also use the corresponding eigenvectors as the eigenfunctions with $f(\mathbf{x}_i) = \mathbf{v}_i$, where i is the node index at location \mathbf{x}_i . Given that the eigenvalue problem is solved only once as a pre-processing step, this methodology provides an efficient way to compute the kernel and the prior in a manifold.

2.4 Multi-fidelity framework

In this work, we will assume that we have 2 information sources of different fidelity. We will call the high fidelity, computationally expensive, and hard to acquire information source with the subscript H and the inexpensive, faster to compute, low fidelity source with the subscript L . For the problem of interest, the high and low fidelity model are as described in Sec. 2.1, but with different spatial resolution of the eikonal solver, respectively with grid size of 0.5 mm and 1 mm. Now, our data set comes from these two sources $\mathcal{D} = \left\{ (\mathbf{x}_{Li}, y_{Li})_{i=1}^{N_L}, (\mathbf{x}_{Hi}, y_{Hi})_{i=1}^{N_H} \right\} = \{\mathbf{X}, \mathbf{y}\}$. We will postulate two latent functions f_H and f_L , respectively, that are related through an auto-regressive prior, see Kennedy and O’Hagan (2000)

$$f_H(\mathbf{x}) = \rho f_L(\mathbf{x}) + \delta(\mathbf{x}). \quad (11)$$

Under this model structure, the high fidelity function is expressed as a combination of the low fidelity function scaled by ρ , corrected with another latent function $\delta(\mathbf{x})$ that explains the difference between the different levels of fidelity. Following Kennedy and O’Hagan (2000), we assume Gaussian process priors on these latent functions

$$f_L \sim \mathcal{GP}(\mathbf{0}, k(\mathbf{x}, \mathbf{x}'; \boldsymbol{\theta}_L)), \quad (12)$$

$$\delta \sim \mathcal{GP}(\mathbf{0}, k(\mathbf{x}, \mathbf{x}'; \boldsymbol{\theta}_H)). \quad (13)$$

The vectors $\boldsymbol{\theta}_L$ and $\boldsymbol{\theta}_H$ contain the kernel hyperparameters of this multi-fidelity Gaussian processes model. The choice of the auto-regressive model leads to a joint prior distribution over the latent functions that can be expressed as

$$\mathbf{f} = \begin{bmatrix} f_L \\ f_H \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \mathbf{0} \\ \mathbf{0} \end{bmatrix}, \begin{bmatrix} \mathbf{K}_{LL} & \mathbf{K}_{LH} \\ \mathbf{K}_{LH} & \mathbf{K}_{HH} \end{bmatrix} \right), \quad (14)$$

with

$$\begin{aligned} \mathbf{K}_{LL} &= k_L(\mathbf{X}_L, \mathbf{X}'_L; \boldsymbol{\theta}_L), \\ \mathbf{K}_{LH} &= \rho k_L(\mathbf{X}_L, \mathbf{X}'_H; \boldsymbol{\theta}_L), \\ \mathbf{K}_{HH} &= \rho^2 k_L(\mathbf{X}_H, \mathbf{X}'_H; \boldsymbol{\theta}_L) + k_H(\mathbf{X}_H, \mathbf{X}'_H; \boldsymbol{\theta}_H). \end{aligned} \quad (15)$$

The global covariance matrix \mathbf{K} of this multi-fidelity Gaussian process model has a block structure corresponding to the different levels of fidelity, where \mathbf{K}_{HH} and \mathbf{K}_{LL} model the spatial correlation of the data observed in

each fidelity level, and \mathbf{K}_{LH} models the cross-correlation between the two levels of fidelity. We also have kernel parameters for the different levels of fidelity. We again use the Matérn as described in Sec. 2.3, which results in parameters $\boldsymbol{\theta}_H = (\eta_H, \ell_H)$, and $\boldsymbol{\theta}_L = (\eta_L, \ell_L)$. Again, we find the optimal parameters by minimizing the negative log-marginal likelihood described in (3).

2.5 Bayesian Optimization

To minimize the mismatch between the signals predicted by our model and the observed ones we will use Gaussian process regression to perform Bayesian optimization. In this setting, we will start by computing the objective function at small number of locations ($N = 10$) in the manifold. With this data, we will train a GP regressor to predict the objective function in the entire manifold. Then, we will select the next point in the manifold to compute the objective function as:

$$\mathbf{x}^{N+1} = \arg \min_{\mathbf{x} \in \mathcal{S}} \left(\mu(\mathbf{x}) - 2\sqrt{\Sigma(\mathbf{x})} \right), \quad (16)$$

where μ and Σ are the posterior mean and variance defined in (4). This is known as the lower confidence bound criterion, see Snoek et al. (2012), and combines the exploitation of regions of high uncertainty (large variance Σ) and exploitation of regions that are potentially close to the minimum (low mean μ). We solve this optimization problem by brute force, which involves evaluating the Gaussian process in all the nodes of the mesh that represents the manifold. Once we have acquired the value of the objective function we add it to our dataset $\{\mathbf{X}, \mathbf{y}\}$ and repeat this process until \mathbf{x} matches the true location.

In the multi-fidelity setting, we employ the same approach. The only difference is that we start with $N_H = 5$ high fidelity samples, which are complemented with $N_L = 35$. (The choice is such that the multi-fidelity training cost matches the low fidelity one.) In this approach, we have a better initial approximation of the objective function of the manifold, which should speed up the optimization process. We only acquire high fidelity samples using the lower bound criterion and keep the amount of low fidelity samples fixed.

3. NUMERICAL EXPERIMENTS

We present here a comparison of the single-fidelity (SF) and multi-fidelity (MF) approach. In both cases, the ground truth \mathbf{x}^* was located in the mid-free wall of the right ventricle. The corresponding ECG $\hat{\mathbf{V}}(t)$ was obtained from the model in Sec. 2.1, paced at \mathbf{x}^* . Figure 2 shows an example of the optimization for SF alone. The total number of iterations was 6, all except the first one very close to the true minimum. The Figure also displays the mean and the standard derivation of the GP.

In the MF case, Pearson’s correlation between the HF and the LF ECG at the ground truth was 0.98, as shown in Figure 3. The computational cost for the LF model was roughly 9-fold lower than the HF model (runtime respectively equal to 0.6 s and 5.6 s.)

The MF framework converged in only 3 iterations, thus highlighting the benefit of LF samples. A specific example

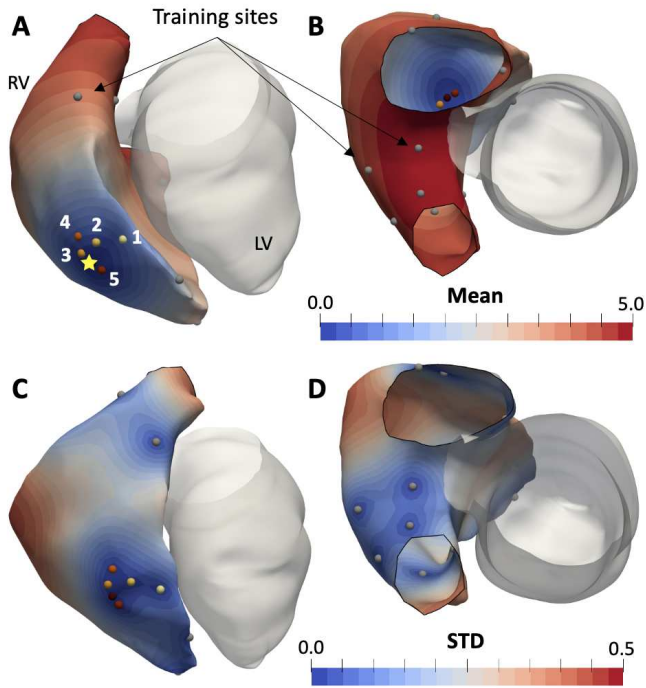


Fig. 2. Single-fidelity experiment. Top row (A-B panels): mean of the GP, with training sites indicated in gray and numbered iterations. The star indicates the true site. Bottom row (C-D, views as above): standard deviation of the GP.

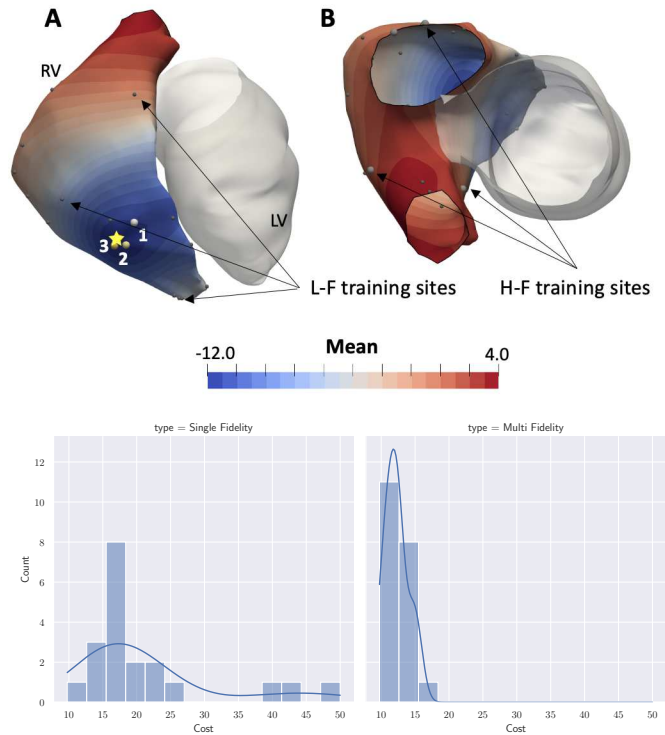


Fig. 4. Multi-fidelity experiment. Top row (A-B): two views of the mean of the GP, with iteration steps and training sites. Bottom row: randomized study on the training set ($N = 20$), showing the distribution of the cost (one unit is an evaluation of the H-F model).




Fig. 3. Simulated ECGs for the HF and LF models (8 derivations). The ECGs correspond to a stimulation from the ground truth site (mid free wall).

is reported in Figure 4. In order to assess the influence the training data, we performed an experiment with randomized training set (20 runs in total). The result is shown in Figure 4, bottom panel. The cost (one unit cost corresponds to a HF evaluation) of MF is about half of SF (median respectively 17 and 11, inter-quartile range 6 and 2), and we observed high cost runs (more than 20 iterations) only the SF case.

4. DISCUSSION

In this work, we have presented a novel methodology to identify an EAS from the ECG. We take advantage of state of the art machine learning tools—Gaussian process on manifolds—to build a surrogate model of the objective function and use Bayesian optimization to minimize it. In our results, we have shown that we need very few evaluations of our model to find the minimum in the SF case, and this is exacerbated in the MF case. In our tests, we were able to solve the minimization problem in less than a minute, which paves the way towards the translation of this tool to a clinical setting. Intriguingly, the MF framework could be applied real-time with a high fidelity model corresponding to the real patient. For instance, our framework could be adapted to solve the problem of localizing ectopic foci during endocardial mapping, a tedious process usually performed by the clinician by manually pacing tentative locations.

We are also able to take advantage of inexpensive approximations of our model by building a MF GP surrogate model. We observed that this approach robustly reduced the median computational cost of 35%, but more importantly the MF scheme reduced the variability in the cost by 3-fold.

One benefit of this method compared to existing ones, is the fact that we rely on global optimization, here possible because the dimension of the parameter space is small. In this way, the global minimum is readily available from the GP representing the loss function. The EAS

localization problem is in fact potentially non-convex and hence prone to multiple local minima. The use of a criterion that balances exploration and exploitation, such as the lower confidence bound, is key to find the global minimum in this problem more robustly.

The proof of concept presented in this work opens some new research directions. So far, we only apply the optimization criterion to the high fidelity model in the MF setting. We could develop new rules that could acquire either the low or the high fidelity model, reducing the computational cost even further. We have also tested our methodology only with synthetic data, and we would need to verify its performance with patient data. Finally, we have only considered one EAS, but in practice there could be multiple, nearly simultaneous, activations at different locations. We plan to extend this framework to identify multiple sites from the same data.

In summary, we have presented a novel methodology to identify the EASs from non-invasive signals in cardiac electrophysiology. We hope this method will enable more precise interventions to treat cardiac arrhythmias.

REFERENCES

- Borovitskiy, V., Terenin, A., Mostowsky, P., and Deisenroth, M. (2020). Matérn gaussian processes on riemannian manifolds. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, 12426–12437. Curran Associates, Inc.
- 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.
- Gillette, K., Gsell, M.A.F., Prassl, A.J., Karabelaş, E., Reiter, U., Reiter, G., Grandits, T., Payer, C., Stern, 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.
- Grandits, T., Efland, A., Pock, T., Krause, R., Plank, G., and Pezzuto, S. (2021a). GEASI: Geodesic-based Earliest Activation Sites Identification in cardiac models. *International Journal for Numerical Methods in Biomedical Engineering*, 37.
- Grandits, T., Pezzuto, S., Costabal, F.S., Perdikaris, P., Pock, T., Plank, G., and Krause, R. (2021b). Learning atrial fiber orientations and conductivity tensors from intracardiac maps using physics-informed neural networks. In *International Conference on Functional Imaging and Modeling of the Heart*, 650–658. Springer.
- Kennedy, M.C. and O’Hagan, A. (2000). Predicting the output from a complex computer code when fast approximations are available. *Biometrika*, 87(1), 1–13.
- Kunisch, K., Neic, A., Plank, G., and Trautmann, P. (2019). Inverse localization of earliest cardiac activation sites from activation maps based on the viscous Eikonal equation. *J. Math. Biol.*, 79(6-7), 2033–2068.
- Lindgren, F., Rue, H., and Lindström, J. (2011). An explicit link between gaussian fields and gaussian markov random fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(4), 423–498.
- Neic, A., Campos, F.O., Prassl, A.J., Niederer, S.A., Bishop, M.J., Vigmond, E.J., and Plank, G. (2017). Efficient computation of electrograms and ECGs in human whole heart simulations using a reaction-eikonal model. *J. Comput. Phys.*, 346, 191–211.
- Pagani, S. and Manzoni, A. (2021). Enabling forward uncertainty quantification and sensitivity analysis in cardiac electrophysiology by reduced order modeling and machine learning. *International Journal for Numerical Methods in Biomedical Engineering*, 37(6), e3450.
- Palamara, S., Vergara, C., Catanzariti, D., Faggiano, E., Pangrazzi, C., Centonze, M., Nobile, F., Maines, M., and Quarteroni, A. (2014). Computational generation of the Purkinje network driven by clinical measurements: the case of pathological propagations. *International journal for numerical methods in biomedical engineering*, 30(12), 1558–1577.
- Peirlinck, M., Costabal, F.S., Yao, J., Guccione, J., Tripathy, S., Wang, Y., Ozturk, D., Segars, P., Morrison, T., Levine, S., et al. (2021). Precision medicine in human heart modeling. *Biomechanics and modeling in mechanobiology*, 1–29.
- Pezzuto, S., Kalavský, 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.
- 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.
- Quaglino, A., Pezzuto, S., and Krause, R. (2019). High-dimensional and higher-order multifidelity monte carlo estimators. *Journal of Computational Physics*, 388, 300–315.
- Rasmussen, C.E. and Williams, C.K.I. (2006). *Gaussian Processes for Machine Learning*. MIT Press, Cambridge.
- Sahli Costabal, F., Hurtado, D., and Kuhl, E. (2016). Generating Purkinje networks in the human heart. *Journal of Biomechanics*, 49, 2455–2465.
- Sahli Costabal, F., Perdikaris, P., Kuhl, E., and Hurtado, D.E. (2019). Multi-fidelity classification using gaussian processes: Accelerating the prediction of large-scale computational models. *Computer Methods in Applied Mechanics and Engineering*, 357, 112602.
- Snoek, J., Larochelle, H., and Adams, R.P. (2012). Practical bayesian optimization of machine learning algorithms. *Advances in neural information processing systems*, 25.
- Whittle, P. (1963). Stochastic-processes in several dimensions. *Bulletin of the International Statistical Institute*, 40(2), 974–994.