BOATMAP: Bayesian Optimization Active Targeting for Monomorphic Arrhythmia Pace-mapping

Casey Meisenzahl[a], Karli Gillette[b,c], Anton J. Prassl[b], Gernot Plank[b,c], John L Sapp[d], Linwei Wang[a,*]

[a] *Rochester Institute of Technology, Rochester, New York, USA*
[b] *Gottfried Schatz Research Center, Division of Medical Physics and Biophysics, Medical University of Graz, Graz, Austria*
[c] *BioTechMed-Graz, Graz, Austria*
[d] *QEII Health Sciences Centre, Dalhousie University, Halifax, Nova Scotia, Canada*

Abstract

Recent advances in machine learning and deep learning have presented new opportunities for learning to localize the origin of ventricular activation from 12-lead electrocardiograms (ECGs), an important step in guiding ablation therapies for ventricular tachycardia. Passively learning from population data is faced with challenges due to significant variations among subjects, and building a patient-specific model raises the open question of where to select pace-mapping data for training. This work introduces BOATMAP, a novel active learning approach designed to provide clinicians with interpretable guidance that progressively assists in locating the origin of ventricular activation from 12-lead ECGs. BOATMAP inverts the input-output relationship in traditional machine learning solutions to this problem and learns the similarity between a target ECG and a paced ECG as a function of the pacing site coordinates. Using Gaussian processes (GP) as a surrogate model, BOATMAP iteratively refines the estimated similarity landscape while providing suggestions to clinicians regarding the next optimal pacing site. Furthermore, it can incorporate constraints to avoid suggesting pacing in non-viable regions such as the core of the myocardial scar. Tested in a realistic simulation environment in various heart geometries and tissue properties, BOATMAP demonstrated the ability to accurately localize the origin of activation, achieving an average localization accuracy of 3.9 ± 3.6 mm with only 8.0 ± 4.0 pacing sites. BOATMAP offers real-time interpretable guidance for accurate localization and enhancing clinical decision-making.

Keywords: Ventricular tachycardia, electrocardiogram, pace-mapping, active learning

1. Introduction

Monomorphic ventricular tachycardia (VT) is an arrhythmia in which abnormal electrical activity causes the lower chambers of the heart (*i.e.*, ventricles) to beat rapidly in a uniform pattern. Sustained VT can be a life-threatening condition, as it can lead to decreased blood flow to the rest of the body and can progress to ventricular fibrillation with a high risk of sudden cardiac death. To terminate and prevent future VT events, treatment may include catheter ablation to destroy the electrical pathway producing the abnormal pulses. A common target is the origin of ventricular activation that can be detected on the 12-lead electrocardiogram (ECG) – this may be the foci of premature ventricular contractions (PVC) or, in the case of scar-related reentrant circuits, the site at which the VT reentrant circuit exits the central isthmus protected by the myocardial scar [1].

The QRS complex is a part of an ECG that represents rapid depolarization of the right and left ventricles. Since the morphology of the QRS complex varies with the origin of ventricular activation, it plays an important role in the localization of VT origin in clinical practice [2]. Specifically, pace-mapping is a common clinical procedure that involves the iterative process of stimulating different locations on the heart (*i.e.*, pacing) and then comparing the resulting ECG with that of a previously observed VT event. When the reference VT morphology is matched on the majority of the ECG leads (*e.g.,* 12 out of 12 match), the corresponding pacing site is considered the target of VT ablation [3]. However, the current pace-mapping practice is of a "trial-and-error" nature that requires rapid qualitative interpretations of the ECG by clinicians, which can be inefficient and prone to errors. To address this gap, an increasing number of studies have investigated the use of machine learning (ML) and deep learning to localize the origin of ventricular activation using 12-lead ECGs [4].

Related Works: Existing works can be loosely divided into two categories: population-based learning [5, 6, 7, 8,

*Corresponding author: Linwei Wang, PhD, 74-1075, 102 Lomb Memorial Drive, Rochester Institute of Technology, Rochester, NY 14623, USA., linwei.wang@rit.edu, Tel: 585-475-4238, Fax: 585-475-5669.

Email addresses: casey.meisenzahl@gmail.com (Casey Meisenzahl), karli.gillette@gmail.com (Karli Gillette), anton.prassl@medunigraz.at (Anton J. Prassl), gernot.plank@medunigraz.at (Gernot Plank), john.sapp@nshealth.ca (John L Sapp), linwei.wang@rit.edu (Linwei Wang)

9, 10], and patient-specific learning [11, 4, 12, 13]. These include popular machine learning techniques such as linear regression [12] and support vector regression (SVR) [11], as well as deep learning architectures such as multilayer perceptions (MLP) [6], convolution neural networks (CNN) [14], and variational autoencoders (VAE) [15].

Population-based models utilize pre-existing pace-mapping data in a cohort of patients to predict the origin of ventricular activation from 12-lead ECGs. When there is a large amount of data, advanced deep neural networks can be used to extract hidden patterns from the QRS morphology [16], and the resulting model can be directly applied to a new patient to predict the origin of ventricular activation. However, inter-subject variability in ECGs still remains a key challenge in population-based learning. Although ECGs may originate from the same anatomical location in the heart among individuals, the QRS morphology can vary significantly due to factors such as heart and thorax geometry, as well as other patient-specific physiological and pathological conditions. [15]. As a result, the localization error of population-based learning has been limited to the order of more than 10 mm on clinical data [16]. This falls short of meeting a clinically relevant accuracy of 5 mm, the diameter of a typical single ablation lesion.

Patient-specific learning provides an attractive alternative to address the issue of inter-subject variability by incorporating pace-mapping data collected directly from the patient for whom localization is needed. These patient-specific approaches have shown the ability to obtain an average localization error close to 5 mm if the pacing sites are well-collected [11]. Unfortunately, this performance is highly dependent on the location of the pacing sites selected for training. Without knowledge of the location of the ablation targets, collecting such training sites remains an open challenge due to the need to balance multiple competing objectives: pacing sites must be carefully selected not only to pinpoint the potential target, but also to provide sufficient information about the landscape of the paced ECGs across the heart. These challenges are further increased due to the variability among individuals in their arrhythmogenic substrates. How to efficiently explore the heart to identify regions of interest while avoiding unnecessary pacing is crucial, as it reduces the risk of complications during this invasive procedure. If a large number of such training sites becomes necessary, it defeats the purpose of the ML model compared to the trial-and-error status quo.

Instead of an ML model *passively* trained on predefined data samples, a model that can *actively* query for needed data to improve training can be desirable in facilitating the clinical workflow of pace-mapping. It would provide an ML-driven tool that can suggest to clinicians where to pace and collect ECG data to progressively narrow down the ablation target. This concept was demonstrated by training an SVR model to produce the origin of ventricular activation from an input of a 12-lead ECG [11]. The SVR was trained with an initial small number of randomly selected pacing sites before entering an iterative process. In each iteration, the predicted SVR site was suggested as the next pacing site, and the obtained ECG data was added to the training data to retrain the model. This iteration continues until a termination criterion is reached, such as the match of a QRS morphology on 12/12 ECG leads from the suggested pacing site [11].

Although impressive performance has been demonstrated [11], there are two key limitations. First, pace-mapping at regions with myocardial scars can be challenging and noninformative, and is often avoided in pace-mapping procedures. However, such information cannot be incorporated into the model described in [11] to avoid nonviable pacing regions. Second, since the model guidance on where to pace is provided only as a predicted location, there is limited interpretable information to help clinicians better decide how to maneuver the heart during the pace-mapping procedure.

Contributions: To overcome these challenges, we present a novel active learning approach, BOATMAP (Bayesian Optimization Active Targeting for Monomorphic Arrhythmia Pace-mapping), to provide interpretable guidance to the clinician to progressively narrow down the ablation target while avoiding non-viable pacing regions during pace-mapping. To achieve this, we reverse the commonly used input-output relation in existing ML models for this problem and instead model the similarity $\mathcal{S}$ between a target ECG and a paced ECG as a function of the coordinate $\mathbf{r}$ of the pacing site. Using Bayesian optimization, we approximate this similarity $\mathcal{S}(QRS_{\text{tgt}}, QRS(\mathbf{r}))$, with a Gaussian process (GP) [17] based on available training data. The GP's mean and uncertainty guide the balancing of exploitation and exploration within the pacing space, informing the selection of optimal pace-mapping sites for improved training.

Compared to earlier work [11], this approach provides a theoretically rigorous formulation to use minimal training data to locate the site whose ECG would best match the target ECG in QRS morphology. More importantly, non-viable pacing regions can be simply incorporated as constraints on the input coordinate space. In addition to suggesting the next pacing site, the progressively refined GP provides an estimate of the similarity landscape between the target ECG and the pacing ECG throughout the heart: as illustrated in Fig. 1, this approximation is iteratively updated and improved as training progresses (A), along with an uncertainty map indicating the confidence of the current approximation (B). Not only do these maps provide clinicians with an explanation of the model's decision-making process (by showing why a location is predicted as a target as estimated by the GP to have the highest similarity to the target ECG, or why a new pacing site is suggested according to the acquisition criterion (C)), they also provide interpretable information that can inform clinicians to use/add their own decisions to collect pace-mapping data and with which to improve the training

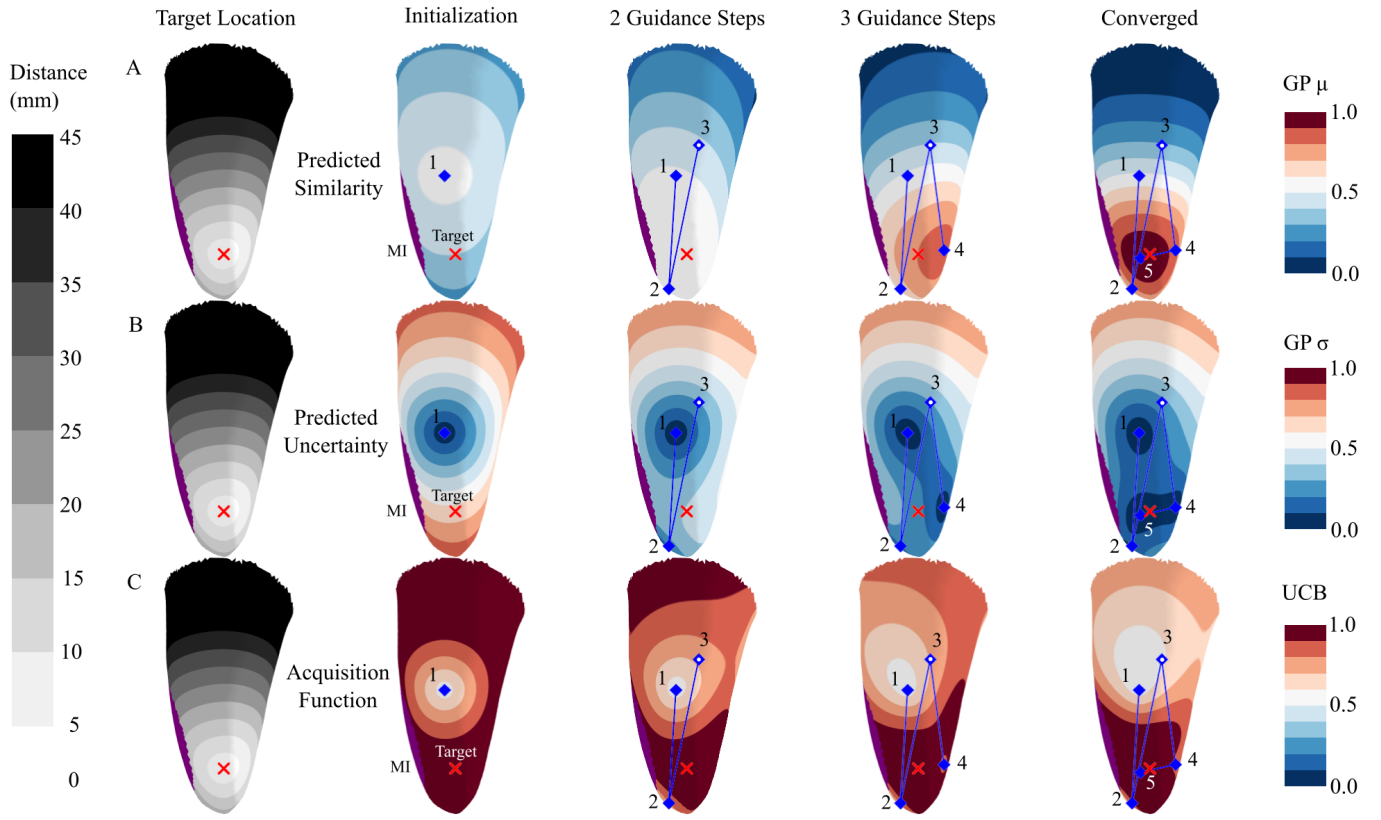

Figure 1: Overview of the BOATMAP procedure in localizing an unknown target origin of ventricular activation with 5 guided pace-mapping sites. (A): The GP mean is iteratively updated as new pace-mapping data are acquired, presenting a landscape of the expected similarity between the target QRS and QRS arising from all other potential pacing sites across the left endocardium. (B): The corresponding GP variance map indicates the uncertainty in the current estimate of the similarity map. (C): Maps of the upper confidence bound (UCB) that is used as an acquisition function to select sites to collect pacing data from, considering a balance between exploring high-uncertainty (*i.e.,* high variance) areas and exploiting most likely targets exhibiting high-similarity (*i.e.,* high mean) with the target QRS.

and prediction of the patient-specific model.

We develop BOATMAP with three alternative GP functions based on different spatial representations of the input coordinate space: a GP defined over the Cartesian 3D $x-y-z$ space, a GP defined over a 2D polar projection of the $x-y-z$ space, and a manifold-GP defined over the triangular mesh of the heart surface following [18].

Due to its active learning nature, the experimental evaluation of BOATMAP inherently requires a *prospective* setting where the heart is paced as suggested by the ML model. This presents significant challenges for *in-vivo* settings, as a result of which existing work has resorted to *retrospective* emulation of pace-mapping procedures using previously existing clinical data [11]: while a location is suggested as the next pace-mapping site, the pre-existing ECG from the nearest location is chosen as the actual location for updating the ML model. The spatial resolution of pre-existing pace-mapping data – which is limited in clinical data, thus constitutes the major bottleneck to how closely the prospective setting can be retrospectively emulated.

In this work, we leverage a rapid high-fidelity ECG simulation pipeline to provide an *in-silico quasi-prospective* testbed for BOATMAP. While an *in-silico* testbed in theory allows us to acquire ECG data wherever needed (by executing the simulation at the suggested pace-mapping site), in practice for computation and logistic ease, we pre-generate a dense grid of ECG simulations at an average resolution of 14 pace-mapping sites per square cm: as illustrated in Fig. 2, this provides a resolution not possible to achieve in either clinical or experimental settings, providing a setting closely emulating prospective active guidance.

We evaluated BOATMAP in this quasi-prospective *in-silico* environment on two different heart geometries, each considering six different configurations of tissue property (one healthy and five myocardial scars with varying sizes and locations), as illustrated in Fig. 2. In each setting, we evaluated BOATMAP for targeting 1,561 origins of ventricular activation, compared to the SVR-based active guidance model described in [11], as well as *passive learning* while varying the amount of training data and also varying the proximity of training data collected in relation to the target site.

We considered evaluation metrics that include the localization accuracy of the target site, the number of pac-

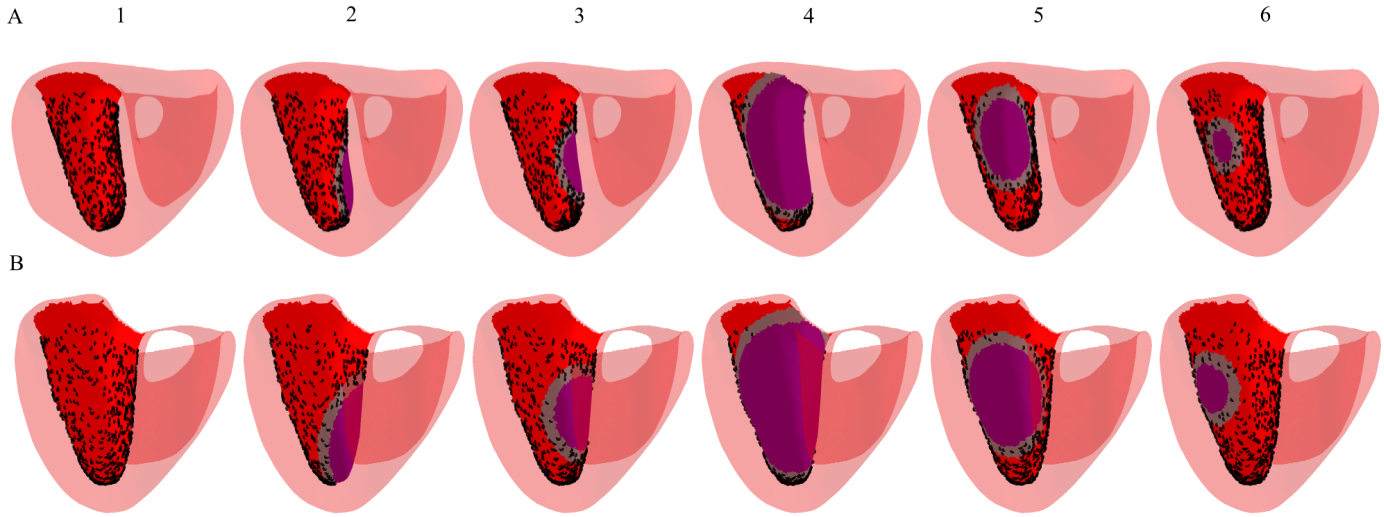

Figure 2: Two ventricular geometry (one per row) and their tissue property settings (1: healthy; 2-6: five different myocardial infarctions) used in the virtual test bed for this study. Purple represents the scar core, and gray represents the scar border. Dots indicate the origins of ventricular activation at which 12-ECGs are simulated.

ing sites needed to reach the target site, and the convergence rate described by the reduction of localization error achieved by each added pacing site. We considered a most common pace-mapping setting that considered searching the space of the left ventricular (LV) endocardium, although the presented methodology is generally applicable to the entire 3D space of the heart. Experimental results demonstrated that, while BOATMAP is comparable to active guidance based on SVR in a healthy myocardium, the latter cannot be performed in the presence of myocardial scars, whereas BOATMAP is able to locate the origin of activation – including those within the border zone – with an accuracy of 3.9 ± 3.6 mm using only 8.0 ± 4.0 pacing sites.

We have made the code for BOATMAP available, and it can be accessed at https://github.com/caseymeiz/boatmap.

2. Methodology

Localizing the origin of ventricular activation, which manifests in the QRS complex, can be reformulated into a mathematical problem of finding the spatial location $\hat{\mathbf{r}}$ where the paced QRS complex $QRS(\mathbf{r})$ will have the maximum similarity with a target QRS complex, QRS_{tgt}, as:

$$\hat{\mathbf{r}} = \arg\max\{\mathcal{S}(QRS_{\text{tgt}}, QRS(\mathbf{r}))\}, \mathbf{r} \in \Omega_h. \quad (1)$$

Here, $\mathbf{r}$ represents the spatial location within a chosen coordinate system and Ω_h denotes the domain of the heart where the search is carried out. As mentioned above, in this study, we consider Ω_h to be the endocardial surface of the heart, while the presented methodology is generally applicable to a broader definition of Ω_h (*e.g.,* the 3D volume of the myocardium).

Using this reformulation, Bayesian optimization facilitates a sequential search over Ω_h to effectively approximate $\mathcal{S}(QRS_{\text{tgt}}, QRS(\mathbf{r}))$ and locate the optimal position $\hat{\mathbf{r}}$ using a minimal number of M carefully selected pacing sites $P = \{(\mathbf{r}_i, QRS(\mathbf{r}_i))\}_i^M$.

To achieve this, starting with the initialization set $P_1 = \{(\mathbf{r}_1, QRS(\mathbf{r}_1))\}$, which is a location on the heart paired with the corresponding ECG, we build a GP surrogate to approximate $\mathcal{S}(QRS_{\text{tgt}}, QRS(\mathbf{r}))$ with the available data pairs. This process generates an approximated *landscape* that illustrates the expected similarity score and the confidence level between the target QRS complex and any other location in Ω_h. In an iterative fashion, based on this current landscape, an *acquisition function*, a, is then used to suggest the next location to collect ECG data,

$$\mathbf{r}_{m+1} = \arg\max\{a(\mathbf{r})\}, \mathbf{r} \in \omega \quad (2)$$

and the new pair of data is added to the training set $P_{m+1} = P_m \cup \{(\mathbf{r}_{m+1}, QRS(\mathbf{r}_{m+1}))\}$. Each additional training pair refines the surrogate GP model, allowing for a more informed collection of additional pace-mapping data. The search is terminated when a QRS complex obtained from a pacing site meets or exceeds a predefined similarity criterion in comparison with the target QRS complex.

Known locations of infarct tissues (or any non-viable regions not suitable for pace-mapping), Ω_{exc}, can be excluded from the search space defined by $\omega = \Omega_h - \Omega_{exc}$. This is a key advantage of BOATMAP over existing formulations that define the target ECG in the input space and the activation origin in the output space, as it is much more difficult to incorporate prior knowledge to constrain the search over the space of ECGs. By swapping the input and output variables in the function being learned, BOATMAP elegantly enables the incorporation of clini-

cal knowledge about the viable pace-mapping space in the heart.

This iterative process of BOATMAP as described above not only provides *online* guidance to the pace-mapping procedure but also a progressively refined probabilistic *landscape* that clinicians can interpret and inform their actions.

Figure 1 outlines this process in which the BOATMAP-guided pace-mapping (in selected snapshots) is shown overlaid with the landscape of the predicted similarity scores (A), along with confidence maps (B) and visuals of the acquisition functions that lead to the suggested pacing sites in the corresponding snapshots (C). In the following, we describe each key component of BOATMAP.

2.1. QRS Complex Similarity Measure

The Pearson correlation coefficient (PCC) is a strong candidate for the QRS complex similarity measure $\mathcal{S}$ in Equation (1), as it is clinically used as a criterion to determine the sites matched during pace-mapping [19]. As a statistical measure, PCC can be calculated as follows:

$$\mathcal{S}(\mathbf{g},\mathbf{h}) = \frac{\sum_{i=1}^{n}(g_i-\bar{g})(h_i-\bar{h})}{\sqrt{\sum_{i=1}^{n}(g_i-\bar{g})^2}\sqrt{\sum_{i=1}^{n}(h_i-\bar{h})^2}} \tag{3}$$

where $\mathbf{g}$ and $\mathbf{h}$ represent the amplitude vectors of the target and paced QRS complexes, respectively, each containing all 12 ECG leads. Here, g_i and h_i denote the amplitude at each time index i, while $\bar{g}$ and $\bar{h}$ denote the mean amplitude of the target ECG and paced ECG, respectively. The PCC value ranges from -1 to 1, with higher values indicating greater similarity. The use of PCC, recognized for its clinical applicability, significantly enhances the interpretability of BOATMAP in clinical settings.

2.2. Learning Similarity Functions as a Gaussian Process

BOATMAP is built on learning to approximate the objective function $\mathcal{S}(QRS_{\text{tgt}}, QRS(\mathbf{r}))$ with a GP surrogate. A GP models a set of random variables, each representing the similarity between a target QRS complex and one from a pacing site in Ω_h, such that any subset of these variables has a joint Gaussian distribution [17]. The GP is thus characterized by a mean similarity function $\mu(\mathbf{r})$, such that any location on the heart has a predicted expected similarity with the target, and a covariance function or kernel $k(\mathbf{r},\mathbf{r}')$ that determines how correlated two QRS complexes should be based on the distance between two pacing sites:

$$\mathcal{S}(QRS_{\text{tgt}}, QRS(\mathbf{r})) \sim \mathcal{GP}(\mu(\mathbf{r}), k(\mathbf{r},\mathbf{r}')). \tag{4}$$

Given a set of training pairs consisting of $R = [\mathbf{r}_1^T;\mathbf{r}_2^T;...;\mathbf{r}_m^T]$ and $\boldsymbol{s} = [s_1, s_2, ..., s_m]$, where each pacing site $\mathbf{r}_i$ has an observed similarity measure s_i such that $s_i = \mathcal{S}(QRS_{\text{tgt}}, QRS(\mathbf{r}_i))$. For simplicity, we assume a zero-mean function in our GP. This is done by finding the average similarity score across the training set and subtracting this value from each sample. In this case, we adopt the commonly used Matérn kernel function for the GP:

$$k(\mathbf{r}_i,\mathbf{r}_j) = (1+\frac{\sqrt{5}d_{ij}}{l}+\frac{5d_{ij}^2}{3l^2})\exp(-\frac{\sqrt{5}d_{ij}}{l}), \tag{5}$$

where l is the length-scale hyperparameter, and d_{ij} is the distance between to pacing sites $\mathbf{r}_i$ and $\mathbf{r}_j$. The Matérn kernel includes a smoothness parameter which, in Equation (5), has been substituted with 5/2 and reduced, simplifying the formulation [17]. This Matérn kernel describes a class of smooth functions that is twice differentiable, which makes it a good choice for approximating the similarity function between the QRS_{tgt} and other QRS complexes which is expected to gradually change based on proximity to the target pacing site.

The length-scale hyperparameter is optimized by maximizing the log marginal likelihood during the iterative active learning process. This separate optimization process adjusts the length-scale hyperparameter to optimally fit the observed data while maintaining model correctness and minimizing complexity.

To make similarity predictions $\mathcal{S}(QRS_{\text{tgt}}, QRS(\mathbf{r}_*))$ at unobserved pacing sites $\mathbf{r}_*$ we can use the predictive posterior mean $\mu(\mathbf{r}_*)$ and the predictive posterior variance $\sigma^2(\mathbf{r}_*)$:

$$\mu(\mathbf{r}_*) : \mathbf{k}^T(\mathbf{K}+\zeta^2\mathbf{I})^{-1}\mathbf{s} \tag{6}$$

$$\sigma^2(\mathbf{r}_*) : k(\mathbf{r}_*,\mathbf{r}_*) - \mathbf{k}^T(\mathbf{K}+\zeta^2\mathbf{I})^{-1}\mathbf{k} \tag{7}$$

such that $\mathbf{k} = [k(\mathbf{r}_*,\mathbf{r}_1); k(\mathbf{r}_*,\mathbf{r}_2); ...k(\mathbf{r}_*,\mathbf{r}_m)]$, $\mathbf{k}$ encodes the amount of influence the existing pacing sites should have on the unobserved locations, $\mathbf{K}$ is the covariance matrix of the observed pacing sites, $\mathbf{K}_{i,j} = k(\mathbf{r}_i,\mathbf{r}_j)$. The observation of the similarity may have noise ϵ, assuming it is normally distributed, $\epsilon = \mathcal{N}(0,\zeta^2)$

These provide probabilistic predictions on how ECGs arising from unexplored areas over the space of Ω_h may resemble the target ECG, which can help make informed decisions about where to collect training data to narrow down the target site. As illustrated in Figure 1, as ECG training data is collected from a location, the predictive uncertainty around that location, in general, is reduced. Therefore, a high level of uncertainty in the model predictions signals a potential gap in our understanding of the target function in that region, suggesting the need for further exploration to gather more data. Conversely, low uncertainty indicates that the model has a high degree of confidence in its predictions based on the existing data, suggesting that additional data points in this region may not significantly improve the model's accuracy.

We consider three different formulations of the GP based on different representations of the input domain Ω_h to capture the spatial characteristics of the heart. The primary distinction between the three formulations lies in how the kernel function k within the GP framework utilizes distance metrics between points $\mathbf{r}_i$ and $\mathbf{r}_j$ on Ω_h.

2.2.1. 3D Euclidean space

First, we consider a simple Euclidean space where Ω_h is defined by the 3D (x, y, z) Cartesian coordinates. This representation is simple and can be directly interface with the clinical pace-mapping practice, where the electroanatomic mapping system used to support pace-mapping is already equipped with abilities to define the $\mathbf{r} = (x, y, z)$ coordinates of any mapping sites.

To incorporate the 3D model as a surrogate, we need to provide the distance metric for the kernel, d_{ij} is the Euclidean distance between two pacing sites $\mathbf{r}_i = \{x_i, y_i, z_i\}$ and $\mathbf{r}_j = \{x_j, y_j, z_j\}$ as:

$$d_{ij} = \sqrt{(x_j - x_i)^2 + (y_j - y_i)^2 + (z_j - z_i)^2} \tag{8}$$

Although this formulation is straightforward, the myocardium Ω_h of interest occupies only a small portion of the 3D (x, y, z) grid. Therefore, the GP may be defined with a greater degree of freedom than is necessary.

2.2.2. 2D Euclidean space

Second, when Ω_h represents only the LV endocardial surface, we consider a 2-dimensional input space by projecting the endocardium into a polar plane using Universal Ventricular Coordinates (UVC) [20]. Each pacing site is described with 4 dimensions (c, ρ, θ, v): apicobasal, transmural, rotational, and transventricular, respectively. Specifically, we used the apicobasal coordinate c to measure the distance from the center of the polar plot and used the rotational coordinate θ to represent the angle. This can be represented in a polar form, $w = (c + \epsilon)e^{i\theta}$, $w \in \mathbb{C}$. We then place the complex number w into a 2D Cartesian plane by extracting the real and imaginary components. An ϵ is added to prevent coordinates from collapsing to $(0, 0)$. For example, two pace-mapping sites may have apicobasal distances of $c = 0$ and different rotational θ values.

In this case, we use the same Matérn kernel function as defined in Equation (5), but with the distance between pacing sites defined by the Euclidean distance in the 2D polar coordinate space.

This approach offers a condensed yet geometrically meaningful representation of Ω_h, although it is less generally applicable when Ω_h is beyond the LV endocardial surface. Furthermore, while the projection forces pacing sites that are neighbors on the 3D surface to be neighbors to each other in the 2D representation, their distances between each other are not preserved.

2.2.3. Manifold space

Lastly, we explore the use of a manifold representation of the endocardial surface to describe the non-Euclidean spatial relationships between pacing sites, incorporating local and global spatial characteristics into the model. Although the earlier two GPs eventually rely on Euclidean distances in the kernel function to calculate the covariance between pacing sites, we explore the notion that a QRS complex may have a high similarity to a pacing site that is closer along the geodesic surface of the heart as opposed to a direct 3D Euclidean distance. Replacing the Euclidean distance directly with the geodesic distance defined by the mesh, however, is not feasible because the resulting covariance matrix may not be positive-definite.

We consider the manifold defined by the vertices Ω_h, edges Ψ_h, and faces F_h of the surface mesh of the LV endocardium. Following [18, 21], we use the kernel spectral density and eigen solutions of the Cotangent Laplacian on Ω_h to express the stationary kernel in the surface manifold. Extracting the Laplacian from the mesh, we are left with a sparse weighted adjacency matrix with Neumann boundary conditions at the mesh edges. The Laplacian operator of an arbitrary scalar field f at a vertex i can be approximated as follows:

$$(\nabla^2 f)_i = \frac{1}{2A_i} \sum_{j \in N(i)} (\cot\, a_{ij} + \cot\, b_{ij})(f_i - f_j) \tag{9}$$

The Laplacian is then encoded into a matrix $\mathcal{L}$ as follows,

$$\mathcal{L}_{ij} = \begin{cases} \frac{1}{2}\sum_{j \in N(i)} \cot\, a_{ij} + \cot\, b_{ij} & \text{if } ij \in \Psi_h,\ j \in N(i), \\ -\sum_{k \in N(i)} \mathcal{L}_{ik} & \text{if } i = j, \\ 0 & \text{otherwise} \end{cases}, \tag{10}$$

where $N(i)$ defines the neighbors of a vertex i in the manifold, and the angles a_{ij} and b_{ij} represent the angles opposite to the edge ij [18]. We solve for the eigenvectors and eigenvalues of the Laplacian matrix as $\mathcal{L}\phi = \lambda \mathbf{A}\phi$, where $\mathbf{A}$ is a diagonal matrix with A_i capturing the area of a Voronoi cell centered at vertex i defined by the vertices $N(i)$ [18].

The covariance between pacing sites can be approximated by utilizing the eigenvectors and the spectral density function D from the Matérn family [21] as follows:

$$k(\mathbf{r}_i, \mathbf{r}_j) \approx \sum_{k=1}^{E} D(\sqrt{\lambda_k})\phi_k(\mathbf{r}_i)\phi_k(\mathbf{r}_j), \tag{11}$$

where E denotes the number of the smallest eigenvalues and their corresponding eigenvectors, and the spectral density function D (by the Wiener–Khinchin theorem) is:

$$D(\tau) = \frac{4\pi\Gamma(\nu + 1)(2\nu)^\nu}{\Gamma(\nu) l^{2\nu}} \left(\frac{2\nu}{l^2} + 4\pi^2\tau^2\right)^{-(\nu+1)}. \tag{12}$$

2.3. Active Guidance to Learn the Gaussian Process

Instead of passively learning a patient-specific GP from given data pairs, our goal is to minimize the number of pacing sites required to build the GP in order to accurately localize the site that can best explain the target QRS complex. Intelligently choosing the pacing sites that explore the heart surface (reducing the uncertainty of GP predictions) and simultaneously localizing the target site (maximizing similarity) is achieved by optimizing an acquisition function.

2.3.1. Initialization

To initialize the GP, a small number of pacing sites can be chosen at random while there is no prior knowledge of where the origin of activation may be, or these sites can be selected by expert knowledge depending on their hypothesis about potential target sites.

2.3.2. Acquisition Function

The acquisition function plays a crucial role in guiding BOATMAP to select the most informative next pacing locations. When QRS complexes are collected at a particular site, the mean value of the GP at that site is fitted to the measured PCC between the paced QRS complex and the target QRS complex, and the standard deviation nears 0 (*i.e.*, low uncertainty given observed data). For regions where QRS complexes have not yet been collected, often a higher standard deviation can be expected, as seen in Figure 1 B. To balance the search for high PCC regions with the exploration of unknown regions, we consider the commonly used upper confidence bound (UCB), which is defined as:

$$a(\mathbf{r}) = \mu(\mathbf{r}) + \lambda\sigma(\mathbf{r}) \tag{13}$$

where λ is a tunable parameter that controls the exploration and exploitation trade-off. A higher value of λ encourages more exploration, while a lower value promotes more exploitation.

2.3.3. Guidance Strategy

The optimal pacing site $\mathbf{r}_{m+1}$ can be obtained by finding $\mathbf{r}$ that maximizes the UCB function defined in Equation (2). This can be done by another optimization procedure on the space of $\Omega_h - \Omega_{exc}$. Alternatively, to simplify the process, we create a dense discretization of Ω_h and find the node with the maximum value of the acquisition function. To exclude non-viable pacing areas from the search space, we can simply eliminate grid points that overlap with these non-viable areas based on expert knowledge or other clinical data obtained prior to pace-mapping procedures.

2.3.4. Convergence Criteria

Each time a pacing site is gathered, it presents a chance to examine and determine whether the target QRS complex has been localized. We develop two quantitative criteria based on the PCC values. The ECG has 12 leads that each report a QRS complex representing different perspectives of the heart. We consider two possible termination criteria: 1) we ensure that all leads attain a similarity threshold t:

$$\mathcal{T}_l(\mathbf{r}, t) = \forall l \in QRS_{12}(\mathbf{r}), S(QRS_{tgt-lead}, l) \geq t \tag{14}$$

or 2) we evaluate the similarity with all leads concatenated and ensure it attains a similarity threshold of t.

$$\mathcal{T}_c(\mathbf{r}, t) = l \in QRS_{12}(\mathbf{r}), S(QRS_{tgt}, (l_1, ..., l_{12})) \geq t \tag{15}$$

Algorithm 1 Localization Algorithm

Require: QRS_{tgt} Ω_h $\Omega_{exclude}$ $init\_size$

// Restrict the search space to viable regions on the heart surface
$\omega \leftarrow \Omega_h - \Omega_{exclude}$
$R \leftarrow init\_size$ random samples from ω
$\mathbf{s} \leftarrow$ initialize empty vector to hold similarity values
for i from 1 to $|\text{rows}(R)|$ **do**
  $\mathbf{s}_{[i]} \leftarrow \mathcal{S}(QRS_{tgt}, QRS(R_{[i,:]}))$
end for
// Continue the search until a pace-mapping site satisfies the termination criterion
while $\forall \mathbf{r} \in \text{rows}(R), \mathcal{T}_c(\mathbf{r}, t)$ is False **do**
  // Estimate the mean and variance on the dense grid
  $\mu(\mathbf{r}_*) : \mathbf{k}^T(\mathbf{K} + \zeta^2\mathbf{I})^{-1}\mathbf{s}$
  $\sigma^2(\mathbf{r}_*) : k(\mathbf{r}_*, \mathbf{r}_*) - \mathbf{k}^T(\mathbf{K} + \zeta^2\mathbf{I})^{-1}\mathbf{k}$
  // Suggest the highest utility pace-mapping site
  $\mathbf{r} \leftarrow \arg\max a(\mathbf{r}_*), \mathbf{r}_* \in \omega$
  $m \leftarrow |\text{rows}(R)|$
  $R_{[m+1,:]} \leftarrow \mathbf{r}$
  // Collect and store the similarity between the target and the pace-mapping site
  $\mathbf{s}_{[m+1]} \leftarrow \mathcal{S}(QRS_{tgt}, QRS(\mathbf{r}))$
end while
return $\mathbf{r}$

As outlined in Algorithm 1, when a pacing site which was suggested by BOATMAP reaches a criterion $\mathcal{T}$, then the suggested pacing site is returned as the predicted target.

3. Experiments and Data

3.1. Virtual ECG Simulation Pipeline

As explained earlier, to allow the generation of ECG data at any pacing site suggested by BOATMAP, we leveraged a high-fidelity ECG simulation pipeline [22] as a *quasi-prospective in-silico* test environment. We considered two human bi-ventricular models made available through the Experimental Data and Geometric Analysis Repository (EDGAR) [23]. On each heart, we considered a setting of healthy myocardium, and five settings of myocardial scars of various locations of the LV and sizes ranging between 1 cm^2 and 29 cm^2, as illustrated in Fig. 2.

For computational efficiency, instead of online execution of the ECG simulation pipeline at any BOATMAP-suggested location, we precomputed a collection of 12-lead ECGs density distributed over the LV-endocardium at an average resolution of 14 pacing sites per cm^2, where the average PCC between ECGs generated from pacing sites within 5 mm is greater than 0.96 PCC. This means that, although the ECGs are precomputed offline, a selected pacing site will always be less than 14.6 mm from what is suggested by BOATMAP and, on average, 1.5 mm away.

Simulations were performed in two ventricular-torso geometries built as described in [24]. To prepare for simulation, each bi-ventricular geometry was up-sampled to a target resolution 1200 µm and then equipped with rule-based fibers [25] and UVCs [20, 22]. UVCs facilitated the insertion of a physiologically detailed His-Purkinje system representing the cardiac conduction system [26] and of five different transmural myocardial infarcts within the LV with various sizes and locations [27]. Within the infarct, the border zone was modeled as 10% of the outer edge. To facilitate the calculation of 12 lead ECGs, each bi-ventricular geometry was registered using the go-ICP algorithm [28] and meshed into a torso. The geometry of the torso had previously been created from MRI of the torso obtained for a single male patient using an automated model generation pipeline [22].

For each geometry, 12-lead ECGs were computed for 1,000 pacing sites under healthy conditions and 7,790 pacing sites in the five configurations of myocardial infarction. Simulations were carried out within the openCARP framework [29]. A *simulation grid* was created by sampling locations in the LV endocardium. The sampling was performed using UVCs implemented in *meshtool*. Under healthy conditions, the cellular dynamics within the heart was modeled using the Mitchell-Schaeffer ionic model with parameter settings of $V_{gate} = 0.13$, $V_{min} = -86.2$, $V_{max} = 40.0$, $\tau_{in} = 0.3$, $\tau_{out} = 5.4$, $\tau_{open} = 80.0$, and $\tau_{close} = 150$ [30]. Heart conductivities were assigned according to [31] and the torso conductivity was assigned a value of $0.22\,\mathrm{S\,m^{-1}}$ [32]. Conduction velocities within the healthy myocardium were assigned $0.6\,\mathrm{m\,s^{-1}}$ with an off-axis ratio of 4:2:1. A general conduction velocity of $2.0\,\mathrm{m\,s^{-1}}$ was prescribed within the His-Purkinje system, with slight deviations assigned to the fascicular branches. An anterograde and retrograde delay of 8 ms and 3 ms was assigned, respectively. Within the infarct, the core was considered electrically inactive. The border zone was assigned an isotropic conduction velocity of $0.15\,\mathrm{m\,s^{-1}}$ [33]. The Mitchell-Schaeffer model in these regions was prescribed with resting membrane parameter settings of $V_{min} = -73.1$, $V_{max} = 12.5$, $\tau_{in} = 0.45$, $\tau_{out} = 3.6$, $\tau_{open} = 44.0$, and $\tau_{close} = 100$. Simulations of transmembrane voltages were performed using a reaction-eikonal model in monodomain formulation without diffusion using CARPentry [34, 35]. 12-lead ECGs were computed from electrode potentials on the torso surface simulated using lead-field projection [36]. Lead-field matrices were calculated for each infarction configuration, as well as under healthy conditions, because of altered conductivity regions.

3.2. Emulation of BOATMAP

To simulate BOATMAP's assistance in pace-mapping, as outlined in Algorithm 1, we initiated the Gaussian process with two pacing sites and their corresponding 12-lead ECGs. The coordinates of the pacing sites were centered at the origin and normalized relative to the largest dimension in Ω_h. The regions that are not suitable for pacing, identified by $\Omega_{exclude}$, were omitted. We cropped the ECGs to isolate the QRS complex, which was then normalized to its maximum amplitude. We evaluated three BOATMAP surrogate models: 3D GP, 2D GP, and manifold GP, as detailed in Section 2.2. Each model used the Matérn kernel with a smoothness parameter of 2.5; length scales were optimized by minimizing the negative log marginal likelihood using the L-BFGS algorithm. The hyperparameter λ for the UCB acquisition function (Eq. 13) was set at 1, balancing exploration and exploitation. The 3D and 2D models used scikit-learn [37], while the GP manifold was implemented using the quLATi package [38].

Table 1: Experiment Setting Summary

Parameter	**Value**
Heart Models	2
MI Settings Per Heart Model	5
No MI Settings Per Heart Model	1
Total Targets	1,561
No MI Targets	18.1%
MI Border Zone Tissue Targets	3.3%
MI Normal Tissue Targets	78.6%

When the acquisition function suggested a pacing site, we selected the nearest corresponding site from the simulation grid along with its 12-lead ECG for updating the GP. If the nearest site in the simulation grid was more than 3 mm away from the suggested site, the search was terminated and removed from the experiments as it would bias the *prospective* evaluation where ECG data should be available whenever asked. The convergence criterion was set to be when a collected 12-lead ECG has a similarity to the target ECG greater than 0.97, with all leads concatenated together.

Emulation was performed in a variety of scenarios, localizing targets near and far from myocardial infarction (MI). In total, 1561 unique targets were included, 18.1% in healthy hearts, 78.6% in healthy tissue of hearts with MI, and 3.3% in tissue of the border zone of hearts with MI. The available data broken down by configuration are summarized in Table 1.

4. Results

4.1. Baselines

We evaluated BOATMAP against two categories of baseline models to locate the origin of ventricular activation: *passive* and *active models*. For *active* baselines, we considered the existing active-SVR model as described in [11]: it is based on an SVR model that independently predicted the x, y, and z coordinates using a radial basis function (RBF) kernel, with a regularization parameter C; the input ECG features include QRS complex time-integrals on each of the 12-lead ECGs; the prediction of the *current*

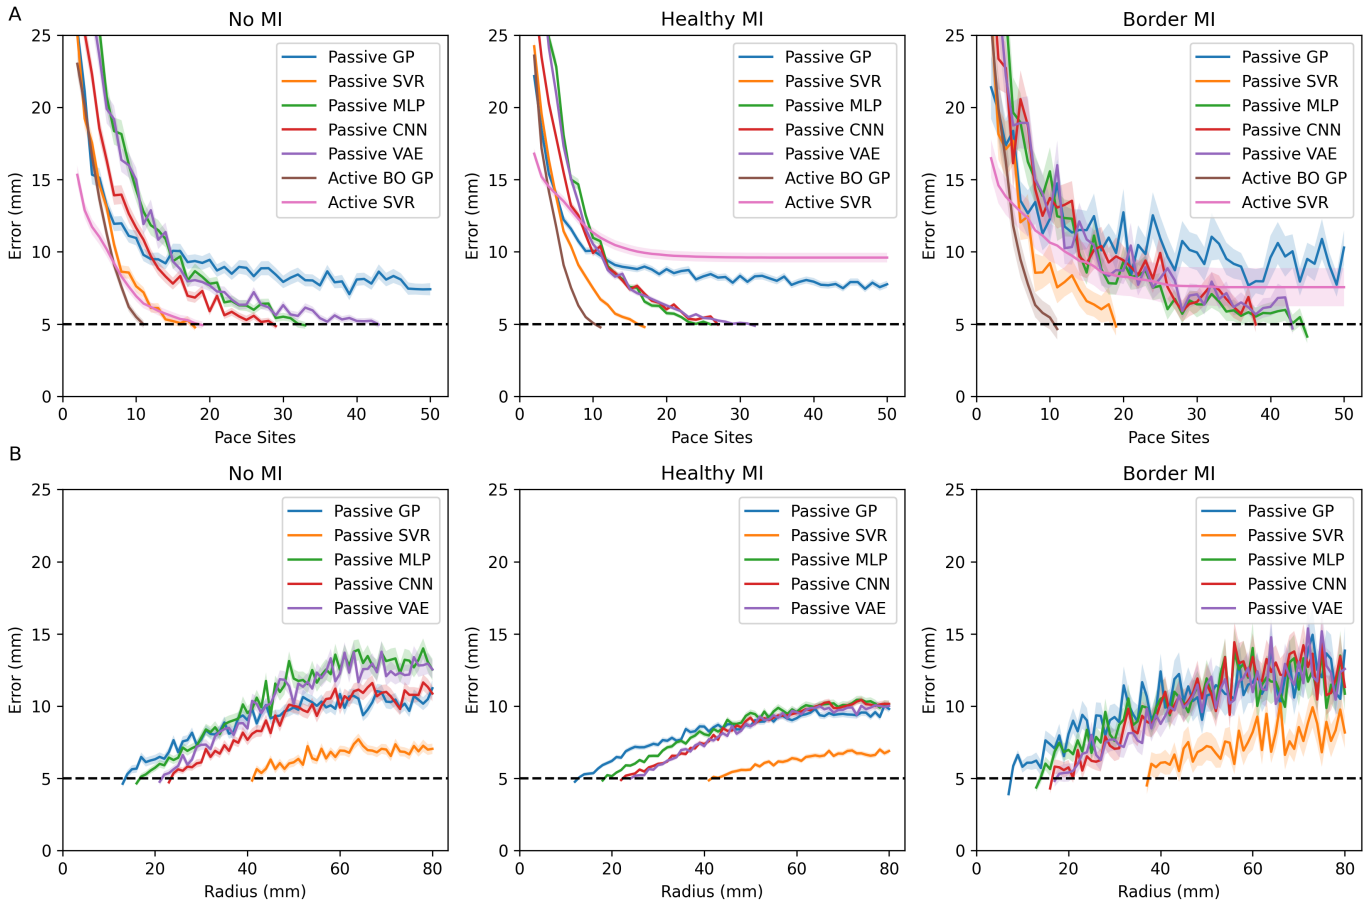

Figure 3: A: Localization error as a function of the number of training sites across various models. The columns represent different sampling settings: the first column is for targets sampled from a heart with no myocardial infarction (No MI), the second column is for targets sampled from healthy regions of a heart with myocardial infarction (Healthy MI), and the third column for targets sampled from the border zone of MI (Border MI). B: Localization error as a function of the radius size from the actual target for passive models using 11 training sites.

SVR model is used directly to suggest the next pacing site and to select training data to improve the SVR model. To evaluate the SVR model, we preprocessed the ECGs as described in [11] including down-sampling, cropping, and normalizing the amplitudes.

By *passive* baselines, we refer to machine learning or deep learning models trained on randomly selected training data. We included the following passive models as baselines:

- Passive GP Model: Serves as the non-active counterpart to the BOATMAP model. It shares the same 3D GP architecture. Fitting the model to the randomly selected pacing sites and then predicting the location with the highest expected similarity to the target ECG.
- Passive SVR Model: Is a non-active version of the active-SVR model.
- Passive MLP Model: A simple neural network baseline was constructed with three fully connected layers, each utilizing rectified linear unit (ReLU) activation functions. The model was trained using mean squared error loss and optimized with the Adam optimizer over a fixed number of epochs. This model serves as a more expressive baseline compared to the simpler SVR and GP models.
- Passive VAE Model: A combination between a VAE and a MLP classifier, such that the VAE learned generative features for recreating ECG signals driven by maximizing the evidence lower bound (ELBO) and minimizing the error in the predictions made by the MLP network, inspired by [15].
- Passive CNN Model: The architecture includes 4 convolutional layers followed by two fully connected layers, similar to [11]. This model aims to capture temporal features within the ECG signals.

4.2. Evaluation Metrics

We considered the following evaluation metrics:

Table 2: Summary of performance by BOATMAP *vs.* active-SVR methods.

Condition	Model	N	Overall		Error > 5			Error ≤ 5		
			Training Sites	Error (mm)	Training Sites	Error (mm)	Percentage(↓)	Training Sites	Error (mm)	Percentage(↑)
No MI	BOATMAP	282	**8.7** ± 4.0	**3.9** ± 3.6	**6.8** ± 2.9	**8.7** ± 3.1	**28**%	9.5 ± 4.2	**2.0** ± 1.4	**72**%
No MI	SVR	282	8.8 ± 5.6	4.5 ± 3.8	7.6 ± 4.7	8.9 ± 3.2	34%	**9.4** ± 6.0	2.2 ± 1.4	66%
Border MI	BOATMAP	52	**8.0** ± 4.5	**4.2** ± 5.0	8.9 ± 4.9	**9.2** ± 6.4	**31**%	**7.6** ± 4.4	**2.0** ± 1.7	**69**%
Border MI	SVR	52	9.5 ± 7.2	7.6 ± 9.5	**6.3** ± 4.8	16.4 ± 10.3	38%	11.5 ± 7.8	**2.0** ± 1.6	62%
Healthy MI	BOATMAP	1227	7.9 ± 4.1	**3.9** ± 3.6	6.5 ± 3.4	**8.6** ± 3.4	**27**%	**8.4** ± 4.2	**2.2** ± 1.5	**73**%
Healthy MI	SVR	1227	**7.2** ± 5.2	9.6 ± 11.7	**5.8** ± 4.4	17.6 ± 12.8	48%	**8.4** ± 5.6	2.4 ± 1.4	52%

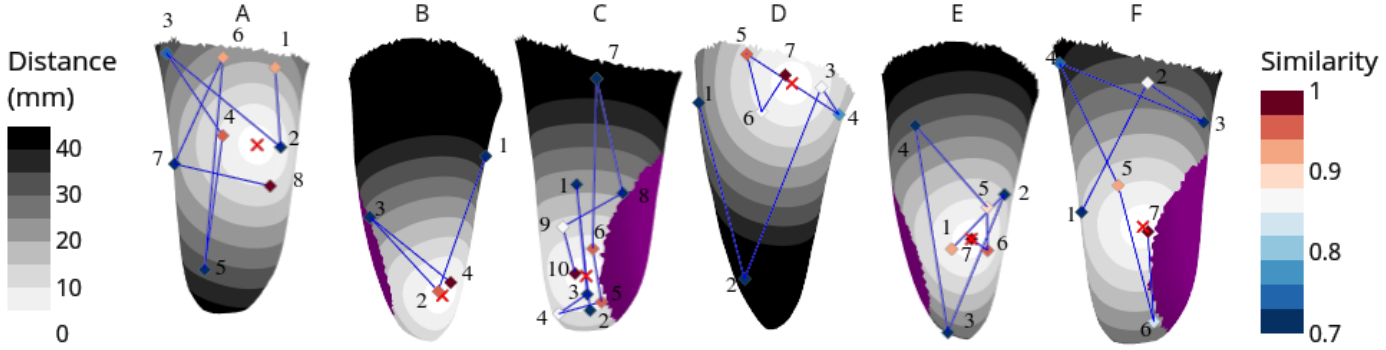

Figure 4: Examples with problematic (A-C) and successful (D-F) localization by BOATMAP. (A): The predicted target is outside the 5 mm goal with an error of 11 mm. (B): A false negative on the 2nd pacing site collected, the pacing site is within 5 mm of the target but continues because the QRS similarity did not reach the convergence threshold. (C): The search took more than the average number of pacing sites to converge. (D-F): Examples of successful localization using 7 pacing sites to converge.

1. Localization error: The Euclidean distance between the target and the prediction. A smaller distance error indicates a more accurate localization of the QRS complex's origin.
2. Pace-mapping sites: The number of pacing sites needed to reach a prediction that satisfied the termination criterion. A lower number of pacing sites indicates a more efficient search.

4.3. BOATMAP vs. passive prediction models:

Experiments in this section were carried out in a total of 1561 unique origins of ventricular activation in healthy hearts (n = 282), healthy myocardium in infarcted hearts (n = 1227), and border zone in infarcted hearts (n = 52). For passive models, we examined 1) how many randomly selected pacing sites are needed, or 2) using a similar number of pacing sites needed by BOATMAP, how close the randomly selected training sites need to be to the target (in terms of a radius to the target), in order to achieve similar accuracy as BOATMAP.

Fig. 3A shows the drop in localization error as the number of training sites increased between the two active models and the three passive models. In healthy hearts, BOATMAP was able to achieve the clinical target of a 5-mm localization error with a smaller number of training sites (11). Both active and passive SVR models need approximately double the number of training data. The passive neural networks MLP, CNN, and VAE need nearly four times the training data in order to achieve similar accuracy, while the passive GP struggles to reach the target accuracy even with 50 training sites.

This trend generally holds in infarcted hearts, with ventricular activation origins in the healthy myocardium or border zones, with two notable differences. First, the gap between BOATMAP and passive baselines in general increased, suggesting increased benefits of BOATMAP when used on infarcted hearts and especially for the origin of activation in the border zone. Second, interestingly, a larger gap between the two active models was also induced, with the active-SVR model not able to achieve the clinical target accuracy on infarcted hearts. We will delve into the possible causes in Section 4.4

Across all conditions, as shown, BOATMAP was able to consistently achieve an average target accuracy of 5-mm localization error with 11 training sites, significantly improving over its passive baselines. Fig. 3B then shows, if the passive models used only 11 training sites, how their localization accuracy will increase as a function of the radius within which the training sites are sampled from the target ECG. As shown, for passive GP and MLP, in order to achieve a localization error of 5 mm using the same number of BOATMAP training sites needed, their training sites must be selected within, respectively, 13 or 16 mm of the target on healthy hearts, and increasingly closer to the healthy myocardium of the infarcted heart (10 and 18 mm, respectively) or the border zone (7 and 6 mm, respectively). For passive CNN and VAE, the training sites need to be within roughly 20 mm of targets throughtout

Table 3: Summary of performance by BOATMAP with different GP-variants.

Condition	Model	N	Overall		Error > 5			Error ≤ 5		
			Training Sites	Error (mm)	Training Sites	Error (mm)	Percentage(↓)	Training Sites	Error (mm)	Percentage(↑)
No MI	BO 3D	1063	8.9 ± 4.3	**3.8** ± 3.6	7.4 ± 3.4	**8.6** ± 3.1	**27**%	9.5 ± 4.4	**2.0** ± 1.5	**73**%
No MI	BO 2D	490	9.4 ± 6.4	5.1 ± 4.0	8.0 ± 4.9	8.8 ± 3.3	42%	10.5 ± 7.1	2.4 ± 1.5	58%
No MI	BO Manifold	529	**8.7** ± 5.1	4.1 ± 3.6	**7.0** ± 3.9	8.7 ± 3.3	28%	**9.4** ± 5.4	2.2 ± 1.5	72%
Border MI	BO 3D	152	8.8 ± 6.0	**3.9** ± 4.5	7.9 ± 4.2	**8.9** ± 4.6	**32**%	9.2 ± 6.6	**1.6** ± 1.6	**68**%
Border MI	BO 2D	83	**7.4** ± 5.4	5.8 ± 4.9	**7.5** ± 5.7	9.9 ± 3.3	51%	**7.2** ± 5.1	1.7 ± 1.7	49%
Border MI	BO Manifold	98	8.7 ± 6.3	4.4 ± 4.8	8.9 ± 6.9	9.5 ± 4.9	34%	8.6 ± 6.1	1.8 ± 1.7	66%
Healthy MI	BO 3D	3285	7.8 ± 4.1	**4.1** ± 3.7	6.4 ± 3.6	**8.7** ± 3.4	**30**%	8.4 ± 4.2	**2.2** ± 1.5	**70**%
Healthy MI	BO 2D	1432	8.0 ± 6.0	5.2 ± 4.1	6.8 ± 4.9	8.9 ± 3.5	43%	8.9 ± 6.5	2.4 ± 1.5	57%
Healthy MI	BO Manifold	1932	**7.3** ± 4.6	4.5 ± 4.0	**6.2** ± 3.8	8.9 ± 3.5	34%	**7.9** ± 4.9	**2.2** ± 1.5	66%

all cases. For passive SVR, training sites must be selected within 41 mm of targets in the healthy myocardium and 34 mm of the border zone of infarcted hearts. Since the actual location of the target is unknown in prospective use, the need to obtain pacing sites within such distances to the target is practically not possible.

This set of experiments provided strong evidence for the advantage of BOATMAP in guiding pace-mapping toward the unknown target with a small number of pacing sites.

4.4. BOATMAP vs. alternative active models:

As Fig. 3A suggests that a larger performance gap was induced between BOATMAP and active-SVR on infarcted hearts, Table 2 provides a more detailed breakdown of the performance of the two models in the three conditions considered in Fig. 3. As shown, while the two models demonstrated comparable performance in localizing targets in healthy hearts, a significant difference was found for infarcted hearts, with BOATMAP showing better accuracy than active-SVR. This was confirmed by a paired t-test on 1279 paired samples, resulting in a significant p-value of less than 0.001. More specifically, on infarcted hearts, active-SVR was not able to achieve the 5-mm localization accuracy in a larger number of targets, increasing from 34% of targets in healthy hearts to 38% within the border zone in infarcted hearts. A closer investigation revealed that because active SVR suggests its current prediction as the next site to pace, it would have to terminate when the predicted site ended up in the infarcted region – this happened in 9% of the targets in the healthy myocardium and 22% of the targets in the border zone of the infarcted hearts. This demonstrated the key motivation for BOATMAP: the ability to exclude non-viable pace-mapping regions – available from other clinical data such as scar imaging/mapping – from the active guidance, and the importance of this ability for effective use on infarcted hearts. Fig. 4 provides several examples of successful and problematic pace-mapping guided by BOATMAP.

Overall, compared to both passive and alternative active baselines, experiments in Sections 4.3-4.4 showed that BOATMAP is able to achieve the clinical target of less than 5mm localization accuracy with a significantly smaller number of pace-mapping sites needed, and this benefit is especially strong on the hearts with scars where mechanisms are needed to steer the actively-suggested pacing sites away from non-viable pacing areas.

4.5. BOATMAP with 3D-, 2D-, and Manifold-GPs:

Table 3 compares the three alternative GP formulations described in Section 2.2, with the breakdown of metrics based on the location of the target in the infarcted heart border zone, the healthy myocardium of the infarcted heart, and the healthy heart.

The targets were randomly selected from the LV endocardium, in total 8807 searches were initiated for all three heart conditions. If the search terminated due to sparsity in the simulations (a suggested pacing site was more than 3 mm from a simulated pacing site), then those targets were dropped.

Using the 3D input domain, BOATMAP was able to consistently localize the target within a 5-mm radius, the most often compared to the 2D and manifold input domains. On average, it incurred the least error for all three heart conditions while using a similar number of pacing sites to perform the search. The 2D input domain excelled in the case where the targets resided in the border zone of the infarcted hearts, on average 7.4±5.4 training sites to localize. The manifold input domain performed well in cases where the target resided in healthy hearts, using 8.7±5.1 pacing sites.

4.6. Ablation studies

We evaluated the effect of initial conditions and convergence criteria on BOATMAP using the 3D-GP. All ablation studies were conducted on a subset of 1,500 targets, 500 each originating from the three distinct conditions: the border zone of infarcted hearts, the healthy myocardium of infarcted hearts, and healthy hearts.

BOATMAP was initialized by a set of randomly selected pacing sites to build the initial GP. Intuitively, a larger initialization would require fewer iterative acquisitions of guided pacing sites to localize the target, although at the expense of a large number of initial pacing sites. We investigated how the size of the initial random sample affected the total number of pacing sites needed to localize the origin of activation. Results in Fig. 5A showed that BOATMAP can work well with minimal initialization, requiring fewer initialization pacing sites than the SVR model to yield similar results.

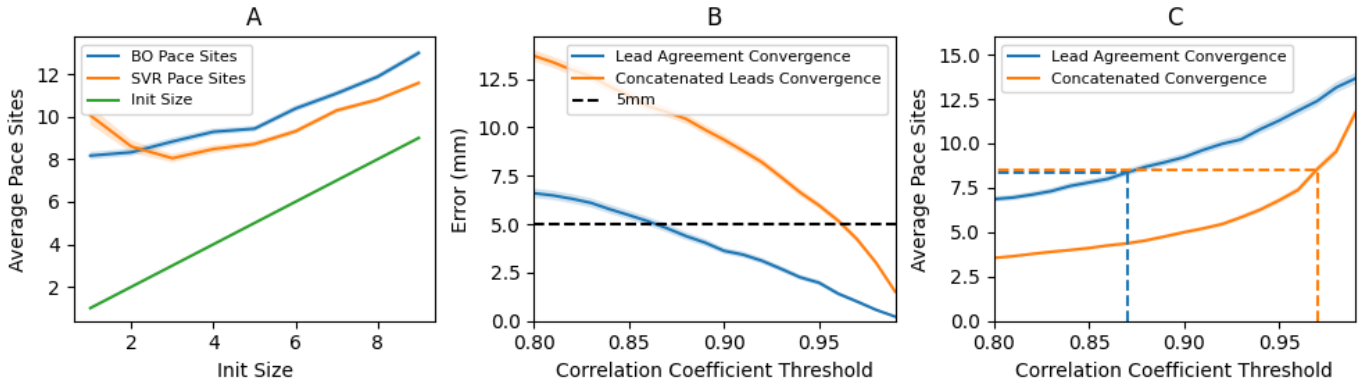

Figure 5: (A) The number of pacing sites to initialize the model (x-axis) versus the total number of pacing sites used (y-axis) to localize the target. (B) The effect of varying the correlation coefficient thresholds on the accuracy of localization. (C) The effect of varying the correlation coefficient thresholds on the total number of pacing sites needed to localize the target.

The convergence criterion indicates when the search should stop, which is determined by the similarity of the 12-lead ECGs. We tested how the change in the convergence criterion affected the final localization accuracy of BOATMAP. As shown in Fig. 5B-C, to be within 5-mm of the target using lead agreement, all leads need to reach a correlation coefficient of 0.87, which would require 8.36 total pacing sites on average. Concatenating the leads and then applying the correlation coefficient required a threshold of 0.97, which corresponds to 8.55 pacing sites used on average.

5. Discussion

5.1. Bound Constraints on GP Output

Since the GP learns the PCC function between two ECGs, we know that its output values should be bounded between -1 and 1. In the current study, no special strategies were used to incorporate this bound on the GP being learned: as observed in our experiments (Fig. 6), it is not uncommon for the learned GP to make predictions outside this range during the process of active guidance. Constraining the GP predictions into a bounded range may improve the GP fitting, alter the search, and thereby reduce the number of pacing sites needed to locate the target. Future works will explore approaches to limit the range of GP predictions, such as using splines as described in [39].

5.2. Relationship between ECG similarity and distance from target

A fundamental assumption underpinning this work (or any works that aim to infer the origin of ventricular activation from ECG) is that the similarity – as measured by PCC – between two ECGs decreases as a smooth function over the distance between the two sites of origins over the myocardium. During some exploratory analysis of our data, we found that this relationship may change along different spatial directions. Fig. 6 shows examples of the

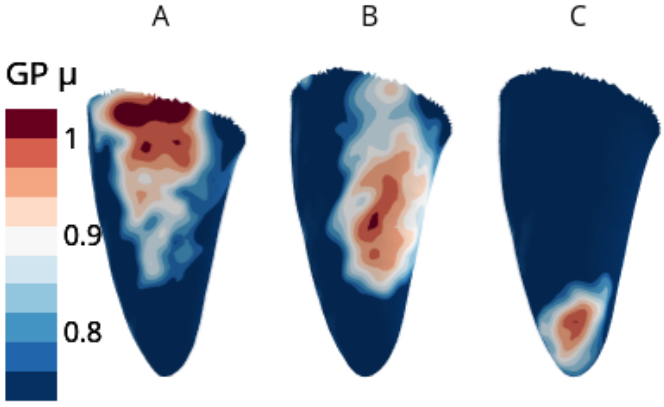

Figure 6: Examples of GP fit to all available training data to show the regions where the similarity threshold would be evaluated to be converged (the darkest two red contours). (A) Target located near the base, (B) target between base and apex, (C) target near the apex.

region of pace-mapping sites that would give rise to ECGs meeting a given PCC threshold with a target ECG: as shown, this region is sometimes elongated or irregular in shape, indicating the similarity between two ECGs may change differently along different spatial dimensions.

Furthermore, this similarity as a function of spatial distance also varies substantially with the location of the target origin producing the target QRS complex. For example, as shown in Fig. 6, in a healthy heart within the apical area of the LV, we observe a tighter region of pacing sites matching the PCC threshold to a target versus a much broader region in the base region of the LV, suggesting that the similarity between two ECGs decreases much more rapidly at the apex, whereas it decreases more slowly in the basal region.

This suggests that to meet the same localization accuracy for a target origin, different PCC thresholds may be needed for different regions of the heart: for instance, stricter criteria (*i.e.,* a higher PCC threshold) may be

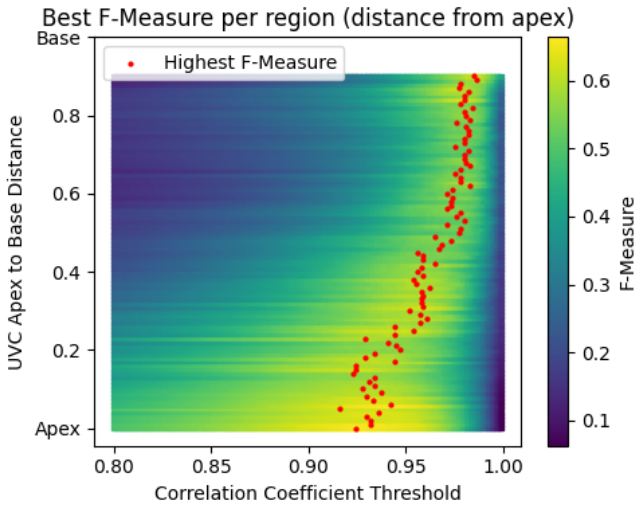

Figure 7: Heatmap visualization of the average F-Measure values obtained by varying correlation coefficient thresholds and apicobasal distances. Color intensity within each cell indicates the F-Measure value, with red markers highlighting the configurations that yield the highest F-Measure for each apicobasal distance. These optimal settings reflect the balance between precision and recall, considering true positives as those instances meeting the similarity threshold and lying within a 5 mm radius from the target site.

needed for the basal region of the heart where ECGs change less rapidly *versus* a lower PCC threshold can be used in the apical region where ECGs change more rapidly over space.

Determining a criterion in individual heart models can be done by finding the best balance between true positives and false negatives characterized by the f-measure. This is exemplified in Fig. 7. To generate these graphs, however, we used full knowledge of where the activation origin is. This method would not be viable for personalized pace-mapping, but it may be possible to learn a relationship from population-based models.

In our evaluation, we chose PCC as a similarity measurement because of its clinical use in localizing the origin of ventricular activation. It is feasible to explore alternative similarity measures, including Dynamic Time Warping (DTW). DTW, known for its ability to account for variations in time sequences, might offer a robust alternative for analyzing cases where the timing of ventricular activation varies.

5.3. Other areas of limitations and future works

5.3.1. Methodological limitations and future works

The proposed BOATMAP builds on the GP as an emulator – while equipped with advantages such as analytical and efficient computation as well as inherent uncertainty measure, the use of the GP has a disadvantage in that the similarity function being approximated has to be re-learnt from scratch each time the target ECG changes, even if we are trying to locate the targets for multiple ECGs on the same patient. An interesting future work is to investigate the use of alternative surrogates, such as a neural network that is able to generate 12-lead ECG given an input pacing site within the BOATMAP framework, such that this surrogate can be continually learned and can be used to accelerate active learning for new targets by re-using the knowledge already learned in the surrogate.

5.3.2. Experimental limitations and future works

General evaluation setup and future works:. In this study, we restricted the BOATMAP method to the left ventricular endocardium to gather initial results. Its core methodology however can be directly extended to the entire heart surface or volumetric mesh by expanding the input domain. The 2D version of BOATMAP does not extend naturally outside the LV endocardium, while the manifold variant does not directly extend outside surface meshes. The 3D BOATMAP, however, can be extended without modification to methodology. Furthermore, while this work focuses on the use of BOATMAP in pace-mapping for localizing the target of the earliest ventricular activation site, the general concept underlying BOATMAP can be extended to a wider variety of clinical applications such as ablation of the Wolff-Parkinson-White (WPW) condition and atrial fibrillation, where active learning can be used to steer the acquisition of data during the procedure and progressively improve the localization of interventional targets.

BOATMAP faces a unique challenge of *in-vivo* validation due to its nature of *active* learning. Because the method is designed to actively select and therefore influence the collection of training data, its use will directly influence the standard of care (regarding how pace-mapping will be performed to arrive at the intervention target) if used in actual clinical settings. Similarly, such prospective testing in animal models will require challenging setups to allow BOATMAP's output to steer the actual pace-mapping process. Alternatively, retrospective testing of BOATMAP on existing *in-vivo* data – as shown in our previous work [11] – faces the challenge that existing pacemapped ECG data are often too sparse to make sure a paced-ECG is always available in locations near where BOATMAP suggests the collection of training data. This creates a gap between the performance we see during retrospective evaluation and what we can expect in prospective deployment. To overcome this limitation, we move to an *in-silico* setting which provides an opportunity to ensure training data is available in the locations where BOATMAP suggests the collection of ECG data: in other words, the *in-silico* test bed provides a setting that is the closest emulation of a prospective setting where 12-lead ECG data can be collected wherever as suggested by BOATMAP. It thus serves as an important step in the evaluation of BOATMAP before we can move it to an expensive prospective evaluation on animal or clinical settings where BOATMAP can influence pace-mapping procedures.

In the particular implementation of this *in-silico*

testbed, we used a dense grid of pre-computed simulations of pace-mapping ECGs to mimic the availability of ECGs wherever BOATMAP suggests. The potential impact is that, if the distance between a BOATMAP-suggested site and the nearest site of pre-computed ECG data is large, the experiment will not reflect the actual performance of BOATMAP. In this work, we mitigated this issue by two strategies. First, we used a dense simulation grid with an average resolution of 1.45 ± 0.66 mm on which the average CC between ECGs simulated on the two closest sites is 0.99 ± 0.03, ensuring that the available ECGs capture the spatial variability of the ECGs and that there is a sufficient density of training ECGs available to be selected by BOATMAP to model such variability. Second, during BOATMAP experimentation, we simply terminate and discard experiments when the distance between a BOATMAP-suggested site and the nearest site of pre-computed ECG data is larger than 3mm. This ensures that the included experimental results closely emulate what can be expected if BOATMAP is tested in a setting where ECG data are available wherever suggested. Because of the large number of test targets available, this simple mitigation plan also has minimal effects on the sample size of our test data ($n = 1561$ after excluding these cases). A natural next step is to replace pre-computed ECGs with on-the-fly ECG simulations integrated into the BOATMAP workflow. In this setting, BOATMAP can suggest pacing sites and generate 12-lead ECGs at that location, which can be obtained by executing a computer simulation using that site as the earliest activation, emulating the collection of pace-mapped ECGs as suggested by BOATMAP in *in-vivo* settings. Once completed, we will have accumulated evidence to support the logistic investment and alleviate the safety concerns associated with *in-vivo* animal and clinical studies to use BOATMAP to prospectively guide and thus influence the pace-mapping procedure to localize interventional targets for VT. These future studies would require us to integrate 3D BOATMAP within the 3D coordinate system used in standard electroanatomical mapping systems such as CARTO (Biosense Webster) and NAVIX (EnSite NavX system by Abbott).

Specific modeling assumptions and future works:. The study utilized two heart geometries; using additional geometries originating from more patients could provide more insight into how BOATMAP performs in a more diverse setting. Additionally, the border zones were constructed to be 10% of the scar volume. However, the border zone morphology is known to vary greatly across the literature, and it contributes to ECG morphology and also potential re-entry genesis. In the future, to overcome this modeling assumption, we plan to consider realistic geometries with realistic scar morphology segmented from images such as late gadolinium-enhanced magnetic resonance imaging.

Variations in the His-Purkinje system (HPS) were not explored within the current work of this study, which may be a limitation. The largest influence on the robustness of BOATMP may be pacing near or far from the HPS that facilitates retrograde activation and potential differences in the ECG morphology over a smaller spatial resolution. In the future, additional investigations can be carried out to test BOATMAP in a model with 1) a different configuration of the HPS and 2) a fast-conducting endocardial layer but no physiological HPS.

6. Conclusion

In this work we present a robust, accurate, and interpretable method for localizing the origin of QRS complex's in the presence of scar tissue. The presented BOATMAP is able to provide real-time guidance to pace-mapping such that the target origin of ventricular activation can be localized with a small number of pace-mapping data. BOATMAP further allows the constraints on non-viable regions for pacing, such as infarct core, to be incorporated into the suggestion while delivering an interpretable similarity map that can inform the clinicians in their pace-mapping decisions.

7. Acknowledgements

This study was supported by the NIH NHLBI grant R01HL145590 (Wang) and NSF OAC-2212548 (Wang). This study was also supported by funds received from BioTechMed-Graz under the ILearnHeart Flagship Project and from the Austrian Science Fund (FWF) under grant I2760-B30 (Plank).

References

[1] B. A. Koplan, W. G. Stevenson, Ventricular tachycardia and sudden cardiac death, Mayo Clinic Proceedings 84 (3) (2009) 289–297. doi:https://doi.org/10.4065/84.3.289.

[2] K. Park, Y. Kim, F. E. Marchlinski, Using the Surface Electrocardiogram to Localize the Origin of Idiopathic Ventricular Tachycardia, Pacing and Clinical Electrophysiology 35 (12) (2012) 1516–1527. doi:10.1111/j.1540-8159.2012.03488.x.

[3] M. E. Josephson, D. J. Callans, Using the twelve-lead electrocardiogram to localize the site of origin of ventricular tachycardia, Heart Rhythm 2 (4) (2005) 443–446. doi:https://doi.org/10.1016/j.hrthm.2004.12.014.

[4] S. Zhou, A. AbdelWahab, J. L. Sapp, J. W. Warren, B. M. Horáček, Localization of ventricular activation origin from the 12-lead ecg: A comparison of linear regression with non-linear methods of machine learning, Annals of Biomedical Engineering 47 (2) (2019) 403–412. doi:10.1007/s10439-018-02168-y.

[5] M. Yokokawa, T.-Y. Liu, K. Yoshida, C. Scott, A. Hero, E. Good, F. Morady, F. Bogun, Automated analysis of the 12-lead electrocardiogram to identify the exit site of postinfarction ventricular tachycardia, Heart rhythm 9 (3) (2012) 330.

[6] P. K. Gyawali, S. Chen, H. Liu, B. M. Horacek, J. L. Sapp, L. Wang, Automatic coordinate prediction of the exit of ventricular tachycardia from 12-lead electrocardiogram, in: 2017 Computing in Cardiology (CinC), IEEE, 2017, pp. 1–4.

[7] A. Pereira, P. van Dam, R. Abächerli, Automated identification and localization of premature ventricle contractions in standard 12-lead ecgs, in: 2019 IEEE Canadian Conference of Electrical and Computer Engineering (CCECE), IEEE, 2019, pp. 1–4.

[8] T.-Y. Chang, K.-W. Chen, C.-M. Liu, S.-L. Chang, Y.-J. Lin, L.-W. Lo, Y.-F. Hu, F.-P. Chung, C.-Y. Lin, L. Kuo, S.-A. Chen, A high-precision deep learning algorithm to localize idiopathic ventricular arrhythmias, Journal of Personalized Medicine 12 (5) (2022) 764, number: 5 Publisher: Multidisciplinary Digital Publishing Institute. doi:10.3390/jpm12050764.

[9] N. Pilia, S. Schuler, M. Rees, G. Moik, D. Potyagaylo, O. Dössel, A. Loewe, Non-invasive localization of the ventricular excitation origin without patient-specific geometries using deep learning, Artificial Intelligence in Medicine 143 (2023) 102619. doi:10.1016/j.artmed.2023.102619.

[10] T. Nakamura, Y. Nagata, G. Nitta, S. Okata, M. Nagase, K. Mitsui, K. Watanabe, R. Miyazaki, M. Kaneko, S. Nagamine, N. Hara, T. Lee, T. Nozato, T. Ashikaga, M. Goya, T. Sasano, Prediction of premature ventricular complex origins using artificial intelligence–enabled algorithms, Cardiovascular Digital Health Journal 2 (1) (2021) 76–83. doi:10.1016/j.cvdhj.2020.11.006.

[11] R. Missel, P. K. Gyawali, J. V. Murkute, Z. Li, S. Zhou, A. AbdelWahab, J. Davis, J. Warren, J. L. Sapp, L. Wang, A hybrid machine learning approach to localizing the origin of ventricular tachycardia using 12-lead electrocardiograms, Computers in biology and medicine 126 (2020) 104013.

[12] J. L. Sapp, M. Bar-Tal, A. J. Howes, J. E. Toma, A. El-Damaty, J. W. Warren, P. J. MacInnis, S. Zhou, B. M. Horáček, Real-time localization of ventricular tachycardia origin from the 12-lead electrocardiogram, JACC: Clinical Electrophysiology 3 (7) (2017) 687–699.

[13] R. Doste, R. Sebastian, J. F. Gomez, D. Soto-Iglesias, A. Alcaine, L. Mont, A. Berruezo, D. Penela, O. Camara, In silico pace-mapping: prediction of left vs. right outflow tract origin in idiopathic ventricular arrhythmias with patient-specific electrophysiological simulations, EP Europace 22 (9) (2020) 1419–1430. doi:10.1093/europace/euaa102.

[14] T. Yang, L. Yu, Q. Jin, L. Wu, B. He, Localization of Origins of Premature Ventricular Contraction by Means of Convolutional Neural Network From 12-Lead ECG, IEEE Transactions on Biomedical Engineering 65 (7) (2018) 1662–1671. doi:10.1109/TBME.2017.2756869.

[15] P. K. Gyawali, J. V. Murkute, M. Toloubidokhti, X. Jiang, B. M. Horacek, J. L. Sapp, L. Wang, Learning to disentangle inter-subject anatomical variations in electrocardiographic data, IEEE Transactions on Bio-Medical Engineering 69 (2) (2022) 860–870. doi:10.1109/TBME.2021.3108164.

[16] P. K. Gyawali, B. M. Horacek, J. L. Sapp, L. Wang, Sequential factorized autoencoder for localizing the origin of ventricular activation from 12-lead electrocardiograms, IEEE Transactions on Biomedical Engineering 67 (5) (2020) 1505–1516. doi:10.1109/TBME.2019.2939138.

[17] C. E. Rasmussen, C. K. I. Williams, P. (Firm), Gaussian processes for machine learning, 1st Edition, MIT Press, Cambridge, Mass, 2006;2005;.

[18] S. Coveney, C. Corrado, C. H. Roney, D. O'Hare, S. E. Williams, M. D. O'Neill, S. A. Niederer, R. H. Clayton, J. E. Oakley, R. D. Wilkinson, Gaussian process manifold interpolation for probabilistic atrial activation maps and uncertain conduction velocity, Philosophical transactions of the Royal Society of London. Series A: Mathematical, physical, and engineering sciences 378 (2173) (2020) 20190345.

[19] G. Széplaki, T. Tahin, S. Szilágyi, I. Osztheimer, T. Bettenbuch, M. Srej, B. Merkely, L. Gellér, Ablation of premature ventricular complexes originating from the left ventricular outflow tract using a novel automated pace-mapping software, Interventional Medicine and Applied Science 2 (2010) 181–183.
URL `https://api.semanticscholar.org/CorpusID:70723919`

[20] J. Bayer, A. J. Prassl, A. Pashaei, J. F. Gomez, A. Frontera, A. Neic, G. Plank, E. J. Vigmond, Universal ventricular coordinates: A generic framework for describing position within the heart and transferring data, Medical image analysis 45 (2018) 83–93.

[21] A. Solin, S. Särkkä, Hilbert space methods for reduced-rank gaussian process regression, Statistics and Computing 30 (2) (2020) 419–446. doi:10.1007/s11222-019-09886-w.

[22] 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, G. Plank, A framework for the generation of digital twins of cardiac electrophysiology from clinical 12-leads ecgs, Medical Image Analysis 71 (2021) 102080.

[23] K. Aras, W. Good, J. Tate, B. Burton, D. Brooks, J. Coll-Font, O. Doessel, W. Schulze, D. Potyagaylo, L. Wang, P. van Dam, R. MacLeod, Experimental data and geometric analysis repository—edgar, Journal of Electrocardiology 48 (6) (2015) 975–981. doi:https://doi.org/10.1016/j.jelectrocard.2015.08.008.

[24] B. Erem, J. Coll-Font, R. M. Orellana, P. Št´ovíček, D. H. Brooks, Using transmural regularization and dynamic modeling for noninvasive cardiac potential imaging of endocardial pacing with imprecise thoracic geometry, IEEE Transactions on Medical Imaging 33 (3) (2014) 726–738. doi:10.1109/TMI.2013.2295220.

[25] J. D. Bayer, R. C. Blake, G. Plank, N. A. Trayanova, A novel rule-based algorithm for assigning myocardial fiber orientation to computational heart models, Annals of biomedical engineering 40 (2012) 2243–2254.

[26] K. Gillette, M. A. Gsell, J. Bouyssier, A. J. Prassl, A. Neic, E. J. Vigmond, G. Plank, Automated framework for the inclusion of a his–purkinje system in cardiac digital twins of ventricular electrophysiology, Annals of Biomedical Engineering 49 (12) (2021) 3143–3153.

[27] K. Gillette, M. A. F. Gsell, C. Nagel, J. Bender, B. Winkler, S. E. Williams, M. Bär, T. Schäffter, O. Dössel, G. Plank, A. Loewe, Medalcare-xl: 16,900 healthy and pathological 12 lead ecgs obtained through electrophysiological simulations, Nature: Scientific Data 10 (1) (2023) 531. doi:10.1038/s41597-023-02416-4.

[28] J. Yang, H. Li, D. Campbell, Y. Jia, Go-icp: A globally optimal solution to 3d icp point-set registration, IEEE transactions on pattern analysis and machine intelligence 38 (11) (2016) 2241–2254.

[29] G. Plank, A. Loewe, A. Neic, C. Augustin, Y.-L. Huang, M. A. Gsell, E. Karabelas, M. Nothstein, A. J. Prassl, J. Sánchez, et al., The openCARP simulation environment for cardiac electrophysiology, Computer Methods and Programs in Biomedicine 208 (2021) 106223. doi:https://doi.org/10.1016/j.cmpb.2021.106223.

[30] C. C. Mitchell, D. G. Schaeffer, A two-current model for the dynamics of cardiac membrane, Bulletin of mathematical biology 65 (5) (2003) 767–793.

[31] D. E. Roberts, A. M. Scher, Effect of tissue anisotropy on extracellular potential fields in canine myocardium in situ., Circulation Research 50 (3) (1982) 342–351.

[32] D. U. Keller, F. M. Weber, G. Seemann, O. Dössel, Ranking the influence of tissue conductivities on forward-calculated ecgs, IEEE Transactions on Biomedical Engineering 57 (7) (2010) 1568–1576.

[33] C. Mendonca Costa, G. Plank, C. A. Rinaldi, S. A. Niederer, M. J. Bishop, Modeling the electrophysiological properties of the infarct border zone, Frontiers in Physiology 9 (2018) 356.

[34] A. Neic, F. O. Campos, A. J. Prassl, S. A. Niederer, M. J. Bishop, E. J. Vigmond, G. Plank, Efficient computation of electrograms and ecgs in human whole heart simulations using a reaction-eikonal model, Journal of computational physics 346 (2017) 191–211.

[35] E. Vigmond, R. W. Dos Santos, A. Prassl, M. Deo, G. Plank, Solvers for the cardiac bidomain equations, Progress in biophysics and molecular biology 96 (1-3) (2008) 3–18.

[36] M. Potse, Scalable and accurate ecg simulation for reaction-diffusion models of the human heart, Frontiers in physiology 9 (2018) 370.

[37] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau,

M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: Machine learning in Python, Journal of Machine Learning Research 12 (2011) 2825–2830.

[38] S. Coveney, samcoveney/qulati: v2.0 (Jun. 2021). doi:10.5281/zenodo.5035230.

[39] L. P. Swiler, M. Gulian, A. L. Frankel, C. Safta, J. D. Jakeman, A survey of constrained gaussian process regression: Approaches and implementation challenges, Journal of Machine Learning for Modeling and Computing 1 (2) (2020).