AI-Assisted Pace-Mapping using Continual-Learning Methods in Bayesian Optimization

Dylan O'Hara[1], Pradeep Bajracharya[1], Casey Meisenzahl[1], Karli Gillette[2], Anton J. Prassl[2], Gernot Plank[2], John L Sapp[3], Linwei Wang[1]

[1] Rochester Institute of Technology, Rochester, NY, USA
[2] Gottfried Schatz Research Center, Division of Medical Physics and Biophysics, Medical University of Graz, Graz, Austria
[3] QEII Health Sciences Centre, Dalhousie University, Halifax, Canada

Abstract

Ventricular tachycardia (VT) is a leading cause of sudden cardiac death, and effective treatment often relies on catheter ablation guided by pace-mapping. Traditional pace-mapping is a labor-intensive process requiring expert interpretation of ECGs. Recent methods using Gaussian process (GP)-based Bayesian optimization (BO) have improved efficiency by reducing the number of stimulation sites needed for localization, but they fail to transfer knowledge across different VT targets, requiring retraining for each new case. This study introduces a novel BO framework that integrates ensemble neural networks (ENN) with continual learning (CL) strategies to enable knowledge transfer across tasks. Allowing increased efficiency without the need for initial data and memory across tasks. Evaluated on one healthy and one infarcted setting of a heart geometry, our proposed method demonstrated an average 90% *reduction in required pace-mapping sites compared to unguided approaches and a* 65% *reduction relative to GP-based BO.*

1. Introduction

Ventricular tachycardia (VT) is a life-threatening cardiac arrhythmia and a major cause of sudden cardiac death [1,2]. A common and effective treatment for VT is catheter ablation, wherein accurate localization of the abnormal heart activity is critical. This localization is often achieved through pace-mapping, a process that involves stimulating various sites on the heart and comparing the resulting ECG signals to those recorded during clinical VT episodes [3]. Traditional pace-mapping methods rely heavily on expert interpretation and involve a time-consuming trial-and-error approach to identify the optimal ablation site.

To improve this process, recent strategies have employed deep learning and other machine learning techniques [4–6] designed to predict the origin of VT signals. The first class of these methods use population-based learning, in which models are constructed using large and diverse datasets of patient ECGs, which are then applied to new patients. However, this requires acquiring a robust dataset that properly captures the distribution of heart geometries and ECG responses, which is difficult, and new patients whose heart geometries lie outside of the training distribution may see poor performance from these models.

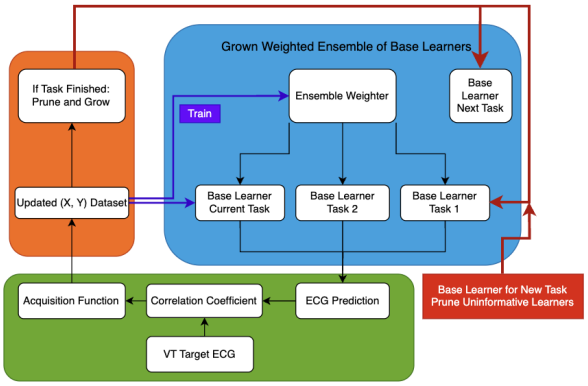

Figure 1. Workflow of the Proposed Method

The second class of methods uses patient-specific learning, where Bayesian Optimization (BO) frameworks guide the selection of stimulation sites more efficiently. Patient specific models, such as the Gaussian Process (GP) [7], can capture the relationship between the pacing site location and target ECG, while also providing a measure of uncertainty in the relationship. The uncertainty estimate is a crucial component of the BO process and allows the algorithm to explore and exploit the search space effectively.

However, there is a key limitation of this method; the GP models the relationship between a pacing-site location and the similarity between that paced ECG and the target ECG.

Therefore, the model cannot transfer information between tasks (as the target ECGs would be different), and must be retrained from scratch for each new target, even within the same patient heart, limiting scalability and clinical utility.

In this work, we present an improvement on existing strategies by incorporating neural networks trained to predict full 12-lead ECG responses at each stimulation site. This allows the model to utilize signal-level information rather than relying solely on spatial coordinates. Furthermore, we introduce a continual learning (CL) framework that enables the model to transfer knowledge across sequential tasks without requiring direct transfer of ECG data. This continual learning approach allows the model to selectively incorporate relevant prior knowledge, significantly improving efficiency across multiple tasks.

We evaluate our proposed method using simulated ECG data generated via a realistic biophysical modeling framework. Performance is assessed using several key metrics, including the number of pace-mapping sites required for localization, localization error, and convergence rate. We benchmark our approach against a state-of-the-art GP-based BO method, analyzing performance within one healthy and one infarcted setting of a heart geometry. Our results demonstrate the effectiveness and generalizability of our approach, highlighting its potential for improving pace-mapping efficiency and clinical applicability.

2. Methods

In this section we define the main components of our study: our base surrogate model, the BO algorithm including the acquisition function, similarity scoring, and convergence criteria, and the continual learning algorithm.

2.1. Base Surrogate Learner

The base surrogate learner is an ensemble of neural networks (ENN), consisting of three independently trained neural networks (MLPs). The models are architecturally identical, using three layers of size 128 followed by ReLU activations, and is optimized using the Adam optimizer. Each neural network models the relationship between pacing site location $X_i = [q_i, r_i, s_i]$ and response $Y_i = [y_{i_1}, y_{i_2}, ..., y_{i_t}]$ for the ith site-ECG pair with t representing the length of the ECG signal.

Given the ENN (N), each MLP ($n \in N$), and a coordinate X_i,

$$N_\mu(X_i) = \frac{1}{3}\sum_{j=1}^{3} n_j(X_i) \tag{1}$$

$$N_\sigma(X_i) = \sqrt{\frac{1}{3}\sum_{i=1}^{3}(n_j(X_i) - N_\mu(X_i))^2} \tag{2}$$

2.2. Bayesian Optimization

2.2.1. Correlation Coefficient and Uncertainty

The acquisition function in BO requires that we generate an objective criteria, including $\mu(x)$ and $\sigma(x)$, to optimize for subsequent suggestions. We use a metric called the correlation coefficient (CC). After obtaining the prediction and uncertainty for the ECG at a given pacing site, we can compute the similarity and uncertainty of similarity between the predicted ECG and the target ECG. They are computed as follows,

$$\begin{aligned} \mu(X_i) &= CC(Y_i, N(X_i)) \\ CC_{upper} &= CC(Y_i, N(X_i) + N_{std}(X_i)) \\ CC_{lower} &= CC(Y_i, N(X_i) + N_{std}(X_i)) \\ \sigma(X_i) &= CC_{upper} - CC_{lower} \end{aligned} \tag{3}$$

2.2.2. Acquisition Function

With the $\mu(x)$ and $\sigma(x)$ functions defined in Section 2.2.1, we can use the Expected Improvement (EI) acquisition function to balance exploration and exploitation in selecting the next stimulation site.

$$\begin{aligned} \mathrm{EI}(\mathbf{x}) &= (\mu(\mathbf{x}) - \mathbf{f}^+)\Phi\left(\tfrac{\mu(\mathbf{x})-\mathbf{f}^+}{\sigma(\mathbf{x})}\right) + \\ &\quad \sigma(\mathbf{x})\phi\left(\tfrac{\mu(\mathbf{x})-\mathbf{f}^+}{\sigma(\mathbf{x})}\right), \end{aligned} \tag{4}$$

where, f^+ is the best observed CC so far, and Φ and ϕ are the CDF and PDF of the standard normal distribution, respectively.

The function is then evaluated at each pacing site in our ECG database (excluding previously paced sites) and the location with the highest expected improvement is chosen as the next stimulation site to sample.

2.2.3. Convergence Criteria

The optimization process terminates when a response ECG with at least 97% CC to the target ECG is identified, this means that on average the paced site is within a distance of $5mm$ from the target; a threshold deemed sufficient for clinical localization accuracy,

2.3. Continual Learning

We propose a novel Grown Weighted Ensemble of Base Learners (GWEBL) framework that mitigates this issue by dynamically assembling an ensemble of base learners (ENNs) for each task. The method balances plasticity and stability by using a learned weighting module, or

"weighter", that selects and combines experts based on both recent task performance and historical utility.

2.3.1. Weighting the Ensemble

A separate neural network is used to weight both the current and prior learners based on their performance in the current task. The weighter is implemented using a bottleneck architecture consisting of three layers of size 64 followed by one of size 32, then 16. Each layer is followed by ReLU activations, except for the last which uses SoftMax to turn the outputs into normalized probabilities.

The weighter also starts with a significantly disruptive dropout rate of $p = 0.5$ at the final weighting step, forcing it to consider many models as possible sources of prior information. As the task progresses the dropout rate decreases by 0.05 each round, allowing the model to give more weight to the best performing model.

The prediction and uncertainty of the GWEBL (G) at any given time is the weighted (W) average of each base learner (B), in the GWEBL ensemble (E).

$$G(x) = \sum_{j=1}^{|E|} W(x)_j \cdot B_j(x) \tag{5}$$

Using the dataset of the current task only, backpropagation is done to update the parameters of the weighter W and the newest base learner $B_{|E|}$, all other $B \in E$ remain unchanged.

2.3.2. Growing and Pruning the Ensemble

When a new task is started, an additional learner is appended to E with its parameters randomly initialized, and its weight $W_{|E|}$ is 0.

If the new task begins and the size of E is equal to 10, then there will be a pruning stage to remove one of B_j in favor of a newly initialized B. To do this, we keep a record of the ensemble weights at the convergent round for each task. Then the historical weight of B_j is the average W_j over all previous tasks, and the B_j with the lowest historical weight (where $j \leq 7$) is removed from E.

$$\bar{W}_j^h = \frac{1}{|W_j^h|} \sum_{i=1}^{7} (W_j^h)_i \tag{6}$$

W_j^h is the weight history of B_j

We put this restriction on j to allows newer models to proliferate in a number of tasks before they can be decided to be of low weight.

2.3.3. Integration with BO

For the first task in BO, there will only be one B so we initialize the first task using 2 randomly selected pacing-site locations. For each subsequent task, the first round of acquisition will use W of the ultimate round in the previous task (after the ensemble has been pruned, if applicable). The full algorithm is shown,

Algorithm 1 GWEBL in BO

Require: Dataset (X, Y), ensemble E, weighter W, weight history W^h, base learner B.
for t in tasks **do**
$E, W^h, W \leftarrow \text{GrowPrune}(W_j, E, B)$
repeat
for e in epochs **do**
$\theta_W \leftarrow \theta_W - \nabla_{\theta_W} \text{Loss}(X, Y)$
$\theta_{B_{|E|}} \leftarrow \theta_{B_{|E|}} - \nabla_{\theta_{B_{|E|}}} \text{Loss}(X, Y) \cdot W_{|E|}$
end for
$(X_{\text{new}}, Y_{\text{new}}) \leftarrow \text{Acquisition}(W, E)$
until $\text{CC}(y_{\text{target}}, Y_{\text{new}}) > 0.97$
end for

3. Experiment

3.1. Data

We consider a single human biventricular model obtained from the Experimental Data and Geometric Analysis Repository (EDGAR) [8]. The model includes a healthy (sinus) condition and an infarcted condition. For each condition, 12-lead ECGs were calculated for 924 pacing sites under healthy conditions and 687 pacing sites under infarcted condition. ECG signals are processes such that each lead signal is of length 170.

3.2. Baselines

We compare our results against three baselines namely: *Random*, *GP-based Bayesian Optimization (GP)*, and *ENN-based Bayesian Optimization (ENN)*, i.e the base learner used in our continual-learning method. In the *Random*, candidate pacing sites are randomly selected from the coordinate search space without considering model uncertainty. In contrast, the *GP* leverages the uncertainty of a task-specific surrogate model, modeled for individual target ECG for each patient, to guide the selection of the next candidate site. The *ENN* baseline works similarly to *GP* but predicts the entire ECG signal rather than just the CC. Importantly, none of these baselines incorporates knowledge transfer, as a new surrogate model is created independently for each task.

Sinus Condition

Methods	Task 1	Task 2	Task 3	Task 4	Task 5	Task 6
Random	53.9 +/- 38.0	23.6 +/- 21.0	85.3 +/- 27.7	75.7 +/- 32.6	69.3 +/- 36.8	76.5 +/- 34.9
GP	16.7 +/- 6.0	9.6 +/- 4.6	14.6 +/- 6.1	16.5 +/- 7.2	16.2 +/- 6.8	11.2 +/- 4.3
ENN	6.3 +/- 1.4	6.9 +/- 1.6	12.8 +/- 3.2	14.6 +/- 5.0	6.9 +/- 1.4	7.5 +/- 1.0
GWEBL	6.4 +/- 2.3	2.94 +/- 1.5	10.2 +/- 3.4	4.4 +/- 2.3	2.4 +/- 1.8	6.1 +/- 3.2
	Task 7	Task 8	Task 9	Task 10	Task 11	Task 12
Random	52.6 +/- 34.4	33.5 +/- 27.3	57.5 +/- 31.4	24.0 +/- 20.4	54.0 +/- 32.0	28.7 +/- 22.9
GP	8.5 +/- 3.4	7.5 +/- 2.83	10.0 +/- 4.2	8.4 +/- 4.4	8.8 +/- 4.0	8.2 +/- 3.4
ENN	14.30 +/- 4.1	5.2 +/- 1.3	5.7 +/- 1.2	5 +/- 0.8	6.5 +/- 0.9	6 +/- 3.0
GWEBL	3.5 +/- 1.2	3.0 +/- 1.0	3.5 +/- 1.1	2.1 +/- 1.1	3.9 +/- 2.2	2.7 +/- 1.1

Infarct Condition

Methods	Task 1	Task 2	Task 3	Task 4	Task 5	Task 6
Random	45.5 +/- 31.8	49.3 +/- 32.8	32.2 +/- 25.8	16.5 +/- 12.7	58.6 +/- 35.0	43.1 +/- 31.3
GP	10.5 +/- 5.0	7.6 +/- 5.7	9.4 +/- 5.3	5.9 +/- 2.8	7.9 +/- 4.1	8.8 +/- 5.7
ENN	8.3 +/- 2.6	3.4 +/- 0.8	5.2 +/- 1.0	5.8 +/- 1.2	7.6 +/- 2.2	8.4 +/- 4.5
GWEBL	7.7 +/- 1.8	1.7 +/- 1.0	3.0 +/- 1.3	2.4 +/- 1.4	5.0 +/- 3.0	4.3 +/- 2.9
	Task 7	Task 8	Task 9	Task 10	Task 11	Task 12
Random	51.5 +/- 35.5	66.6 +/- 36.3	43.2 +/- 19.7	39.1 +/- 31.2	96.5 +/- 14.5	35.2 +/- 30.6
GP	9.9 +/- 4.1	10.3 +/- 4.4	12.9 +/- 5.8	8.0 +/- 3.8	26.1 +/- 3.8	7.5 +/- 4.7
ENN	6.7 +/- 1.0	5.7 +/- 0.9	8.0 +/- 3.8	7 +/- 1.5	15.0 +/- 3.7	8.9 +/- 3.6
GWEBL	3.4 +/- 1.1	4.0 +/- 1.4	4.3 +/- 1.1	2.9 +/- 1.1	8.3 +/- 4.5	3.0 +/- 1.7

Figure 2. Localization steps (mean +/- std) for Sinus and Infarct. Task 1 uses no prior information.

3.3. Experiment Setup

On the first task we used 2 initial pacing-site ECG pairs to start our *GWEBL* model, then we generated a set of 11 subsequent tasks (so 12 in total) for our continual-learning strategy to optimize without additional seeded data. The sequence of tasks remained the same in each run, but we conducted 50 random restarts with different initial data on the first task.

4. Results

Figure 2 summarizes the results of the steps to localize each task for both the healthy and infarcted heart condition, expressed as mean $\pm$ std. The *Random* method shows an extremely large number of steps to localize and high variability. Whereas the *GP* converges much quicker and more consistently, requiring on average only 10.9 rounds compared to 50.5 using *Random*.

The *ENN* (without continual-learning) shows an improvement over the *GP* as well. The ability to predict the ECG signal itself allows the model to include more information in its representation, thus it only requires on average only 7.8 rounds to converge. Finally, we examine our proposed *GWEBL* method that is able to dynamically transfer knowledge across tasks without any actual data being kept, and using this method results in convergence after only 4.2 steps on average.

Statistical analysis shows that with 95% confidence, our method localizes ECG targets in significantly fewer steps, with an improvement of 65% over *GP* and 40% over *ENN*.

5. Conclusion

In conclusion, we present a novel framework employing continual-learning techniques to effectively transfer knowledge across pace-mapping tasks without the need for copying of data or memory. Using this our method in the Bayesian Optimization context allows us to localize the origin of VT episodes using significantly fewer steps, enabling a faster and more reliable procedure.

Acknowledgments

This study was supported by funding provided by NIH Award No: R01HL145590 and NSF Award No: OAC-2212548.

References

[1] Di Cesare M, Perel P, Taylor S, Kabudula C, Bixby H, Gaziano TA, McGhie DV, Mwangi J, Pervan B, Narula J, et al. The heart of the world. Global heart 2024;19(1).

[2] Koplan BA, Stevenson WG. Ventricular tachycardia and sudden cardiac death. In Mayo clinic proceedings, volume 84. Elsevier, 2009; 289–297.

[3] Josephson ME, Callans DJ. Using the twelve-lead electrocardiogram to localize the site of origin of ventricular tachycardia. Heart Rhythm 2005;2(4):443–446.

[4] Sapp JL, Bar-Tal M, Howes AJ, Toma JE, El-Damaty A, Warren JW, MacInnis PJ, Zhou S, Horáček BM. Real-time localization of ventricular tachycardia origin from the 12-lead electrocardiogram. JACC Clinical Electrophysiology 2017;3(7):687–699.

[5] Zhou S, AbdelWahab A, Sapp JL, Warren JW, Horáček BM. 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 2019;47:403–412.

[6] Missel R, Gyawali PK, Murkute JV, Li Z, Zhou S, AbdelWahab A, Davis J, Warren J, Sapp JL, Wang L. A hybrid machine learning approach to localizing the origin of ventricular tachycardia using 12-lead electrocardiograms. Computers in biology and medicine 2020;126:104013.

[7] Meisenzahl C, Gillette K, Prassl AJ, Plank G, Sapp JL, Wang L. Boatmap: Bayesian optimization active targeting for monomorphic arrhythmia pace-mapping. Computers in Biology and Medicine 2024;182:109201.

[8] Aras K, Good W, Tate J, Burton B, Brooks D, Coll-Font J, Doessel O, Schulze W, Potyagaylo D, Wang L, et al. Experimental data and geometric analysis repository—edgar. Journal of electrocardiology 2015;48(6):975–981.

Address for correspondence:

Dylan O'Hara
1 Lomb Memorial Drive, Rochester, NY, USA
dbo8671@rit.edu