



High-Dimensional Bayesian Optimization of Personalized Cardiac Model Parameters via an Embedded Generative Model

Jwala Dhamala¹(✉), Sandesh Ghimire¹, John L. Sapp², B. Milan Horáček², and Linwei Wang¹

¹ Rochester Institute of Technology, New York, USA
jd1336@rit.edu

² Dalhousie University, Halifax, Canada

Abstract. The estimation of patient-specific tissue properties in the form of model parameters is important for personalized physiological models. However, these tissue properties are spatially varying across the underlying anatomical model, presenting a significance challenge of high-dimensional (HD) optimization at the presence of limited measurement data. A common solution to reduce the dimension of the parameter space is to explicitly partition the anatomical mesh, either into a fixed small number of segments or a multi-scale hierarchy. This anatomy-based reduction of parameter space presents a fundamental bottleneck to parameter estimation, resulting in solutions that are either too low in resolution to reflect tissue heterogeneity, or too high in dimension to be reliably estimated within feasible computation. In this paper, we present a novel concept that embeds a generative variational auto-encoder (VAE) into the objective function of Bayesian optimization, providing an implicit low-dimensional (LD) search space that represents the generative code of the HD spatially-varying tissue properties. In addition, the VAE-encoded knowledge about the generative code is further used to guide the exploration of the search space. The presented method is applied to estimating tissue excitability in a cardiac electrophysiological model. Synthetic and real-data experiments demonstrate its ability to improve the accuracy of parameter estimation with more than 10x gain in efficiency.

Keywords: Parameter estimation · Model personalization
Cardiac electrophysiology · Variational auto-encoder
Bayesian optimization

1 Introduction

Patient-specific cardiac models have shown increasing potential in personalized treatment of heart diseases [9]. A significant challenge in personalizing these

models arises from the estimation of patient-specific tissue properties that vary spatially across the myocardium. To estimate these high-dimensional (HD) tissue properties (in the form of model parameters) is not only algorithmically difficult given indirect and sparse measurements, but also computationally intractable in the presence of computing-intensive simulation models.

Numerous efforts have been made to circumvent the challenge of HD parameter estimation. Many works assume homogeneous tissue property that can be represented by a single global model parameter [7]. To preserve local tissue properties, a common approach is to reduce the parameter space through an explicit partitioning of the cardiac mesh. These efforts can be generally summarized in two categories. In one approach, the cardiac mesh is pre-divided into 3–26 segments, each represented by a uniform parameter value [11]. Naturally, this artificial low-resolution division has a limited ability to represent tissue heterogeneity that is not known *a priori*. It has also been shown that the initialization of model parameters becomes increasingly more critical as the number of segments grows [11]. Alternatively, a multi-scale hierarchy of the cardiac mesh can be defined for a coarse-to-fine optimization, which allows spatially-adaptive resolution that is higher in certain regions than the others [3,4]. However, the representation ability of the final partition is limited by the inflexibility of the multi-scale hierarchy: homogeneous regions distributed across different scales cannot be grouped into the same partition, while the resolution of heterogeneous regions can be limited by the level of scale the optimization can reach [4]. In addition, because these methods involve a cascade of optimizations along the hierarchy of the cardiac mesh, they are computationally expensive. In the presence of models that could require hours or days for a single simulation, these methods could quickly become computationally prohibitive.

In this paper, we present a novel HD parameter optimization approach that replaces the explicit anatomy-based reduction of the parameter space, with an implicit low-dimensional (LD) manifold that represents the generative code for HD spatially-varying tissue properties. This is achieved by embedding within the optimization a generative variational auto-encoder (VAE) model, trained from a large set of spatially-varying tissue properties reflecting regional tissue abnormality with various locations, sizes, and distributions. The VAE decoder is utilized within the objective function of the Bayesian optimization [2] to provide an implicit LD search space for HD parameter estimation. Meanwhile, the VAE-encoded posterior distribution of the generative code is used to guide an efficient exploration of the LD manifold. The presented method is applied to estimating tissue excitability of a cardiac electrophysiological model using non-invasive electrocardiogram (ECG) data. On both synthetic and real data experiments, the presented method is compared against the use of anatomy-based LD [11] or multi-scale representation of the parameter space [4]. Experiments demonstrate that the presented method can achieve a drastic reduction in computational cost while improving the accuracy of the estimated parameters. To the best of our knowledge, this is the first work that utilizes a probabilistic generative model

within an optimization framework for estimating HD model parameters. It provides an efficient and general solution to personalizing HD model parameters.

2 Background: Cardiac Electrophysiological System

Cardiac Electrophysiology Model: Among the different types of cardiac electrophysiological models, phenomenological models such as the Aliev-Panfilov (AP) model [1] can explain the macroscopic process of cardiac excitation with a small number of model parameters and reasonable computation. Therefore, the AP model given below is chosen to test the feasibility of the presented method:

$$\begin{aligned}\partial u/\partial t &= \nabla(\mathbf{D}\nabla u) - cu(u - \theta)(u - 1) - uv, \\ \partial v/\partial t &= \varepsilon(u, v)(-v - cu(u - \theta - 1)).\end{aligned}\tag{1}$$

Here, u is the transmembrane action potential and v is the recovery current. The transmural action potential is computed by solving the AP model (1) on a 3D myocardium discretized using the meshfree method [10]. Because u is most sensitive to the value of the parameter θ [4], we focus on its estimation in this study.

Body-Surface ECG Model: The propagation of the spatio-temporal transmural action potential \mathbf{U} to the potentials measured on the body surface \mathbf{Y} can be described by the quasi-static approximation of the electromagnetic theory [8]. Solving the governing equations on a discrete heart-torso mesh, a linear relationship between \mathbf{U} and \mathbf{Y} can be obtained as: $\mathbf{Y} = \mathbf{H}(\mathbf{U}(\boldsymbol{\theta}))$, where $\boldsymbol{\theta}$ is the vector of local parameters θ at the resolution of the cardiac mesh.

3 HD Parameter Estimation

To estimate $\boldsymbol{\theta}$, we maximize the similarity between the measured ECG and those simulated by the combined electrophysiological and ECG model $M(\boldsymbol{\theta})$:

$$\hat{\boldsymbol{\theta}} = \arg \max_{\boldsymbol{\theta}} -\|\mathbf{Y} - M(\boldsymbol{\theta})\|^2.\tag{2}$$

To enable the estimation of $\boldsymbol{\theta}$ at the resolution of the cardiac mesh, the presented method embeds within the Bayesian optimization framework a stochastic generative model that generates $\boldsymbol{\theta}$ from a LD manifold. It includes two major components as outlined in Fig. 1: (1) the construction of a generative model of HD spatially-varying tissue properties at the resolution of the cardiac mesh, and (2) a novel Bayesian optimization method utilizing the embedded generative model.




Fig. 1. Outline of the presented method, with the dimension of each VAE layer labeled.

3.1 LD-to-HD Parameter Generation via VAE

Generative VAE Model: We assume that the spatially varying tissue properties at the resolution of a cardiac mesh θ is generated by a small number of unobserved continuous random variables \mathbf{z} in a LD manifold. To obtain the generative process from \mathbf{z} to θ , the VAE consists of two modules: a probabilistic deep encoder network with parameters α that approximates the intractable true posterior density as $q_\alpha(\mathbf{z}|\theta)$; and a probabilistic deep decoder network with parameters β that can probabilistically reconstruct θ given \mathbf{z} as $p_\beta(\theta|\mathbf{z})$. Both networks consist of three fully-connected layers as shown in Fig. 1.

To train the VAE, we generate $\Theta = \{\theta^{(i)}\}_{i=1}^N$ consisting of N configurations of heterogeneous tissue properties in a patient-specific cardiac mesh. The training involves optimizing the variational lower bound on the marginal likelihood of each training data $\theta^{(i)}$ with respect to network parameters α and β :

$$\mathcal{L}(\alpha; \beta; \theta^{(i)}) = -D_{\text{KL}}(q_\alpha(\mathbf{z}|\theta^{(i)})||p_\beta(\mathbf{z})) + E_{q_\alpha(\mathbf{z}|\theta^{(i)})}[\log p_\beta(\theta^{(i)}|\mathbf{z})], \quad (3)$$

where we model $p_\beta(\theta|\mathbf{z})$ with a Bernoulli distribution. To optimize Eq. (3), stochastic gradient descent with standard backpropagation can be utilized. Assuming the approximate posterior $q_\alpha(\mathbf{z}|\theta)$ as a Gaussian density and the prior $p_\beta(\mathbf{z}) \sim \mathcal{N}(0, 1)$, their KL divergence can be derived analytically as:

$$D_{\text{KL}}(q_\alpha(\mathbf{z}|\theta^{(i)})||p_\beta(\mathbf{z})) = -\frac{1}{2} \sum (1 + \log(\sigma_j^2) - \mu_j^2 - \sigma_j^2), \quad (4)$$

where j is along the dimensions of \mathbf{z} , and μ and σ^2 are mean and variance from $q_\alpha(\mathbf{z}|\theta^{(i)})$. Because stochastic latent variables are utilized, the gradient of the expected negative reconstruction term during backpropagation cannot be directly obtained. The popular re-parameterization trick is utilized to express \mathbf{z} as a deterministic variable as $\mathbf{z}^{(i)} = \mu^{(i)} + \sigma^{(i)}\epsilon$, where $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ is noise [6].

Probabilistic Modeling of the Latent Code: The trained encoder provides an approximated posterior density of the LD latent code $q_\alpha(\mathbf{z}|\theta)$. This represents valuable knowledge about the probabilistic distribution of \mathbf{z} learned from a large training dataset. To utilize this in the subsequent optimization, we integrate $q_\alpha(\mathbf{z}|\theta)$ over the training data Θ to obtain the density $q_\alpha(\mathbf{z})$ as a mixture of Gaussians $1/N \sum_i^N \mathcal{N}(\mu^{(i)}, \Sigma^{(i)})$, where $\mu^{(i)}$ and $\Sigma^{(i)}$ are mean and covariance from $q_\alpha(\mathbf{z}|\theta^{(i)})$. Because the number of mixture components in $q_\alpha(\mathbf{z})$ scales

linearly with the number of training data, we approximate $q_{\alpha}(\mathbf{z})$ with a single Gaussian density as $\mathcal{N}(1/N \sum_i^N \boldsymbol{\mu}^{(i)}, 1/N \sum_i^N (\boldsymbol{\Sigma}^{(i)} + \boldsymbol{\mu}^{(i)} \boldsymbol{\mu}^{(i)T}) - \boldsymbol{\mu} \boldsymbol{\mu}^T)$. Alternatively, we approximate $q_{\alpha}(\mathbf{z})$ with a mixture of Gaussians with $K \ll N$ components, where k-means clustering with the Bregman divergence [5] as a similarity metric is used to reduce the number of mixture components.

In this way, we obtain a generative model $p_{\beta}(\boldsymbol{\theta}|\mathbf{z})$ of HD tissue properties from an implicit LD manifold, and prior knowledge of the LD manifold $q_{\alpha}(\mathbf{z})$ from the probabilistic encoder. Both will be embedded into Bayesian optimization to enable efficient and accurate HD parameter estimation.

3.2 Bayesian Optimization with Embedded Generative Model

Representing $\boldsymbol{\theta}$ with the expectation of the trained decoder $p_{\beta}(\boldsymbol{\theta}|\mathbf{z})$, we obtain:

$$\hat{\mathbf{z}} = \arg \max_{\mathbf{z}} -\|\mathbf{Y} - M(\mathbb{E}[p_{\beta}(\boldsymbol{\theta}|\mathbf{z})])\|^2, \quad (5)$$

which allow us to optimize the HD parameter $\boldsymbol{\theta}$ in an implicit LD manifold of \mathbf{z} . For Bayesian optimization, we assume a zero mean Gaussian process (GP) with an anisotropic Matérn 5/2 kernel as a prior over the objective function (5). The optimization then consists of two iterative steps: (1) select point in the LD manifold that allows the GP to globally approximate Eq. (5) (exploration) while locally refining the area of optimum (exploitation); and (2) update the GP.

VAE-Informed Acquisition Function: To select points on LD manifold, we adopt the expected improvement (EI) function that picks a point with maximum expectation of improvement over the current best objective function value f_m [2]. For a GP posterior $\sim \mathcal{N}(\mu(\cdot), \sigma(\cdot))$, it can be obtained as:

$$\text{EI}(\mathbf{z}) = (\mu(\mathbf{z}) - f_m) \Phi\left(\frac{\mu(\mathbf{z}) - f_m}{\sigma(\mathbf{z})}\right) + \sigma(\mathbf{z}) \phi\left(\frac{\mu(\mathbf{z}) - f_m}{\sigma(\mathbf{z})}\right), \quad (6)$$

where Φ and ϕ are the cumulative distribution function and density function of the standard normal distribution. Here, the first term controls exploitation (through high μ) and the second term controls exploration (through high σ). Because using only f_m can lead to excessive exploitation, it is common to augment f_m with a constant trade-off parameter ε as: $f_m + \varepsilon$ [2]. Here, we utilize the VAE-encoded knowledge about the LD manifold $q_{\alpha}(\mathbf{z})$ to enforce higher exploitation in the areas of high probability density for \mathbf{z} , and lower elsewhere. In specific, we define $\varepsilon(\mathbf{z}) = -f_m \sum_i w_i (\mathbf{z} - \boldsymbol{\mu}_i) \boldsymbol{\Sigma}_i^{-1} (\mathbf{z} - \boldsymbol{\mu}_i)$, where w , $\boldsymbol{\mu}$, and $\boldsymbol{\Sigma}$ are the weight, mean, and variance of the K Gaussian mixture components in $q_{\alpha}(\mathbf{z})$.

GP Update: After a new point $\mathbf{z}^{(n)}$ is selected by maximizing the modified EI, the objective function (5) is evaluated at the HD parameter given by the mean of the generative model $p_{\beta}(\boldsymbol{\theta}|\mathbf{z}^{(n)})$. The GP is then updated by adding the new pair of $\mathbf{z}^{(n)}$ and objective function value, and maximizing the log marginal likelihood with respect to kernel parameters: length scales and kernel amplitude.

4 Experiments

Synthetic Experiments: We include 27 synthetic experiments on three CT-derived human heart-torso models. In each case, an infarct sized 2%–40% of the heart was placed at differing locations using various combinations of the AHA segments. The value of the parameter θ in the infarcted and the healthy region is set to 0.5 and 0.15, respectively. 120-lead ECG is simulated and corrupted with 20 dB Gaussian noise as measurement data. We evaluate the accuracy in estimated parameters with two metrics: (1) root mean square error (RMSE) between the true and estimated parameters; and (2) dice coefficient (DC) = $\frac{2(|S_1 \cap S_2|)}{|S_1| + |S_2|}$, where S_1 and S_2 are the sets of nodes in the true and estimated regions of infarct; these regions are determined by *Otsu’s thresholding* method.

VAE Architecture and Training: For each heart, we generate a training dataset of tissue properties with various heterogeneous infarcts. Each infarct is generated by random region growing in which, starting with one infarct node, one out of the five closest neighbors of the present infarct is randomly added to the infarct until an infarct of desired size is obtained. It is then added to training data. Because infarcts thus generated tend to be very irregular, we also include infarcts generated by growing the infarct with the node closest to its center. For each heart, we extract 123,896, 155,099, and 116,459 data. The training of VAE with an architecture as shown in Fig. 1 with the *Adam* optimizer on Titan X GPU took 9.77, 13.96, and 9.0 min for each dataset.




Fig. 2. Comparison of BO-VAE EI Post-1 (blue bar) with: (1) FH and FS (green bars); and (2) BO-VAE using standard EI, EI Isotropic, and EI Post-K (yellow bars) in terms of DC, RMSE, and number of model evaluations (from left to right).

Comparison with Existing Methods: The presented method (termed as BO-VAE) is compared against two common approaches based on explicit LD representation: (1) optimization over fixed 18 segments (fixed-segment (FS) method); and (2) coarse-to-fine optimization along a fixed multi-scale hierarchy (fixed-hierarchy (FH) method). As summarized in Fig. 2(a)(b), BO-VAE (blue bar) is more accurate than the other two methods (green bars) in both DC and RMSE (paired t-tests, $p < 0.012$). This is achieved at a reduction of the computational cost by: 87.57% for the FS method and 98.73% for the FH method (Fig. 2(c)).

The FS method shows the lowest accuracy with some estimated parameters either missing the infarct or including large false positives (Fig. 3) left. The FH method overcomes this issue, although to a limited extent. In the LD representation obtained by the FH method as shown in Fig. 3 right several dimensions




Fig. 3. Left: Examples of estimated parameters with BO-VAE, FH, and FS. Right: Progression of FH on the multi-scale hierarchy for parameter estimation of (a) (green leaf: homogeneous tissues; red leaf: heterogeneous tissues).

are wasted at representing homogeneous healthy regions (green) across different scales, which limits its ability to optimize deeper along the tree. BO-VAE is not limited by such explicitly-imposed anatomy-based structure, allowing it to attain higher accuracy with only 2 latent dimensions and 1–10% of the computation.

The Effect of VAE-Encoded Knowledge About the LD Manifold: To study the effect of incorporating the VAE-encoded $q_{\alpha}(\mathbf{z})$ in the EI, we compare the standard EI with EI augmented with three types of distributions on \mathbf{z} : (1) $p_{\beta}(\mathbf{z}) \sim \mathcal{N}(0, 1)$ (EI Isotropic), (2) approximated $q_{\alpha}(\mathbf{z})$ with a single Gaussian density (EI Post-1); and (3) approximated $q_{\alpha}(\mathbf{z})$ with a mixture of 10 Gaussian densities (EI Post-K). As shown in Fig. 2, the accuracy using all three distributions is higher than that without using any, among which EI Post-1 has the highest accuracy. Figure 4(b) illustrates that, when $q_{\alpha}(\mathbf{z})$ is utilized, the exploration proceeds from the region of high probability density to the region of low probability density. In comparison, with standard EI, the points are spread to reduce overall variance (Fig. 4(a)); this could result in incorrect (Fig. 4(c)) or suboptimal (Fig. 4(d)) solutions.




Fig. 4. Comparison of training points selected by EI (a) and EI Post-1 (b), and examples of the estimated parameters by the two acquisition functions (c–d).

We also experimented with HD latent code \mathbf{z} . As shown in Fig. 5(a)(b), there was only a marginal improvement in accuracy with a five *vs.* a two dimensional (2d) latent code. It suggests that, given the focus of the training data on local




Fig. 5. (a–b): Examples of estimated parameters using five *vs.* two dimensional latent codes. (c–d): Latent code manifold based on (c) infarct location, and (d) infarct size.

infarcts, a 2d latent code may be sufficient to capture the necessary generative factors. The plot of these 2d latent codes in Fig. 5(c)(d) show that they cluster by infarct location and their radial direction accounts for the infarct size.

Real Data Experiments: Real-data studies are conducted on two patients with previous myocardial infraction. Patient-specific heart and torso meshes are constructed from axial CT images. Tissue excitability is estimated from 120-lead ECG data. The results are evaluated by *in-vivo* bipolar voltage data which, although not a direct measure of tissue excitability, provides a reasonable reference about the region of infarcts. The first two columns of Fig. 6 show the original voltage data (red: dense infarct; purple: healthy tissue; green: infarct border) and the same data registered to cardiac meshes.

The voltage map in case 1 (Fig. 6(a)) shows a highly heterogeneous infarct spread over a large region in the lateral LV. The estimated parameters by all methods capture this region of infarct. For this accuracy, the FH and FS methods required 4056 and 1058 model evaluations, whereas BO-VAE required only 105 model evaluations. By contrast, as shown in Fig. 6(b), case 2 has a smaller region of dense scar in the lateral LV. The estimated parameters by BO-VAE and FH correctly reveal this region of scar, whereas the FS method is less accurate. In this case, BO-VAE, FH method, and FS method required 105, 5798, and 1501 model evaluations respectively.




Fig. 6. Model parameter estimated with BO-VAE, FH, and FS on real-data study.

Conclusion: We present a novel approach to estimating HD model parameters, achieved by embedding within the Bayesian optimization a generative model of HD tissue properties from a LD manifold. Experiments show a gain in accuracy

with drastically reduced computation. Future works include the incorporation of training data from high resolution 3D imaging and study of alternatives to incorporate the knowledge of latent manifold in Bayesian optimization.

Acknowledgment. This work is supported by the National Science Foundation under CAREER Award ACI-1350374 and the National Institute of Heart, Lung, and Blood of the National Institutes of Health under Award R21HL125998.

References

1. Aliev, R.R., Panfilov, A.V.: A simple two-variable model of cardiac excitation. *Chaos Solitons Fractals* **7**(3), 293–301 (1996)
2. Brochu, E., Cora, V.M., De Freitas, N.: A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. arXiv preprint [arXiv:1012.2599](https://arxiv.org/abs/1012.2599) (2010)
3. Chinchapatnam, P., Rhode, K.S., Ginks, M.: Model-based imaging of cardiac apparent conductivity and local conduction velocity for diagnosis and planning of therapy. *IEEE Trans. Med. Imaging* **27**(11), 1631–1642 (2008)
4. Dhamala, J., Sapp, J.L., Horacek, M., Wang, L.: Spatially-adaptive multi-scale optimization for local parameter estimation: application in cardiac electrophysiological models. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) *MICCAI 2016*. LNCS, vol. 9902, pp. 282–290. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46726-9_33
5. Garcia, V., Nielsen, F., Nock, R.: Levels of details for Gaussian mixture models. In: Zha, H., Taniguchi, R., Maybank, S. (eds.) *ACCV 2009*. LNCS, vol. 5995, pp. 514–525. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-12304-7_48
6. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint [arXiv:1312.6114](https://arxiv.org/abs/1312.6114) (2013)
7. Lê, M., Delingette, H., Kalpathy-Cramer, J.: MRI based Bayesian personalization of a tumor growth model. *IEEE TMI* **35**(10), 2329–2339 (2016)
8. Plonsey, R.: *Bioelectric Phenomena*. Wiley Online Library, Hoboken (1969)
9. Sermesant, M., Chabiniok, R., Chinchapatnam, P., Mansi, T., et al.: Patient-specific electromechanical models of the heart for the prediction of pacing acute effects in CRT: a preliminary clinical validation. *Med. Image Anal.* **16**(1), 201–215 (2012)
10. Wang, L., Zhang, H., Wong, K.C., Liu, H., Shi, P.: Physiological-model-constrained noninvasive reconstruction of volumetric myocardial transmembrane potentials. *IEEE Trans. Biomed. Eng.* **57**(2), 296–315 (2010)
11. Wong, K.C., Sermesant, M., Rhode, K.: Velocity-based cardiac contractility personalization from images using derivative-free optimization. *J. Mech. Behav. Biomed. Mater.* **43**, 35–52 (2015)