Spatially-Adaptive Multi-scale Optimization for Local Parameter Estimation: Application in Cardiac Electrophysiological Models

Jwala Dhamala[1(✉)], John L. Sapp[2], Milan Horacek[2], and Linwei Wang[1]

[1] Rochester Institute of Technology, Rochester, NY 14623, USA
jd1336@rit.edu
[2] Dalhousie University, Halifax, Canada

Abstract. The estimation of local parameter values for a 3D cardiac model is important for revealing abnormal tissues with altered material properties and for building patient-specific models. Existing works in local parameter estimation typically represent the heart with a small number of pre-defined segments to reduce the dimension of unknowns. Such low-resolution approaches have limited ability to estimate tissues with varying sizes, locations, and distributions. We present a novel optimization framework to achieve a higher-resolution parameter estimation without using a high number of unknowns. It has two central elements: (1) a multi-scale coarse-to-fine optimization that uses low-resolution solutions to facilitate the higher-resolution optimization; and (2) a spatially-adaptive scheme that dedicates higher resolution to regions of heterogeneous tissue properties whereas retaining low resolution in homogeneous regions. Synthetic and real-data experiments demonstrate the ability of the presented framework to improve the accuracy of local parameter estimation in comparison to optimization based on fixed-segment models.

Keywords: Parameter estimation · Cardiac electrophysiological model · Multi-scale optimization · Gaussian process

1 Introduction

Many cardiac diseases stem from abnormal myocardial tissues with altered material properties. The quantitative knowledge about these abnormal tissues is paramount to the diagnosis, treatment, and prevention of relevant cardiac diseases. Since it is difficult to directly measure the material property of cardiac tissues, one effective way to quantify pathological tissue properties is to estimate the three-dimensionally distributed parameters of a cardiac model using indirect measurement data. This will in addition provide a patient-specific model useful for personalized treatment planning and prognosis [4].

Much effort has been reported on parameter estimation for complex physiological models. For example, derivative-free optimization methods [9] are shown

S. Ourselin et al. (Eds.): MICCAI 2016, Part III, LNCS 9902, pp. 282–290, 2016.
DOI: 10.1007/978-3-319-46726-9_33

to be effective in handling the complex objective functions. Alternatively, surrogate models such as spectral representation based on polynomial chaos [4], multivariate polynomial regression [10], and Gaussian processes [5] have gained increasing interest in recent years.

Nevertheless, progress has been limited in estimating local parameters that are three-dimensionally distributed in space. Many previous works focus on the estimation of global parameters [5] by assuming uniform tissue property throughout the myocardium. Although it provides a fast calibration of a model, global parameter estimation does not reveal the local change of tissue properties. Toward local parameter estimation, a commonly-used approach is to divide the myocardium into a set of pre-defined segments and assume the parameter to be uniform within each segment. This substantially reduces the dimension of unknowns (to a range of 3 to 27) [9,10], yet the resolution is too low to capture abnormal tissues with different sizes, locations, and distributions. Moreover, as the number of segments increases, a good initialization becomes critical [9] which typically requires additional data to delineate diseased regions *a priori*. A critical gap remains between the need for a high-resolution local parameter estimation and the difficulty to accommodate high-dimensional optimization.

To bridge this gap, we propose a novel framework that goes beyond fixed low-resolution parameter estimation without invoking an infeasible number of unknowns. This is achieved via two primary elements. First, a multi-scale hierarchy is used to progressively use low-resolution results to facilitate higher-resolution optimization, thereby alleviating the issue of identifiability. Second, instead of uniform resolution, an adaptive scheme is used to selectively allocate higher resolution in heterogeneous regions whereas retaining lower resolution in homogeneous regions. It shares an important intuition with [1] where non-uniform mesh is used, although with fundamental differences in the coarser-to-finer transition of information and the motivation (heterogeneity) for adaptive resolution.

The proposed framework is applied to local parameter estimation for a 3D cardiac electrophysiological (EP) model using non-invasive electrocardiographic (ECG) data. It is noteworthy that the remote global ECG data increase the difficulty in identifying local tissue properties in comparison to local direct mapping data. The presented method is tested on a set of synthetic and real-data experiments. In comparison to the derivative-free BOBYQA [6] method carried out on a predefined 18-segment model, the presented method demonstrates higher accuracy using a similar or even fewer unknowns. While this framework is reported with GP based optimization, it can be used with other optimization methods. It is also applicable to local parameter estimation beyond cardiac EP models.

2 Cardiac Electrophysiology and ECG

2.1 Cardiac Electrophysiological Model

The spatiotemporal evolution of cardiac action potential can be described by a set of differential equations, ranging from complex ionic models with tens of

hundreds of parameters to simpler models with a few parameters [2]. As an initial demonstration of feasibility for the proposed framework, we consider parameter estimation for the *Aliev-Panfilov* (AP) [2] model because of its ability to simulate electrical dynamics with fewer parameters and reasonable computation.

$$\begin{aligned} \partial u/\partial t &= \partial/\partial x_i d_{ij} \partial u/\partial x_j - ku(u-a)(u-1) - uv \\ \partial v/\partial t &= \varepsilon(u,v)(-v - ku(u-a-1)). \end{aligned} \tag{1}$$

where, u is the transmembrane action potential and v is the recovery current. Parameter d_{ij} is the spatial conductivity, ε controls the coupling between the recovery current and action potential, k controls the repolarization, and a controls the excitability of a cell. In this study, we focus on a as it is one of the most sensitive model parameter and its value is associated to the ischemic severity of the myocardial tissue [2]. The meshfree method is used to discretize and solve (1) on the 3D myocardium [8], with a resolution of ∼6-mm (∼10^3 nodes).

2.2 ECG Measurement Model

Cardiac action potential produces potential on the body surface that is measured as time-varying ECG signals. This measurement process can be described by the quasi-static approximation of the electromagnetic theory [8]. Solving this governing equation on the discrete mesh of heart and torso models specific to an individual, a linear model between ECG data $\mathbf{\Phi}$ and transmural action potential $\mathbf{u}$ can be obtained as: $\mathbf{\Phi} = \mathbf{H}\mathbf{u}$ [8].

3 Spatially-Adaptive, Multi-scale Optimization

Estimation of the three-dimensionally distributed tissue excitability $\boldsymbol{\theta}$ from ECG data $\mathbf{y}$ can be formulated as a bounded global maximization problem:

$$\max_{\mathbf{l}\le\boldsymbol{\theta}\le\mathbf{u}} G(\boldsymbol{\theta}) = \max_{\mathbf{l}\le\boldsymbol{\theta}\le\mathbf{u}} \sum_{i=1}^{L} \Big(\frac{\sum_{t=1}^{M}(y_{it}-\bar{y_i})(\Phi_{it}-\bar{\Phi}_i)}{L\sqrt{\sum_{t=1}^{M}(y_{it}-\bar{y_i})^2 \sum_{t=1}^{M}(\Phi_{it}-\bar{\Phi}_i)^2}} - \lambda \sum_{t=1}^{M}(\Phi_{it}-y_{it})^2 \Big). \tag{2}$$

where $\mathbf{\Phi} = f(\boldsymbol{\theta})$ is a composite of the measurement an AP models (see Sect. 2). The objective function (2) includes both a correlation coefficient and a squared error to balance the morphology and magnitude similarity between the ECG data and the model output. L is the number of ECG leads.

The direct estimation of $\boldsymbol{\theta}$ requires a high-dimensional optimization (order of 10^3) that is not feasible due to both un-identifiability and high computation. To achieve higher-resolution local parameter estimation, the optimization framework described below includes two key components: (1) a hierarchical coarse-to-fine estimation, and (2) a spatially-adaptive resolution that is refined at regions of heterogeneity. This framework is developed with GP-based optimization.

3.1 Multi-scale Hierarchy

A coarse-to-fine optimization has the advantage to use lower-resolution solution to reduce the search during higher-resolution optimizations. To facilitate this, we construct a multi-scale representation of the cardiac mesh using Agglomerative Hierarchical Clustering [3], exploiting the spatial smoothness of tissue properties. A partial tree structure of this multi-scale model can be seen in Fig. 1. The clustering starts with each node in the cardiac mesh as a separate cluster. Every two closest clusters, based on the Euclidean distance and average linkage metric, are then merged until the entire ventricular mesh belongs to a single cluster. On this hierarchy model, the optimization starts at the root as a global estimation, and progressively moves to a higher level of resolution. Each level of optimization consists of two primary tasks: (1) optimization exploiting the lower-level solution (Sect. 3.3); and (2) determination of the spatial resolution for the next level of optimization (Sect. 3.2).

3.2 Adaptive Spatial Refinement

Instead of uniform resolution, we aim for a spatially-adaptive resolution so that higher resolution is used at regions of heterogeneity. In other words, after each level of optimization, instead of splitting all leaf nodes selective splitting and retraction is done to generate a skewed tree.

The key task is to identify the heterogeneous versus homogeneous clusters in tissue properties. Intuitively, if a cluster is homogeneous, its split is expected to yield children clusters with similar parameter values; *i.e.*, there will be minimal gain in the objective function (2). The contrary is true for heterogeneous clusters. Therefore, we propose a criterion based on gains in the objective function value.

Specifically, after obtaining an optimal solution $\boldsymbol{\theta}^k$ at level k, we examine two types of leaf nodes. First, we examine each pair of sibling nodes $(\theta^k_{i,c1}, \theta^k_{i,c2})$ that share same parent node $\boldsymbol{\theta}^{k-1}_i$. For each pair of $(\theta^k_{i,c1}, \theta^k_{i,c2})$, we evaluate the gain of splitting them from their parent as the difference in objective function evaluated on $\boldsymbol{\theta}^k$ versus replacing $(\theta^k_{i,c1}, \theta^k_{i,c2})$ with their parent θ^{k-1}_i:

$$r_{k,i} = G(\boldsymbol{\theta}^k) - G(\mathbf{s}^k), \quad \text{where} \quad \mathbf{s}^k = (\boldsymbol{\theta}^k \setminus (\theta^k_{i,c1}, \theta^k_{i,c2}), \boldsymbol{\theta}^{k-1}_i) \tag{3}$$

Second, for leaf nodes that do not any sibling, no resolution change has occurred but their values may have been changed as a result of resolution change elsewhere. For such a node $\boldsymbol{\theta}^k_i$, the gain $r_{k,i}$ equals the change in the objective function due to the change in $\boldsymbol{\theta}^k_i$ before and after the optimization.

Based on $r_{k,i}$, we take two actions on tree structure before the next level of optimization: (1) for a leaf node or a pair of leaf nodes with maximum gain $r_{k,i}$, we consider them to be most heterogeneous and warrant a higher-resolution representation (*i.e.*, a split); and (2) for those that bring negligible or negative gain ($r_{k,i} < \delta$, δ is the same tolerance in the improvement of global optimum used for the convergence of overall framework), the split suggested by the previous level was not beneficial and retract it. The rest of nodes are unchanged.

3.3 Optimization via Surrogate Models

The proposed framework can be used in combination with any optimization method suitable for handling a complex objective function like (2). In this paper, a GP surrogate model based method is used [7]. In brief, the optimization assumes a prior distribution, in the form of a GP $\sim \mathcal{N}(\mu(\boldsymbol{\theta}), \sigma(\boldsymbol{\theta}))$, to denote the belief over the objective function (2) and sequentially updates the prior based on new data to better approximate the objective function, especially in the region of global optimum. Here, we elaborate the three main steps of the optimization at each level of resolution:

1. Initialize the GP: We start with a GP with zero mean and "Matern 5/2" covariance function [7] to impose a minimal assumption of smoothness over the objective function (2). While the GP-based optimization is gaining increasing attention for optimizing highly expensive cost functions, it suffers from an inability to scale in high dimension ($\leq$15) [7]. Here, we utilize the low-resolution optimum to facilitate the higher-resolution GP optimization. In the proposed framework, a set of higher-resolution points are generated from the previous lower-resolution optimum through a convolution operator and parameter bounds. These points serve to quickly obtain an initial higher-resolution GP surrogate.

2. Determine the next query point: To update the GP, the best point to query should both exploit the solution space of the current GP where the predictive mean $\mu(\boldsymbol{\theta})$ is high and explore the solution space where the predictive uncertainty $\sigma(\boldsymbol{\theta})$ is high. This is done by finding the point that maximizes the upper confidence bound $\mu(\boldsymbol{\theta}) + \kappa\sigma(\boldsymbol{\theta})$ of the current GP [7], using the BOBYQA [6] optimization. The parameter κ balances the exploitation and exploration.

3. Update the GP: On the new query point obtained from step 2, the objective function (2) is evaluated and the posterior distribution of the GP is updated [7]. Steps 2 and 3 run in iteration until convergence of the GP based optimization.

4 Experiments

Synthetic Experiments: In a set of 22 synthetic experiments conducted on 3 image-derived realistic human heart-torso models, we test the proposed method in estimating the excitability of cardiac tissue in presence of infarct of varying locations and sizes. The parameter "a" of the AP model (1) is set to be 0.15 ± 0.01 and 0.45 ± 0.01, respectively, for normal and infarction tissues. 120-lead ECG are simulated and corrupted with 20 dB Gaussian noise. Infarct covering 1 % to 40 % of the LV/RV is set at different locations using various combinations of the AHA segments and random initializations with sizes smaller than one segment.

The presented method is compared with the BOBYQA method [6] carried out on a fixed 18-segment model (17 LV AHA segments + 1 RV segment) [9]. Because GP-based optimization did not scale well to 18-dimensional optimization in our

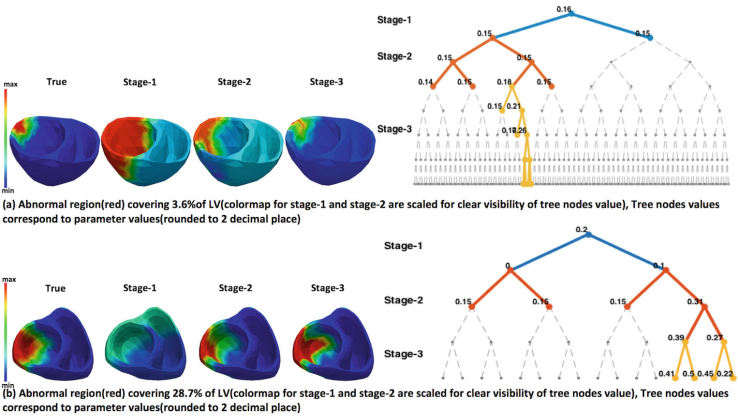

Fig. 1. Examples of the progression of the multi-scale optimization. Left: true parameter settings *vs.* estimation results over 3 successive stages of the optimization. Right: the corresponding growth of the tree at each stage. The gray, dotted structure shows the full hierarchy, whereas the colored line show the path taken by the presented method. (Color figure online)

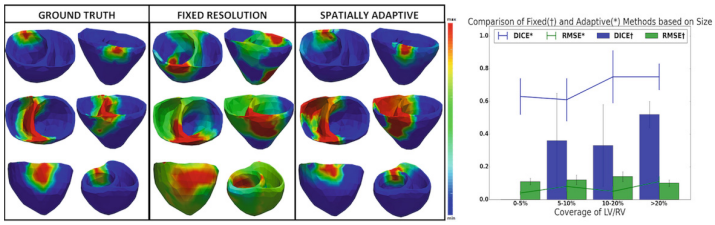

Fig. 2. Comparison of the presented method with BOBYQA on a 18-segment model. Left: examples of different infarcts. Right: quantitative comparison in DC and RMSE.

experiments, it is not included in this paper. We evaluate the estimated parameters using two metrics: (1) root mean square error (RMSE) between the true and estimated parameters; and (2) dice coefficient DC $= \frac{2(S_1 \cap S_2)}{S_1 \cup S_2}$, where S_1 and S_2 are the sets of cardiac nodes in the true and estimated regions of infarct; these regions are defined from the final tree in the presented method, and by thresholding parameter values in the BOBYQA method. Both metrics are evaluated at the resolution of the cardiac mesh.

Figure 1 demonstrates how the presented adaptive coarse-to-fine optimization progresses: the left panel shows the improvement in estimation at 3 successive stages, with the corresponding growth of the tree in the right panel. Figure 1a shows an example on a small infarct (3 %). The tree shows that since stage 1, the optimization split only along the heterogeneous region that contains the abnormality. It continued narrowing down the infarct with higher resolution, generating a narrow yet deep tree. The estimation shown in stage 3 of Fig. 1b

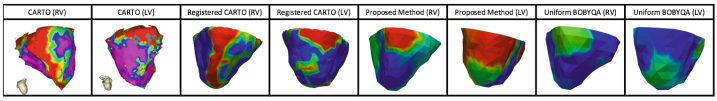

Fig. 3. Real-data experiment: comparison with *in-vivo* voltage maps of scar. (Color figure online)

was achieved with only 13 unknowns. In comparison, if a uniform resolution is used, a dimension of 128 is needed to achieve estimation at the same resolution. Figure 1b shows an example with a larger infarct (29 %). Because abnormal tissues span a larger number of clusters compared to a small infarct, it is not until stage 2 before the tree can be split along major branch. In addition, because both normal and abnormal tissues are large enough to be represented by low-resolution, homogeneous regions, an overall lower-resolution solution is obtained with a wider yet shallow tree. Similarly, in this case the presented method converged at a dimension of 7 whereas a uniform resolution of 16 is required.

Figure 2 summarizes the comparison between the presented method and the BOBYQA method directly on the fixed 18-segment. The improvement of the presented method is statistically significant in both DC and RMSE (paired-t tests, $p < 0.001$). More specifically, the performance of the presented method is much more robust to the size and shape of the infarct. While optimization on fixed segments has trouble handling infarcts of size equal to a single AHA segment, the presented method could provide an accurate estimation using only 12–14 unknowns. Furthermore, optimization on fixed segments tends to show false-positives across multiple segments, falling short to reveal the spatial distribution of an infarct. The presented method improves this accuracy by adaptively allocating higher-resolution on the heterogeneous regions. Depending on the type of the infarct, the computational cost of the presented method is comparable or higher than direct optimization on 18 uniform divisions. For medium sized infarcts (5–25 % of LV), the tree is shallower (Fig. 1b) and requires fewer coarse to fine optimizations. For small infarcts ($\leq$5 % of LV), the tree is deeper with many leaf nodes (Fig. 1a) and requires larger coarse to fine optimizations. In such cases, the number of model evaluations for presented method was at most 1.5 times that needed for direct optimization on 18 segments. Although for such size, direct optimization has trouble in estimation mainly due to un-identifiability.

Real-Data Experiment: As a feasibility test, we conducted a case study on a patient who underwent catheter ablation of ventricular tachycardia due to prior tissue infarction. Tissue excitability was estimated from 120-lead ECG on the patient-specific heart-torso geometry obtained from CT images. Bipolar voltage data from *in-vivo* CARTO mapping were used as reference: as illustrated in Fig. 3, they reveal low-voltage regions at both lateral LV and RV (red: dense scar $\leq$0.5 mV; green: scar border = 0.5–1.5 mV; blue: normal >1.5 mV). Excitability estimated from the presented framework successfully captured abnormal tissues

at both locations whereas estimation from direct BOBYQA optimization using pre-defined 23 segments (17-LV, 6-RV) captured the abnormal tissue located on RV only. Interestingly, during post-processing of CARTO, the clinician marked that the low-voltage region at middle-apical lateral RV was caused by poor catheter contact rather than scar tissue. The estimated excitability values on RV, reflected this tissue property on RV. Overall, the core and border of abnormal tissues as revealed by the estimated excitability appear to co-locate with CARTO maps. It should be noted that, because CARTO maps show voltage data whereas the estimated parameter map shows tissue excitability, they are not expected to appear identical; further caution is needed in interpreting the results.

5 Conclusion

This paper presents a novel framework to achieve a higher-resolution local parameter estimation using a small number of unknowns. This is enabled by a multi-scale optimization, and a spatially-adaptive scheme that allocates higher resolution only at heterogeneous regions. Theoretically, the proposed method has the potential to reach the resolution of mesh. Experiments show that at the current stage, the accuracy is limited around the infarct border. One main future work is to improve the ability to go deeper along the tree, overcoming the issues of computation and observability. Additionally, it is desired to incorporate probabilistic estimation to handle the uncertainties in real data.

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. Chinchapatnam, P., Rhode, K.S., Ginks, M., Rinaldi, C.A., Lambiase, P., Razavi, R., Arridge, S., Sermesant, 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)
2. Clayton, R., Panfilov, A.: A guide to modelling cardiac electrical activity in anatomically detailed ventricles. Prog. Biophys. Mol. Biol. **96**(1), 19–43 (2008)
3. Hastie, T., Tibshirani, R., Friedman, J.: Unsupervised Learning. Springer, New York (2009)
4. Konukoglu, E., Relan, J., Cilingir, U., Menze, B.H., Chinchapatnam, P., Jadidi, A., Cochet, H., Hocini, M., Delingette, H., Jaïs, P., et al.: Efficient probabilistic model personalization integrating uncertainty on data and parameters: application to eikonal-diffusion models in cardiac electrophysiology. Prog. Biophys. Mol. Biol. **107**(1), 134–146 (2011)
5. Lê, M., Delingette, H., Kalpathy-Cramer, J., Gerstner, E.R., Batchelor, T., Unkelbach, J., Ayache, N.: Bayesian personalization of brain tumor growth model. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015, Part II. LNCS, vol. 9350, pp. 424–432. Springer, Heidelberg (2015). doi:10.1007/978-3-319-24571-3_51

6. Powell, M.J.: The bobyqa algorithm for bound constrained optimization without derivatives. Cambridge NA report NA2009/06, University of Cambridge, Cambridge (2009)
7. Shahriari, B., Swersky, K., Wang, Z., Adams, R.P., Freitas, N.: Taking the human out of the loop: a review of bayesian optimization. Proc. IEEE **104**(1), 148–175 (2016)
8. 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)
9. Wong, K.C., Sermesant, M., Rhode, K., Ginks, M., Rinaldi, C.A., Razavi, R., Delingette, H., Ayache, N.: Velocity-based cardiac contractility personalization from images using derivative-free optimization. J. Mech. Behav. Biomed. Mater. **43**, 35–52 (2015)
10. Zettinig, O., et al.: Fast data-driven calibration of a cardiac electrophysiology model from images and ECG. In: Mori, K., Sakuma, I., Sato, Y., Barillot, C., Navab, N. (eds.) MICCAI 2013, Part I. LNCS, vol. 8149, pp. 1–8. Springer, Heidelberg (2013)