

ECG-Based Reconstruction of Heart Position and Orientation with Bayesian Optimization

Jaume Coll-Font, Setareh Ariaifar, Dana H Brooks

SPIRAL Group, ECE Dept. Northeastern University, Boston (MA), USA

Abstract

Respiratory motion is known to cause beat-to-beat variation of the ECG. This observation suggests that it may be possible to use this variation to track position and orientation of the heart. Electrocardiographic Imaging (ECGI) would benefit from such a reconstruction since one contribution to errors in its solutions is respiratory motion of the heart. ECGI solutions generally rely on prior computation of a “forward” model that relates cardiac electrical activity to ECGs. However, the ill-posed nature of the inverse solution leads to large errors in ECGI even for small amounts of error in the forward model. The current work is a first step towards reducing those errors using a nominal forward model and the ECG itself. We describe a method that can reconstruct cardiac position / orientation using known potentials on both the heart and torso. Our current implementation is based on Bayesian Optimization and efficiently optimizes for the position / orientation of the heart to minimize error between measured and forward-computed torso potentials. We evaluated our approach with synthesized torso potentials under a model of respiratory motion and also using potentials recorded in a tank experiment on a canine epicardium and the tank surfaces. Our results show that our method performs accurately in synthetic experiments and can account for part of the error between forward-computed and measured ECGs in the tank experiments.

1. Introduction

Electrocardiographic Imaging (ECGI) is a technology whose objective is to non-invasively image the electrical function of the heart. That is, it reconstructs the electrical potentials on the heart from the body surface potential measurements and a mathematical model of the volume conductor that relates them. Its purpose is to detect abnormal behavior of the myocardial tissue, such as points of re-entry or ectopic beats, and has great potential as a tool for pre-interventional planning in ablation procedures.

ECGI methods rely on the availability of a “forward”

model that relates the electrical activity on the heart to the body surface ECG measurements. Inverting the relationship described by this model provides solutions to ECGI. However, its intrinsic ill-posedness causes large errors in the solutions with even small variations in the measurements or the model itself. In order to provide stable solutions to ECGI, it is necessary to regularize the inverse solutions or reduce the noise present in the measurements and the model. In this work, we show a proof-of-concept approach to reduce effects of model errors that stem from respiratory movement of the heart.

The relationship between electrical activity on the heart and body surface ECG measurements can be approximated using discretized geometries of, at least, the heart and torso, along with the conductivities of the tissues involved [1]. Moreover, when the electrical sources on the heart are characterized by the distribution of potentials on a surface enclosing it, the relation between heart and body surface potentials at any time instant t —which we denote as x_t and y_t respectively—can be modeled as multiplication by a “forward” matrix A :

$$y_t = Ax_t \quad (1)$$

This forward matrix is numerically computed from the heart and torso geometries and thus changes with the movement of the heart. Hence, to characterize the changes in the measured ECG produced by movements of the heart, it is necessary to create a forward model that accounts for the underlying position and orientation of the heart, whose parameters are collected in the variable q , to generate forward matrices $A(q)$. This model can be used to solve the problem of estimating the position of the heart using the electrical measurements on the torso and on the heart:

$$\min_q \sum_{t=1}^T \|y_t - A(q)x_t\|_2^2 \quad (2)$$

Unfortunately, solving this optimization problem can be computationally too expensive due to the cost of numerically computing a forward matrix at every iteration (and in addition running a simulation of the torso potentials when the heart potentials are not known).

There has been little work addressing this challenge. Our group previously presented an algorithm that interpolated the manifold of forward matrices $A(q)$ with a multi-dimensional function and then optimized Equation 2 using this function as a surrogate [2]. In parallel Rodrigo *et. al.* developed a method that uses an alternative measure to determine the true underlying forward model. They used the curvature of the L-curve from Tikhonov regularization to find the geometry that best described the observed data [3].

Here we address this challenge employing a family of optimization methods known as Bayesian Optimization (BO). These black-box optimization methods were designed to optimize objective functions that are too costly to evaluate [4]. They minimize the number of function evaluations needed by choosing the sequentially “best” point at which to evaluate the objective function at each iteration.

Here we present a proof of concept of the application of BO methods to finding the true position of the heart. We assume we know both torso and heart potentials to evaluate our method and report tests on a series of synthetic experiments simulating the movement of the heart due to respiration and evaluate if this method can correct the discrepancies between measured and synthesized potentials in a canine tank experiment.

2. Methods

2.1. Bayesian Optimization

The basic assumptions in BO methods is that the objective function to be optimized ($f(q) : \mathbb{R}^D \rightarrow \mathbb{R}$) is not known and that its evaluations are costly. BO methods maintain a probabilistic model of the objective function as well as a “utility function” (UF) that estimates the “benefit” of evaluating any new point in the domain. The basic procedure of any BO method consists of following these steps iteratively until convergence:

1. Update probability model of objective function with the available samples.
2. Update the UF.
3. Determine the most beneficial point to sample (maximize the UF).
4. Evaluate the objective function at that point.

The key elements in a BO method are the probabilistic model and the UF. Any probability distribution over functions can be used as the probabilistic model. Gaussian Processes (GPs) are the most commonly used priors since they characterize smooth functions and offer easy-to-compute closed form solutions. The UFs are measures of “evaluation benefit” derived from this probability model. There are many utility functions proposed in the literature, including probability of improvement, entropy search, and expected improvement. However the UF that has attracted

most attention is expected improvement (EI)[4]. EI measures the expectation that the function value of a new point will improve with respect to the best point observed (f'):

$$EI(q) = \int_{-\inf}^{f'} (f' - f(q))p(f(q))df(q) \quad (3)$$

A common approach to combining a probabilistic model and a UF is to use a GP with EI. The advantage of this combination is that the equation for the EI simplifies to:

$$EI(q) = (f' - \mu(q))\Phi(f'; \mu(q), \sigma(q)) + \sigma(q)\mathcal{N}(f'; \mu(q), \sigma(q)), \quad (4)$$

where $\mu(q)$ is the mean and $\sigma(q)$ is the variance the GP model and $\Phi(\cdot)$ the cumulative distribution function for a normal distribution. This UF can be cheaply evaluated at every point of the domain q and expressions to compute its gradient are known. One limitation of BO is that all UF are highly non-linear and, when the dimensionality of the problem is high, its optimization becomes computationally challenging in itself. This limitation introduces an inherent trade-off between the computational cost of evaluating the objective function and the time spent optimizing the UF. The general formulation of BO does not provide strategies to find solutions to this trade-off, but it is possible to use domain-specific knowledge about the application to facilitate it.

2.2. Optimizing the Position of the Heart

In this work, we use BO method to reconstruct the position of the heart from ECG measurements. The objective function that we use to solve this problem is Equation 2 and, in practice, it is evaluated as follows:

1. Move the heart to the coordinates $q \in \mathbb{R}^6$ determined by EI (here, 3 coordinates for translation and 3 for rotation).
2. Compute a forward model ($A(q)$) using a numerical method (in this work we used the Boundary Element Method (BEM)).
3. Evaluate the fitting error from Equation 2

Although our objective function is non-linear and non-convex, there is a dominant basin of attraction that contains the global minimum. Thus we take advantage of this characteristic and combine BO with conjugate gradient descent. The conjugate gradient approach reduces the search space dimension at each step from 6D to 1D, and the BO reduces the number of evaluations needed in each iteration.

The complete algorithm, described in **Algorithm 1**, sequentially solves a BO problem in each dimension using a GP model combined with EI. The algorithm is considered to have converged when the best observed heart position q^{opt} is stable between consecutive iterations or a pre-determined maximum number of function evaluations is reached.

Algorithm 1 Heart position / orientation reconstruction algorithm

```

1: Initialize samples  $Q = \{q\}$  and  $Y = \{f(q)\}$ 
2: while  $\sim$ convergence do
3:   for  $d \in [1 \dots D]$ : do
4:      $q_d^{opt} = \min_q Y(q)$ 
5:     Update GP model with samples  $X$  and  $Y$ 
6:     Update EI with  $\mu(q)$  and  $\sigma(q)$  from GP
7:      $q_d^{new} = \min_{q_d} EI(q_d)$ , where  $q_d \in \mathfrak{R}$  is one
       dimension of  $q^{opt}$ 
8:     Update  $d^{th}$  dimension of  $q^{opt}$  with  $q_d^{new}$ 
9:      $X = [QQ^{opt}]$ 
  
```




Figure 1: Left: Relative error between synthesized and true ECG versus the respiration phase in the synthetic experiment. The blue curve shows the relative error between synthesized and true ECG. For comparison, the red curve shows the relative error between the ECG generated with the heart placed at the (uncorrected) at maximum exhale and the true ECG. The yellow curve shows the difference, that is the improvement using BO. Right: Same as on left but showing RMS error of the heart position.

3. Experiments and Results

To evaluate the algorithm we used the geometry and electrical measurements on both the heart and tank surface in a torso-tank experiment. It was originally acquired at the Cardiovascular Research and Training Institute (CVRTI), University of Utah, with applicable IACUC approval. The experiment consisted of the placement of an excised canine heart within a torso shaped tank filled with conducting medium, while electrical recording were taken on the surface of the tank (192 electrodes) and the heart ventricles (247 electrodes). In this work, we used this dataset in two different ways. First, we took the geometries and the heart potentials measured for one heartbeat and generated a series of synthetic ECG examples mimicking the effects of respiration. Second, we used the real electrical measurements on both the heart and tank surface to estimate the unknown position of the heart within the tank.

To simulate a realistic movement of the heart, we used a model of a respiration process obtained in previous work, where we tracked the movement of the heart from a series of MRI scans [5]. With this model, we moved the heart

to positions and orientations that correspond to multiple phases of the respiration cycle. Specifically, we simulated half of a respiratory cycle by positioning the heart at 10 equally spaced respiration phases between maximum inhale and maximum exhale. For each step of the cycle, we computed a forward matrix using the BEM code in the SCIRun software package [6] and multiplied the heart potentials by it to synthesize the corresponding ECG per each beat. Finally, we added white Gaussian noise with sufficient power to create synthetic measurements with $SNR = 30dB$ ¹.

The torso-tank experiments were provided to us with a geometry that included a nominal position of the heart. However, the ECG synthesized with this geometry and that measured on the tank surface have a relative error between 0.4 and 1.3. Our hypothesis is that a significant part of this error is due to a misplacement of the heart in the nominal geometry. To evaluate this hypothesis, we applied the geometry reconstruction method to 35 heartbeats with the heart paced at 6 different locations.

For both synthetic and real experiments, we applied the BO algorithm separately to each heartbeat and estimated the position of the heart. We computed the relative error between simulated and measured ECG and —when the ground truth was available— we also computed the RMS error between all nodes in the reconstructed versus the true hearts.

The relative error in the synthetic experiment is shown in Figure 1(left). The red and blue curves compare the relative error of the ECG generated after moving the heart to the reconstructed position (blue curve) to the error obtained with an ECG generated with the heart fixed at the maximum exhale position (red curve). As expected, when the heart moves with respiration, the reference error increases until reaching its maximum at the maximum inhale position —when the heart is furthest from the fixed exhale position. On the other hand, reconstruction obtained with our method is capable of reducing the error to almost the levels of the SNR introduced with additive noise. These results match the observations of the error in position of the heart shown in Figure 1(right). Again, the blue curve corresponds to the RMSE for the heart reconstructed with our algorithm and the red curve to the RMSE for the heart fixed at the maximum exhale position. The method is capable of finding the position of the heart with an RMSE below 10mm.

To evaluate if this method can account for the errors observed between synthesized ECG and measured on the tank surface, we computed the difference between relative error using the nominal geometry and the geometry after “correcting” the position of the heart. Figure 2 shows

¹Here SNR is computed as $SNR(dB) = 20\log_{10}(\frac{V_{signal}}{V_{noise}})$, where V is variance




Figure 2: Histogram of the differences between the error of the ECG synthesized with the heart fixed at the original position and that of the ECG synthesized with the reconstructed heart position

the histogram of this difference for all 35 measured heartbeats. As expected, all heartbeats show some improvement in ECG fit. In some cases, the change in position of the heart can only account for 0.125 of relative error, which suggests that the nominal placement of the heart was correct. On the other hand, many measured heartbeats show a reduction in relative error that accounts for up to 0.5, indicating that the nominal geometry has considerable error.

4. Discussion

In the synthetic experiments, the position of the heart was correctly estimated with an error below 10mm and the algorithm was capable of explaining up to 0.5 of the relative error observed in the torso-tank experiments. The fact that simple translation and rotation of the heart can account for part of the error in the real experiments is encouraging and needs to be further explored with the other potential sources of error. All these results combined suggest that the position of the heart can be reliably reconstructed from ECG measurements by solving the optimization problem in Equation 2

The computational costs of solving this optimization problem could be daunting for forward models with densely discretized geometries, multiple organs or including cardiac simulations. BO algorithms can overcome this limitation by smartly sampling the objective function. This smart sampling reduces the number of function evaluations needed and, with it, the computational costs of this problem. Moreover, BO algorithms can be adapted to include domain-specific knowledge of the problem being solved and thus further decrease the computational costs at no loss of accuracy.

This work is a proof-of-concept of BO applied to reconstruct the position of the heart. There are various limitations of this study that need further research. The current approach is restricted to experimental settings where

the potentials of the heart are being measured. In the real clinical environment that is not the case and this algorithm should be combined with inverse methods or simulations of cardiac activation. An alternative approach to solving the clinical problem is to apply this optimization method to other objective functions as in [3].

5. Conclusions

Errors in positioning of the heart introduce a considerable amount of noise in ECGI solutions. We presented an algorithm that takes advantage of novel Bayesian Optimization methods to efficiently reconstruct the position of the heart from the ECG measurements. Our approach combines classic optimization methods and BO to improve the computational efficiency while maintaining the accuracy of the solutions. The results obtained suggest that this method will yield accurate solutions in a feasible time.

Acknowledgements

This work was supported in part by NIH grant P41GM103545. The authors would like to thank the CVRTI for providing the data.

References

- [1] MacLeod R, Buist M. The forward problem of electrocardiography. In *Comp. Elec.* Springer. ISBN 9781848820456, 2010; .
- [2] Coll-Font J. Model Based Approaches to Incorporate Recordings of Multiple Heartbeats into the Inverse Problem of Electrocardiography. Ph.D. thesis, Northeastern University, 2016.
- [3] Rodrigo M, *et.al.* Solving inaccuracies in anatomical models for electrocardiographic inverse problem resolution by maximizing reconstruction quality. *IEEE Transactions on Medical Imaging* 2017;25:1–1. ISSN 0278-0062.
- [4] Snoek J, Larochelle H, Adams RP. Practical Bayesian Optimization of Machine Learning Algorithms. In *NIPS*. 2012; 2951–2959.
- [5] Coll-Font J, Brooks DH, Tadmor G. Comparison of low-order 4d dynamical models for cardiac respiratory motion using mri. Ph.D. thesis, Universitat Politècnica de Catalunya, Barcelona, 2011.
- [6] SCI-Institute, 2014. SCIRun: A Scientific Computing Problem Solving Environment, Scientific Computing and Imaging Institute (SCI), Download from: <http://www.scirun.org>.

Address for correspondence:

Jaume Coll-Font - jcollfont@ece.neu.edu,
Spiral group in the ECE dept. at Northeastern University
360 Huntington Ave., Boston, Massachusetts 02115