Published in final edited form as:
Conf Proc IEEE Eng Med Biol Soc. 2011 ; 2011: 3913–3916. doi:10.1109/IEMBS.2011.6090972.

Analysis of the Criteria of Activation-Based Inverse Electrocardiography using Convex Optimization

Burak Erem,
CDSP Center, Dept. of ECE, Northeastern Univ., Boston, MA USA

Peter M. van Dam, and
Donders Institute, Radboud Univ. Nijmegen, The Netherlands

Dana H. Brooks
CDSP Center, Dept. of ECE, Northeastern Univ., Boston, MA USA

Burak Erem: berem@ece.neu.edu; Peter M. van Dam: peter.van.dam@peacs.nl; Dana H. Brooks: brooks@ece.neu.edu

Abstract

In inverse electrocardiography (ECG), the problem of finding activation times on the heart noninvasively from body surface potentials is typically formulated as a nonlinear least squares optimization problem. Current solutions rely on iterative algorithms which are sensitive to the presence of local minima. As a result, improved initialization approaches for this problem have been of considerable interest. However, in experiments conducted on a subject with Wolff-Parkinson-White syndrome, we have observed that there may be a mismatch between favorable solutions of the optimization problem and solutions with the desired physiological characteristics. In this work, we use a method based on a convex optimization framework to explore the solution space and analyze whether the optimization criteria target their intended objective.

I. Introduction

The inverse problem of electrocardiography (ECG) is to estimate source parameters on or in the heart given a geometric and conductivity model of the torso volume and observed electric potentials on the body surface. Activation-based inverse ECG models the functional sources of the heart at any location on the heart surface (synthesized by closing the ventricular endocardium to its epicardium, for example) as having two sequential states: off, then on [1]. One can model the on/off waveform behavior of each source during the QRS complex of a heart beat as a phase-shifted step function, a parameterization that reduces the temporal behavior of each source to a single unknown variable: the activation time [2]. Because this waveform parameterization is nonlinear in the unknown activation time, the inverse problem is typically formulated as a nonlinear least squares (NLLS) minimization problem. It turns out that this problem is not convex and that its objective function tends to have many suboptimal local minima.

This has led to several attempts to incorporate the observed data and prior physiological knowledge into initialization methods, with the belief that local minima found near these initializations are likely to be close to being optimal. An important example is the Fastest Route Algorithm (FRA), an initialization method that employs a simplified wavefront propagation model based on finding the shortest path on a graph. In this method, each edge that connects two nodes of the graph that represents the heart surface is weighted by an assumed propagation velocity [3], [4]. Each node on the surface is considered as a candidate for the site of first activation. An initial earliest node is selected by choosing the resulting wavefront, as determined by the estimated propagation pattern which would follow from a signal triggering activation at that node, whose predicted body surface potentials (BSPs) have the highest correlation with the data. Wavefronts arising from later breakthroughs are then also considered and combined node-wise by retaining the activation time of the first wavefront to arrive [5], [6]. The emphasis of this method is on choosing an initialization with physiologically-plausible propagation behavior that is consistent with the measured BSPs.

More recently, we introduced a method of initialization that reformulates the NLLS problem as an optimization problem with non-convex constraints, relaxes the constraints to be convex, and then solves the new problem for a globally optimal solution. This solution is typically infeasible for the NLLS problem because it does not satisfy the original non-convex constraints, but we suggested in [7] one possible method that solved for the nearest neighbor to the convex relaxation solution in the original non-convex constraint set and subsequently used this point as initialization for the NLLS problem. We found that this method of initialization yielded solutions that were close to ground truth in experiments conducted on data simulated from known activation times and perturbed with pseudorandom "measurement" noise.

However, when conducting experiments with clinical data recorded from a subject with Wolff-Parkinson-White (WPW) syndrome (published previously in [8]–[10]), we observed what appeared to be a mismatch between the optimization criteria and our intended objective. The measured BSPs exhibited typical WPW behavior during the QRS complex, with extra activations initiated from the Kent bundle (whose location was determined invasively). It was previously reported that FRA-initialized candidate solutions to the NLLS problem were able to localize the initial activations [10]. These candidate solutions had activation patterns exhibiting propagation consistent with cardiac electrophysiology. Surprisingly, we were able to find physiologically-inadequate candidate solutions that outperformed the FRA-initialized solutions in terms of objective value. Furthermore, we found that the duration of the upstroke in this optimization setting can have a significant effect on the objective value. In this paper, we use the convex relaxation to explore the solution space and examine whether the NLLS objective function targets the solutions we wish to find.

We begin the rest of the paper by presenting the NLLS problem, its convex relaxation, and briefly explaining the nearest neighbor initialization method in the Background section. In the Methods section, we introduce an optimization procedure that, together with the convex relaxation, we use to explore the solution space of the NLLS problem. In the Experiments section, we report the results we obtained by applying FRA and our newer methods to two clinical WPW datasets. We analyze the results and their implications in the Discussion section below.

II. Background

In this section, we review a framework for the activation-based problem as an optimization problem and use it to show how the original NLLS problem can be equivalently expressed as a non-convex constrained optimization problem. We explain how a convex relaxation can be obtained and one method of converting its global solution into a feasible initialization for the NLLS problem.

For the remainder of the paper we assume that the data is regularly sampled in time and we only consider those samples that correspond to the QRS complex of a single heartbeat. At any given time, the linear relationship between a vector of body surface potentials, $y \in \mathbb{R}^M$, and a vector of on/off sources on the heart, $x \in \mathbb{R}^N$, is $y = Ax$, where A is the forward matrix that results from solving the forward problem on spatially discretized heart and body surface domains. Furthermore, we assume that the waveforms for the sources are unit step functions whose true amplitudes are known (a vector v) and have been multiplicatively absorbed into the linear forward model ($A \leftarrow A\,\mathrm{diag}(v)$) for notational simplicity.

In order to reformulate the NLLS problem as a constrained optimization problem, we define an alternative set of constraints to describe the nonlinearly parameterized waveforms. The original nonlinear parameterization is that every source, x_n, has the waveform

$$x_n(t)=u(t-\tau_n)=\begin{cases} 0 & , t-\tau_n<0 \\ 1 & , t-\tau_n \geq 0 \end{cases}$$

where τ_n is the activation time. Using this parameterization, the original NLLS problem is

$$\text{minimize} \quad \sum_t \left\|y(t)-Ax(t)\right\|_2^2+\lambda\left\|Lx(t)\right\|_2^2$$

where L is a Tikhonov regularization matrix, λ is the regularization parameter, and the optimization variables are the activation times. The Gauss-Newton algorithm and similar nonlinear least squares solvers require that the objective function is differentiable, so a smoothed step function (with a specified upstroke duration) is typically used and the subsequent approximate version of the original NLLS problem is solved instead [2], [11].

If we let QRS correspond to the sample times $t = 1, \dots T$ then we can define a source matrix X that contains all of the temporal samples of each spatial source such that $X_{n,t} = x_n(t)$. Key characteristics of this matrix are that its values are either 0 or 1, are nondecreasing as the column index increases, and that they always increase from 0 to 1 between column indices 1 and T.

Let D be a first-order temporal differencing matrix (i.e. D is $T \times T$ with 1 on the diagonal and −1 on the subdiagonal). If we define the sets $\mathcal{R}$ and $\mathcal{E}$ as

$$\begin{aligned} \mathcal{R}&=\{X \in \mathbb{R}^{(N\times T)} | 0 \leq X \leq 1,\ XD^{\mathrm{T}} \geq 0, XD^{\mathrm{T}} 1_{(T\times 1)}=1_{(N\times 1)}\} \\ \mathcal{E}&=\{X \in \mathbb{R}^{(N\times T)} | \mathrm{tr}(X^{\mathrm{T}}X)=1^{\mathrm{T}}_{(N\times 1)}\, X 1_{(T\times 1)}\} \end{aligned}$$

(where $1_{(i\times j)}$ denotes a $i \times j$ matrix of ones) then $X \in \mathcal{R} \cap \mathcal{E}$. Thus we can express the original NLLS problem as a constrained optimization problem

$$\begin{array}{ll} \text{minimize} & \left\|Y-AX\right\|_{\mathrm{F}}^2+\lambda\left\|LX\right\|_{\mathrm{F}}^2 \\ \text{subject to} & X \in \mathcal{R} \cap \mathcal{E} \end{array}$$

where the optimization variable is the matrix X and $\|\cdot\|_{\mathrm{F}}$ denotes the Frobenius norm. We showed this problem was non-convex in [7] and formulated a convex relaxation by relaxing the domain from $\mathcal{R} \cap \mathcal{E}$ to $\mathcal{R}$. Theoretically, if the solution to the convex relaxation satisfies $X \in \mathcal{E}$, it is the global solution to the non-convex problem as well. In general, the convex relaxation does not solve the original problem because of infeasibility. In this case, the objective value $f(X)$ for the convex relaxation is simply a lower bound on the objective value of the original constrained problem (which follows from $\mathcal{R} \cap \mathcal{E} \subset \mathcal{R}$ and convex f). If the solution is infeasible, one can find its nearest neighbor in the feasible set and use that as initialization for the NLLS problem [7].

III. Methods

In this section we describe an optimization procedure that solves a sequence of alternate NLLS problems that approaches the original NLLS problem. This is done by using a sequence of parameter values that effectively transforms a subset of the convex domain $\mathcal{R}$ into $\mathcal{R} \cap \mathcal{E}$ in its limit.

We start by modeling the values of a source matrix X as samples taken at regular intervals from an underlying continuous-time function. Rather than limit the function for each row to phase-shifted unit step functions, we use an arbitrary function h_n that is non-decreasing with values between 0 and 1. Therefore we have

$$X_{n,t}=x_n(t)=h_n(\beta(t-\tau_n))$$

where β is a time scaling parameter and τ_n is the phase shift. Given these conditions, as β goes from $1 \rightarrow \infty$, X approaches the set $\mathcal{R} \cap \mathcal{E}$. Thus, as an alternate NLLS problem, we minimize the least squares objective function from the original NLLS with this new parameterization (for fixed β and variable τ_n). We note that this alternate NLLS problem and the original NLLS problem have the same convex relaxation.

We initialize the sequence of optimization problems with the convex relaxation, X_c, such that $\beta = \beta_0 = 1$, the initial phase shifts τ_n are the activation times from the nearest neighbor initialization method, and the waveforms h_n are chosen such that $X = X_c$. We use a Gauss-Newton solver to search for a local minimizer, using this solution to initialize the next problem in the sequence with $\beta = \beta_1 > \beta_0$. The procedure continues in this manner until X is within some specified precision of its nearest neighbor in $\mathcal{R} \cap \mathcal{E}$.

At the conclusion of every iteration, we store the minimizing activation times and source matrix. When the sequence terminates, we find the nearest feasible neighbor to all of the stored source matrices and keep the one with the best objective value as the “best” feasible candidate solution. We also keep the final source matrix of the sequence as the “last” candidate solution (not necessarily feasible).

IV. Experiments

For our experiments, we used measured body surface potentials during the QRS complex of beats from a subject diagnosed with Wolff-Parkinson-White (WPW) syndrome. Subjects with WPW have an additional conduction system called the Kent bundle that leads to additional initial activations to those caused by the regular conduction system. In this case, the location of the subject’s Kent bundle was determined invasively. We use two different beats for our experiments: first a fusion beat (i.e. ventricular activation initiated by both the Kent bundle and regular conduction system), and then a Kent-bundle-only beat. The fusion beat is considered a typical WPW beat, whereas the Kent-bundle-only beat was induced by administration of adenosine to block the AV-node. Thus, although exact activation patterns may not be known, much is known about the beats that can be used to evaluate whether candidate solutions to the inverse problem capture the correct behavior.

For each beat, we computed candidate solutions by solving the NLLS problem initialized by FRA, applying the optimization procedure described in Methods for the “last” and “best” candidate solutions, and solving the convex relaxation. Optimization problems were solved for a fixed value of the regularization parameter ($\lambda = 1$) in each case, but the objective values of resulting candidate solutions were compared over a wider range of regularization parameters.

V. Discussion

In this paper, we described the results of a procedure for activation-based inverse electrocardiography as applied to data measured from a patient with WPW syndrome. In these experiments, we explored the solution space of the NLLS minimization problem of activation-based inverse ECG with a number of optimization procedures: FRA-initialized NLLS, a convex relaxation of the NLLS problem, and a sequence of NLLS problems.

The first experiment was conducted on a fusion beat that combined both the Kent bundle and the normal conduction system. The results of the optimization procedures can be seen as isochrone maps in Figure 1. The aforementioned sites of initial activation can be seen in the candidate solution found by FRA initialization. This candidate solution has its earliest activation near the base of the heart where the Kent bundle is known to be located, and a second activation on the epicardium of the other ventricle. Furthermore, the spacing of the isochronal contours for this activation pattern are consistent with the behavior of propagating wavefronts. The second experiment was conducted on an induced Kent-bundle-only beat and its results can be seen in Figure 2. Again, in this case the FRA-initialized candidate solution accurately localizes the site of initial activation and captures the expected behavior of propagating wavefronts.

For the fusion beat experiment, the other three candidate solutions that were obtained using the optimization procedures described in the Background and Methods sections loosely resemble the FRA-initialized activation pattern but have isochronal contours that are inconsistent with propagating wavefronts. In the case of the Kent-bundle-only beat, the candidate solutions for the same three methods find erroneous endocardial sites of initial activation in addition to those of the Kent bundle. In terms of their qualitative physiological characteristics and the degree to which they adhere to the known behavior of the beats in these experiments, these three sets of candidate solutions are clearly unfavorable in comparison to the comparable FRA-initialized candidate solutions.

On the other hand, in Figures 1 & 2, we also plot the objective values for these candidate solutions evaluated for a range of regularization parameters. The green and red curves in this plot correspond to the FRA-initialized candidate solution and its smoothed approximation with a 10ms upstroke duration, respectively. The other three curves correspond to the unfavorable candidate solutions described in the previous paragraph. As the plot shows, the unfavorable candidate solutions have a lower objective value than FRA for all of the regularization parameter values.

In theory, the objective function of an optimization problem serves as a way of ranking the suitability of candidate solutions. The implication of the results we have presented here is that there is a mismatch between the objective of activation-based inverse ECG and the objective function of the corresponding optimization problem. This suggests that the optimization criteria for this problem need to be carefully reviewed and modified, if necessary, to achieve the intended goal. Another conclusion we draw from this work is that the convex relaxation method, which compared favorably to FRA on simulated data, does not compare favorably on this clinical case. We are currently investigating the possible causes of this discrepancy with an eye to improving both methods.

NIH-PA Author Manuscript

NIH-PA Author Manuscript

Acknowledgments

Support for the work of BE and DHB was provided in part by the NIH/NCRR Center for Integrative Biomedical Computing (CIBC), 2P41 RR0112553-12. The authors would like to thank Dr. Gerald Fischer for use of the data.

References

1. Cuppen, J.; Van Oosterom, A. Biomedical Engineering, IEEE Transactions on. 1984. Model Studies with the Inversely Calculated lsochrones of Ventricular Depolarization; p. 652-659.
2. Huiskamp G, Van Oosterom A. The depolarization sequence of the human heart surface computed from measured body surface potentials. IEEE Transactions on Biomedical Engineering. 1988; 35(12):1047–1058. [PubMed: 3220498]
3. van Dam P, van Oosterom A. Volume conductor effects involved in the genesis of the P wave. Europace. 2005; 7(s 2):30–38. [PubMed: 16102501]
4. van Dam P, Oostendorp T, van Oosterom A. Simulating ECG Changes during Acute Myocardial Ischemia. Computers in Cardiology, 2007. 2007:325–328.
5. van Dam P, Oostendorp T, van Oosterom A. Application of the fastest route algorithm in the interactive simulation of the effect of local ischemia on the ECG. Medical and Biological Engineering and Computing. 2009; 47(1):11–20. [PubMed: 18766396]
6. van Dam P, Oostendorp T, Linnenbank A, van Oosterom A. Non-invasive imaging of cardiac activation and recovery. Annals of Biomedical Engineering. 2009:1–18. [PubMed: 18982451]
7. Erem B, van Dam P, Brooks D. A convex relaxation framework for initialization of activation-based inverse electrocardiography. NFSI. 2011
8. Fischer G, Hanser F, Pfeifer B, Seger M, Hintermüller C, Modre R, Tilg B, Trieb T, Berger T, Roithinger F, et al. A Signal Processing Pipeline for Noninvasive Imaging of Ventricular Preexcitation. Methods Inf Med. 2005; 44:508–15. [PubMed: 16342917]
9. Berger T, Fischer G, Pfeifer B, Modre R, Hanser F, Trieb T, Roithinger F, Stuehlinger M, Pachinger O, Tilg B, et al. Single-beat noninvasive imaging of cardiac electrophysiology of ventricular pre-excitation. Journal of the American College of Cardiology. 2006; 48(10):2045–2052. [PubMed: 17112994]
10. van Dam, PM. PhD dissertation. Radboud University; 2010. The shortest path to cardiac activation.
11. Modre R, Tilg B, Fischer G, Wach P. Noninvasive myocardial activation time imaging: a novel inverse algorithm applied to clinical ECG mapping data. Biomedical Engineering, IEEE Transactions on. 2002; 49(10):1153–1161.

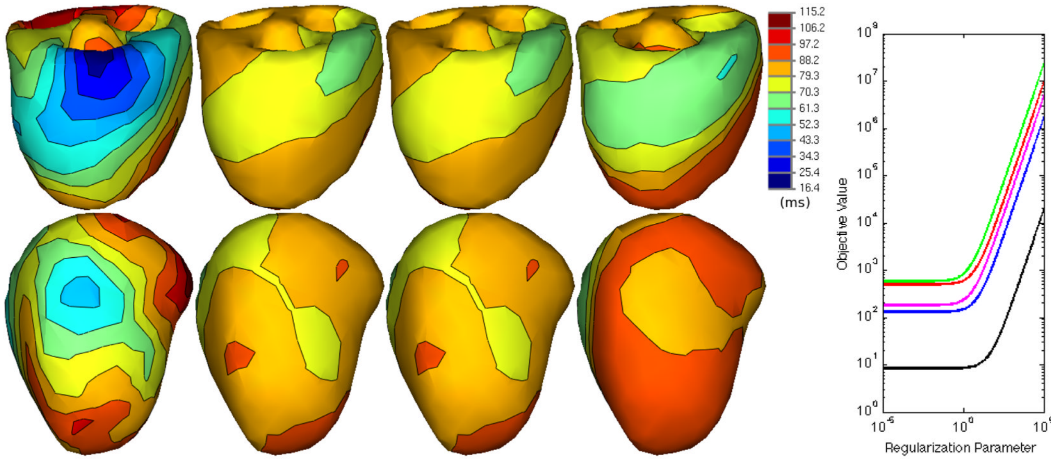

Fig. 1.

Isochronal activation time maps of candidate solutions (left to right: FRA-initialized, “last”, “best”, and convex relaxation) for the WPW fusion heart beat, rows showing alternate views. Plot of objective values (black=convex relaxation, blue=“last”, magenta=“best”, red=smooth FRA-initialized (upstroke=10ms), green=FRA-initialized) for a range of regularization parameters.

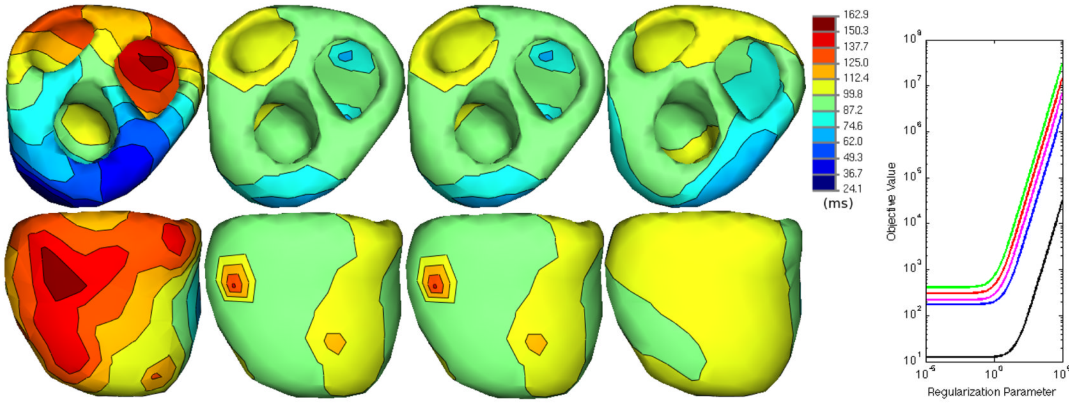

Fig. 2.

Isochronal activation time maps of candidate solutions (left to right: FRA-initialized, "last", "best", and convex relaxation) for the WPW Kent-bundle-only heart beat, with rows showing alternate views. Plot of objective values for a range of regularization parameters (colors of lines same as in Fig. 1).