Identifying Model Inaccuracies and Solution Uncertainties in Noninvasive Activation-Based Imaging of Cardiac Excitation Using Convex Relaxation

Burak Erem*, *Member, IEEE*, Peter M. van Dam, and Dana H. Brooks, *Senior Member, IEEE*

***Abstract*—Noninvasive imaging of cardiac electrical function has begun to move towards clinical adoption. Here, we consider one common formulation of the problem, in which the goal is to estimate the spatial distribution of electrical activation times during a cardiac cycle. We address the challenge of understanding the robustness and uncertainty of solutions to this formulation. This formulation poses a nonconvex, nonlinear least squares optimization problem. We show that it can be relaxed to be convex, at the cost of some degree of physiological realism of the solution set, and that this relaxation can be used as a framework to study model inaccuracy and solution uncertainty. We present two examples, one using data from a healthy human subject and the other synthesized with the ECGSIM software package. In the first case, we consider uncertainty in the initial guess and regularization parameter. In the second case, we mimic the presence of an ischemic zone in the heart in a way which violates a model assumption. We show that the convex relaxation allows understanding of spatial distribution of parameter sensitivity in the first case, and identification of model violation in the second.**

***Index Terms*—Activation time imaging, biomedical imaging, biomedical signal processing, convex relaxation, electrocardiography (ECG), inverse problems.**

I. Introduction

AFTER many years of research interest, imaging cardiac electrical function from noninvasive measurements of electric potentials on the body surface, via solutions of the inverse problem of electrocardiography (ECG), has recently begun to move towards clinical adoption [1]–[3][1]. Clinical utility would be greatly enhanced by an understanding of the robustness of these solutions. We address the problem of uncertainty in this work. Specifically we treat one of the two most common formulations of the inverse ECG problem, activation-based inverse ECG, in which cardiac electrical function during the QRS complex is parameterized by the activation times of equivalent electrical sources. This formulation requires that we numerically solve a nonconvex optimization problem that has many suboptimal local solutions. In this paper we describe a theoretical framework to relax the problem to a convex one. We then illustrate that this convex relaxation allows us to study parameter sensitivity and model inaccuracy in the original problem. We want to emphasize that we are not presenting here a new and improved inverse solution method. As will become clear below, the method we present finds solutions that are either not guaranteed (and are very unlikely) to be valid solutions in the sense of the activation-based problem formulation, or when modified to be valid, are not optimal solutions in any meaningful way. Our goal rather is to provide a tool which can help evaluate confidence in results of an important existing method, both in terms of spatial location on the heart surface and in terms of presence of possible pathologic conditions which violate the underlying method's basic assumption.

As is described in the sequel, the assumptions behind this method and the nonlinear and nonconvex optimization it requires make uncertainty quantification both particularly important and particularly challenging. We consider here sensitivity to two types of model assumptions. One we address in the context of measured data obtained from a human subject, and the other by means of simulated measurements and models, using the ECGSIM software package [5].

In activation-based ECG, the reconstructed sources are typically distributed spatially on a model heart surface comprising a union of the endocardial (inner) and epicardial (outer) surfaces of (typically) the ventricles [2], [6], [7]. (A 3-D formulation has also been reported [8]–[10] but here we only treat the more common surface-based formulation.) The physiological assumption underlying this parameterization is that the sub-

Manuscript received October 22, 2013; revised December 26, 2013; accepted December 30, 2013. Date of publication January 09, 2014; date of current version March 31, 2014. This work was supported in part by the National Center for Research Resources under Grant 5P41RR012553–14 and in part by the National Institute of General Medical Sciences under Grant 8 P41 GM103545–14 from the National Institutes of Health. *Asterisk indicates corresponding author.*

*B. Erem is with the Department of Radiology, Children's Hospital Boston/Harvard Medical School, Boston, MA 02115 USA (e-mail: burak.erem@childrens.harvard.edu).

P. M. van Dam is with the Donders Institute for Brain Cognition Behaviour, Radboud University Medical Center, 6500 HB Nijmegen, The Netherlands (e-mail: peter.m.vandam@xs4all.nl).

D. H. Brooks is with the Department of Electrical/Computer Engineering, Northeastern University, Boston, MA 02115 USA (e-mail: brooks@ece.neu.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TMI.2014.2297952

[1]We note that a clinical device that images the electrical activity of the *inner* (endocardial) heart surface from measurements with a catheter-based probe introduced into the chamber has been in use for years [4]. We are not directly concerned with such systems here.

ject's transmembrane potential (TMP) waveform morphology is essentially uniform across this heart surface, and that propagation can be encoded in differences in TMP phase, referred to as activation time. The advantage of activation-based formulations is that the number of unknown parameters to be estimated is reduced to one per modeled heart surface location, and that this reduction is achieved by imposition of a strong electrophysiologically-inspired constraint. It has been reported to lead to greater robustness to geometric errors and measurement noise in the face of the sensitivity of this notoriously ill-posed problem [7]. The tradeoff is the uncertainty introduced by these model assumptions. Its robustness to one key assumption was discussed in [11]. Efforts to increase clinical acceptance of this technology could be greatly aided by a method of identifying and characterizing the uncertainty in a given solution. There has been other work in characterizing uncertainty in cardiac electrophysiology [12], [13], but to the best of our knowledge, there have not been any published attempts to characterize the uncertainty in inverse activation time estimates. The ability to characterize the spatial distribution of this uncertainty could be especially useful, so that clinicians could apply inverse solution results to guide diagnosis or treatment with knowledge of where on the myocardial surface these results can be expected to be more reliable and where they might be more uncertain. Here, we present a general method for exploring this robustness and give two specific examples.

In particular, as a consequence of its underlying model assumptions, the activation-based formulation leads to a nonlinear, and nonconvex, optimization problem, which is solved iteratively. These iterative solutions have been found to depend strongly on a suitable initialization [2]. The existence of many local minima and the resulting sensitivity to initialization complicates achieving the goal of spatial uncertainty quantification. In this work, we address the difficulties posed by the local minima by formulating an inverse problem which is derived from, and closely related to, the original problem but which is convex and has a single global minimum. In particular we present a convex relaxation of activation-based inverse ECG obtained by relaxation of a single constraint. The relaxation has a clear interpretation in terms of the electrophysiological assumptions of the activation-based model, and its solution is closely related to that of the original problem, but guaranteed to be a unique and independent of initialization. We illustrate how this convex relaxation can be used as a framework to explore model inaccuracy and solution uncertainty.

More specifically, the two model assumptions we treat are typical of activation-based methods. To describe these particular assumptions, first we quickly summarize some aspects of the forward model upon which inverse solutions depend; more mathematical detail is given below in Section II. The forward model for activation-based formulations is an explicit mapping of timing parameters (wavefront arrival, i.e., activation time) to body surface potentials. This mapping is generally constructed in two steps: first activation time parameters are mapped to time-shifted fixed-morphology TMPs, then the TMPs are mapped to body surface potentials. Since the first of these mappings is inherently nonlinear, iterative solutions typically minimize a regularized nonlinear least-squares (NLLS) fit of the activation time parameters to the measured body surface potentials, mediated by the forward model, which "predicts" body surface potentials from a given candidate set of activation times. The regularization involved typically requires the choice of a scalar regularization parameter.

NLLS algorithms typically depend on an initialization, and for the problem at hand it has been found that the initialization strongly influences the specific locally-optimal minimum to which the iterations converge. Thus, there has been considerable recent effort towards devising more effective initializations [2], [14]–[18]. The fastest route algorithm (FRA) [2] is one such initialization method that has been reported to be particularly effective. It chooses an initialization from a collection of candidate activation patterns that simulate the spread of activation wavefronts from one or more foci, based on the similarity of their predicted body surface potentials to the measured ones. Like all such initialization algorithms, FRA depends on some parameter choices; for FRA this includes assumptions about propagation speed in the myocardium. In this work, we look at the spatial distribution of solution sensitivity across the heart surface to this parameter choice. We show that our convex relaxation approach predicts and illuminates this sensitivity. We compare this sensitivity to sensitivity of the convex solution to regularization parameter choice to elucidate these relationships. Thus, this example is a study of parameter robustness.

A second model assumption inherent in the activation-based method is the uniformity of TMP morphologies. In settings where pathology may be present, however, regions of the heart with tissue damage or genetic abnormalities may have TMPs with different (typically lower) amplitudes, for example due to ischemia, Brugada syndrome, or arrhythmogenic right ventricular dysplasia/cardiomyopathy [18]–[24]. When available for segmentation, magnetic resonance imaging (MRI) or X-ray computed tomography (CT) images of the heart may allow identification of such regions, but this imaging may not always be available, or it may not be possible to determine the true amplitudes, or such regions may be missed in the images or appear only after the time of imaging. It turns out that our convex relaxation specifically relaxes this TMP uniformity assumption. As a result, as we show, it can lead to direct identification of regions with lower TMP amplitude, and thus to *localization* of model error in the activation-based results. Thus, this second part of our study looks at detection of violations of model assumptions.

We derive the convex relaxation by first reformulating the NLLS problem in a way that allows us to identify mathematically the precise source of the nonconvexity. Indeed we show that there is a single nonconvex constraint implicit in the method. This identification leads to a simple relaxation of that constraint, resulting in a "close-by" convex problem which can be easily solved using modern convex optimization algorithms, leading to a unique minimizer of the relaxed problem. We then suggest a simple method to find a "feasible" solution to the original problem (one that obeys the original nonconvex constraint) that is close to the global minimum of the relaxed problem. We explore the relationships between the global minimum, the nearby feasible solution, the FRA initialization, and the resulting NLLS solution, by varying first one and then the other of the two parameters described above—propagation speed in

FRA initialization and local TMP amplitude. We also look at sensitivity of the convex relaxation to regularization parameter choice, which helps us to uncover solution uncertainties.

The optimization problem we propose in this work takes a similar form to the one introduced by Messnarz *et al.* for TMP imaging [25]. Although the authors of that work motivated their approach as an attempt to ease the restrictions of the activation-based inverse problem, they did not elaborate on an explicit mathematical relationship between the two problems. Since here we derive the convex relaxation directly from the activation-based problem, we are able to study the activation-based problem directly.

In Section II, we review the activation-based inverse problem and its initialization with the FRA. We also show how the NLLS problem can be restated within a constrained optimization framework and that it has a single nonconvex constraint. In Section III, we use that framework to formulate a relaxation of the NLLS problem with a convex objective function and convex constraints, and then present one approach for extracting activation times from its solution. Experiments are presented in Section IV, along with a brief description of results, which are then discussed in greater detail in Section V.

II. Background

In this section, we review how the activation-based inverse is posed as an unconstrained optimization problem. We establish our framework by showing how the original NLLS problem can be equivalently expressed as a constrained optimization problem. We then use this framework to show that the optimization problem is nonconvex, leading to the convex relaxation proposed in Section III.

For the remainder of the paper we assume that our electrocardiographic measurements are regularly sampled in time and we only consider those samples that correspond to the QRS complex of a single heartbeat. At any of these sample times, the linear relationship between a vector of body surface potentials, $y \in \mathbb{R}^M$, and a vector of TMP sources on the heart, $x \in \mathbb{R}^N$, is $y = Ax$, where A is the forward matrix that results from assuming that the volume conductor model of both the torso and myocardium are homogeneous and isotropic in conductivity [26], [27], and then solving a quasi-static approximation of the linear part of the forward problem on spatially discretized heart and body surface domains [16], [17], [28]. When TMP amplitudes are assumed to be known (represented here by a vector v), we assume that these have been multiplicatively absorbed into the linear forward model by the substitution $A \leftarrow A\mathrm{diag}(v)$ (where $\mathrm{diag}(v)$ places the elements of the vector v in the main diagonal of a square matrix). As a result, all "normal" TMP waveforms are assumed to have unit amplitude.

A. Constrained Optimization Problem

To reformulate the NLLS problem as a constrained optimization problem, first we state the NLLS problem itself, and then reformulate it in terms of an alternative set of constraints to describe the nonlinearly parameterized waveforms. To start, a discrete time unit step function, $u(t)$, is defined piecewise as

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

The nonlinear TMP source parameterization is such that every equivalent source, x_n, has the waveform

$$x_n(t) = u(t - \tau_n)$$

where τ_n is the activation time. We use $x(t;\tau)$ to denote the parameterization of TMP sources $x(t) = [x_1(t) \ \cdots \ x_N(t)]^{\mathrm{T}}$ by activation times $\tau = [\tau_1 \ \cdots \ \tau_N]^{\mathrm{T}}$. Then the original NLLS problem is

$$\underset{\tau}{\text{minimize}} \ \sum_t \|y(t) - Ax(t;\tau)\|_2^2 + \lambda \|Lx(t;\tau)\|_2^2$$

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 width for the smoothed step transition) is typically used and the resulting approximate version of the original NLLS problem is solved instead [11], [29].

We can also define the set of discrete-time step functions by mathematically describing the set of constraints that they must obey. To accomplish this, let QRS correspond to the sample times $t = 1, \ldots, T$, and define a source matrix X that contains all of the temporal samples of each TMP, such that $X_{n,t} = x_n(t)$. Key characteristics of this matrix, which reduce its degrees of freedom to specification of the activation times, are that its element values are either 0 or 1, are nondecreasing as the column index increases, are 0 in the first column, and are 1 in column T.

To define constraints for the reformulated optimization problem, we create sets whose elements must obey these key characteristics. In anticipation of the convex relaxation that we will describe in the sequel, let us define the two sets, $\mathcal{R}$ and $\mathcal{E}$, such that $\mathcal{R}$ will be known as the relaxed constraint set, and $\mathcal{R} \cap \mathcal{E}$ will be known as the full or exact constraint set. Let D be a first-order temporal differencing matrix (i.e., D is $T \times T$ with 1s on the diagonal and -1's on the subdiagonal). If we define the sets $\mathcal{R}$ and $\mathcal{E}$ as

$$\mathcal{R} = \Big\{ X \in \mathbb{R}^{(N\times T)} \mid 0 \leq X \leq 1, \ XD^{\mathrm{T}} \geq 0, \ XD^{\mathrm{T}} 1_{(T\times 1)} = 1_{(N\times 1)} \Big\}$$

$$\mathcal{E} = \{ X \in \mathbb{R}^{(N\times T)} \mid \mathrm{tr}(X^{\mathrm{T}}X) = 1_{(N\times 1)}^{\mathrm{T}} X 1_{(T\times 1)} \}$$

(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

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

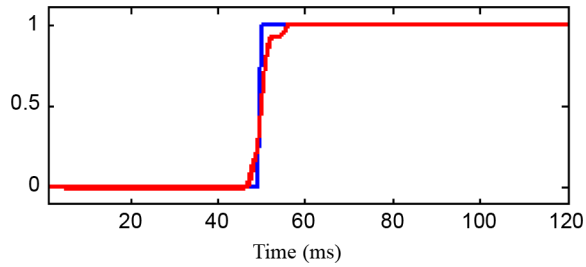

Fig. 1. Example TMP waveforms: the NLLS problem can be equivalently restated as a constrained optimization problem whose constraint set, $\mathcal{R} \cap \mathcal{E}$, consists of TMPs resembling the blue waveform shown here. On the other hand, an example of a similar waveform in the set $\mathcal{R}$ (but *not* in $\mathcal{R} \cap \mathcal{E}$) is shown in red. In general, waveforms in $\mathcal{R}$ may increase from 0 to 1 more gradually than those in $\mathcal{R} \cap \mathcal{E}$.

where the optimization variable is the matrix X and $\|\cdot\|_{\mathrm{F}}$ denotes the Frobenius norm. As shown in the Appendix, with these constraints, the only feasible waveforms have the shape of a TMP and thus the only free variables are the timing of their $0-1$ transition, that is the activation times. In other words, the intersection of $\mathcal{R}$ and $\mathcal{E}$ contains only valid TMP waveforms, and all valid TMP waveforms are in that same intersection. In Fig. 1, we show an example of one such waveform from the exact constraint set, i.e., $X \in \mathcal{R} \cap \mathcal{E}$, in blue, and, to illustrate the relationship between the sets, an example from the relaxed constraint set, $\mathcal{R}$, in red. The red curve is in fact the result of solving the activation based problem over a domain consisting of only the set $\mathcal{R}$, as we describe in the sequel. Note that it has all the characteristics of a valid model TMP except the jump discontinuity. The set $\mathcal{E}$ imposes that constraint, so that its intersection with $\mathcal{R}$ contains only TMP waveforms like the blue curve in Fig. 1.

B. Nonconvex Optimization Problem

Now we can show precisely how the constrained optimization problem in the previous section is nonconvex. A convex optimization problem is one that can be expressed as the minimization of a convex function whose domain is a convex set [30]. The objective function in our problem

$$f(X) = \|Y - AX\|_{\mathrm{F}}^2 + \lambda \|LX\|_{\mathrm{F}}^2$$

is the sum of two convex quadratic functions, which is also convex. The intersection of the constraints of the optimization problem define the domain of the objective function. In the problem at hand, the set $\mathcal{R}$ is convex because it is defined as the nonempty intersection of linear equality and inequality constraints which are all convex. However, the set $\mathcal{E}$ is defined by a nonconvex quadratic equality constraint and the constraint set, $\mathcal{R} \cap \mathcal{E}$, is discrete and therefore also nonconvex [30]. Consequently, the constrained optimization problem, which is equivalent to the NLLS problem, is nonconvex.

C. Fastest Route Algorithm

The FRA is an initialization method for the NLLS problem which we use in our experiments to explain the results of the robustness analysis method proposed in the sequel [18]. The FRA produces an initialization for the NLLS problem by synthesizing wavefront propagation patterns originating from one or more foci and choosing the most suitable one. We only describe the most basic form of the method here, which is sufficient to understand the algorithm's dependence on the parameters which we will vary in Section IV, and refer the reader to the literature for a more extensive description [2].

Given a tessellated surface representation for the epicardial and endocardial surfaces, the FRA forms a graph data structure whose vertices are the corners of the triangles in the surface, with vertices connected by edges if they are part of the same triangle, and whose edge weights are the Euclidean distances between connected triangle corners. The FRA supplements the set of edges of this graph with transmural edges that connect vertices on the epicardial surface to nearby vertices on the endocardial surface. Existing edge weights are then replaced by "propagation times," computed by dividing the edge distances by the propagation velocities along each edge. These propagation velocities are a set of parameters that must be specified *a priori*. The shortest path algorithm is then applied to the graph, which computes the minimum travel time between every pair of vertices in the graph, effectively yielding activation times for a set of candidate propagation patterns, each originating from one of the vertices in the tessellated surface.

In the most basic form of the algorithm, the forward problem is solved for each such propagation pattern and its correlation with measured body surface potentials is computed. The FRA then chooses as its initialization the activation times for the propagation pattern with the highest correlation. In the full algorithm, the set of propagation patterns is extended to simulate propagation originating from more than one site.

III. Convex Relaxation of the Activation-based Formulation

We formulate a convex relaxation of the original optimization problem by relaxing the domain of the problem to a convex set. Once the resulting convex optimization is solved, we will have found a unique global solution to the relaxed problem but, in general, it will not be feasible according to the original problem specification. In other words, it will not be a set of valid model TMP waveforms. However, one can always look for a set of valid waveforms which is "close" to that solution. We describe one approach to find such a feasible solution, which we will use in the sequel. This approach may be used to extract activation times or, if desired, used to initialize the original NLLS problem.

A. Convex Relaxation

A convex relaxation of an optimization problem removes nonconvex constraints or terms in the objective function or approximates them with convex ones [30]. In Section II-B, we narrowed the nonconvex part of the activation-based inverse problem to a single equality constraint. Our convex relaxation of this problem simply removes the restriction to the nonconvex constraint, $\mathcal{E}$, and keeps the objective function, $f(X)$, along with the remaining convex constraint set, $\mathcal{R}$. Therefore, the convex relaxation is

$$\begin{aligned} \text{minimize} \quad & f(X) \\ \text{subject to} \quad & X \in \mathcal{R}. \end{aligned}$$

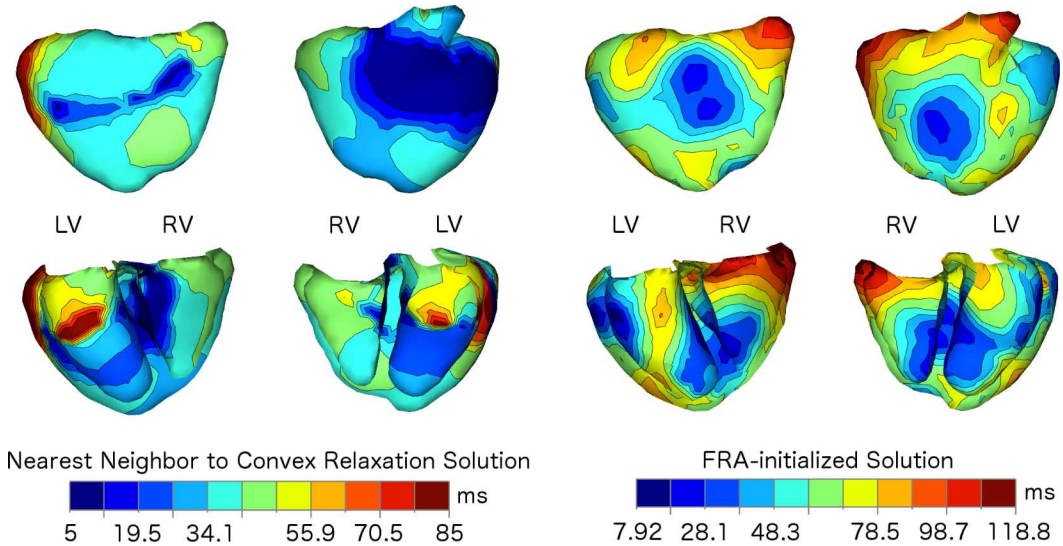

Fig. 2. Healthy male subject: nearest neighbor activation times obtained from the convex relaxation solution (left), and activation-based inverse solution obtained by initializing with the fastest route algorithm (right). Each column shows an isochronal map of the full ventricular heart surfaces (top row) along with a cutaway view to better show the endocardial surfaces (bottom row). We show each solution from two views, rotated approximately 180°, and have labeled the locations of the LV and RV, respectively, to indicate the orientation of the heart in each view. Color shows activation time relative to the start of the QRS complex.

This problem can be solved globally by any number of numerical optimization methods regardless of initialization. For this work, we have developed our own custom solver based on the alternating directions method of multipliers (ADMM) [31] which we have validated against the well-known CVX modeling environment [32]. We note that, in our experience, our method is two orders of magnitude faster than straightforward application of CVX for this problem. A detailed convergence proof and analysis of the ADMM can be found in [31].

Theoretically, if the solution X^* to the convex relaxation is feasible for the original problem (i.e., it satisfies $X^* \in \mathcal{E}$), it is the global solution to the nonconvex problem as well. However, as noted, in general this is not the case, and the value of the objective function, $f(X^*)$ for the convex relaxation is instead a lower bound on the objective function values over the nonconvex constraint set, $\mathcal{R} \cap \mathcal{E}$. That is, if X^* is infeasible, the objective value for any feasible point in the nonconvex set is guaranteed to be greater than $f(X^*)$.

B. Nearest Feasible Neighbor to Convex Relaxation Solution

When the convex relaxation solution is infeasible for the original problem, we can quantify its infeasibility by measuring the degree to which the nonconvex constraint has been violated. Going one step further, we can apply the violated constraint after solving the convex relaxation. This can be done, for example, by finding the nearest neighbor to the convex relaxation solution, X^*, in the original feasible set, $\mathcal{R} \cap \mathcal{E}$. Such an approach does not guarantee an optimal solution to the NLLS problem, in general, but provides a point of comparison for solutions of the original NLLS problem.

The nearest feasible neighbor is found by solving

$$\begin{aligned} \underset{X}{\text{minimize}} \quad & \|X^* - X\|_{\text{F}}^2 \\ \text{subject to} \quad & X \in \mathcal{R} \cap \mathcal{E}. \end{aligned}$$

This problem is row-wise separable, meaning that we can solve for each row of X independently. In other words, we can find the closest valid TMP to the convex solution at each heart surface node separately. Doing so involves testing a small set of possible waveforms to find the minimizer. Specifically, for each row of X, we consider each of the T possible step functions shifted to sample times and choose the one that minimizes the sum of squared differences from the corresponding row of X^*. Then the activation time for each row is the phase shift from which the minimizing waveform was generated.

IV. Experiments

In this section, we describe our experiments on both measured and simulated electrocardiographic data and present results. Discussion of those results is deferred to the following section. First, we present a case of data measured from the body surface of a healthy male subject, whose heartbeats are expected to be normal. Second, we consider a synthetic case created using ECGSIM where the model assumptions of the original NLLS problem have been violated by introducing heterogeneous TMP waveform amplitudes.

A. Healthy Male Subject

Our first experiment uses data recorded from a healthy male, 40 years old, whose heartbeats are expected to be normal. Body surface potentials were measured from 65 electrodes, sampled at a rate of 1000 samples/s. Segmentations of MRI images of the subject's torso were used to discretize the surfaces of the heart, lungs, and torso, and numerically solve the linear part of the forward problem using a boundary element method [33]. We first computed nominal solutions, that is solutions using a "standard" set of parameters used to estimate normal activation in previous studies, as specified in the next few sentences, to both the FRA-initialized NLLS problem and the nearest-neighbor

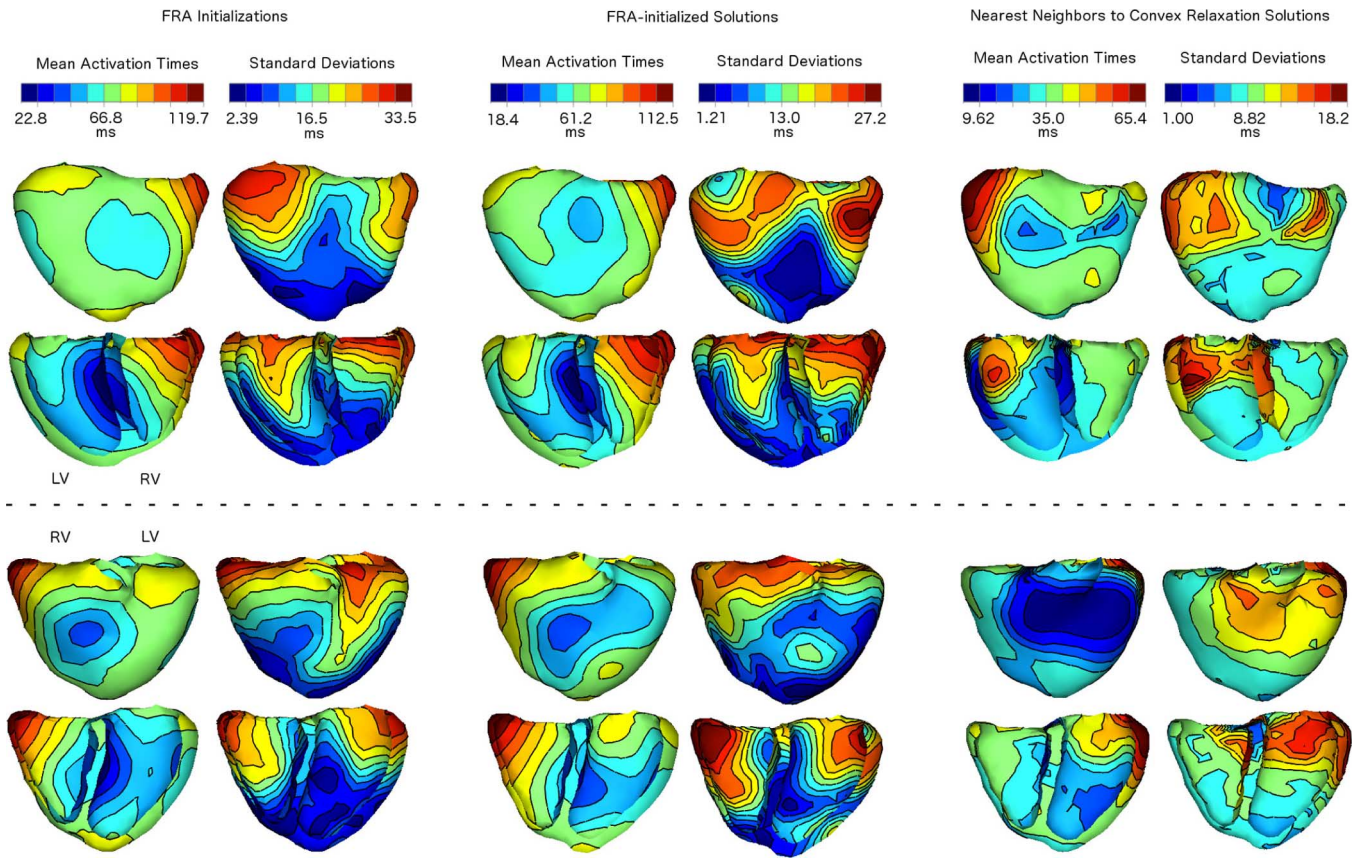

Fig. 3. Isochronal maps of means and standard deviations of activation times: for a number of assumed surface propagation velocities, 0.6–1.4 m/s, the FRA was used to initialize and then solve the NLLS optimization problem. The mean and standard deviation of the resulting sets of activation times (FRA initializations and FRA-initialized solutions) for each spatial location were computed. The spatial distribution of these means and standard deviations are shown as isochronal maps (columns 1 and 2: FRA initializations, columns 3 and 4: FRA-initialized solutions). Additionally, the convex relaxation was solved for a number of regularization parameters, $\lambda = 0.1 \pm 0.01$, and nearest neighbor activation times were computed. The means and standard deviations of these activation times were computed and are also shown as isochronal maps in columns 5 and 6. The top two rows of isochronal maps show the ventricles from a view that is rotated approximately 180° from the view shown in the bottom two rows. Labels indicate the orientation of the LV and RV above and below the dashed line. Rows 1 and 3 show isochronal maps of the epicardial surface, and rows 2 and 4 show cutaway views of the endocardial surfaces.

convex relaxation problem. Specifically, for FRA we used typical parameter values for propagation speed: 0.8 m/s for surface propagation and 40% of that value, 0.32 m/s, for transmural propagation (i.e., in the direction normal to the surface) [2]. The convex relaxation was solved using our custom solver and activation times were obtained using the method in Section III-B. Both methods used $\lambda = 0.1$ for the regularization parameter. In Fig. 2, we show isochronal maps of the activation times resulting from both methods. We show four maps for each case: the top two show epicardial surfaces and the bottom two a cutaway view revealing the chamber walls. For each surface we show two views, rotated approximately 180° from each other, so that both left ventricle (LV) and right ventricle (RV) epicardial freewalls, and the corresponding chamber walls, can be seen. The LV and RV are noted on the plots for greater clarity. Color shows activation time relative to the start of the QRS complex, which was determined manually from the waveform of the rms values of the body surface potential maps.

To explore the dependence of the NLLS problem on initializations, we computed nine different FRA initializations with surface propagation velocities chosen in increments of 0.1 m/s from the interval $[0.6, 1.4]$ m/s, while keeping the surface/transmural velocity ratio constant. The means and standard deviations of the activation times of the equivalent sources were computed over the set of FRA initializations, and are shown as isochronal maps in the first two columns of Fig. 3. Each initialization was supplied to the NLLS solver to find its corresponding local minimum, and again means and standard deviations over the initializations were computed, as shown in the middle two columns of Fig. 3. As in Fig. 2, every epicardial view is accompanied by a cutaway view below it. Results shown above and below the dashed line are rotated approximately 180° from each other, with the orientation of the heart indicated by the placement of the "LV" and "RV" labels.

To induce variability in the convex relaxation solutions, we computed solutions using 50 regularization parameters regularly sampled from the interval $[0.09, 0.11]$, centered on the nominal value of 0.1 used initially and also in all NLLS solutions. Nearest neighbor activation times were computed from these solutions, and maps of their means and standard deviations over the regularization parameter values are shown in the last two columns of Fig. 3. We also computed NLLS solutions using each of those regularization parameters, all starting from the same FRA initialization (using 0.8 m/s for surface propagation and 0.32 m/s for transmural propagation), and computed their means and standard deviations. In this case, of the 1200

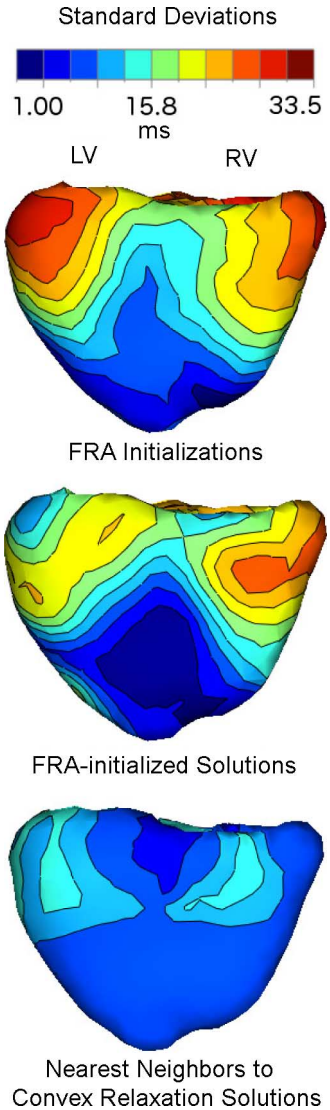

Fig. 4. Standard deviations of activation times with a global colormap: standard deviation maps from the first row of Fig. 3 are shown here again, but rescaled to a global colormap.

modeled sources on the heart surfaces, 1198 of them had standard deviations below 0.04 ms, and the remaining two had standard deviations of 0.83 and 3.8 ms. Maps of these results were omitted from Fig. 3 because these values were much smaller than the other induced variabilities and thus did not provide any insights into the sensitivity of the problem to this type of perturbation.

Finally to emphasize the difference in scales between the different standard deviation maps, we have selected a subset of the standard deviation maps to show in a global colormap in Fig. 4.

B. Synthesized Low-Amplitude Transmembrane Potentials

Our second case consists of simulated electrocardiographic data using a computerized model of a 22-year-old healthy male, distributed with the ECGSIM software package [5]. We reduced the TMP amplitude parameter of this model by 40% of the default values in a region on the epicardial surface of the right ventricle near the apex. This region can be clearly seen in the "Ground Truth" isopotential map in Fig. 5, which shows the end of QRS when all sources have activated. The goal was to mimic one electrocardiographic effect of an ischemic zone [34]. Obviously the heterogeneity of the TMP waveform amplitudes after this modification violates one of the default assumptions of the original NLLS inverse problem formulation. Using the forward model provided with ECGSIM, we computed potentials at 65 body surface electrodes from the TMPs and added pseudorandom noise with an average signal-to-noise ratio of 30 dB with reference to the QRS complex. The convex relaxation was solved using this data and then compared to the ground truth from which it was synthesized. Fig. 5 shows isopotential maps of the convex relaxation solution and ground truth at the same sample time ($\mathrm{t} = 103$ ms) at the end of the QRS complex as well as examples of normal (blue, left) and low-amplitude (red, right) waveforms from the ground truth and the convex relaxation solution. We also designed a simple classifier of "low-amplitude" and "normal" sources for this particular experiment by thresholding the reconstructed source amplitudes of the convex relaxation solution at $\mathrm{t} = 103$ ms, the same time as the snapshot shown in Fig. 5. We tested thresholds in the interval, $[0, 1]$, of normalized TMP values, and compared the classification results with the ground truth on a source-by-source (or, equivalently, node-by-node) basis. The results of using this approach as a "low-amplitude" source detector over the range of possible thresholds are shown in Fig. 6 as a receiver operator characteristic (ROC) curve, which shows the performance trade-off in terms of the true and false positive rates.

V. Discussion

Our results suggest that solutions to the convex relaxation formulation are closely related to those of the original NLLS problem and can be used to study its solution uncertainties and to identify violations of its underlying model assumptions.

For example, in Section IV-A, we compared the NLLS solution to the convex relaxation in the case of a healthy male subject who can be expected to exhibit normal sinus rhythm propagation. In Fig. 2, we show isochronal maps of the activation times extracted from the convex relaxation solution, by finding its nearest feasible neighbor, along with the NLLS solution, initialized by the FRA, where both methods used the same regularization parameter ($\lambda = 0.1$). Both the nearest neighbor to the convex relaxation solution and NLLS solution showed earliest activation sites on the endocardial surface of the septum in the LV, closely followed by the RV septum and the endocardial LV free wall (see bottom row of isopotential maps in Fig. 2). These sites are in agreement with those reported in the classic study of Durrer *et al.* [35]. In addition, both methods find early epicardial breakthroughs, but they differ substantially in their sizes and shapes. Although the NLLS and nearest neighbor convex relaxation solutions are not the same, their similarities suggest that they are mathematically "close" and therefore also close to the convex relaxation solution.

We then considered the dependence of this result on the sensitivity of the FRA initialization to assumed propagation velocities. To try to separate sensitivity of the initializations themselves from sensitivity of the NLLS algorithm to initialization, we presented the activation time means and standard deviations of both the initializations (Fig. 3, columns 1 and 2, respectively)

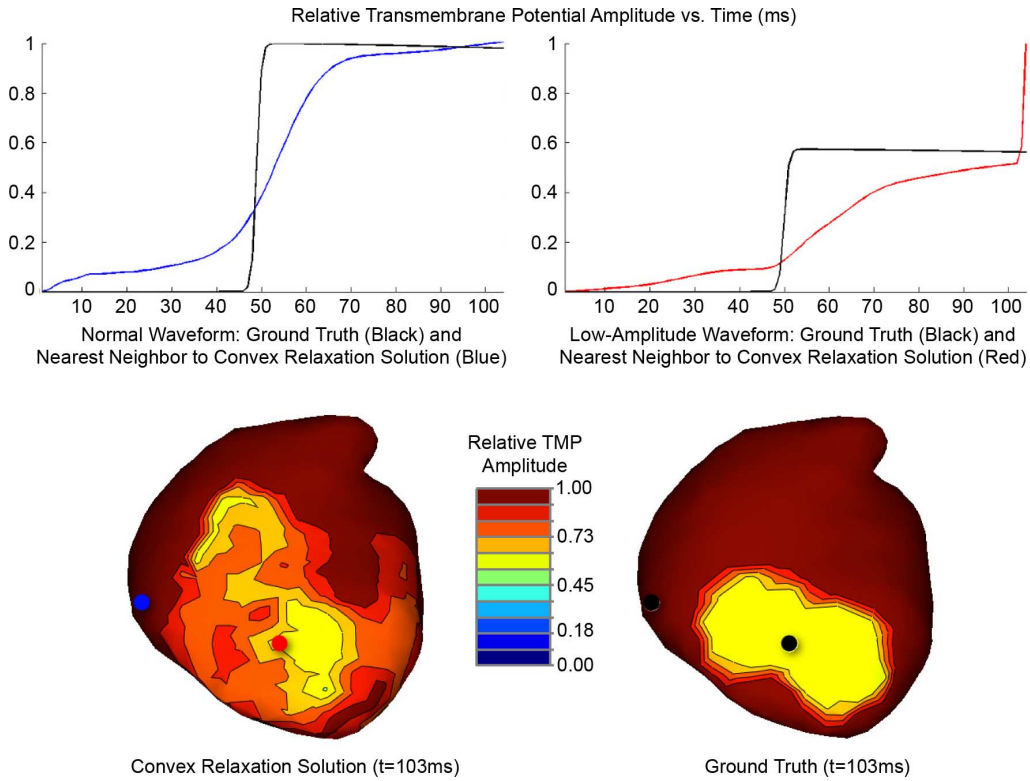

Fig. 5. Synthesized low-amplitude transmembrane potentials: A comparison of the convex relaxation solution with the ground truth, simulated with ECGSIM. Example TMP waveforms are shown in the top row, and isopotential maps from the end of the QRS complex ($t = 103$ ms) are shown in the bottom row. The example TMP waveforms from the convex relaxation solution, plotted as blue (normal) and red (low-amplitude) curves, are from sources located at the blue and red dots, respectively, on the isopotential map of the convex relaxation solution. The equivalent locations on the ground truth isopotential map have been marked using black dots, with their corresponding normal and low-amplitude TMP plots shown in black.

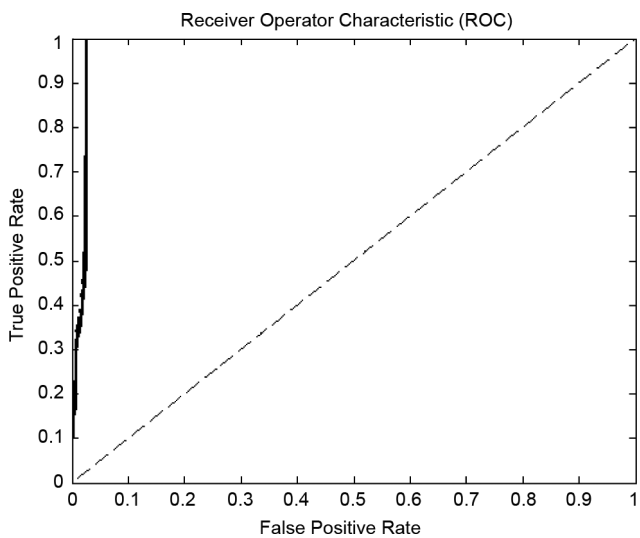

Fig. 6. Detection of synthesized low-amplitude transmembrane potentials: the convex relaxation solution at $t = 103$ ms, shown in Fig. 5, was thresholded at different values within the range, $[0, 1]$, of normalized TMP values to classify reconstructed sources as either "low-amplitude" or "normal." Comparing these results to the synthesized ground truth, also shown in Fig. 5, the ROC curve was computed by varying the threshold of the detector over the full range, and shows the resulting relationship between the false positive and true positive rates for detecting "low-amplitude" sources.

and NLLS solutions (Fig. 3, columns 3 and 4). If the NLLS method were robust to the range of initialization variation tested, the standard deviations shown in column 4 of this figure would be quite small, but this is clearly not the case. Indeed, the similarities between columns 1 and 3 suggest that the means of the solutions remain largely unchanged from those of the initializations, both in terms of their spatial distributions and their values. Furthermore, columns 2 and 4 reveal that the standard deviations also change very little after NLLS optimization, with the maximum standard deviation reduced only moderately, from 33.5 ms to 27.2 ms. This is direct evidence of the sensitivity of the NLLS formulation to initialization, and therefore to the FRA model's assumed propagation velocity.

Another interesting finding is that this sensitivity is highly variable over the heart surface, since some spatial regions of both the FRA initializations and the FRA-initialized solutions exhibited much smaller standard deviations than others. On both the endocardium and epicardium, the standard deviations near the apex were relatively low, whereas those near the base were higher.

Because the convex relaxation solution is independent of initialization it may be more useful for characterizing the intrinsic uncertainty of the NLLS solution, without compounding variation in initialization with intrinsic uncertainty in the rest of the model. To simulate this uncertainty, we studied its dependence on the regularization parameter value. The means and standard deviation maps of the nearest neighbor activation times, over the range of regularization parameters described, are shown in columns 5 and 6 of Fig. 3. We note that regions near the apex

have lower standard deviations than those near the base. Therefore, it is not surprising that the earliest activations in the septum and LV free wall which are near the apex are similar in the nominal solution for $\lambda = 0.1$ (shown in Fig. 2) and the mean solution in column 5 of Fig. 3). Interestingly, the FRA-initialized results (in both Figs. 2 and 3) are also similar to those nearest neighbor means. This reinforces the relative stability of the underlying activation-based formulation in these regions. On the other hand, the high standard deviations of the nearest neighbor solutions near the base suggest considerably greater uncertainty in *any* solution to this problem in that region (e.g., the early activations near the base on the epicardium). This stems from the observation that the two different perturbation analyses were performed independently on different sets of variables, yet resulted in similar spatial patterns of solution variability. Thus, using the convex relaxation we are able to predict and better understand this variability in the NLLS problem. We note that here the NLLS problem is not sensitive to regularization parameters, because there is so little change from initializations to solutions. As reported in Section IV-A, the perturbation analysis of the NLLS solution over regularization parameters yielded standard deviations below 0.04 ms for all but two of the modeled heart surface sources. Our results indicate that this sensitivity is absorbed into that of the initialization instead.

In the experiment described in Section IV-B, we also constructed a simulation in which the NLLS model assumptions about TMP amplitudes were violated. We chose this particular modification to mimic one electrocardiographic effect of ischemic zones. We solved the convex relaxation using these simulated body surface measurements without prior knowledge of the nonuniformity of the TMP amplitudes. We note that, despite the violation of its underlying assumption, it may be possible for a NLLS solution to correctly estimate activation times under such circumstances. However, such a solution may not be clinically useful because it is not sensitive to the physiological irregularity due to the strict constraints it places on true waveform amplitudes. When we solved the convex relaxation optimization, we were able to detect evidence of the presence and location of the lowered TMP amplitudes. Fig. 5 reveals that, although the constraints force the signal to reach an incorrect maximum amplitude at the very end of the QRS complex, the reconstructed waveforms show a sustained nonzero amplitude lower than that maximum until just before the QRS endpoint. Therefore, it may be possible to identify and localize an ischemic zone by inspecting the convex relaxation solution at just a few samples *before* the end of the QRS complex (when most tissue is expected to be activated) for regions of low amplitude. The isopotential map of the convex relaxation solution in Fig. 5 is an example of what one would see using such a strategy. Although our result does not exactly identify the boundaries and values of the ground truth ischemic zone, the broad region with lowered TMP amplitudes in the convex relaxation is roughly collocated with the ground truth region at its center, and therefore may be useful for localization of these inaccurate model assumptions. To test this, we built a simple "low-amplitude" source detector based on thresholding the reconstructed values, which we described in Section IV-B, and whose ROC curve over the range of possible thresholds is shown in Fig. 6. A perfect true positive rate could be achieved with a false positive rate of less than 2.5% by classifying sources with normalized TMPs below a threshold of 0.5926 as "low-amplitude" (as compared to the ground truth value of 0.6). This suggests that localization of low-amplitude TMPs indeed may be possible, although a more extensive evaluation on a larger number of cases would be necessary before making any stronger claims. An interesting implication of these results is that it may be possible to estimate ischemic zones using only data acquired during the QRS complex, whereas typical analysis of ischemia is predominantly performed using the ST segment. It may also be possible to improve localization by modeling the amplitudes of the roughly-localized ischemic zone in this solution, adjusting the assumed waveform amplitude endpoints accordingly, and then rerunning the convex relaxation solution.

The strong spatial structure of solution uncertainty identified by our first set of results was not necessarily expected, and may bear further investigation. It is not clear whether this is a result of loss of geometric accuracy in the anatomic model, or change in relative complexity of the propagation pattern towards the end of QRS, or some combination of both, or if indeed other factors are involved as well.

Also, to date we have only applied the convex relaxation approach to analysis of depolarization during QRS. Since NLLS activation-based methods, with minor modifications, have been used to identify recovery (repolarization) timing as well [2], we intend to attempt to extend the convex relaxation method for the recovery problem as well. We also plan to investigate possible connections between the spatial dependence of solution uncertainty, as characterized by our method, and the structural accuracy of our geometric representations of the heart and body surfaces.

Finally, activation-based methods have recently been shown to be reasonably robust even when much smaller sets of electrocardiographic leads are used [36]. Whether the convex relaxation is similarly able to extract activation timing information from smaller lead sets is still an open question.

Appendix
Proof of Equivalence of Constraint Sets

Let $\mathcal{R}$ and $\mathcal{E}$ be sets as defined in Section II-A. In addition, let $\mathcal{X}$ be a set defined as

$$\mathcal{X} = \Big\{ X \in \mathbb{R}^{N\times T} \mid X_{n,t} = u(t-\tau_n),\ \ldots \\ \ldots \forall 1 < \tau_n \leq T,\ n = 1,\ldots,N \Big\}.$$

Then to prove that $\mathcal{X} = \mathcal{R}\cap\mathcal{E}$, we must show that $\mathcal{X} \subseteq \mathcal{R}\cap\mathcal{E}$ and $\mathcal{X} \supseteq \mathcal{R}\cap\mathcal{E}$.

$\mathcal{X} \subseteq \mathcal{R}\cap\mathcal{E}$*:* Let $X \in \mathcal{X}$. Then, by definition, $\exists\tau = [\tau_1,\ldots,\tau_N]^{\mathrm{T}}$ s.t. $X_{n,t} = u(t-\tau_n)$. Because every value of $X_{n,t}$ is either 0 or 1, then $0 \leq X \leq 1$. In addition, from its definition, we have that $\forall\tau'$ and $t_0 \leq t_1$, $u(t_0-\tau') \leq u(t_1-\tau')$, and it is straightforward to verify that $XD^{\mathrm{T}} \geq 0$. We also know that $\forall 1 < \tau' \leq T$, $u(1-\tau') = 0$ and $u(T-\tau') = 1$, and therefore $XD^{\mathrm{T}}1_{(T\times 1)} = 1_{(N\times 1)}$. Finally, to show that

$\mathrm{tr}\left(X^{\mathrm{T}}X\right) = 1^{\mathrm{T}}_{(T\times 1)}X1_{(N\times 1)}$, we begin by restating this condition using summation notation

$$\sum_{n=1}^{N}\sum_{t=1}^{T} X_{n,t}^2 = \sum_{n=1}^{N}\sum_{t=1}^{T} X_{n,t}$$
$$\sum_{n=1}^{N}\sum_{t=1}^{T} X_{n,t}^2 - X_{n,t} = 0.$$

To verify this condition, we consider the two applicable cases

$$X_{n,t} = 0 \Rightarrow 0^2 - 0 = 0$$
$$X_{n,t} = 1 \Rightarrow 1^2 - 1 = 0.$$

Then because $\sum_{n,t} 0 = 0$, the condition is verified. Finally, because we have verified that an arbitrary $X \in \mathcal{X}$ also satisfies $X \in \mathcal{R} \cap \mathcal{E}$, we have shown that $\mathcal{X} \subseteq \mathcal{R} \cap \mathcal{E}$.

$\mathcal{X} \supseteq \mathcal{R} \cap \mathcal{E}$*:* Let $X \in \mathcal{R} \cap \mathcal{E}$. Then we must show that $\exists \tau$ s.t. $X_{n,t} = u(t - \tau_n)$. Given $0 \leq X \leq 1$ and $\sum_{n,t} X_{n,t}^2 - X_{n,t} = 0$, we have already shown that $X_{n,t} = 0$ or $X_{n,t} = 1$ will satisfy these conditions. In all other cases, $0 < X_{n,t} < 1$, and $X_{n,t}^2 < X_{n,t}$, so it follows that $X_{n,t}^2 - X_{n,t} < 0$. If $\exists n', t'$ s.t. $0 < X_{n',t'} < 1$, then $\sum_{n,t} X_{n,t}^2 - X_{n,t} < 0$, and because this cannot be the case, we have shown that it is necessary for $X_{n,t} \in \{0, 1\}$. This is also sufficient to show that $\mathcal{R} \cap \mathcal{E}$ is a discrete set.

Now we consider the condition that $XD^{\mathrm{T}} \geq 0$, which means that $\forall t_0 \leq t_1$, $X_{n,t_0} \leq X_{n,t_1}$. Then if $X_{n,t_1} = 0$, our previous result requires that $X_{n,t_0} = 0$. Similarly, if $X_{n,t_0} = 1$, it is required that $X_{n,t_1} = 1$. In other words, we have shown that every row of X must satisfy one of three conditions.

1) Constant row of zeros: all of its values are equal to 0.
2) Constant row of ones: all of its values are equal to 1.
3) Phase-shifted unit step function: $\exists \tau_n$ such that for $t < \tau_n$ all values are equal to 0, and for $t \geq \tau_n$ all values are equal to 1.

The last of these conditions may be equivalently restated as $X_{n,t} = u(t - \tau_n)$.

Also consider that $XD^{\mathrm{T}}1_{(T\times 1)} = 1_{(N\times 1)}$, which may be equivalently restated as $\forall n$, $X_{n,T} - X_{n,1} = 1$. Clearly, again given the previous result, the only way this can be satisfied is if $X_{n,T} = 1$ and $X_{n,1} = 0$. In other words, this says that every row of X must start from 0 in the first column and end with 1 in the last column. This rules out constant-valued rows of zeros or ones, and the only remaining possibility is the condition that the values of X may be parameterized as $X_{n,t} = u(t - \tau_n)$ with $1 < \tau_n \leq T$.

This shows that $X \in \mathcal{X}$ and therefore $\mathcal{X} \supseteq \mathcal{R} \cap \mathcal{E}$. ■

REFERENCES

[1] Y. Rudy, "Noninvasive electrocardiographic imaging of arrhythmogenic substrates in humans," *Circ. Res.*, vol. 112, no. 5, pp. 863–874, 2013.

[2] P. M. V. Dam, T. F. Oostendorp, A. C. Linnenbank, and A. V. Oosterom, "Non-invasive imaging of cardiac activation and recovery," *Ann. Biomed. Eng.*, vol. 37, no. 9, pp. 1739–1756, 2009.

[3] T. Berger *et al.*, "Single-beat noninvasive imaging of ventricular endocardial and epicardial activation in patients undergoing CRT," *PloS One*, vol. 6, no. 1, p. E16255, 2011.

[4] C. C. Gornick, S. W. Adler, B. Pederson, J. Hauck, J. Budd, and J. Schweitzer, "Validation of a new noncontact catheter system for electroanatomic mapping of left ventricular endocardium," *Circulation*, vol. 99, no. 6, pp. 829–835, 1999.

[5] A. van Oosterom and T. Oostendorp, "ECGSIM: an interactive tool for studying the genesis of QRST waveforms," *Heart*, vol. 90, no. 2, p. 165, 2004.

[6] B. Tilg, G. Fischer, R. Modre, F. Hanser, B. Messnarz, M. Schocke, C. Kremser, T. Berger, F. Hintringer, and F. X. Roithinger, "Model-based imaging of cardiac electrical excitation in humans," *IEEE Trans. Med. Imag.*, vol. 21, no. 9, pp. 1031–1039, Sep. 2002.

[7] L. K. Cheng, J. M. Bodley, and A. J. Pullan, "Comparison of potential-and activation-based formulations for the inverse problem of electrocardiology," *IEEE Trans. Biomed. Eng.*, vol. 50, no. 1, pp. 11–22, Jan. 2003.

[8] Z. Liu, C. Liu, and B. He, "Noninvasive reconstruction of three-dimensional ventricular activation sequence from the inverse solution of distributed equivalent current density," *IEEE Trans. Med. Imag.*, vol. 25, no. 10, pp. 1307–1318, Oct. 2006.

[9] C. Liu, M. D. Eggen, C. M. Swingen, P. A. Iaizzo, and B. He, "Noninvasive mapping of transmural potentials during activation in swine hearts from body surface electrocardiograms," *IEEE Trans. Med. Imag.*, vol. 31, no. 9, pp. 1777–1785, Sep. 2012.

[10] L. Wang, F. Dawoud, S.-K. Yeung, P. Shi, K. Wong, and A. Lardo, "Transmural imaging of ventricular action potentials and post-infarction scars in swine hearts," *IEEE Trans. Med. Imag.*, vol. 32, no. 4, pp. 731–747, Apr. 2013.

[11] R. Modre, B. Tilg, G. Fischer, and P. Wach, "Noninvasive myocardial activation time imaging: A novel inverse algorithm applied to clinical ECG mapping data," *IEEE Trans. Biomed. Eng.*, vol. 49, no. 10, pp. 1153–1161, Oct. 2002.

[12] S. E. Geneser, R. M. Kirby, and R. S. MacLeod, "Application of stochastic finite element methods to study the sensitivity of ECG forward modeling to organ conductivity," *IEEE Trans. Biomed. Eng.*, vol. 55, no. 1, pp. 31–40, Jan. 2008.

[13] E. Konukoglu *et al.*, "Efficient probabilistic model personalization integrating uncertainty on data and parameters: Application to eikonal-diffusion models in cardiac electrophysiology," *Prog. Biophys. Mol. Biol.*, vol. 107, no. 1, pp. 134–146, 2011.

[14] B. Erem, P. M. van Dam, and D. H. Brooks, "Analysis of the criteria of activation-based inverse electrocardiography using convex optimization," in *Proc. Annu. Int. Conf. IEEE EMBS*, 2011, pp. 3913–3916.

[15] B. Erem, P. M. van Dam, and D. H. Brooks, "A convex relaxation framework for initialization of activation-based inverse electrocardiography," in *Proc. IEEE Int. Symp. Noninvasive Funct. Source Imag. Brain Heart Int. Conf. Bioelectromagn.*, 2011, pp. 12–17.

[16] A. Pullan, L. Cheng, M. Nash, C. Bradley, and D. Paterson, "Noninvasive electrical imaging of the heart: Theory and model development," *Ann. Biomed. Eng.*, vol. 29, no. 10, pp. 817–836, 2001.

[17] A. Pullan, L. Cheng, and M. Buist, *Mathematically Modelling the Electrical Activity of the Heart: From Cell to Body surface and back again*. Singapore: World Scientific, 2005.

[18] P. M. van Dam, T. F. Oostendorp, and A. van Oosterom, "Application of the fastest route algorithm in the interactive simulation of the effect of local ischemia on the ECG," *Med. Biol. Eng. Comput.*, vol. 47, no. 1, pp. 11–20, 2009.

[19] R. Coronel, J. Fiolet, F. Wilms-Schopman, A. Schaapherder, T. Johnson, L. Gettes, and M. Janse, "Distribution of extracellular potassium and its relation to electrophysiologic changes during acute myocardial ischemia in the isolated perfused porcine heart," *Circulation*, vol. 77, no. 5, pp. 1125–1138, 1988.

[20] R. Coronel *et al.*, "Right ventricular fibrosis and conduction delay in a patient with clinical signs of Brugada syndrome," *Circulation*, vol. 112, no. 18, pp. 2769–2777, 2005.

[21] P. G. Meregalli, A. A. Wilde, and H. L. Tan, "Pathophysiological mechanisms of Brugada syndrome: Depolarization disorder, repolarization disorder, or more?," *Cardiovasc. Res.*, vol. 67, no. 3, pp. 367–378, 2005.

[22] A. Wilde and C. Antzelevitch, "The continuing story: The aetiology of Brugada syndrome: Functional or structural basis?," *Eur. Heart J.*, vol. 24, no. 22, p. 2073, 2003.

[23] C. Wolpert *et al.*, "Intravenous drug challenge using flecainide and ajmaline in patients with Brugada syndrome," *Heart Rhythm*, vol. 2, no. 3, pp. 254–260, 2005.

[24] F. I. Marcus *et al.*, "Diagnosis of arrhythmogenic right ventricular cardiomyopathy/dysplasia proposed modification of the task force criteria," *Circulation*, vol. 121, no. 13, pp. 1533–1541, 2010.

[25] B. Messnarz, B. Tilg, R. Modre, G. Fischer, and F. Hanser, "A new spatiotemporal regularization approach for reconstruction of cardiac transmembrane potential patterns," *IEEE Trans. Biomed. Eng.*, vol. 51, no. 2, pp. 273–281, Feb. 2004.
[26] D. B. Geselowitz, "On the theory of the electrocardiogram," *Proc. IEEE*, vol. 77, no. 6, pp. 857–876, Jun. 1989.
[27] D. B. Geselowitz, "Description of cardiac sources in anisotropic cardiac muscle: Application of bidomain model," *J. Electrocardiol.*, vol. 25, pp. 65–67, 1992.
[28] R. M. Gulrajani, *Bioelectricity and biomagnetism*. New York: Wiley, 1998.
[29] G. Huiskamp and A. van Oosterom, "The depolarization sequence of the human heart surface computed from measured body surface potentials," *IEEE Trans. Biomed. Eng.*, vol. 35, no. 12, pp. 1047–1058, Dec. 1988.
[30] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge, MA: Cambridge Univ. Press, 2004.
[31] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, 2011.
[32] CVX: MATLAB software for disciplined convex programming. ver. 2.0 beta, CVX Research, Inc., Sep. 2012.
[33] J. W. Meijs, O. W. Weier, M. J. Peters, and A. van Oosterom, "On the numerical accuracy of the boundary element method," *IEEE Trans. Biomed. Eng.*, vol. 36, no. 10, pp. 1038–1049, Oct. 1989.
[34] A. G. Kleber, M. J. Janse, F. Wilms-Schopmann, A. Wilde, and R. Coronel, "Changes in conduction velocity during acute ischemia in ventricular myocardium of the isolated porcine heart," *Circulation*, vol. 73, no. 1, pp. 189–198, 1986.
[35] D. Durrer, R. T. Van Dam, G. Freud, M. Janse, F. Meijler, and R. Arzbaecher, "Total excitation of the isolated human heart," *Circulation*, vol. 41, no. 6, pp. 899–912, 1970.
[36] P. M. van Dam, T. Oostendorp, and A. van Oosterom, "Non-invasive cardiac imaging based on just the standard 12-lead signals?," *IEEE Comput. Cardiol.*, pp. 173–176, 2009.