Annals of Biomedical Engineering, Vol. 26, pp. 278–292, 1998
Printed in the USA. All rights reserved.

An Admissible Solution Approach to Inverse Electrocardiography

GHANDI F. AHMAD,* DANA H. BROOKS,† and ROBERT S. MACLEOD‡

*Electronics Engineering Department, College of Science and Technology Jerusalem, West Bank, †Communications and Digital Signal Processing (CDSP) Center, Electrical and Computer Engineering Department, Northeastern University, Boston, MA, and ‡Cardiovascular Research and Training Institute (CVRTI), University of Utah, Salt Lake City, UT

(Received 8 July 1996; accepted 22 August 1997)

Abstract—The goal of the inverse problem of electrocardiography is noninvasive discrimination and characterization of normal/abnormal cardiac electrical activity from measurements of body surface potentials. Smoothing and attenuation in the torso volume conductor cause the problem to be ill posed. Standard regularized solutions employ an *a priori* constraint to achieve reliability and may be biased by the constraint chosen as well as the regularization parameter used to weight it. In this paper, we describe an approach that reformulates this inverse problem as the search for a solution that is a member of an *admissible solution set*; admissibility is defined in terms of the available constraints. In principle, this approach can utilize as many constraints as may be available, unlike standard techniques which do not easily permit the use of multiple constraints. No regularization parameter is required; instead, we need to choose the nature and size of the constraint sets. Constraints described include several spatial constraints, weighted constraints, and temporal constraints. We describe a solution approach based on iterative convex optimization, and the algorithm—the ellipsoid algorithm—which we have used. Accuracy and feasibility of the method are illustrated with simulation results using dipole sources and measured epicardial potentials. © *1998 Biomedical Engineering Society.*
[S0090-6964(98)01002-9]

Keywords—Inverse problem of electrocardiography, Convex optimization, Multiple constraints, Set theoretic estimation, Regularization.

INTRODUCTION

The ability to understand and characterize the heart's electrical activity, and to discriminate normal from abnormal activity, is a goal of great interest to physicians and researchers. The electrocardiogram (ECG) offers a safe, cheap, noninvasive means of measuring this activity and is the basis of many diagnostic techniques. The standard ECG permits only a relatively coarse description of the spatial complexity of cardiac electrical activity—one result is that interpretation of the standard ECG cannot be based on a rigorous biophysical model, but rather must depend on a heuristic match between waveforms and disease state. Enhanced forms of the ECG permit a more complete representation of the heart and can be coupled with visualization, signal processing, and mathematical modeling techniques to create powerful tools. This manuscript deals with one such enhancement—the inverse problem of electrocardiography—and describes a novel solution method.

Since the standard 3–10 ECG electrodes significantly undersample the distribution of electrical potentials on the torso surface, the ECG provides only a limited view of underlying cardiac activity. Additionally, information about cardiac electrical activity is blurred and attenuated in the torso volume conductor, the medium between the heart and the body surface. Researchers (and some clinicians) have used electrode arrays to achieve denser spatial sampling, along with geometrical models of the thorax, to establish mathematical relationships with which one can estimate cardiac electrical activity from torso surface potentials. This approach—explicitly solving for the underlying cardiac electrical activity of the heart from body surface electric potentials—is known as the *inverse problem of electrocardiography*. A successful inverse solution would allow us to detect, quantify, and localize cardiac electrical activity from noninvasive torso measurements. Contributions by many researchers (see, for a sample, Refs. 17, 27, 31, and 33), have established that inverse solutions posed in terms of reconstructing epicardial potentials (the electrical potential distribution on the outer surface of the heart) offer the advantages of a concise mathematical formulation, uniqueness, and direct links to measurements. To obtain such an inverse solution, we must first solve the associated *forward problem*, in which torso surface potentials are estimated based on epicardial potentials and a geometric model (including electric conductivity) of the thorax.[16]

The inverse problem is ill posed,[17,18,33] meaning that small perturbations in the measured data (due to noise, errors in the forward model, discretization effects, etc.) can result in unbounded errors in the inverse solution. The ill posedness is a mathematical reflection of physical phenomena that include the attenuation and smoothing

Address correspondence to Robert S. MacLeod, N.E. Harrision CVRTI, Building 500, University of Utah, Salt Lake City, Utah 84112. Electronic mail: macleod@cvrti.utah.edu

effects of the volume conductor. Sharp spatial variations in the epicardial potentials are blurred or smoothed in the body surface potentials. Thus, a small amount of noise in the body surface potential will tend to be magnified in epicardial potentials estimated through an inverse procedure. Inverse solutions need to be stabilized by the incorporation of additional physiological information; the best-known approach to this type of ill-posed inverse problem is known as *regularization*. In regularized solutions, a compromise is sought between a solution that matches the data but may be unrealistic and unstable, on the one hand, and fidelity to a constraint based on *a priori* knowledge (or assumptions) about a realistic, stable solution, on the other. We wish to estimate epicardial potential distributions that, when put through the forward model, generate body surface potentials similar to those we measure, and yet which also respect certain reasonable constraints. Typical constraints include the size (two-norm) of the solution or of its gradient or Laplacian (for details and references, see Ref. 33). An important issue in regularized inverse solutions is the choice of a weight, called the regularization parameter, which controls the tradeoff between the data fit and the constraint. The accuracy of the inverse solution is quite sensitive to this parameter, and much research has been devoted to developing good methods for determining appropriate values.[18,33]

The main features of epicardial potential distributions can often be roughly reconstructed by regularization approaches. However, due to the complex nature of epicardial potentials, the imposition of any single constraint often fails to produce globally satisfactory solutions. For instance, regularization with a two-norm amplitude constraint may capture areas of large gradients and accurately constrain regions of large amplitude (such as near activation wave fronts) but will tend to be noisy where the signal level is low. On the other hand, derivative-based constraints will be less noisy but may smear wave fronts.[5,6,24]

In response to the limited success of these standard methods, techniques which attempt to incorporate more *a priori* information into the solution have recently been investigated.[25] Spatially local regularization,[22] orthogonality constraints,[36] combined truncated and weighted singular value constraints,[35] a constraint on the normal component of current on the epicardial surface,[23] simultaneous imposition of two distinct spatial constraints,[5,6] and constraints on individual solution elements based on over-regularized and under-regularized two-norm solutions[21] have all been reported. In addition to these spatial constraints, some researchers have recently reported results achieved by using *temporal* constraints. These constraints include on–off constraints in the form of activation source models,[10] more recently restricted to the reconstruction of first the breakthrough events and then the full activation sequences,[13,14,20] solutions using power spectrum techniques,[7] two-step postfiltering techniques[29] (equivalent to temporally low-pass filtering the spatially regularized solution), joint temporal/spatial multiple constraint solutions,[5,6] and solutions based on Kalman filtering.[12]

In our own recent work,[5,6] we found encouraging results using two constraints simultaneously. In particular, using two constraints rather than one increased robustness to error in the particular values of each regularization parameter and preserved quantitative and qualitative accuracy in ways difficult to achieve with a single constraint (e.g., preserving sharp gradients around wave fronts while suppressing noise in low-amplitude regions, or combining temporal and spatial smoothness). In general, this approach increased both the reliability of the result and our ability to impose physiologically reasonable prior information on the solution. However, the problem of selecting a regularization parameter became more complicated as this approach required *two* regularization parameters. Although we solved this problem in the two-constraint case[6] by extending a well-known single constraint technique,[18] it is difficult to see how one could choose regularization parameters for more than two regularization constraints, even though additional constraints might result in further improvements to the inverse solution. It is in response to the potential advantage of employing a flexible combination of diverse constraints that we have adopted and tested the *admissible solution* method described in this paper.

This admissible solution method reflects a different philosophy from the methods referred to so far, which all feature the selection of one or two constraints and the calculation of a unique minimum of an error measure (such as the 2-norm), which is chosen as much for mathematical convenience as physical reasonableness. Instead, we start with the assumption that we have available a number of appropriate constraints on the solution. These constraints can be based on physical/physiological principles or empirical results obtained from experimental measurements. For instance, we may have reasonable bounds on spatial norms or maxima/minima of the solution at each point in time, or temporal norms at each point in space, or spatial or temporal gradients, or temporal frequency content, or breakthrough locations and timing, etc., or combinations of these bounds. Moreover, the forward model along with the measured data is itself a constraint on the solution, since any epicardial solution projected to the torso by the forward model should be ''similar enough'' to the measured body surface potentials (where ''similar'' can be defined in a number of ways); this is equivalent to the data fit term in the error that is minimized in regularized solutions.

We can interpret each constraint as defining a set of possible solutions in the solution space. The intersection

of these sets of constrained solutions then represents a region of acceptability for a solution—the set of all solutions which are feasible or *admissible*, that is, consistent with our knowledge about the problem as encoded by the constraints. The philosophy we have adopted is to search for a solution which meets this collection of constraints. We can find such an admissible solution by applying an iterative algorithm which locates one element within the intersection of the constraint sets.[8,9] This is in clear contrast to the notion of a unique optimal solution in standard regularization approaches, which locate a point in the solution space by minimizing an *a priori* cost function.

The main advantages of this approach are the following:

(1) There is no need for regularization parameters as required in standard regularization techniques. The regularization parameters are replaced by direct bounds on the solution, which may be easier to derive, test, and interpret based on biophysics and/or experimental measurements. Any bias imposed on the solution is applied directly in terms of the bounds on the constraints rather than indirectly in terms of regularization parameters.

(2) We can incorporate a large number and wide variety of mixed constraints in a flexible manner, with no essential change in the underlying solution method.

(3) As the solution methods are iterative and employ only one constraint at a time, we can monitor how the constraints are employed and thus gain understanding both of whether a particular constraint is effective in restricting the solution set, and if so, how it interacts with other constraints.

We note that, in general, a single constraint will not suffice—one needs two constraints at a minimum to hope to achieve reasonable results. This is the dual to the need to regularize in standard approaches. Moreover, at least one of the constraints must involve the data and the forward model. One disadvantage of the admissible solution approach is that there are no closed-form expressions for the solution. Another disadvantage is that there is no guarantee that the resulting solution is optimal in any sense, although there is no clear rationale that standard currently used optimality criteria are, in fact, optimal in any *physiological* sense. A third disadvantage is that the algorithms that are used to find an admissible solution are computationally intensive and may have slow convergence rates, especially if the size of the feasible set is small and the dimension of the space defining the sets is large. However, if we restrict the constraint sets to the class of convex functions, we can employ one of several effective numerical convex optimization methods, of which projection onto convex sets (POCS)[37] is, perhaps, the best known. These algorithms have tradeoffs in terms of prior knowledge required, flexibility of problem formulation, convergence speed, and algorithmic and computational complexity. For instance, POCS is relatively simple to understand and implement, but requires exact projection operators for each constraint, does not converge as quickly as some other algorithms, and requires a ''stepsize-''type parameter. We have employed an algorithm known as the ellipsoid algorithm,[4] which offers a reasonable compromise to the tradeoffs mentioned above.

There have been a few previous attempts to incorporate constraints of this type into a regularization scenario for the inverse problem of electrocardiography, most notably in Refs. 21 and 27. However, only one or two constraints were used and a global objective function was minimized in a traditional fashion. Techniques similar to our admissible solution method have been used in image restoration problems, which are formally related to the inverse problem of electrocardiography. For example, adaptive regularized methods based on POCS have been used to restore images in Refs. 30 and 37.

In the next section, we formulate the admissible solution approach as a convex optimization problem, describe one algorithm which we have employed to solve it, and discuss some aspects of implementation. We describe our experimental procedures and present supporting simulations in the section on results. We discuss our results and conclusions and outline our future work in the last section.

ADMISSIBLE SOLUTION APPROACH

We use a standard discretized formulation of the forward equation,

$$\mathbf{y}(k)=\mathbf{A}\cdot\mathbf{h}(k)+\mathbf{e}(k), \quad k=1,2,...,L, \tag{1}$$

where $\mathbf{y}$ is an $M\times 1$ vector of torso potentials at time instant k, $\mathbf{h}(k)$ is the $N\times 1$ vector of epicardial potentials, $\mathbf{A}$ is the $M\times N$ matrix representing the forward solution, $\mathbf{e}$ is measurement noise of the same dimensions as the body surface vector, and k and L are a discrete time index and the number of time samples, respectively. Thus, we assume that we have a forward model, represented by the matrix $\mathbf{A}$, which expresses the potentials at each body surface node as a linear combination of the potentials at all the nodes of the epicardial surface at a given time instant. This implies, in turn, assumptions of linearity and quasistatic propagation, both of which are reasonable physical approximations.[32] In addition, we assume that we have some *a priori* bounds, or constraints, on the solution, each of which can be described as a convex set in the solution space. Although the convexity condition is not necessary for the problem to be well formulated, it makes tractable solutions possible while retaining a very large class of potential constraints—for instance, there is no requirement that the constraint functions be differentiable. Figure 1 shows a schematic illus-

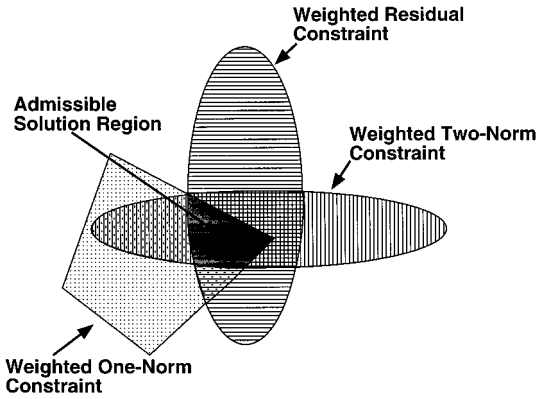

FIGURE 1. Three convex sets with their intersection.

tration of the approach. Each shaded shape represents a particular constraint. In the diagram, the ellipse marked ''Weighted Residual Constraint'' would reflect a bound on the noise power (more precisely, a weighted two-norm of the residual between the measured torso potentials and those predicted by a candidate solution together with the forward model **A**). The set labeled ''Weighted two-Norm Constraint'' shows a bound on a weighted two-norm of the solution, while the ''Weighted One-Norm Constraint'' limits a weighted one-norm of the solution. This constraint is nondifferentiable but convex. The intersection of these constraint sets is the region of admissible solutions.

In the rest of this section, we describe how the problem of finding an admissible solution with convex constraints can be posed as a convex optimization problem. We then describe some examples of convex constraints that might be useful in the electrocardiographic inverse problem. Finally, we describe the ellipsoid algorithm we have implemented to find admissible solutions.

Admissible Solutions and Convex Optimization

Initially, for simplicity, we will treat the problem at only one time instant, and in the sequel, will expand our formulation to include many time instants. The unknown epicardial potentials at a given time instant are assumed to be an element of an appropriate Hilbert space $\mathscr{H}$. After sampling in space and time, the desired solution will be a vector $\mathbf{x}\in\mathscr{R}^N$, where N is the number of nodes in the model of the epicardial surface. We use each piece of *a priori* information, i.e., each constraint, to restrict the solution to a closed convex set in $\mathscr{H}$ (or equivalently $\mathscr{R}^N$); for each constraint there is a corresponding function $\phi(\mathbf{x})$: $\mathscr{R}^N\rightarrow\mathscr{R}$ and constraint bound ϵ, so that the constraint can be written as $\phi(\mathbf{x})-\epsilon<0$. In addition, we require that each ϕ satisfy the convexity condition

$$\phi(\alpha\mathbf{x}+(1-\alpha)\mathbf{y})\leqslant\alpha\phi(\mathbf{x})+(1-\alpha)\phi(\mathbf{y}),\quad\forall\alpha\in[0,1]. \tag{2}$$

Thus, with m such constraints ϕ_i, $i=1,2,...,m$, each will correspond to a closed convex set $(\phi_i(\mathbf{x})-\epsilon_i)\in\mathscr{H}$, for $i=1,2,...,m$. Since the region in $\mathscr{R}^N$ enclosed by the intersection of the ϕ_i is itself convex, we can define a new convex function ϕ as

$$\phi=\max_i(\phi_i-\epsilon_i). \tag{3}$$

Thus, the set $\{\mathbf{x}:\phi(\mathbf{x})<0\}$ represents the intersection of all the constraints and a point $\mathbf{x}\in\mathscr{R}^N$ such that $\phi(\mathbf{x})\leqslant 0$ is an admissible solution.

If we assume we have an iterative algorithm available which can perform convex optimization, i.e., minimize a convex function, then we might apply it to *minimize* one or more of the constraints. However, this would not be appropriate here as none of the constraints is a traditional objective to minimize—not even the residual error, since forcing it to be too small implies an under-regularized solution. The appropriate approach here, then, is to run the convex optimization algorithm only until all the constraints are satisfied, i.e., $\phi(\mathbf{x})\leqslant 0$, and then stop and declare the resulting solution to be admissible. The bounds on the constraints, the ϵ_i, play an equivalent role to regularization parameters. However, unlike regularization parameters, they represent direct constraints on solutions, and as such can be developed and tested based on measured data, and generally are subject to direct physiological interpretation.

To consider L time instants simultaneously, in order to use temporal constraints or temporal frequency constraints, one can consider an augmented problem in a higher dimensional space $\mathscr{R}^{NL}$.[6] The vectors in Eq. (1) become block vectors (for instance, $\overline{\mathbf{y}}\overset{\text{def}}{=}[\mathbf{y}^T(1),\ldots,\mathbf{y}^T(L)]^T$), and the forward matrix becomes a block diagonal matrix $\underline{\mathbf{A}}\overset{\text{def}}{=}\text{diag}(\mathbf{A},\mathbf{A},\ldots,\mathbf{A})=\mathbf{I_L}\otimes\mathbf{A}$, where $\otimes$ is the Kronecker product. As we will see in the next section, constraints can be defined to operate separately on any one time instant or spatial node or globally on all L time instants and/or N nodes. The obvious costs, to be weighed against the benefits, are that

(1) the computational complexity will increase as the dimensionality of the problem does;
(2) the algorithm will take longer to converge if constraints on each spatial node or time instant are included, as it will have to find a solution satisfying all such constraints.

Examples of Convex Constraints

In this section, we describe some examples of the types of convex constraints one might employ to con-

strain inverse solutions. To clarify the discussion, we have classified these examples into four categories; standard, weighted (both of which operate at each time instant separately), spatiotemporal, and novel. The constraint formulations are based on the measured data $\mathbf{y}(k)$, a candidate solution $\mathbf{x}(k)$, and the forward transfer matrix $\mathbf{A}$. As above, each constraint is specified by a constraint *function* of the estimate, denoted as $\phi(\mathbf{x})$, and a constraint *bound* ϵ, so that the ith constraint is written as $\phi_i(\mathbf{x}) \leqslant \epsilon_i$. Determination of each constraint bound requires *a priori* knowledge or an estimate of some particular aspect of the desired solution, the measurement noise, or the model error.

Standard Constraints. Here we illustrate four types of standard constraints: residual constraints, regularization constraints (i.e., 2-norms of solutions or of differentiated or filtered solutions), Tikhonov-type error measure constraints, and nondifferentiable constraints:

(1) Residual constraints: One can constrain the solution to fit reasonably to the measurements and the forward model by using the function $\phi(\mathbf{x}) = \|\mathbf{Ax} - \mathbf{y}\|_2^2$ and the constraint $\phi \leqslant \epsilon$, where ϵ is related to the power (variance) of the measurement noise and any error in the forward model $\mathbf{A}$.

(2) Regularization constraints: One can easily employ most of the constraints used in standard inverse electrocardiography, as described in the introduction, using the convex function $\phi(\mathbf{x}) = \|\mathbf{Rx}\|_2^2$ and the constraint $\phi \leqslant \epsilon$, where $\mathbf{R}$ is, for instance, the identity matrix or a discrete spatial differentiator, or the output of a discrete spatial filter, and ϵ is based on previous measurements, physiological assumptions, or an initial solution (for instance, using the L curve[18]).

(3) Tikhonov-type error constraints: One can include as a constraint a function similar to the error that is minimized in Tikhonov regularization by using the function

$$\phi_\lambda(\mathbf{x}) = \left\| \begin{pmatrix} \mathbf{A} \\ \sqrt{\lambda}\mathbf{R} \end{pmatrix} x - \begin{pmatrix} \mathbf{b} \\ \mathbf{0} \end{pmatrix} \right\|_2^2,$$

and the constraint $\phi_\lambda \leqslant \epsilon_\lambda$. We note that the unique minimizer of $\phi_\lambda(\mathbf{x})$ is the corresponding Tikhonov solution with regularization parameter λ and Tikhonov error ϵ_λ^*. Thus, we must use this constraint with $\epsilon_\lambda > \epsilon_\lambda^*$ (for instance, $\epsilon_\lambda = k\epsilon_\lambda^*$ for some reasonable k). This constraint can be used to ensure that the iterative solution comes close to the minimum error for some Tikhonov solution.

(4) Nondifferentiable constraints: One example would be to constrain the maximum amplitude of the solution or its spatial derivatives with the function $\phi(\mathbf{x}) = \|\mathbf{Rx}\|_\infty$ and the constraint bound $\phi \leqslant \epsilon$.

We note that we can combine multiple instances of these constraints, constraining, for example, the two-norm and Laplacian of the solution, the error in both corresponding Tikhonov solutions for several values of λ, and the max norm of both the solution and its Laplacian.

Weighted Constraints. In addition to constraints of the type illustrated above, we can also use weighted norm constraints, which will effectively be applied with different strength in different regions of the epicardium. For example, if we have prior knowledge of locations of large amplitude and small amplitude regions (e.g., from preliminary over-regularized and under-regularized solutions as suggested in Ref. 21), we may want to constrain amplitudes and smoothness more stringently in small amplitude regions and less so in large amplitude regions. If we have prior knowledge of regions of sharp spatial gradients (for instance, from preliminary solutions which locate breakthroughs),[14,19,20] we may want to constrain gradients differently at different epicardial locations. Weighting the constraints may reduce bias to the noise in the data [by effectively adjusting to a local signal to noise ratio (SNR)] and thus help sharpen the solution. To accomplish this, one can write the spatial constraints above using weighted norms rather than standard norms. The constraint on the weighting matrix for norm properties to be preserved is that it be symmetric and positive definite.

Spatiotemporal Constraints. If we use the augmented model described earlier to include multiple time instants in the problem formulation, we may wish to define constraints in the augmented solution space that operate pointwise in space, pointwise in time, globally in space or time, or jointly in space and time.

We can define global spatial constraints simply by extending the single time instant formulations above. In the case of derivative operators and weighting matrices, it may be necessary to premultiply the block data vector $\overline{\mathbf{y}}$ or the block solution vector $\overline{\mathbf{x}}$ by the appropriate block diagonal matrix. For instance, to constrain the two-norm of the Laplacian over L time instants, we would define

$$\phi_{2,L}(\overline{\mathbf{x}}) = \|\overline{\mathbf{L}}\overline{\mathbf{x}}\|_2^2, \tag{4}$$

and the constraint $\phi_{2,L} \leqslant \epsilon_{2,L}$, where $\overline{\mathbf{L}}$ is an $(NL) \times (NL)$ block diagonal matrix, $\mathrm{diag}(\mathbf{L}, \mathbf{L}, \ldots, \mathbf{L})$ with $\mathbf{L}$ a Laplacian operator matrix as in the section on standard constraints. Global temporal constraints such as temporal derivative approximators or high-pass filters[6] can be defined in a similar fashion except that the block constraint matrix will have both diagonal and off-diagonal blocks, each a diagonal submatrix. Constraints on points over both space and time can be obtained by generalizations. We can form pointwise constraints in space and time by premultiplying the solution vector $\overline{\mathbf{x}}$ either by a block

diagonal matrix with only one nonzero block—selecting all locations at a particular time—or by a block diagonal matrix each of whose blocks is nonzero only at the same diagonal location—selecting a particular point in space over all time instants. By using max norms over collections of such localized constraints, we can also constrain the behavior of the worst time instant or spatial mode.

Novel Constraints. We give only a few examples of constraints which would be difficult to employ in conventional approaches but would not cause any conceptual difficulties here (although some might cause computational difficulties in practice). We can, for example, consider constraints on spatial norms such as the $l1$ or $l4$ norm, or on the standard deviation of the potentials, over the epicardium. Using the augmented model, we can constrain not only temporal norms over electrograms, similarly to the spatial norms, but also the temporal frequency behavior of the solution electrograms. For instance, we can constrain the location of the peak or median frequency, or the size of the bandwidth, or a variation measure such as the standard deviation of peak frequencies. We can also employ entropy measures[4] such as minimum relative entropy to constrain the multichannel temporal frequency behavior of the solution.[7] We note that, as described in Ref. 11, we have begun to explore the elucidation of such constraints using recordings from isolated canine hearts suspended in a torso-shaped electrolytic tank (see the section on results).

The Ellipsoid Algorithm

In this subsection, we describe the ellipsoid algorithm as we used it. In the description that follows, we will use (sub)gradients, which define a supporting hyperplane. The main idea here is to use the gradient $\nabla_{\mathbf{x}}$ or subgradient $\mathbf{g}_{\mathbf{x}}$ of the convex function ϕ at a point $\mathbf{x}$ to define a hyperplane normal to the gradient or subgradient. This hyperplane divides $\mathscr{R}^N$ into two parts, one of which is ''above'' the level set of ϕ at $\mathbf{x}$. Formally, for convex functionals of the form $\Phi : \mathscr{R}^N \rightarrow \mathscr{R}$ that are differentiable and satisfy the convexity condition in Eq. (2), we have

$$\phi(\mathbf{y}) \geqslant \phi(\mathbf{x}) + \nabla_x \phi(\mathbf{x})^T(\mathbf{y}-\mathbf{x}), \quad \forall \mathbf{y} \in \mathscr{R}^N. \tag{5}$$

In other words, the hyperplane tangent to the graph of ϕ at $\mathbf{x}$ never lies below the level set of ϕ at $\mathbf{x}$.

If ϕ is convex but not necessarily differentiable, then $\mathbf{g} \in \mathscr{R}^N$ is by definition a subgradient of ϕ at $\mathbf{x}$ if

$$\phi(\mathbf{y}) \geqslant \phi(\mathbf{x}) + \mathbf{g}^T(\mathbf{y}-\mathbf{x}), \quad \forall \mathbf{y} \in \mathscr{R}^N. \tag{6}$$

Thus the subgradient, like the gradient, divides $\mathscr{R}^N$ by its normal hyperplane into two sets, on one of which $\phi \geqslant \phi(\mathbf{x})$. For example, in the diagram in Fig. 1, a subgradient of the function ϕ_3 at a corner (where it is not differentiable) would be any vector pointing *away from* the polygon, whose normal hyperplane, the line normal to it through the corner point, did not intersect the interior. There may be an infinite number of subgradients for nondifferentiable convex functions at a given point while there is only one subgradient, the gradient, for differentiable ones. More discussion of the properties of convex optimization tools, along with other preliminaries, can be found in Ref. 15.

In general, an *ellipsoid* is defined to be a set of the form

$$\{\mathbf{x} \in \mathscr{R}^N : (\mathbf{x}-\mathbf{c})^T\mathbf{B}^{-1}(\mathbf{x}-\mathbf{c}) \leqslant 1\}, \tag{7}$$

where $\mathbf{B}$ is an $N \times N$ symmetric positive definite matrix that gives the ''size'' and orientation of the ellipsoid (the square roots of the eigenvalues of $\mathbf{B}$ are the lengths of the semiaxis), and $\mathbf{c} \in \mathscr{R}^N$ is its center. Thus the ellipsoid is parameterized by the matrix $\mathbf{B}$ and the vector $\mathbf{c}$.

The ellipsoid algorithm is initialized by choosing an ellipsoid that is large enough to contain the entire feasible region. The initial ellipsoid, for instance, can be chosen with center $\mathbf{c}=\mathbf{0}$ and with $\mathbf{B}=\alpha\mathbf{I}_N$, where α is a constant. Thus all points $\mathbf{x} \in \mathscr{R}^N$ satisfying

$$\mathbf{x}^T\mathbf{x} \leqslant \alpha, \tag{8}$$

are contained in this initial ellipsoid.

The central idea of the ellipsoid algorithm is to iteratively eliminate regions in $\mathscr{R}^N$ found to contain no feasible points. We use the hyperplanes associated with (sub)gradients to shrink the set that contains the points of interest. This shrinking process can be used to arrive at a feasible solution.

Geometrically, we can describe this key idea in the following way: we treat the center of the ellipsoid $\mathbf{c}$ as a candidate solution and evaluate the convex constraint $\phi(\mathbf{c})$. Either the constraint is satisfied, in which case $\mathbf{c}$ is a feasible solution, or it is not satisfied. In the latter case, if we evaluate the (sub)gradient of ϕ at $\mathbf{c}$, the associated hyperplane described above divides $\mathscr{R}^N$ into a set of points for which the constraint is ''less satisfied'' than at $\mathbf{c}$ and a set of points at which it is ''at least as satisfied.'' The goal is to find a new ellipsoid that contains all the points in the original ellipsoid which are on the ''at least as satisfied'' side of the hyperplane. Moreover, we want to find such an ellipsoid in a manner that ensures that it gets smaller. Finding this ellipsoid, in practice, means computing its associated matrix and center. Once the new center has been computed, it becomes the new candidate solution and we then repeat the constraint evaluation and ellipsoid update steps.

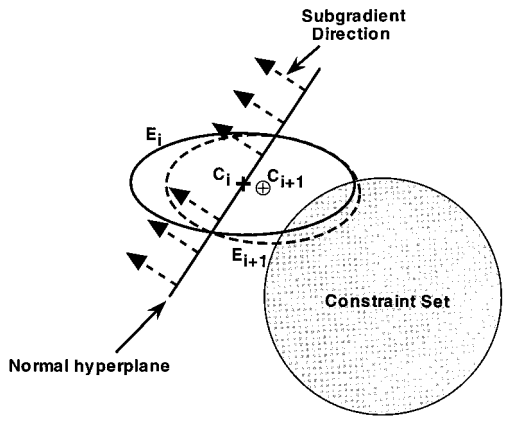

FIGURE 2. Illustration of one iteration of the ellipsoid algorithm. The solid ellipse E_i represents the ellipsoid at one iteration, with center C_i as shown by the "+." The shaded circle represents a constraint set. The hyperplane is normal to the gradient at the ellipse center. The dashed ellipse E_{i+1} is the new ellipsoid, with center C_{i+1} the new candidate solution for the next iteration.

This process is illustrated in Fig. 2. The solid ellipse labeled $\mathbf{E_i}$ is our original "ellipsoid," and the "+" sign labeled $\mathbf{C_i}$ marks its center. The shaded circle is the boundary of the constraint set, the arrows show the direction of the subgradient of the associated function at $\mathbf{C_i}$, and the normal hyperplane defines the undesirable half-space. The dashed ellipse $\mathbf{E_{i+1}}$ is the new, smaller ellipsoid, and the new candidate solution is its center $\mathbf{C_{i+1}}$. This process continues until the center of the ellipsoid enters the constraint set.

More precisely, for a convex function ϕ with $\mathbf{x}$ and $\mathbf{y}$ as two of its elements, with $\mathbf{g}$ a (sub)gradient ($\mathbf{g}$ is also used here for $\nabla_{\mathbf{x}}$ for simplicity), if $\phi(\mathbf{y}) \leqslant \phi(\mathbf{x})$, then

$$\mathbf{g}^T(\mathbf{y}-\mathbf{x}) \leqslant 0. \tag{9}$$

In particular, if a minimizer $\mathbf{x}^*$ exists, then

$$\mathbf{g}^T(\mathbf{x}^*-\mathbf{x}) \leqslant 0, \quad \forall \mathbf{x}. \tag{10}$$

By evaluating $\mathbf{g}$, we can construct a half-space that is guaranteed to contain any minimizer. The new half-space can be denoted as

$$\{\mathbf{y} \in \mathscr{R}^N : \mathbf{g}^T(\mathbf{y}-\mathbf{x}) \leqslant 0\}. \tag{11}$$

Now assume we have an initial ellipsoid with matrix $\mathbf{B}$ and center $\mathbf{c}$. If the constraint ϕ is not satisfied at $\mathbf{c}$, we can then identify a half-space guaranteed to contain the feasible points:

$$\{\mathbf{x} \in \mathscr{R}^N : (\mathbf{x}-\mathbf{c})^T\mathbf{B}^{-1}(\mathbf{x}-\mathbf{c}) \leqslant 1, \quad \mathbf{g}^T(\mathbf{x}-\mathbf{c}) \leqslant 0\}. \tag{12}$$

Now we need to find a smaller volume ellipsoid, represented by the matrix $\tilde{\mathbf{B}}$ and the vector $\tilde{\mathbf{c}}$,

$$\{\mathbf{x} \in \mathscr{R}^N : (\mathbf{x}-\tilde{\mathbf{c}})^T\tilde{\mathbf{B}}^{-1}(\mathbf{x}-\tilde{\mathbf{c}}) \leqslant 1\}, \tag{13}$$

which covers this half-space. It turns out[4] that an appropriate pair $(\tilde{\mathbf{B}}, \tilde{\mathbf{c}})$ can be computed as

$$\tilde{\mathbf{c}} = \mathbf{c} - \frac{\mathbf{B}\tilde{\mathbf{g}}}{N+1}, \tag{14}$$

and

$$\tilde{\mathbf{B}} = \frac{N^2}{N^2-1}\left(\mathbf{B} - \frac{2}{N+1}\mathbf{B}\tilde{\mathbf{g}}\tilde{\mathbf{g}}^T\mathbf{B}\right), \tag{15}$$

with

$$\tilde{\mathbf{g}} = \frac{\mathbf{g}}{\sqrt{\mathbf{g}^T\mathbf{B}\mathbf{g}}} \tag{16}$$

being a normalized subgradient.

Thus the basic ellipsoid algorithm can be summarized as the following.

Step (1), initialization:

(1) Set iteration counter $k = 0$.
(2) Choose the initial ellipsoid and center, such as $\mathbf{B}^{(0)} = \alpha\mathbf{I}_N$ and $\mathbf{c}^{(0)} = 0$.

Step (2), constraint evaluation: At iteration k, evaluate $\phi(\mathbf{c}^{(k)})$ and $\mathbf{g}^{(k)}$, any (sub)gradient of ϕ at the center $\mathbf{c}^{(k)}$ of the k^{th} ellipsoid defined by matrix $\mathbf{B}^{(k)}$. With multiple constraints, this step involves finding the ϕ_i that maximizes ϕ.

Step (3), branch: Check for the termination condition $\phi < 0$. If met, then stop; otherwise go to step (4). Thus, the process continues until an ellipsoid is found whose center meets all the conditions summarized in ϕ.

Step (4), update: Update the new ellipsoid by updating the matrix $\mathbf{B}^{(k+1)}$ and the center $\mathbf{c}^{(k+1)}$ as

$$\tilde{\mathbf{g}} = \frac{\mathbf{g}^{(k)}}{\sqrt{\mathbf{g}^{(k)^T}\mathbf{B}^{(k)}\mathbf{g}^{(k)}}}, \tag{17}$$

$$\mathbf{c}^{(k+1)} = \mathbf{c}^{(k)} - \frac{\mathbf{B}^{(k)}\tilde{g}}{N+1}, \tag{18}$$

$$\mathbf{B}^{(k+1)} = \frac{N^2}{N^2-1}\left(\mathbf{B}^{(k)} - \frac{2}{N+1}\mathbf{B}^{(k)}\tilde{\mathbf{g}}\tilde{\mathbf{g}}^T\mathbf{B}^{(k)}\right). \tag{19}$$

Step (5), loop: Increment k and go to step (2).

This sequence of ellipsoids is guaranteed to have decreasing volume, so that the algorithm will converge as long as the intersection of the constraint sets $\neq\varnothing$. The volume of an ellipsoid defined in Eq. (7) depends on the determinant of **B** and on the dimension of the space $\mathscr{R}^N$. More precisely, we have

$$V(E)=\sqrt{\det \mathbf{B}}\cdot V_N, \tag{20}$$

where V_N is the volume of the unit ball in $\mathscr{R}^N$. Even though the ellipsoid $E^{(k+1)}$ can be larger than the ellipsoid E^k in the sense of maximum semiaxis [$\lambda_{\max}(\mathbf{B}^{(k+1)})>\lambda_{\max}(\mathbf{B}^{(k)})$ is possible], it is shown in Ref. 15 that the volumes are related as follows:

$$V(E^{(k+1)})=\left(\frac{N}{N+1}\right)^{(N+1)/2}\left(\frac{N}{N-1}\right)^{(N-1)/2}V(E^{(k)}). \tag{21}$$

The last equation reveals, in addition, that the shrinking rate of the volume is rather slow and a large number of iterations may be required for convergence. However, since this volume relation only depends on the dimension N, and the initial ellipsoid contains a minimizing point in its interior, the ellipsoid algorithm converges to a feasible point. If there is no feasible point (e.g., the admissible set is empty), the algorithm will continue to iterate indefinitely, alternating between two or more constraints.

We note that the update of the ellipsoid matrix **B** in each iteration equation[19] is computationally expensive and can be avoided by computing the update of the center **c** recursively, as shown in Ref. 1. There is a large computational savings for small k but both computational and memory requirements grow with the iteration index. Combined recursive/block implementations are also possible.

RESULTS

In this section, we present a few examples of the application of the admissible solution approach to inverse electrocardiography. These examples include both dipole source simulations and simulations based on data recorded from experiments using canine hearts. The purpose of this section is to show that the approach gives reasonable results despite the lack of a global objective function, and to give a few illustrations of how it can be used to evaluate constraints and assumptions. Specifically, we will present

(1) results of a feasibility study showing that the method gives reasonable results (first reported in Ref. 2);
(2) an illustration of how the iterative solution can be used to reveal efficacy of constraints (first reported in Ref. 3); and
(3) an illustration of how the method can be used to study the interaction of modeling assumptions and constraints on inverse solutions.

Experimental Techniques and Error Evaluation

We include results based on two simulation scenarios, which we will denote DS, for ''dipole simulation,'' and TS, for ''tank simulation.'' Model DS had two different implementations, denoted DSa and DSb.

In both cases, simulated torso data were computed and then noise was added at specified signal-to-noise ratios, before inverse computing epicardial potentials. Epicardial potentials, either computed from dipole sources or measured from the dog hearts, were the true solution against which different inverse solutions were evaluated. These scenarios, with known epicardial data and computed torso data, allowed us to explicitly compare various inverse solutions.

Dipole Source Simulation, DS. The simulation model DS employed

(1) a single fixed dipole as the equivalent cardiac source;
(2) a forward solution for a homogeneous (i.e., epicardium and body surface only) geometric model based on a human subject (the Dalhousie torso[24]); and
(3) body surface potentials computed in two different ways as described below.

The geometric model used with the dipole source was based on a single subject and consisted of 352 torso nodes connected to form 700 triangles and 98 epicardial nodes connected to form 192 triangles.[24] The cardiac source was a single fixed current dipole located near the center of the left ventricle. We aligned the dipole with the X, Y, and Z axes of the torso geometry to generate different potential distributions, and also generated linear combinations of these three orientations. Epicardial and torso surface potentials were then calculated by two different numerical pathways, producing two distinct implementations of simulation DS:

(1) DSa, dipole computed: both epicardial and torso potentials were computed directly from the dipole source without using the forward model in the computation of the torso potentials.
(2) DSb, forward computed: epicardial potentials were computed from the dipole source. Then, using these computed epicardial potentials as an equivalent source, torso potentials were calculated using the forward solution matrix **A**.

Figure 3 shows isopotential maps for both sets of torso potentials. The potentials produced by these two

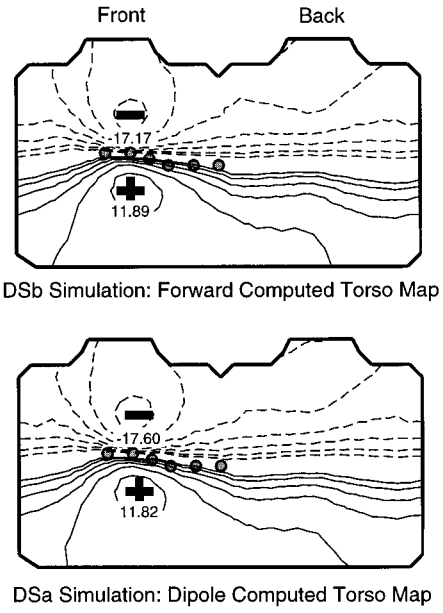

FIGURE 3. Forward computed and dipole computed torso potentials, as described in the text. The isopotential maps were calculated after projecting the potentials onto a two-dimensional surface. The left (right) side of each map represents the projection of the anterior (posterior) torso surface. The contour levels are arranged in logarithmic steps and the location and magnitude of the maximum and minimum are shown by a (+) and (−), respectively. Solid lines show positive potentials and dashed lines negative potentials.

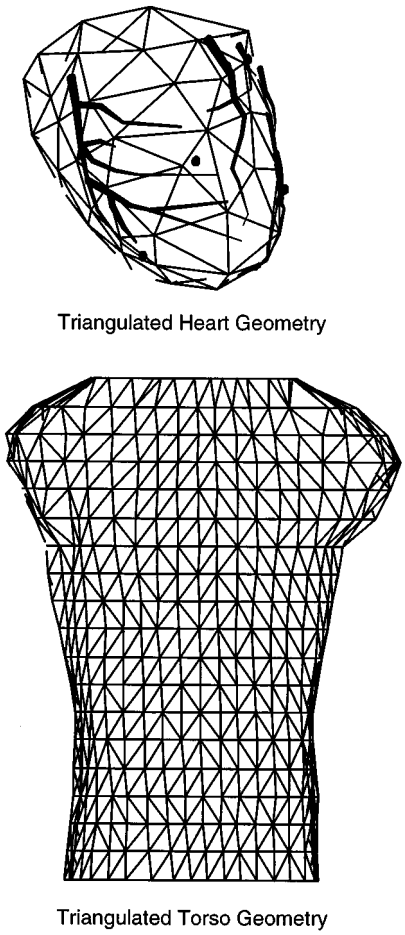

FIGURE 4. Triangulated surface of isolated canine heart geometry (top) and of the torso tank in which it was suspended (bottom). Thick lines in the heart geometry indicate the locations of major coronary arteries.

techniques are similar in shape and slightly different in amplitude. We used these two distinct sets of dipole-based torso surface data to study the relationship between modeling assumptions and multiple constraints. Since the second of these data sets, DSb, uses the same forward model in inverse solutions as was used to generate the data, we refer to DSb as the ''exact model'' case. In contrast, for the scenario DSa, the forward model used in the inverse solution does *not* match exactly the way the data were generated, so we refer to this as the model error, or model mismatch, case.

Two major limitations of the dipole simulation experiments are the difficulty of reproducing realistic epicardial potential distributions (due to the simplicity of the source model) and the lack of a natural way to produce the realistic time-varying epicardial distributions needed to test temporal constraints.

Torso Tank Preparations, TS. To test our inverse methods with time-varying measured data, we used the tank model, TS, which included

(1) epicardial data recorded from an isolated canine heart preparation in a fiberglass tank molded in the shape of an adolescent human torso;
(2) a homogeneous forward solution based on the geometry of the tank; and
(3) body surface potentials computed from the epicardial data and the forward model.

The perfused, isolated canine heart was suspended in a tank molded in the shape of an adolescent thorax (Fig. 4 contains an illustration of the torso tank and heart surface geometry). The tank was filled with electrolytic solution at a resistivity representative of a typical human thorax (500 Ω cm). Epicardial potentials were recorded at a sampling rate of 1000 Hz per channel via 64 electrodes sewn into a nylon sock placed over the isolated heart.[26] Within a 4–7 s recording window, individual beats were either selected or averaged to obtain a representative beat for that window.

Evaluation of Results. We evaluated our results using two error measures that are standard in the literature, relative rms error (RE) and correlation coefficients

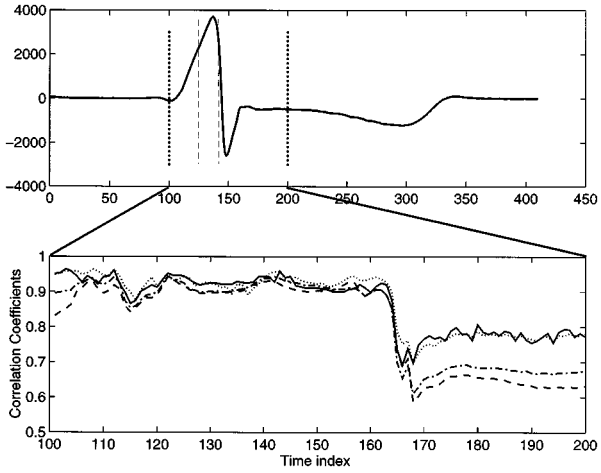

FIGURE 5. Top: An epicardial electrogram. Results are shown in the bottom figure for the portion of the cycle between the dotted vertical lines. Dashed vertical lines show time instants illustrated in the next figure. Bottom: correlation coefficients (CC) for various reconstructions at a SNR of 40 dB. The solid line shows the CC's for convex optimization results and the dotted, dashed, and dash-dot lines show the CC's for energy, Laplacian, and joint energy/Laplacian regularization, respectively.

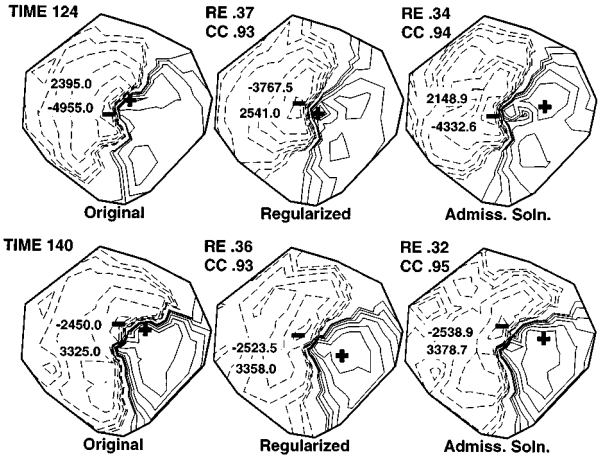

FIGURE 6. This figure displays the inverse computed isopotential maps for two time instants. The time instants are 124 and 140 ms after the stimulus that initiated the heart beat, as indicated by the dashed vertical lines in the previous figure.

(CC),[33] together with visual inspection of epicardial time signals (electrograms) and isopotential maps (IPM). On each IPM, the location and magnitude of the maximum and the minimum are indicated by (+) and (−) signs, respectively, and the contours are drawn as interpolated equipotential lines over a polar two-dimensional projection of the epicardial surface with the apex at the center and the A/V ring around the outside. Solid lines represent positive potentials and dashed lines negative potentials. The contours are drawn in logarithmic steps based on the largest absolute extremum in each map. One map is drawn for each time instant, and where appropriate, the maps are then presented as an IPM sequence.

Feasibility Study

In the feasibility study, we used a small number of constraints, assumed we knew the exact value of these constraints, and then compared the results of the admissible solution inverse method with some standard methods in a simple test scenario. We used five typical ''standard'' constraints as described in the section on standard constraints:

(1) the square of the 2-norm of the residual;
(2) the square of the 2-norm of the Laplacian of the solution;
(3) the square of the 2-norm of the solution;
(4) the square of the 2-norm of the Tikhonov error using an identity matrix as the regularizer; and
(5) the max norm of the solution.

The constraint values (the ϵ's in the section on standard constraints) were calculated based on the measured epicardial potentials. We note that the true epicardial solutions were not used *during* the inverse procedure, but only to obtain the constraint values before beginning. This is analogous to comparing standard regularization methods using their respective optimal regularization parameters, as in Ref. 29.

Thus, we used a very limited subset of available constraints, but with an unrealistically accurate foreknowledge of the correct value of the constraints. The goal here was simply to establish whether such a procedure, which does not pose any overall objective function, can arrive at solutions that are comparable in accuracy to those using traditional methods. The initial ellipsoid was taken as $\mathbf{B}=\alpha\mathbf{I}_N$ with $\alpha=15\|\mathbf{y}\|_2^2$ and initial center $\mathbf{c^0}=0$.

We performed numerical experiments using both dipole and tank scenarios, but will illustrate only our results with tank data (TS) and a SNR of 40 dB. The bottom panel in Fig. 5 shows CC's as a function of time for four different inverse methods over the portion of the cycle marked by the dotted vertical lines in the upper panel. The four inverse methods were: admissible solution approach (solid line), two-norm regularization (dotted line), Laplacian regularization (dashed line), and joint two-norm/Laplacian regularization (dash-dot). To illustrate the spatial behavior of the solution, in Fig. 6, we show isopotential maps of the original measured potentials (on the left) and two inverse reconstructions, using two-norm regularization (middle) and the admissible solution method (right), at two time instants during QRS. The RE and CC values are shown in Fig. 6 for each reconstruction.

Evaluation of Constraint Efficacy

As described in the section on the admissible solution approach, at each iteration we selected the ''worst'' constraint (i.e., the one that achieved the maximum of ϕ) to control the next iteration. Thus, by monitoring which constraint was chosen by the algorithm at each iteration for a particular simulation, we were able to observe which constraints were active and how they interacted as the ellipsoid shrank. We performed many tests using this approach,[1] and here present some illustrations.

Comparing Unweighted and Weighted Norm Constraints. To illustrate this procedure, we present results of two tests using the dipole model DSb at 30 dB SNR. The constraints used in the first test were the same as in the feasibility test. In a second simulation, we replaced the standard two-norm and two-norm-of-Laplacian constraints with *locally weighted* two-norm and two-norm-of-Laplacian constraints. (This is similar in principle to the approach reported in Ref. 21.) The spatial weighting was computed, based on the true solutions, to emphasize each constraint in spatial regions where the amplitude (Laplacian) was small, and deemphasize it in regions where the amplitude (Laplacian) was large. This was accomplished by using a diagonal weighting matrix $\mathbf{W}$ to replace a constraint based on $\phi(\mathbf{x})=\|\mathbf{Rx}\|_2^2$ with one based on $\phi(\mathbf{x})=\|\mathbf{WRx}\|_2^2$. The diagonal elements of $\mathbf{W}$ were weight factors calculated at each node as the ratio of the two-norm of the amplitude (Laplacian) over the whole surface to the two-norm of the amplitude (Laplacian) over the neighborhood of the node. We noted that this is the same as using a *weighted* two-norm where the weighting matrix is diagonal with the square roots of the elements of $\mathbf{W}$ as its nonzero elements.

In Fig. 7, we illustrate a typical result: the top figure shows the correlation coefficient for both methods as a function of the iteration number. The bottom figures show which constraint the algorithm chose at each iteration, coded according to the numbering in the feasibility study section. The left pair of figures shows results with standard constraints and the right pair with weighted constraints.

Efficacy of Temporal Constraints. Here, we present a sample of our results using tank data (TS) and the augmented forward model with ten time samples ($L=10$) and global spatial and global temporal constraints (the latter based on a simple high-pass temporal filter[6]). The temporal constraint used was on the two-norm of the output of this filter, and the spatial constraints were on the two-norm of the residual, amplitude, etc., where all norms are over all time instants and spatial locations included in the augmented model. In general, we observed that at relatively high SNR (40 dB), the temporal

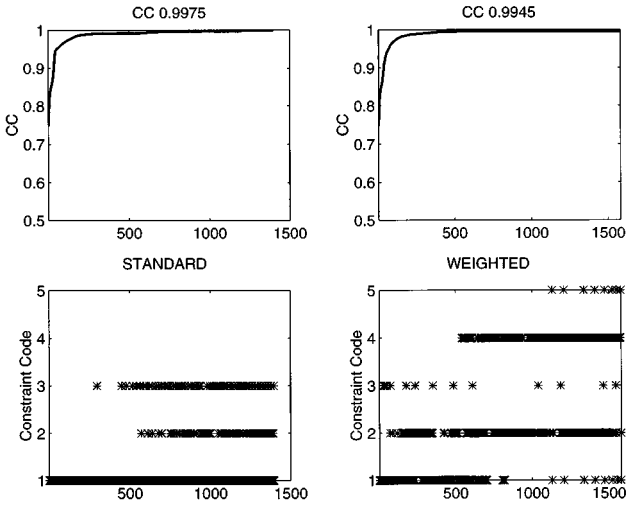

FIGURE 7. Correlation coefficients (top) and active constraint (bottom) as functions of iteration index for DSb simulations, with the dipole source in the *Y* direction at 30 dB SNR using ''exact model'' data. The left two figures show results with standard constraints and the right two with weighted constraints. Constraint coding is described in the text.

constraints were used by the algorithm only at the very end, just before convergence, while at a somewhat lower SNR (30 dB), the temporal constraints were used much earlier. In Figs. 8 and 9, we illustrate these results at 30 dB for an interval of 10 ms (ten time samples) early in QRS. All constraints used were global over all time instants and spatial locations. In addition to the five constraints used in the previous example, we used a global temporal constraint in the form of the two-norm of a simple high-pass filter, as described in the spatiotemporal constraints section.

In Fig. 8, we show examples of time signals from four epicardial locations. The inclusion of the temporal constraint tends to improve reconstructed signal shape, and as can be seen in the signal at electrode No. 18, even when the result is less accurate, it is generally less noisy. The bottom of Fig. 8 contains comparisons of CC and RE and shows which constraints were employed as a function of iteration number during the ellipsoid algorithm. We note that the coding is the same as in Fig. 7, except that the temporal constraint has been inserted as No. 2 and the codes of the other constraints incremented by one. Figure 9 compares an isopotential map sequence from this experiment: the left column shows the original data (shown every 2 ms), while the middle column shows the reconstruction with temporal constraints and the right column without. To make the comparison clearer, the same isopotential contour spacing has been used on all maps. Most notable is the increased density of the isopotential lines in a number of the temporally constrained maps and the somewhat more accurate shape

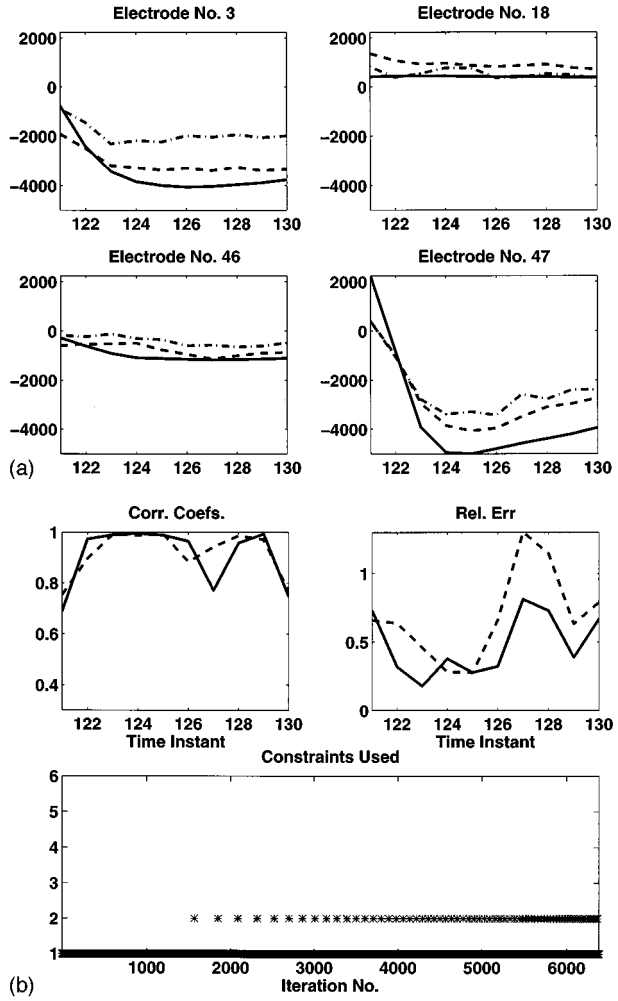

FIGURE 8. Comparison of admissible solutions with and without a global temporal constraint for a 10 ms interval early in QRS: The four panels show time signals at four epicardial locations. The solid lines are the originals and the dashed (dot-dash) lines reconstructions with (without) temporal constraints. The next pair of panels show correlation coefficients and relative error over the 10 ms interval with (solid line) and without (dashed line) the temporal constraints. Constraint usage is shown in the bottom panel: constraint No. 1 is the residual error two-norm and No. 2 is the temporal constraint.

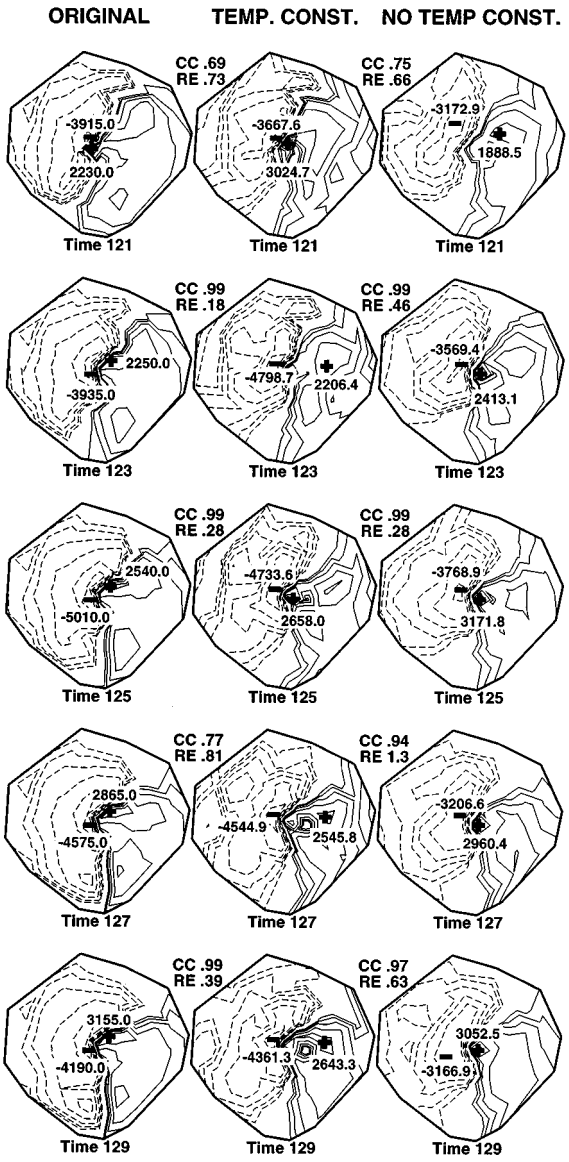

FIGURE 9. Same scenario as in previous figure, but here, comparing isopotential maps every 2 ms.

of the wave front in the later maps in the sequence.

Interaction of Model Assumptions with Constraints

We can also observe how the efficacy of various constraints may change depending on our model assumptions by using the model mismatch dipole body surface data DSa. In the DSb test described above, the true residual error norm is proportional to the noise standard deviation. If we use this as the residual constraint bound, then it is a tight and accurate bound. In the DSa simulation, however, the true residual error between the noisy simulated torso data and the forward-projected epicardial solution is larger than in DSb; it includes both geometric model error and noise. If we use a constraint bound based on the noise only, then it will be ''too tight'' and the algorithm will tend to overfit the data. Some effects of these different model assumptions are shown in Figs. 10 and 11. Figure 10 shows results with the mismatched DSa model when the residual error constraint bound included both the noise and the model error and was, thus, the ''true'' residual error. Figure 11 shows the results with the same DSa model but when the bound was based only on the noise—in other words, the error in the forward model was ignored, and the residual constraint was too tight. The format for both Figs. 10 and 11 is the same as in Fig. 7, showing results for standard and weighted constraints. As can be seen from Fig. 7, in the

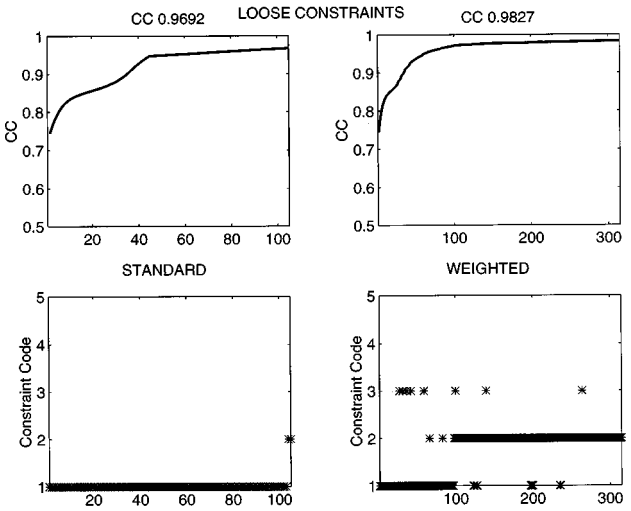

FIGURE 10. Correlation coefficients (top) and active constraint (bottom) as functions of iteration index for DSa simulations, with the dipole source in the *Y* direction at 30 dB SNR using model mismatch data and accurate (loose) residual constraint. Note that the number of iterations is different in the two cases shown. Layout as in the previous figure.

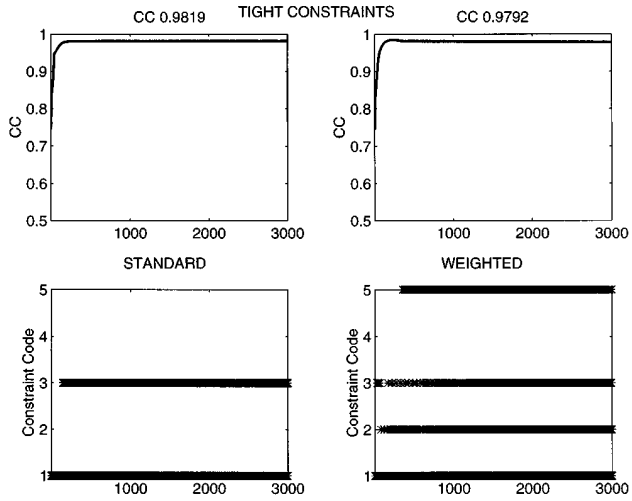

FIGURE 11. Same as in the previous figure, but with the residual constraint bound based only on additive noise variance, and therefore unrealistically tight.

accurate bound case the algorithm stopped after a relatively small number of iterations because the residual bound was so loose and none of the other constraints were effective in further constraining the solution. By contrast, applying the overly tight residual constraint caused more constraints to be employed and the algorithm continued for over ten times as many iterations.

DISCUSSION AND CONCLUSIONS

In this section, we will discuss some of the important features of the admissible solution approach presented in this paper, point out some implications of the simulation results, and present conclusions and current and future research directions.

Discussion of the Method

The admissible solution method proposed in this paper for the inverse problem of electrocardiography has certain advantages over traditional regularization approaches. It is flexible in the number and type of constraints that can be incorporated, requiring only that each constraint be defined by a convex set in the solution space. Nonstandard constraints such as locally weighted constraints, frequency domain constraints, nondifferentiable constraints, etc., can be incorporated without any particular difficulty. The need for regularization parameters is replaced by a need for bounds on the constraints; while this involves the use of *a priori* information, the bounds can be related to measurable physiological quantities. One can ensure ''almost-Tikhonov'' performance by incorporating the Tikhonov error as a constraint in the procedure. In fact, one could incorporate Tikhonov constraints for more than one regularization function and more than one value of the regularization parameter simultaneously. Because of the iterative nature of numerical solutions to the resulting convex optimization problem, one can monitor which constraints are employed by the algorithm as it iterates in order to gain understanding of the effect (or lack of effect) of a particular constraint or constraints on the solution. For instance, one can find a solution with a given set of constraints and then repeat the solution while leaving out or adding one constraint at a time, in order to closely study the efficacy of a particular constraint.

As noted in the introduction, disadvantages of the approach include the lack of an objective function with a unique solution and increased computational load. Running our simulations in MATLAB on a midlevel SGI workstation with no particular attempt to optimize our code, the algorithm takes about 6 min to run for each time instant. Using the temporal constraint on an interval of ten time instants as described earlier, it took 24 h or more to run.

One issue that could surface, in principle, is what happens if the admissible set was empty, in which case the algorithm would tend to run without stopping. In our experience, this only happened when there was a coding or mathematical error, and has been easy to spot as the number of iterations was clearly excessive. Thus, in practice, after a little experience one can put a ceiling on the number of iterations the algorithm is allowed to run. If this ceiling is reached, it can be taken as a warning sign that something is wrong. At a minimum, the bounds can be relaxed on the constraints on which the algorithm is stuck and the algorithm restarted.

Discussion of Simulation Results

We do not wish to overemphasize the significance of the simulation results presented here; as we state in the results section, our purpose in this paper is to present the concepts behind the method and illustrate that it gives reasonable results. However, we offer the following conjectures based on our present results:

(1) The results of the feasibility study show that the method produces reasonable inverse solutions using only a small number of constraints.

(2) When we compare results obtained using weighted constraints to those obtained with standard constraints, two phenomena that occur are as follows: (a) Due to the type of weighting used, as explained in the section on the evaluation of constraint efficacy, constraints on weighted features of the solution are more likely to be used by the algorithm (see Figs. 7, 10, and 11). We presume that this is because weighted constraints, especially when derived from the actual correct answer, fit the candidate solutions better and are thus more likely to dominate the overall constraint and thereby actively direct the iteration path. (b) Although correlation coefficients are usually comparable, the maxima and minima of reconstructed isopotential maps are generally closer in both amplitude and location to the correct values, i.e., less over-regularized, with weighted than with unweighted constraints.

(3) When there is a mismatch between the forward and inverse models, there seems to be a complicated interaction between the constraints used, the values chosen for the constraint bounds (related to the degree of *a priori* knowledge about the error), and the resulting inverse solutions. For instance, in Fig. 10, where the residual constraint is quite loose (since it takes into account the known model error) the advantage of using weighted constraints over standard ones is considerable (isopotential maps bear this out).

(4) When constraints are too tight (as for example in Fig. 11), results are reasonably accurate as long as enough constraints are used (this was borne out by other simulations, not reproduced here, in which constraints were added or removed one by one[1]).

Conclusions and Future Work

We are encouraged by the preliminary results presented here to continue to develop several aspects of this method:

(1) One major area needing further study is how the method behaves when the constraints used are based on less *a priori* knowledge than in this study. We have done some preliminary testing indicating that, at least for inaccuracy in the bounds of less than one order of magnitude, the deterioration in accuracy and reliability of the reconstructions is gradual and ameliorated by the inclusion of more constraints. However, this behavior is of obvious importance and will be the subject of a separate study.

(2) As mentioned previously, we have begun the process of studying recorded epicardial data to develop a library of realistically derived constraints. We will then apply these constraints to inverse solutions using the approach presented here.

(3) We are currently working on a more computationally efficient, parallel implementation of convex optimization algorithms, which will allow us to incorporate a longer time interval into the augmented model, and thus permit the use of promising frequency-domain constraints.

(4) We are exploring analytical methods to measure changes in the ''size'' of the intersection of the constraint sets when different combinations of constraints are used, by studying the problem in a more tractable geometric setting.

(5) We are exploring the use of more efficient optimization algorithms.[28,34]

(6) We feel that the issue of interaction between model error and inverse solutions is a very important one for any practical application of inverse electrocardiography. We tried to address some of the implications of the problem of unmodeled error in the forward solution in a simplified way in some of the dipole simulations described in this paper. We hope to be able to use torso tank experimental results to study the effects of such modeling error in a more realistic setting in the near future.

ACKNOWLEDGMENTS

The authors wish to express their gratitude to Clas Jacobson for many fruitful and insight-producing discussions of this work, to Kadagattur Srinidhi for help with some of the simulations and figures, and Dr. Bruno Taccardi for assistance in obtaining measured data. This material is based upon work supported by the National Science Foundation under Grant No. BCS-9309359, Grant No. HL52338 from the National Institutes of Health, and awards from the Nora Eccles Treadwell Foundation and the Richard A. and Nora Eccles Harrison Fund for Cardiovascular Research.

REFERENCES

[1] Ahmad, G. F. Inverse electrocardiography by simultaneous and iterative imposition of multiple constraints. Ph.D. Thesis, Northeastern University, 1995.

[2] Ahmad, G. F., D. H. Brooks, C. Jacobson, and R. S. MacLeod. A feasibility study of inverse electrocardiography by convex optimization. In: Proceedings of the 21st N. E. Bioengineering Conference, 1995, pp. 245–246.

[3] Ahmad, G. F., D. H. Brooks, C. A. Jacobson, and R. S. MacLeod. Constraint evaluation in inverse electrocardiography using convex optimization. In: Proceedings of the IEEE Engineering in Medicine and Biology Society 17th Annual International Conference, 1995, pp. 209–210.

[4] Boyd, S. P., and C. H. Barret. Linear Controller Design. Englewood Cliffs, NJ: Prentice-Hall, 1991.

[5] Brooks, D. H., G. F. Ahmad, and R. S. MacLeod. Multiply constrained inverse electrocardiography: Combining temporal, multiple spatial, and iterative regularization. In: Proceedings of the IEEE Engineering in Medicine and Biology Society 16th Annual International Conference, 1994.

[6] Brooks, D. H., G. F. Ahmad, R. S. MacLeod, and G. M. Maratos. Inverse electrocardiography by simultaneous imposition of multiple constraints. *IEEE Trans. Biomed. Eng.* (submitted).

[7] Brooks, D. H., C. L. Nikias, and J. H. Siegel. An inverse solution in electrocardiography in the frequency domain. In: IEEE Engineering in Medicine and Biology Society 10th Annual International Conference, 1988, pp. 970–971.

[8] Censor, Y. Iterative methods for the convex feasibility problem. *Ann. Discrete Math.* 20:83–91, 1984.

[9] Combettes, P. L. The foundation of set theoretic estimation. *Proc. IEEE* 81:182, 1993.

[10] Cuppen, J. J. M., and A. van Oosterom. Model studies with the inversely calculated isochrones of ventricular depolarization. *IEEE Trans. Biomed. Eng.* BME-31:652–659, 1984.

[11] Davenport, D. M., D. H. Brooks, and R. S. MacLeod. Experimentally derived realistic constraints on epicardial potential distributions. In: Proceedings of the 21st N.E. Bioengineering Conference, March 1995.

[12] El-Jakl, J., F. Champagnat, and Y. Goussard. Time–space regularization of the inverse problem of electrocardiography. In: Proceedings of the IEEE Engineering in Medicine and Biology Society 17th Annual International Conference, 1995.

[13] Greensite, F. Well-Posed formulation of the inverse problem of electrocardiography. *Ann. Biomed. Eng.* 22:172–183, 1994.

[14] Greensite, F., Y. J. Qian, and G. Huiskamp. Myocardial activation imaging: A new theorem and its implications. In: Proceedings of the IEEE Engineering in Medicine and Biology Society 17th Annual International Conference, 1995, pp. 205–206.

[15] Grotschel, M., L. Lovasz, and A. Schrijver. Geometric Algorithms and Combinatorial Optimization. Berlin, Germany: Springer-Verlag, 1988.

[16] Gulrajani, R. M., F. A. Roberge, and G. E. Mailloux. The forward problem of electrocardiography. In: Comprehensive Electrocardiology, edited by P. W. Macfarlane and T. D. Veitch Lawrie. Oxford, England: Pergamon, 1989, pp. 197–236.

[17] Gulrajani, R. M., F. A. Roberge, and P. Savard. The inverse problem of electrocardiography. In: Comprehensive Electrocardiology, edited by P. W. Macfarlane and T. D. Veitch Lawrie. Oxford, England: Pergamon, 1989, pp. 237–288.

[18] Hansen, P. C. Regularization tools, a MATLAB package for analysis and solution of discrete ill-posed problems. Technical report, UNI.C, DK-2800 Lyngby, Denmark, June, 1992.

[19] Huiskamp, G., and A. van Oosterom. The depolarization sequence of the human heart surface computed from measured body surface potentials. *IEEE Trans. Biomed. Eng.* BME-35:1047–1059, 1988.

[20] Huiskamp, G., A. van Oosterom, and F. Greensite. Physiology based constraints in the inverse problem of electrocardiography. In: Proceedings of the IEEE Engineering in Medicine and Biology Society 17th Annual International Conference, 1995, pp. 207–208.

[21] Iakovidis, I., and R. M. Gulrajani. Improving Tikhonov regularization with linearly constrained optimization: Application to the inverse epicardial potential solution. *Math. Biosci.* 112:55–80, 1992.

[22] Johnson, C. R., and R. S. Macleod. Local regularization and adaptive methods for the inverse Laplace problem. In: 2nd Gauss Symposium: Medical Mathematics and Physics, edited by D. N. Ghista. Wiesbaden, Germany: Vieweg, 1996.

[23] Khoury, D. Use of current density in the regularization of the inverse problem of electrocardiography. In: Proceedings of the IEEE Engineering in Medicine and Biology Society 16th Annual International Conference, 1994, pp. 133 and 134.

[24] MacLeod, R. S. Percutaneous Transluminal Coronary Angioplasty as a Model of Cardiac Ischemia: Clinical and Modelling Studies. Ph.D. Thesis, Dalhousie University, Halifax, N.S., Canada, 1990.

[25] MacLeod, R. S., and D. H. Brooks. Recent progress in inverse problems in electrocardiology. *IEEE Eng. Med. Biol. Mag.* (in press).

[26] MacLeod, R. S., B. Taccardi, and R. L. Lux. Electrocardiographic mapping in a realistic torso tank preparation. In: Proceedings of the IEEE Engineering in Medicine and Biology Society 17th Annual International Conference, 1995, pp. 245 and 246.

[27] Messinger-Raport, B. J., and Y. Rudy. Noninvasive recovery of epicardial potentials in a realistic heart–torso geometry. *Circ. Res.* 66:1023–1039, 1990.

[28] Nesterov, Y., and A. Nemirovksy. Interior point polynomial algorithms in convex programming. Philadelphia, PA: SIAM, 1994.

[29] Oster, H. S., and Y. Rudy. The use of temporal information in the regularization of the inverse problem of electrocardiography. *IEEE Trans. Biomed. Eng.* 39:65–75, 1992.

[30] Ozkan, M. K., A. M. Tekalp, and M. I. Sezan. POCS-based restoration of space-varying blurred images. *IEEE Trans. Image Process.* 3:450–454, 1994.

[31] Pilkington, T. C., and R. Plonsey. Engineering Contributions to Biophysical Electrocardiography, New York, NY: IEEE, 1982.

[32] Plonsey, R. Bioelectric Phenomena, New York: McGraw-Hill, 1969.

[33] Rudy, Y., and B. J. Messinger-Rapport. The inverse solution in electrocardiography: Solutions in terms of epicardial potentials. *Crit. Rev. Biomed. Eng.* 16:215–268, 1988.

[34] Sabharwal, A. S., and L. C. Potter. Set estimation via ellipsoidal approximations. *IEEE Trans. Signal Process* 2:897–900, 1995.

[35] Shahidi, A. V., P. Savard, and R. Nadeau. Forward and inverse problems of electrocardiography: Modeling and recovery of epicardial potentials in humans. *IEEE Trans. Biomed. Eng.* BME-41:249–256, 1994.

[36] Throne, R. D., L. G. Olson, T. J. Hrabik, and J. R. Windle. Generalized eigensystem techniques for the inverse problem of electrocardiography applied to a realistic heart–torso geometry. *IEEE Trans. Biomed. Eng.* 44:447, 1997.

[37] Youla, D. C., and H. Webb. Image restoration by the method of convex projection: Part 1-Theory. *IEEE Trans. Med. Imaging* 1:81–94, 1982.