

Resolution Strategies for the Finite-Element-Based Solution of the ECG Inverse Problem

Dafang Wang*, *Student Member, IEEE*, Robert M. Kirby, *Member, IEEE*, and Chris R. Johnson, *Senior Member, IEEE*

Abstract—Successful employment of numerical techniques for the solution of forward and inverse ECG problems requires the ability to both quantify and minimize approximation errors introduced as part of the discretization process. Our objective is to develop discretization and refinement strategies involving hybrid-shaped finite elements so as to minimize approximation errors for the ECG inverse problem. We examine both the ill-posedness of the mathematical inverse problem and the ill-conditioning of the discretized system in order to propose strategies specifically designed for the ECG inverse problem. We demonstrate that previous discretization and approximation strategies may worsen the properties of the inverse problem approximation. We then demonstrate the efficacy of our strategies on both a simplified and a realistic 2-D torso model.

Index Terms—Adaptive refinement, ECG forward problem, ECG inverse problem, finite-element method (FEM), resolution studies.

I. INTRODUCTION

THE numerical solution of ECG forward and inverse problems has received considerable attention as a means of providing insight into the connection between observable data (ECG signals) and the underlying biophysical phenomena that they represent. Both forward and inverse modeling follow a similar simulation science pipeline. First, a set of model equations is selected that mathematically articulates the biophysical process of electric potential propagation within the torso volume. Second, geometric models are generated that capture the domain over which the mathematical models operate. In the case of the ECG problem, these models normally consist of a geometric discretization of the heart surface, the torso surface, and relevant organs and critical structures. Third, the model equations are then discretized to form a numerical system, the solution of which approximates the “true” solution of the model system. Given the motivating goal of understanding complex biophysical phenomena, it is natural to inquire as to how accurately this pipeline reflects the process of interest. The means by which one answers this question is the final step within the

simulation science pipeline, which is a step that is essential to providing scientific confidence in the results obtained. Within the engineering literature, this final step is known as the process of *validation and verification (V&V)* [1].

As its name implies, V&V partitions the evaluation of the simulation science pipeline into two parts. The validation process is concerned with how faithfully the mathematical and geometric models represent the phenomena of interest. The verification process, on the other hand, is concerned with how accurately one’s numerical and geometric discretizations approximate the aforementioned models. Whenever numerical methods are employed for the solution of science and engineering problems, it is important to understand the impact of one’s discretization choice on how well the process of interest is being approximated. Most numerical method practitioners tackle this issue through the development of “refinement” strategies, which specify how one decreases errors by increasing the resolution (or fidelity) of the numerical approximation at the cost of increased computational work. With such strategies in place, one can specify an acceptable discrepancy between the true and approximate solutions and can tune (or refine) the numerical and geometric approximations accordingly. Most of the literature on adaptive refinement strategies has been targeted toward the solution of “forward” simulations, which are mostly well-posed problems [2], [3]. This paper intends to develop discretization and refinement strategies to be employed when solving the inverse ECG problem with the finite-element method (FEM).

Continuous inverse problems are often ill-posed in the Hadamard sense, in that the existence, uniqueness, and/or stability of the solution may not be guaranteed. Because of this ill-posedness, the discrete version of an inverse problem can be severely ill-conditioned, requiring numerical techniques to address the ill-conditioned systems. The term “regularization” is used to denote a class of techniques for constraining the original ill-posed problem so as to yield a somewhat better-posed problem. Solutions to the better-posed problem are interpreted as constrained approximations of the original problem of interest. The past few decades have seen the creation of a wide range of regularization methods, both theoretical and practical [4]–[9]. In the context of our studies, it is worth noting that one form of regularization to overcome the ill-posedness of inverse problems is the discretization itself [10], which is a method referred to as “self-regularization” or “regularization by projection.” Theoretical optimal discretization formulations have been proposed from both functional analysis [11], [12] and Bayesian statistics [13], [14] viewpoints. The goal of our work is to study regularization by discretization and refinement at a practical level and to develop optimal discretization strategies

Manuscript received February 14, 2009; revised April 30, 2009. First published June 16, 2009; current version published January 20, 2010. This work was supported in part by the National Science Foundation (NSF) Career Award NSF-CCF0347791 and in part by the National Institutes of Health (NIH) National Center for Research Resources (NCRR) Center for Integrative Biomedical Computing under NIH NCRR Grant 5P41RR012553-10. Asterisk indicates corresponding author.

*D. Wang is with the Scientific Computing and Imaging (SCI) Institute and the School of Computing, University of Utah, Salt Lake City, UT 84112 USA (e-mail: dfwang@sci.utah.edu).

R. M. Kirby and C. R. Johnson are with the Scientific Computing and Imaging (SCI) Institute and the School of Computing, University of Utah, Salt Lake City, UT 84112 USA (e-mail: kirby@cs.utah.edu; crj@sci.utah.edu).

Digital Object Identifier 10.1109/TBME.2009.2024928

that can be employed for computational inverse problems. The discretization techniques we have developed can also be used in combination with other classical regularization methods, such as the widely used class of Tikhonov-based methods [4], [5], so as to provide additional improvement to the inverse solution accuracy.

A. Related Work

Two commonly used numerical methods employed for the solution of ECG problems are the boundary element method (BEM) and the FEM. Both are categorized as variational methods: they seek to satisfy the differential (model) equations in a variational (or weak) sense. Each method has strengths and weaknesses in terms of accuracy, fidelity, and computational cost. Practitioners must balance these in making decisions as to which method to employ. No matter which method is chosen, however, one must deal with the verification process. Both methods require the practitioner to understand the impact of resolution on the problem of interest. We now provide an overview of some relevant research literature.

A comparative study applying both the BEM and the FEM to ECG problems has indicated that under similar discretization levels, the BEM yields smaller errors and consumes less computation time but requires more memory than does the FEM [15]. On the other hand, the FEM better accounts for the anisotropic conductivities of human bodies. Discretizing the potential field explicitly, the FEM allows a flexible investigation of the impact of resolution in the verification process. In particular, adaptive finite-element refinement strategies that are widely used in many engineering fields can be applied to the ECG problem [2], [16], [17]. Such strategies are normally based on certain element-wise error estimators and tend to refine the regions where the spatial gradients of the field are high. The adaptation process in such schemes is usually determined by a physiologically based stopping criterion, e.g., a minimax condition [17]. The efficacy and efficiency of such spatial adaptive refinements in simulating forward ECG problems compared to conventional uniform refinements have been demonstrated in [2], supported by similar results reported by an adaptive BEM study [18].

Another somewhat indirect spatial adaption method for ECG problems is to partition the transfer matrix resulting from the FEM into several submatrices, each of which is relevant to a local region of the biophysical potential field, and then, to apply a local regularization procedure specifically to each submatrix [19]. Since discretization is one form of regularization, the rationale of localized regional regularization coincides with the adaptive spatial refinement.

In addition to the aforementioned h-type refinements, which spatially refine the elements, p-type refinements, which approximate the problem using higher order basis polynomials, have also been explored for forward/inverse ECG problems. For example, a finite-element/boundary element model based on cubic Hermite interpolation was proposed in [20], and it was shown in [21] that high-order quadrilateral elements could significantly improve the numerical quality of inverse ECG solution.

Most of the referenced refinement studies report that while increasing the numerical resolution beyond a certain point may still further improve the accuracy of the forward ECG problem, the ill-conditioning of the inverse problem is worsened, and thus, the solution accuracy is diminished. Multilevel methods have been proposed as one possible approach to alleviate this issue, both in theory [22] and in practice [23]. An algebraic multilevel method (AMG) has been proposed that allows more automatic refinements and better stability in the presence of discontinuous coefficients and boundary conditions than ordinary multigrid methods [24]. However, multilevel studies have not addressed the issue of optimization of discretization specifically for inverse ECG problems.

To summarize, there is a notable gap in the current ECG literature concerning the impact of resolution on the practical forward and inverse problems for the FEM. Although the impact of resolution on the epicardium and the body surface has been previously investigated [25], it still remains an open question as to how discretization is related to the ill-posedness of the inverse ECG problem both qualitatively and quantitatively, and correspondingly, how one should adapt one's volumetric mesh to maximize the operator conditioning while at the same time minimizing error. Because the FEM is ubiquitously applied in other fields of engineering, such as in fluid and solid mechanics, especially for forward simulations, it is often assumed that strategies developed in these contexts will naturally carry over to both the solutions of the forward *and* inverse ECG problems. Although many of the refinement strategies developed for mechanics simulations are effective when solving the ECG forward problems, the ill-posedness of the ECG inverse problem undermines the transfer and extension of such refinement strategies toward the inverse solution. To emphasize this point, in the next section, we present a vivid example that illustrates how employing refinement strategies appropriate for the forward ECG problem can undermine the solution accuracy when solving the inverse ECG problem.

B. Motivation

To motivate this study, we consider an example of a finite-element discretization of the forward and inverse ECG problems. The details will be discussed in Sections II and III, respectively; here, we distill for presentation only the salient features that help motivate our study.

The mathematical model used for our forward ECG problem consists of a well-posed elliptic boundary value problem in which one seeks the electric potential on the surface of the torso given as input the potential on the surface of the heart. The discretization by finite elements of boundary value problems of this type is well studied, with very clear theoretical and empirical guidelines as to how and where to place resolution in the form of adding additional, smaller elements or using higher order basis functions to decrease the approximation error. For the ECG forward problem, a noticeable decrease in the error can be obtained by increasing the number of elements at and around the heart surface. The physical rationale for this strategy is that the accuracy of the forward problem approximation is jointly




Fig. 1. Effects of uniform refinement on the forward solution error and singular values of the transfer matrix. (Left) Increasing the resolution both on and normal to the surface of the heart consistently reduces the error in the forward ECG simulation. $|u_H|$ is the resolution on the heart surface and $|\bar{u}_T - u_T|$ is the forward solution error on the torso surface. Four meshes are labeled as A–D in the ascending order of their resolutions. For simplicity, only mesh B and C are displayed. (Right) Increase in resolution worsens the conditioning of the transfer matrix to be inverted in the inverse problem. Curves A–D are singular values in their original length; these singular values are normalized to the length of A, as shown by curves B'–D'.

determined by the discretization's ability to capture the electric potential on the surface of the heart, as well as the ability to capture the strong gradients of the potential moving away from the heart into the torso volume.

Fig. 1 (left) presents a convergence plot showing the decrease in the error between the true and approximate solutions of a forward ECG problem where the resolution is being added at and around the heart surface—precisely what traditional FEM theory would dictate.

In Fig. 1 (right), we present the singular values (a measure of numerical conditioning) of the transform matrix—the discretized operator that “transforms” potentials on the heart surface to potentials on the torso surface. The inverse problem consists of “inverting” the transform matrix so that, given potentials on the torso surface, one can obtain potentials on the heart surface. The magnitudes of the singular values provide a measure of the invertability of the system. As can be seen in this figure, the increase in resolution on and around the heart surface actually increases the ill-conditioned nature of the inverse problem. A refinement strategy developed solely based on considerations within the forward problem leads to an inappropriate discretization for the inverse problem. Thus, in this paper, we seek to develop discretization and refinement strategies that specifically consider the mathematical and numerical constraints imposed by the inverse ECG problem.

C. Outline

The paper is organized as follows. Section II presents the mathematical model for the forward ECG problem and describe how one discretizes the model with the FEM. In Section III, we present the mathematical model for the ECG inverse problem and describe how one modifies the forward-problem finite-element discretization methodology to solve the

inverse problem. We then present a mathematical discussion of the ill-posedness of the inverse problem and the corresponding ill-conditioning of the numerical system. Based upon these discussions, we conjecture finite-element discretization strategies appropriate for the ECG inverse problem. In Section IV, we test our conjectures in two example problems: a simplified 2-D offset annulus problem and a 2-D realistic torso geometry problem. Strategies arising as a consequence of the data presented are summarized in Section V. Section VI includes the conclusion and discussion of future work.

II. DISCRETIZATION OF FORWARD ECG PROBLEM

In this section, we present the mathematical model used to articulate the ECG forward problem, and present the details of how one generates a finite-element approximation of the problem.

A. ECG Forward Problem

The ECG forward problem seeks the potential field on the torso surface induced by the heart's electrical activity either from current sources within the heart or from a potential field on the heart surface. For this study, we use the formulation in terms of epicardial potentials. Given the epicardial potential as a (input) boundary condition, the potential within the torso volume is dictated by a quasi-static approximation of Maxwell's equations expressed as follows:

$$\nabla \cdot (\sigma(\mathbf{x}) \nabla u(\mathbf{x})) = 0, \quad \mathbf{x} \in \Omega \quad (1)$$

$$u(\mathbf{x}) = u_0(\mathbf{x}), \quad \mathbf{x} \in \Gamma_D \quad (2)$$

$$\vec{n} \cdot \sigma(\mathbf{x}) \nabla u(\mathbf{x}) = 0, \quad \mathbf{x} \in \Gamma_N \quad (3)$$

where Ω denotes the torso domain, Γ_D and Γ_N denote the epicardial (Dirichlet) and torso (Neumann) boundaries, respectively,

$u(\mathbf{x})$ is the potential field on the domain Ω , $u_0(\mathbf{x})$ is the known epicardial potential boundary function, $\sigma(\mathbf{x})$ is the symmetric positive-definite conductivity tensor, and \vec{n} denotes the outward facing vector normal with respect to the torso surface.

B. Finite-Element Discretization of Forward Problem

Solving (1)–(3) for any but the simplest geometries requires the use of numerical approximations such as boundary elements or finite elements [26]. We now review how one discretizes the ECG forward problem using the FEM [27].

From the mathematical perspective, solving (1)–(3) is accomplished by assuming that the solution $u(\mathbf{x}) = v(\mathbf{x}) + w(\mathbf{x})$ can be decomposed into homogeneous and heterogeneous parts ($v(\mathbf{x})$ and $w(\mathbf{x})$, respectively). The function $w(\mathbf{x})$ is chosen to satisfy both the Dirichlet boundary and Neumann boundary conditions, and the function $v(\mathbf{x})$ is chosen to satisfy the following system:

$$\nabla \cdot (\sigma \nabla v(\mathbf{x})) = -\nabla \cdot (\sigma \nabla w(\mathbf{x})), \quad \mathbf{x} \in \Omega \quad (4)$$

$$v(\mathbf{x}) = 0, \quad \mathbf{x} \in \Gamma_D \quad (5)$$

$$\vec{n} \cdot \sigma \nabla v(\mathbf{x}) = 0, \quad \mathbf{x} \in \Gamma_N. \quad (6)$$

A mathematical interpretation of this procedure is that one first finds a “lifting” of the boundary conditions onto the space of functions living over the entire domain, and then, solves a homogeneous problem whose forcing function involves the heterogeneous term. By such interpretation, one can immediately see three approximation issues to be encountered when solving the ECG forward problem: 1) how accurately one represents the boundary conditions (expressed in how well $w(\mathbf{x})$ captures $u_0(\mathbf{x})$, $\mathbf{x} \in \Gamma_D$); 2) how accurately one computes the right-hand side forcing term involving $w(\mathbf{x})$ and its interaction with the discretization of the function v over the volume (i.e., how to choose an optimal “lifting” operator); and 3) how accurately one computes the solution of the homogeneous problem over the volume.

Given a domain Ω and a partial differential equation (PDE) operating on a solution u that lives over Ω , the standard FEM attempts to construct a geometric approximation $\tilde{\Omega} = \mathcal{T}(\Omega)$ consisting of a tessellation of polygonal shapes (e.g., triangles and quadrilaterals for 2-D surfaces) of the domain Ω , and to build an approximating function space $\tilde{\mathcal{V}}$ consisting of piecewise linear functions based upon the tessellation [27]. Building on these two things, the goal of a finite-element analysis is to find an approximation $\tilde{u} \in \tilde{\mathcal{V}}$ that satisfies the PDE operator in the Galerkin sense. We utilized a hybrid triangular and quadrilateral tessellation $\mathcal{T}(\Omega)$ of the domain with the set \mathcal{N} denoting indexes of the mesh nodes as the geometric basis for the finite-element computations. For the case of linear/bilinear finite elements, this set consisted of the indexes for the triangle and quadrilateral vertices. We then decomposed the set \mathcal{N} into three nonintersecting sets \mathcal{H} , \mathcal{I} , and \mathcal{T} , which represent nodal indexes that lie on the heart surface denoted with H (at which the heart potentials are known, and hence the Dirichlet boundary), nodal indexes within the interior of the domain for which the solution is sought (denoted with I), and nodal indexes on the torso surface for which

the solution is sought (i.e., the Neumann boundary, denoted with T), respectively.

Let $\phi_i(\mathbf{x})$ denote the global finite-element interpolating basis functions, which have the property that $\phi_i(\mathbf{x}_j) = \delta_{ij}$, where \mathbf{x}_j denotes a node of the mesh for $j \in \mathcal{N}$. Solutions are then of the form

$$u(\mathbf{x}) = \sum_{k \in \mathcal{N}} \hat{u}_k \phi_k(\mathbf{x}) \quad (7)$$

$$= \sum_{k \in \mathcal{I}} \hat{u}_k \phi_k(\mathbf{x}) + \sum_{k \in \mathcal{T}} \hat{u}_k \phi_k(\mathbf{x}) + \sum_{k \in \mathcal{H}} \hat{u}_k \phi_k(\mathbf{x}) \quad (8)$$

where the first two right-hand side terms of (8) denote the sum over the degrees of freedom of the problem consisting of the unknown potential values at the vertices weighted by the basis functions, and the third term denotes the same sum for the (known) Dirichlet boundary conditions of the solution.

Substituting the expansion (8) into the differential equation (1), multiplying by a function from the test space $\{\phi_j(\mathbf{x}); j \in \mathcal{I} \cup \mathcal{T}\}$, taking inner products, and integrating by parts yield a linear system of the form

$$\begin{pmatrix} \mathbf{A}_{II} & \mathbf{A}_{IT} \\ \mathbf{A}_{TI} & \mathbf{A}_{TT} \end{pmatrix} \begin{pmatrix} \mathbf{u}_I \\ \mathbf{u}_T \end{pmatrix} = \begin{pmatrix} -\mathbf{A}_{IH} \\ -\mathbf{A}_{TH} \end{pmatrix} \mathbf{u}_H \quad (9)$$

where $\mathbf{u}_I = (\hat{u}_k)^T, k \in \mathcal{I}$, and $\mathbf{u}_T = (\hat{u}_k)^T, k \in \mathcal{T}$, denote the vectors containing the solution of the system (i.e., potential values at the nodal positions in $\mathcal{I} \cup \mathcal{T}$), $\mathbf{u}_H = (\hat{u}_k)^T, k \in \mathcal{H}$, denotes the vector of known potentials on the surface of the heart, and the matrices are given by

$$\mathbf{A}_{II} = (\nabla \phi_j, \sigma \nabla \phi_k), \quad j \in \mathcal{I}, \quad k \in \mathcal{I} \quad (10)$$

$$\mathbf{A}_{IT} = (\nabla \phi_j, \sigma \nabla \phi_k), \quad j \in \mathcal{I}, \quad k \in \mathcal{T} \quad (11)$$

$$\mathbf{A}_{TI} = (\nabla \phi_j, \sigma \nabla \phi_k), \quad j \in \mathcal{T}, \quad k \in \mathcal{I} \quad (12)$$

$$\mathbf{A}_{TT} = (\nabla \phi_j, \sigma \nabla \phi_k), \quad j \in \mathcal{T}, \quad k \in \mathcal{T} \quad (13)$$

$$\mathbf{A}_{IH} = (\nabla \phi_j, \sigma \nabla \phi_k), \quad j \in \mathcal{I}, \quad k \in \mathcal{H} \quad (14)$$

$$\mathbf{A}_{TH} = (\nabla \phi_j, \sigma \nabla \phi_k), \quad j \in \mathcal{T}, \quad k \in \mathcal{H}. \quad (15)$$

In the previous expressions, (\cdot, \cdot) denotes the inner product taken over the entire spatial domain Ω . Assuming that all elements touching the heart surface do not also touch the torso surface, the aforementioned system can be written as

$$\mathbf{S} \begin{pmatrix} \mathbf{u}_I \\ \mathbf{u}_T \end{pmatrix} = \begin{pmatrix} -\mathbf{A}_{IH} \\ \mathbf{0} \end{pmatrix} \mathbf{u}_H \quad (16)$$

where \mathbf{S} is the so-called stiffness matrix and the right-hand side vector is the “forcing term” induced by the known Dirichlet boundary conditions. Because the stiffness matrix is positive-definite, the solution of the linear system is amenable to iterative methods such as the preconditioned conjugate gradient method [28].

Based upon (16), we recapitulate the observations we made concerning (4). Once an acceptable choice of resolution is made for capturing the boundary conditions via \mathbf{u}_H , spatial accuracy is dictated by how well one captures the heart surface/volume interaction (via \mathbf{A}_{IH}) and how accurate the volume conductor is (via \mathbf{S}).

The forward ECG problem seeks to obtain the potential on the torso surface. Utilizing the fact that \mathbf{A}_{II} is the stiffness matrix resulting from using FEM to solve a zero Dirichlet boundary condition Laplace problem (and hence it is invertible), we can rewrite the first line of (9) to obtain an expression for the potentials in the interior as a function of both the torso and heart potentials: $\mathbf{u}_I = -\mathbf{A}_{II}^{-1}(\mathbf{A}_{IH}\mathbf{u}_H + \mathbf{A}_{IT}\mathbf{u}_T)$. Applying the Schur complement procedure [29], one can then substitute this expression into the second row of (9) to obtain an expression for the torso potentials as a function of the heart potentials

$$\mathbf{M}\mathbf{u}_T = \mathbf{N}\mathbf{u}_H \quad (17)$$

$$\mathbf{M} = \mathbf{A}_{TT} - \mathbf{A}_{TI}\mathbf{A}_{II}^{-1}\mathbf{A}_{IT} \quad (18)$$

$$\mathbf{N} = \mathbf{A}_{TI}\mathbf{A}_{II}^{-1}\mathbf{A}_{IH}. \quad (19)$$

In addition, $\mathbf{K} = \mathbf{M}^{-1}\mathbf{N}$ is often referred to as the transfer matrix, since it “transfers” the potential information from the heart surface to the torso surface. \mathbf{M} is a well-conditioned and invertible matrix; however, \mathbf{N} is severely ill-conditioned. The ECG forward problem given by (1)–(3) is a well-posed problem such that when it is discretized as outlined before (under assumptions of good element quality), it produces a transfer matrix that as a forward linear operator is well-behaved.

III. INVERSE ECG PROBLEM AND DISCRETIZATION

In this section, we present the mathematical model used to articulate the ECG inverse problem and present the details of how one generates a finite-element approximation of the problem. We then discuss the ill-posedness and ill-conditioning issues one should consider when solving the FEM-based ECG inverse problem.

A. ECG Inverse Problem

In the ECG inverse problem, one seeks to describe the electrical activity of the heart either as a potential field on the heart surface or as current sources that induce the potential field measured on the torso surface. In our study, we use the formulation that seeks to find the potentials on the epicardial surface given the measured potentials on the torso surface.

Given the torso surface potential as a (input) boundary condition, the potential within the volume is dictated by a Laplace’s equation expressed as follows:

$$\nabla \cdot (\sigma(\mathbf{x})\nabla u(\mathbf{x})) = 0, \quad \mathbf{x} \in \Omega \quad (20)$$

$$u(\mathbf{x}) = g(\mathbf{x}), \quad \mathbf{x} \in \Gamma_N \quad (21)$$

$$\vec{n} \cdot \sigma(\mathbf{x})\nabla u(\mathbf{x}) = 0, \quad \mathbf{x} \in \Gamma_N \quad (22)$$

where Ω denotes a torso domain bounded by the torso boundary Γ_N and the heart boundary Γ_H , $u(\mathbf{x})$ is the potential field defined on Ω , $g(\mathbf{x})$ is the measured torso boundary potential function, $\sigma(\mathbf{x})$ is the symmetric positive-definite conductivity tensor, and \vec{n} denotes the outward facing vector normal to the torso surface. With both the Dirichlet and Neumann boundary conditions available on Γ_N , one seeks the potentials on Γ_H .

Solutions to (20)–(22) are unique [30], and the physical quantities involved are all measurable. However, these solutions share

a characteristic of all ECG inverse problems in that they are ill-posed in the Hadamard sense, i.e., because the solution does not depend continuously on the data, small errors in the measurement of the torso potentials or thorax geometry can yield unbounded errors in the solution. The origins of this ill-posedness are biophysical; it arises both from the attenuation of potentials as one moves away from the source and the fact that the potential at any point on the torso surface is a weighted superposition of all the individual sources within the heart. Hence, the ECG represents an integration of many sources, the influence of which decreases sharply with distance. To find the inverse solution, we must perform the complementary operations of amplification and differentiation not only on the ECG but also on the inevitable noise that accompanies it. The result is highly sensitive to any fluctuations in the input signals or geometric models.

B. Finite-Element Discretization of Inverse Problem

Instead of attempting to discretize (20)–(22) directly, most practitioners form the finite-element approximation of the inverse problem indirectly by first developing the approximate solution of the forward problem as given in (17). The numerical solution procedure for the inverse problem can then be stated as given \mathbf{u}_T , find \mathbf{u}_H so as to minimize the functional $\|\mathbf{M}\mathbf{u}_T - \mathbf{N}\mathbf{u}_H\|$ in some appropriate vector norm. Many practitioners go one step farther and minimize the functional $\mathbf{u}_T - (\mathbf{M}^{-1}\mathbf{N})\mathbf{u}_H$, as traditional minimization techniques from linear algebra can then be employed. Note that these forms are mathematically equivalent since \mathbf{M} is nonsingular. In either form, the ill-posedness of the continuum problem translates into an ill-conditioning of the discretized problem, thus requiring constraints (via regularization) to ensure a stable solution.

C. Ill-Posedness Considerations

Acting as a volume conductor, the human body is known to respond differently to different frequency components of the electrical source. As such, Fourier analysis, which allows a frequency decomposition of the solution space in infinite dimensions, is an effective analytic tool for understanding the effect of the ill-posedness in inverse problems [31]–[33], and thus, provides a guideline for discretization considerations. We now use the Fourier analysis to study the ill-posedness of the problem.

Assume that $\Omega \subset \mathbb{R}^2$, and without loss of generality, assume $\sigma = 1.0$ in (1). We consider the equation in polar coordinates with the origin point being set within the interior boundary Γ_I . The variable u_0 , the Dirichlet condition on the heart boundary, is a univariate function to the azimuthal variable θ and can be expanded into a Fourier series

$$u_0(\theta) = \frac{A_0}{2} + \sum_m^{\infty} (A_m \cos(m\theta) + B_m \sin(m\theta)) \quad (23)$$

where $\theta \in [-\pi, \pi)$, A_m and B_m are Fourier coefficients, and m is the spatial frequency.

The general solution of (1) can be regarded as a superposition of the solutions stemming from each frequency component of

the source u_0 in (23). By separating the variables, we can derive the general solution

$$u(r, \theta) = \sum_{m=1}^{\infty} \left[\left(c_m r^m + \frac{d_m}{r^m} \right) (a_m \cos(m\theta) + b_m \sin(m\theta)) \right] + a_0 \ln r + b_0 \quad (24)$$

where a_m, b_m, c_m , and d_m are coefficients determined by boundary conditions and domain geometries. To simplify our discussion without losing generality, we consider an annulus model with the outer radius being of unit length and the inner radius being r_0 , i.e., $r \in [r_0, 1]$, $0 < r_0 < 1$. The zero Neumann boundary condition requires $c_m \equiv 1$ and $d_m \equiv 1$; then, (24) is reduced to

$$u(r, \theta) = b_0 + \sum_{m=1}^{\infty} \left[\left(r^m + \frac{1}{r^m} \right) (a_m \cos(m\theta) + b_m \sin(m\theta)) \right]. \quad (25)$$

Combining (23) and (25) yields the algebraic relation between the Dirichlet condition u_0 and the solution u

$$u(r, \theta) = \sum_{m=1}^{\infty} \left(\frac{r^m + r^{-m}}{r_0^m + r_0^{-m}} \right) (A_m \cos(m\theta) + B_m \sin(m\theta)) + \frac{A_0}{2}, \quad \theta \in [-\pi, \pi), \quad r \in [r_0, 1]. \quad (26)$$

In particular, the forward operator \mathbf{K} from the source space (i.e., the heart surface) to the measurement space (i.e., the torso surface) is made concrete by setting $r = 1$

$$u_T = u(r = 1, \theta) = \frac{A_0}{2} + 2 \sum_{m=1}^{\infty} \frac{A_m \cos(m\theta) + B_m \sin(m\theta)}{r_0^m + r_0^{-m}}. \quad (27)$$

The term $2/(r_0^m + r_0^{-m})$ describes the attenuation properties of K , indicating that the magnitude of attenuation is an exponential function with respect to the spatial frequency m . The magnitude of attenuation also depends on r_0 , which characterizes the ratio of the interior/exterior radius of the domain. The ill-posedness of the system can be understood in terms of $(r_0^m + r_0^{-m})/2$ (the reciprocal term of the one mentioned before); it is due to the “physical nature” of the problem being considered, and therefore, cannot be alleviated by any discretizations.

Discretizing a function space is analogous to sampling a continuous signal, and the resolution is similar to the sampling rate. With respect to discretization, spatial frequencies of a continuous function can be approximated by the number of sign changes in the corresponding discrete vector (a measure of variation). According to the sampling theorem, discretization resolution is proportional to the band limit in terms of the spatial frequency m . In this sense, given a discretization of the domain, the heart boundary resolution determines the spatial frequency band limit of the epicardial potentials to be recovered, and hence, provides a sense of the intrinsic ill-conditioning of the discrete inverse problem. This ill-conditioning increases approximately as an exponential function with respect to the heart boundary resolution.

In the forward ECG problem, the heart potential information propagates through the volume conductor to reach the body

surface where the information is recorded. The recorded information then forms the basis for solving the inverse problem. The amount of recorded information not only depends on the source but also on how much information the volume conductor allows to pass. If a frequency component of the discretized heart potentials cannot pass through the volume conductor, this component can never be recovered in the inverse problem. However, since the heart boundary discretization already assumes this frequency component, the resulting numerical system is still required to attempt to resolve this unrecoverable component, leading to what can be considered as “extra” (or supplementary) ill-conditioning. This ill-conditioning is due to discrepancies (or mismatches) in discretization rather than the physical nature of the inverse ECG problem. In conclusion, the recoverable band limit of the heart potential Φ_I is bounded by the minimum between the band limit implied by the heart boundary resolution and the band limit specified by volume discretization.

Key observation: Increasing the desired resolution on the heart surface *increases* the ill-posedness of the continuous system. Requiring more fidelity in the heart-potential reconstruction increases the sensitivity of the system in the Hadamard sense. We would consequently expect that arbitrarily increasing the resolution on the heart surface will increase the ill-conditioning of the numerical inverse problem. The resolution in the volume mesh should be maintained to be not less than the required fidelity of the heart-potential reconstruction.

D. Ill-Conditioning of Numerical System

In this section, we present one traditional means of evaluating the conditioning of the discretized problem—examining the singular value spectrum of the forward operator. We will use this method to demonstrate how the conditioning of the discretized system affects the accuracy of the inverse solution.

After discretization via finite elements, the approximated solutions on the heart and torso boundaries live in finite-dimensional vector spaces, and the relationship connecting them is given by

$$\mathbf{u}_T = \mathbf{K} \mathbf{u}_H \quad (28)$$

where $\mathbf{K} = \mathbf{M}^{-1} \mathbf{N}$ as given in (17). Recall that the number of degrees of freedom on the heart surface and torso surface need not be the same; in general, \mathbf{K} is an $m \times n$ matrix where $m > n$. Here, m and n denote the dimension of the heart potential vector and torso potential vector, respectively. Because of the ill-posed nature of the problem, the corresponding discretization embodied in the transfer matrix \mathbf{K} admits a large condition number. To assess the ill-conditioning of \mathbf{K} , we utilize the concept of valid and null spaces of \mathbf{K} based upon its singular value spectrum, which was introduced into the inverse ECG problem in [34].

To explore the spectrum of \mathbf{K} , we first perform a singular value decomposition (SVD) of the transfer matrix \mathbf{K}

$$\mathbf{K} = \mathbf{U} \cdot \Sigma \cdot \mathbf{V}^T \quad (29)$$

where $\mathbf{U} = (u_1, \dots, u_m) \in \mathbb{R}^{m \times n}$ and $\mathbf{V} = (v_1, \dots, v_n) \in \mathbb{R}^{n \times n}$ are matrices consisting of the left and right singular vectors, respectively, and $\Sigma \in \mathbb{R}^{n \times n}$ is a diagonal matrix with

positive singular values $\sigma_i, i = 1, 2, \dots, n$. We obtain

$$\mathbf{u}_T = \mathbf{U}\Sigma\mathbf{V}^T\mathbf{u}_H = \sum_{i=1}^n u_i \sigma_i (v_i^T \mathbf{u}_H) = \sum_{i=1}^n \alpha_i u_i$$

where $\alpha_i = \sigma_i \cdot (v_i^T \mathbf{u}_H)$ is a scalar value. The variable \mathbf{u}_T , the vector of potentials measured on the torso surface, is a linear combination of u_i with coefficients α_i , which is derived as the product of the singular value σ_i and the projection of \mathbf{u}_H onto the i th right eigenvector v_i . As the singular value σ_i decreases rapidly to zero (or, in practical computations, single-precision or double-precision machine zero), the right eigenspace of \mathbf{K} can be decomposed into a valid subspace spanned by low-indexed eigenvectors and a null subspace spanned by high-indexed eigenvectors. The fraction of the source \mathbf{u}_H that falls in the null space is smoothed out by the zero-valued σ_i . Only the fraction of \mathbf{u}_H in the valid space has a nontrivial contribution to the observable \mathbf{u}_T , and thus, can be recovered.

Accordingly, a slowly descending singular value spectrum with a broader range of nonzero singular values indicates a better conditioning of the discretized inverse problem. The fraction of \mathbf{u}_H in the valid subspace estimates the best solution that is a recoverable problem, regardless of regularization methods, regularization parameters, error measurements, input noise, or other factors that depend on algorithms or numerical solvers.

Key observation: Examination of the singular values of the forward operator is a valuable means of determining the ill-conditioning of the discretized system. Different discretization choices, in terms of number of elements and placement of elements, will impact the formulation of the transfer matrix. This will, in turn, impact its singular value spectrum. In order to understand the impact of discretization choices on the numerical inverse problem, we can investigate the conditioning of the resulting systems by assessing their singular values.

E. Regularization Methods and Parameter Selection

In general, it is extremely difficult to directly solve the inverse problem in (28), especially when the measured torso potentials are contaminated with noise, because the matrix \mathbf{K} is rank-deficient characterized by a cluster of close-to-zero singular values. Regularization introduces extra constraints on the desired solution, and then, seeks a solution that provides reasonable balance between minimizing the residual error and satisfying the constraints. The regularized problem is well posed and provides a stable solution, which should be not too far from the desired solution of the original ill-posed problem.

Perhaps, the best-known regularization technique is the Tikhonov regularization, which is expressed as

$$\mathbf{u}_H(\lambda) = \operatorname{argmin}\{\|\mathbf{K}\mathbf{u}_H - \mathbf{u}_T\|_2^2 + \lambda^2 \|\mathbf{L}\mathbf{u}_H\|_2^2\} \quad (30)$$

where $\mathbf{L}\mathbf{u}_H$ is the regularization term. The matrix \mathbf{L} is typically either the identity matrix or the discrete approximation of the first or second derivative operator, indicating that the regularization term constrains the seminorm of the solution itself or of its derivatives. λ is a regularization parameter that controls the weight placed on the regularization relative to that placed

on solving the original problem. Obviously, a large λ implies a large amount of regularization, placing strict constraints on the solution seminorm without much consideration for the residual error, while a small λ implies a small amount of regularization, focusing on minimizing the residual error without much consideration for the solution constraints. Therefore, λ is an important quantity and should be chosen with care. Many studies have been devoted to finding the optimal parameter [35].

On the other hand, the “optimal” value of λ (however it is selected) somewhat reflects the ill-conditioned severity of the inverse problem being solved. A small λ means that the problem is not severely ill-conditioned, and therefore, does not require much regularization, whereas a large λ offers the opposite implication. In the extreme case of solving a well-conditioned problem, the solution is obtained by minimizing the residual error (a standard least-square problem). No regularization is required at all, and $\lambda = 0$. Therefore, inspecting the value of λ in the regularization process provides a means to assess the ill-conditioning of the problem in addition to that of evaluating singular values as we proposed earlier.

In this study, we obtained the inverse solution by adopting the popular Tikhonov regularization and choosing \mathbf{L} as the discrete operator of the second-order derivative. The value of λ was determined through an exhaustive search. Since the goals of our study were to investigate optimal discretization techniques, we did not intend to propose new, nonrefinement-based regularization schemes or new methods to locate the regularization parameter. Instead, we focused on the influence of various discretization strategies in the regularization process.

IV. RESULTS

A. Simulation Setup

Herein, we present a study of the approximation and conditioning properties of the ECG inverse problem solved by the FEM. We will present three main categories of results: 1) examples on a 2-D simplified geometry with homogeneous conductivities consisting of a circle within a circle, the so-called “offset-annulus” problem; 2) examples on a 2-D realistic torso geometry with homogeneous conductivities; and 3) a 2-D realistic geometry with heterogeneous conductivities. As our goal is to generate guidelines for placing resolution (in the form of the number of elements and the type of elements) in the discretized ECG inverse problem, we examine the following scenarios:

- 1) increasing the resolution of the interior (heart) boundary (tangential direction resolution);
- 2) increasing the resolution of the interior (heart) boundary (normal direction resolution);
- 3) increasing the resolution of the volume conductor;
- 4) increasing the resolution of the exterior (no-flux) torso boundary.

The objective of our experiments is to decipher which of these scenarios leads to better-conditioned (or worse-conditioned) systems. Throughout this study, a first-order finite-element approximation is applied, and h-refinement is employed to adjust the resolution. To facilitate comparison and generalization, we ensure that the annulus and the realistic geometry have the same




Fig. 2. Effects of uniform h-refinement on the transfer matrix K and its components. (A) Mesh with 60 nodes on heart. (B) Mesh with 80 nodes on heart. (C) Mesh with 100 nodes on heart. (D) Singular values of N and A_{IH} plotted against the original index (A, B, C) and the normalized index (A, B', C'). (E) Singular values of K plotted against the original index (A, B, C) and the normalized index (A, B', C').

resolution both on their exterior boundaries and on their interior boundaries. They also have approximately the same number of elements and nodes.

After the heart-to-torso transfer matrix K , given by (17)–(19), is formulated by the FEM, the inverse problem is solved by the standard Tikhonov regularization, in which the optimal regularization parameter is determined via an exhaustive search. Since the choice of regularization methods and parameter values can significantly influence the inverse solution accuracy, we apply the classic Tikhonov regularization to provide consistency and minimize the impact from different regularization techniques. This choice allows us to isolate the influence of discretization on the inverse solution.

The inverse problem is solved in the presence of various levels of noise on the input—the measured potentials on the torso surface. To take into account the randomness of input noise, each experiment is repeated 50 times and the arithmetic average of the results is analyzed.

In each study mentioned later, we present the resulting numerical systems A_{IH} , N , and K in the form of their singular value spectra. We also show how the optimal regularization parameter reflects the numerical conditioning of the system solved in the inverse process. Finally, we evaluate the inverse solution

using two metrics, the relative error (RE) and the so-called max-gradient location error.

The RE between the calculated solution \hat{u}_H and the exact solution u_H is defined as follows:

$$RE = \frac{\|\hat{u}_H - u_H\|_2}{\|u_H\|_2}. \quad (31)$$

The max-gradient location error measures the error in locating the maximum gradient position in the inverse solution. The rationale for this metric is that it normally captures the activation wavefront. The activation time (the time of depolarization) in the ECG is usually estimated by the most negative time derivative, maximum spatial gradient, and zero surface Laplacian [36]. At one time instant, the field between the activated epicardial region and the nonactivated region has the largest potential gradient. This high spatial gradient is also an expression of the high extracellular potential gradient associated with the depolarization phase. ECG researchers use this information to estimate the propagation of activation potential wavefronts, which, combined with some wavefront potential models, provides physiologically based *a priori* estimation of epicardial potentials. The estimation can be directly incorporated into the regularization to improve the inverse solution [37].

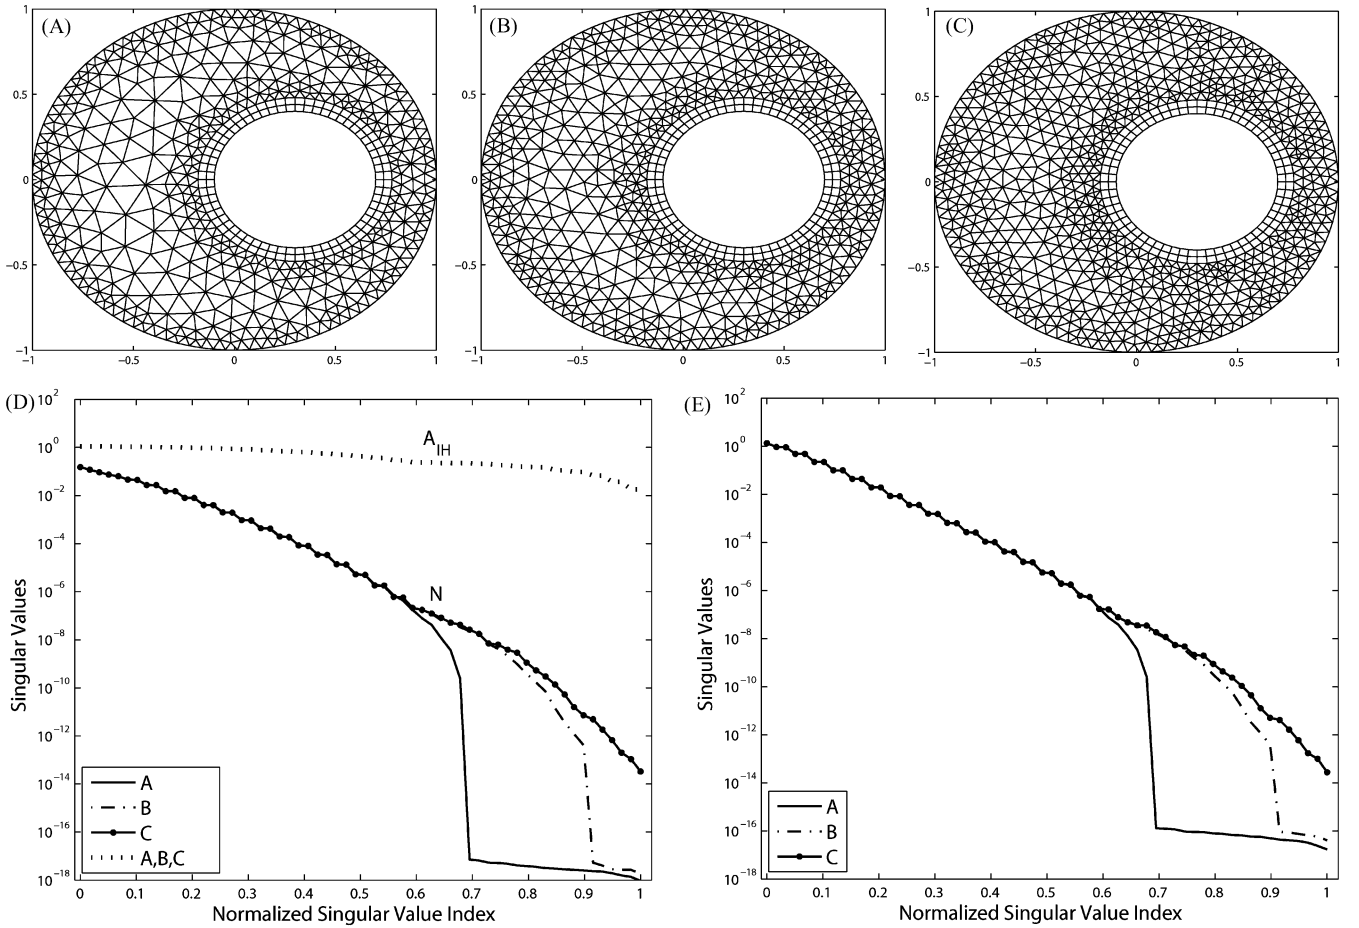


Fig. 3. Refining the volume conductor eliminates the extra ill-conditioning induced by the discretization. The three meshes have the same boundary resolutions: 105 nodes on the torso and 60 nodes on the heart. (A) Mesh with 1045 elements, 665 nodes. (B) 1325 elements, 805 nodes. (C) 1629 elements, 957 nodes. (D) Singular values of N and A_{IH} . (E) Singular values of K .

TABLE I
EVALUATION OF INVERSE SOLUTIONS OF ANNULUS
SIMULATION SHOWN IN FIG. 3

	Noise Level	Mesh A	Mesh B	Mesh C
Mean RE	Noise Free	14.80%	10.99%	9.65%
	30dB	15.10%	12.38%	12.20%
	20dB	20.29%	19.75%	19.09%
λ	Noise Free	0.0019	0.0007	0.0003
	30dB	0.0027	0.0015	0.0014
	20dB	0.0144	0.0128	0.0140
MGL Error	Noise Free	1	1	1
	30dB	0.9	0.54	0.42
	20dB	0.70	0.66	0.56

The data represent the arithmetic average of 50 repeated simulations. White noise of 30 and 20 dB SNR is added to the torso-surface measurements before solving the inverse problem. MGL: maximum gradient location in the vector of the inverse solution.

B. 2-D Simplified Geometry

1) *Uniform Resolution Refinement*: We present in Fig. 2 a uniform refinement typically taken in forward problems. Panels (A)–(C) show three discretization levels of an annulus. The ratio of the outer radius to the inner radius of the annulus is set to be 1:0.4, approximating the ratio in the real human torso. Panel (D) shows the normalized singular value spectra of the resulting matrix A_{IH} and N , while panel (E) illustrates the singular value spectrum of the transfer matrix K . In order to maintain good aspect ratios of the triangle elements, the signif-




Fig. 4. Plots of the inverse solutions from the volume-refinement simulation shown in Fig 3. Evaluations of the solutions are presented in Table I. Mesh C results in a better solution than does the mesh A.

icance of which in the approximation accuracy has been discussed in [1] and [38], the resolution on the interior (heart) boundary is inevitably increased from 60 nodes in mesh A to

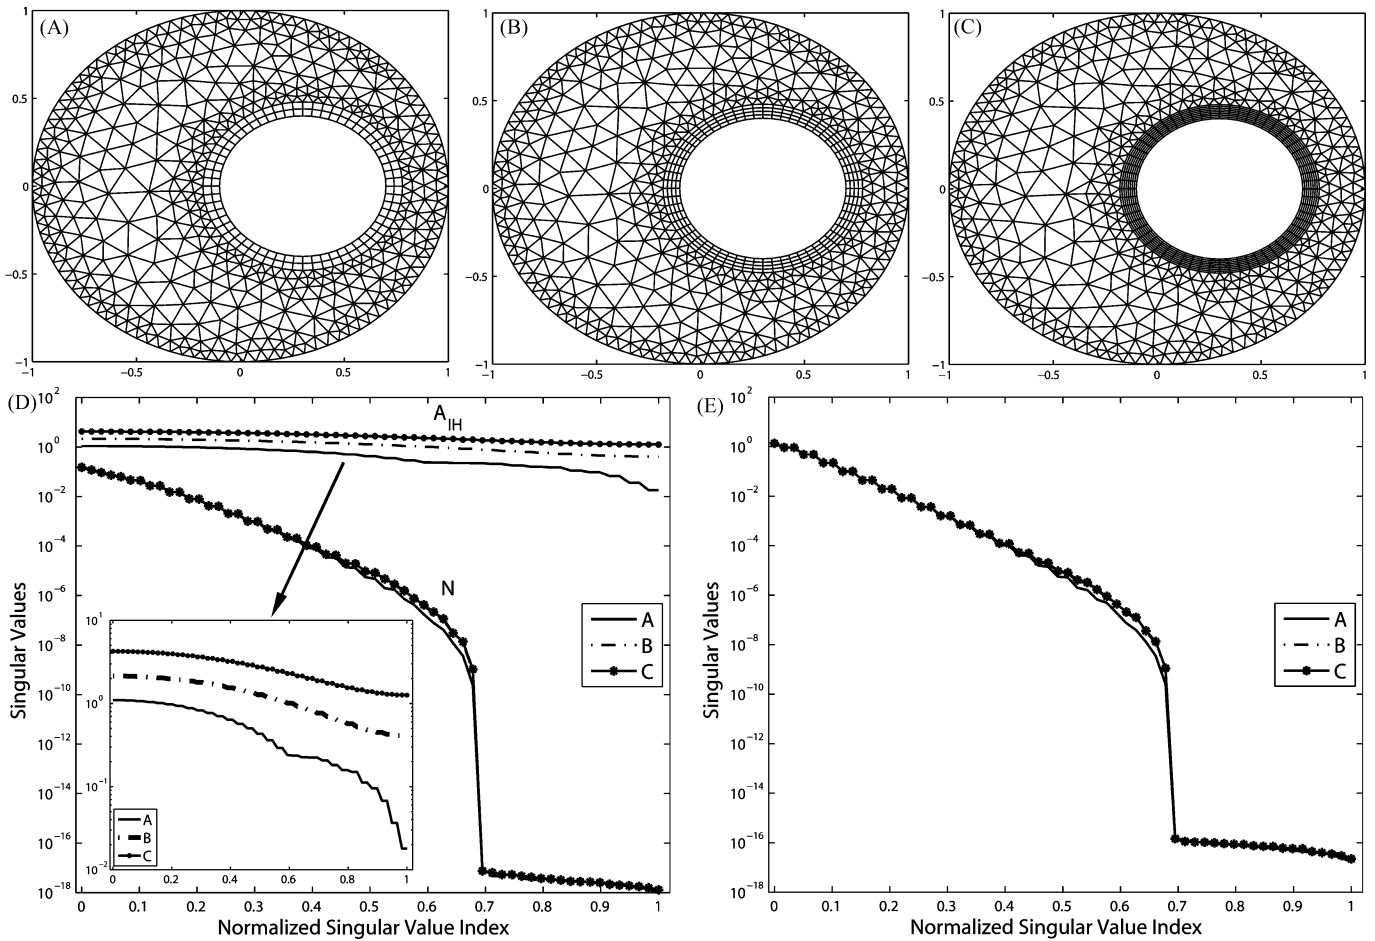


Fig. 5. Refining the resolution normal to the heart under a coarse volume discretization. (A) Two layers of quadrilaterals around the heart. (B) Four layers of quadrilaterals. (C) Eight layers of quadrilaterals. (D) Singular values of N and A_{IH} . (E) Singular values of K .

80 nodes in mesh B and then to 100 nodes in mesh C. Our results show that such refinement worsens the ill-conditioning of N and K .

2) *Volume Conductor Resolution:* In this test, we explore the impact of the azimuthal resolution in the volume conductor. Fig. 3 shows annulus at three refinement levels [panels (A)–(C)], the normalized singular value spectra of A_{IH} and N [panel (D)], and the normalized singular value spectrum of K [panel (E)]. To avoid increasing the resolution on the heart boundary, quadrilateral elements are placed around the heart to decouple the tangential resolution and normal direction resolution. Quadrilateral elements also ensure the same discretization of the high-gradient field near the heart.

Fig. 3(D) shows that such volume refinement improves the singular value spectrum of N , whereas Fig. 3(E) shows that the singular value spectrum of K is also improved. This is consistent because $K = M^{-1}N$ where M is well conditioned. Note that the matrix A_{IH} remains constant in this test because neither the mesh interface between the heart boundary and the volume nor the basis functions spanning the interface changes. In addition, A_{IH} is much better conditioned than either N or K .

An important observation is that the proportion of nontrivial singular values in the entire eigenspace of both N and K is determined by the resolution in the volume of the polygon that

encloses the interior boundary, and that polygon has the least number of nodes. This is clearly manifested in Fig. 3(A): The coarsest polygon in its volume has 42 nodes compared to 60 nodes on the interior boundary. Consequently, singular values of N and K of mesh A suddenly drop to 10^{-16} at the position $42/60 = 0.7$ in the normalized SVD scale. As the volume is refined in meshes B and C, singular value spectra are smoothed and the proportion of trivial singular values (indicating the null space) diminishes. The gap between the singular value spectrum of A and that of C is the additional (supplementary) ill-conditioning caused by the discretization but not associated with the ill-posed nature of the continuum problem. This fact can be inferred from the Fourier analysis in Section III-C, where (23) sets the frequency band limit of the epicardial potential one seeks to recover, and where (27) describes the band limit of the solvable potential field allowed by K . When the former exceeds the latter, one could consider the discretization to be “insufficient,” i.e., not sufficient to capture the aforementioned resolution relationship. Note that this observation is also manifested by our simulation on the realistic torso model, as will be presented in the later section. We will further discuss this issue in detail in Section V. Table I summarizes the evaluation of inverse solutions and the regularization parameter λ in the simulations shown in Fig. 3. Volume refinements consistently

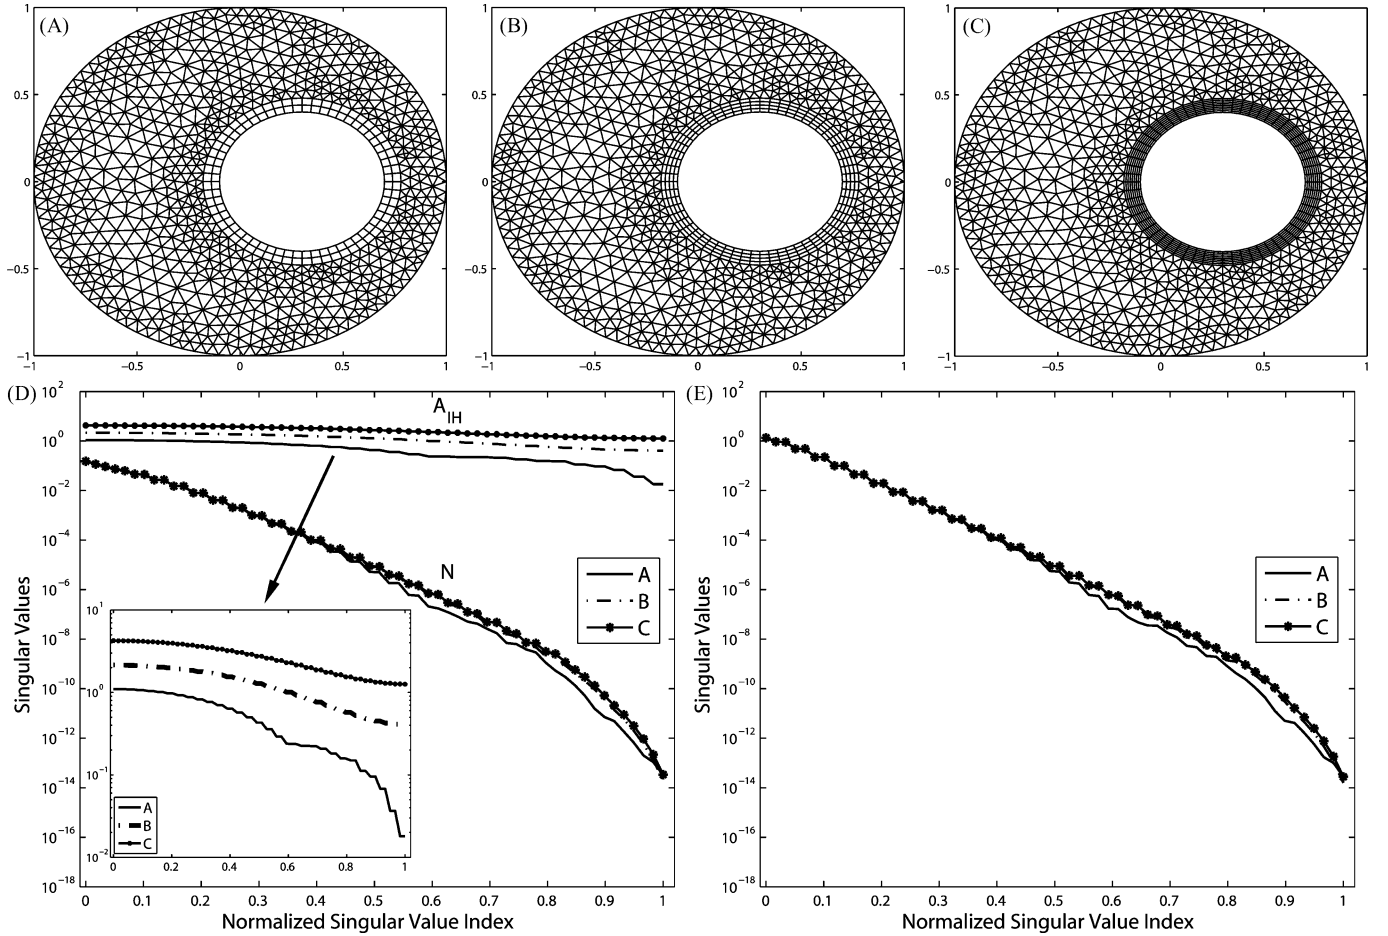


Fig. 6. Refining the resolution normal to the heart under a refined volume discretization. (A) Two layers of quadrilaterals around the heart. (B) Four layers of quadrilaterals. (C) Eight layers of quadrilaterals. (D) Singular values of N and A_{IH} . (E) Singular values of K .

reduce the maximum gradient location error regardless of the presence of input noise. Without input noise, mesh refinements reduce both the optimal value of λ from 0.0019 to 0.0003 and the inverse solution error from 14.80% to 9.65%, indicating that the numerical conditioning of the transfer matrix K is improved from the regularization viewpoint. This is consistent with the improvement of the singular value spectrum of A_{IH} and K , as shown in Fig. 3. While such improvement can still be observed in the case of 30 dB input noise, it is not evident in the case of 20 dB noise. We conjecture that when the input noise goes beyond a certain level, its amplification effect will overwhelm the improvement brought by discretization refinement. In this case, this noise threshold is located between 30 and 20 dB. However, although the quantitative metrics do not indicate apparent improvement in the case of 20 dB input noise, a visual inspection of the inverse solution shows that the solution obtained from mesh C indeed better captures the features of the exact solution than did the solution from mesh A, as shown in Fig. 4.

3) *Normal Direction Resolution*: In this experiment, we study the impact of resolution in the direction normal to the heart boundary. This resolution captures high gradients in the potential field. Its refinement is achieved by increasing the number of quad layers while still keeping the resolution in the tangential direction. Meanwhile, the discretization of the rest volume is

kept constant. The test is performed in two situations: with a coarsely refined volume conductor (shown in Fig. 5) and with a well-refined volume conductor (shown in Fig. 6). Results from both tests consistently indicate that increasing the resolution in the normal direction improves the boundary-to-interior “lifting” matrix A_{IH} . Given that N and K are obtained by multiplying several matrices onto A_{IH} , their singular value spectra are also “lifted” slightly. The basic quality of the singular value spectrum is still dominated by the tangential resolution in the volume: note the abrupt drop of the singular value spectrum in the coarse-volume case in contrast with their gradual decline in the refined-volume case. This implies that only when the azimuthal resolution is reasonably refined does the normal direction resolution matter to the condition of N and K .

4) *Resolution on Torso Boundary*: Measurements of the potential field on the torso boundary constitute the input of the inverse ECG problem. While theoretically one can keep the same resolution on the torso boundary as on the heart boundary so as to derive a square transfer matrix, researchers in practice usually take more measurements (by placing more detecting electrodes) in the belief that improving the numerical approximation of the Cauchy condition will improve the inverse solution [19]. The intuition here is that an overdetermined (though still rank-deficient) system may provide some sort of regularization.




Fig. 7. Refining the resolution on the torso boundary. The heart boundary is fixed to be 60 nodes. (A) 60 nodes on the torso. (B) 100 nodes on the torso. (C) 120 nodes on the torso. (D) Singular values of \mathbf{N} and \mathbf{A}_{IH} . (E) Singular values of \mathbf{K} .

Fig. 7 illustrates our experiment in which the resolution on the torso boundary is adjusted while the heart boundary is fixed. To minimize perturbing effects induced by insufficient volume discretization, as observed in our previous discussion, we refine the volume properly in each case. However, the resolutions on both boundaries are still kept unchanged during the refinement. Quadrilateral elements are again employed to ensure the same discretization of the high-gradient field around the heart. As extra torso nodes are added in the transition from Fig. 7(A)–(C), it is assumed that the real measured data were available on these additional nodes, rather than simple interpolation from the data on the existing nodes. Our results indicate that both \mathbf{N} and \mathbf{K} are improved by the refinement on the torso.

C. 2-D Homogeneous Realistic Geometry

1) *Volume Conductor Resolution*: The geometric model for this study consists of a single 2-D slice of the Utah torso model [39], [40]. We assume that the volume conductor is homogeneous with isotropic conductivities, ignoring the lung, muscle, and other tissue conductivities. Fig. 8 shows our test as well as properties of the resulting matrices. Analogous to the procedure taken in the volume test of the annulus model, we employ the same type of quadrilateral elements around the

heart in order to decouple the tangential resolution and normal direction resolution, as well as to ensure the same discretization of the high-gradient field near the heart. The mesh refinement is hierarchical: mesh B in Fig. 8 includes all nodes of mesh A, while mesh C includes all nodes of mesh B. This hierarchy has clear implications in the first-order FEM: The stiffness matrix \mathbf{A}_{II} is also hierarchical, because all of its degrees of freedom are located on nodes. Given that inverting \mathbf{A}_{II} is the most time-consuming step in the inverse problem computation, it is of practical interest to weigh the marginal gain in approximation accuracy against the extra computational cost. Fig. 8 shows that such volume refinement improves \mathbf{A}_{IH} , \mathbf{N} , and \mathbf{K} , judging by their singular value spectra. Also note that the proportion of nontrivial singular values in the eigenspace of either \mathbf{N} or \mathbf{K} is determined by the resolution of the polygon in the volume that encloses the interior boundary and that has the least number of nodes, an observation consistent with the discovery in the annulus model. Table II presents the regularization parameter λ and the evaluation of inverse solutions, which further confirms that the improvement of singular value spectra leads to better inverse solution. Volume refinement consistently reduces the maximum gradient location error regardless of the presence of input noise. In the noise-free case, mesh refinement reduces λ from 0.0077

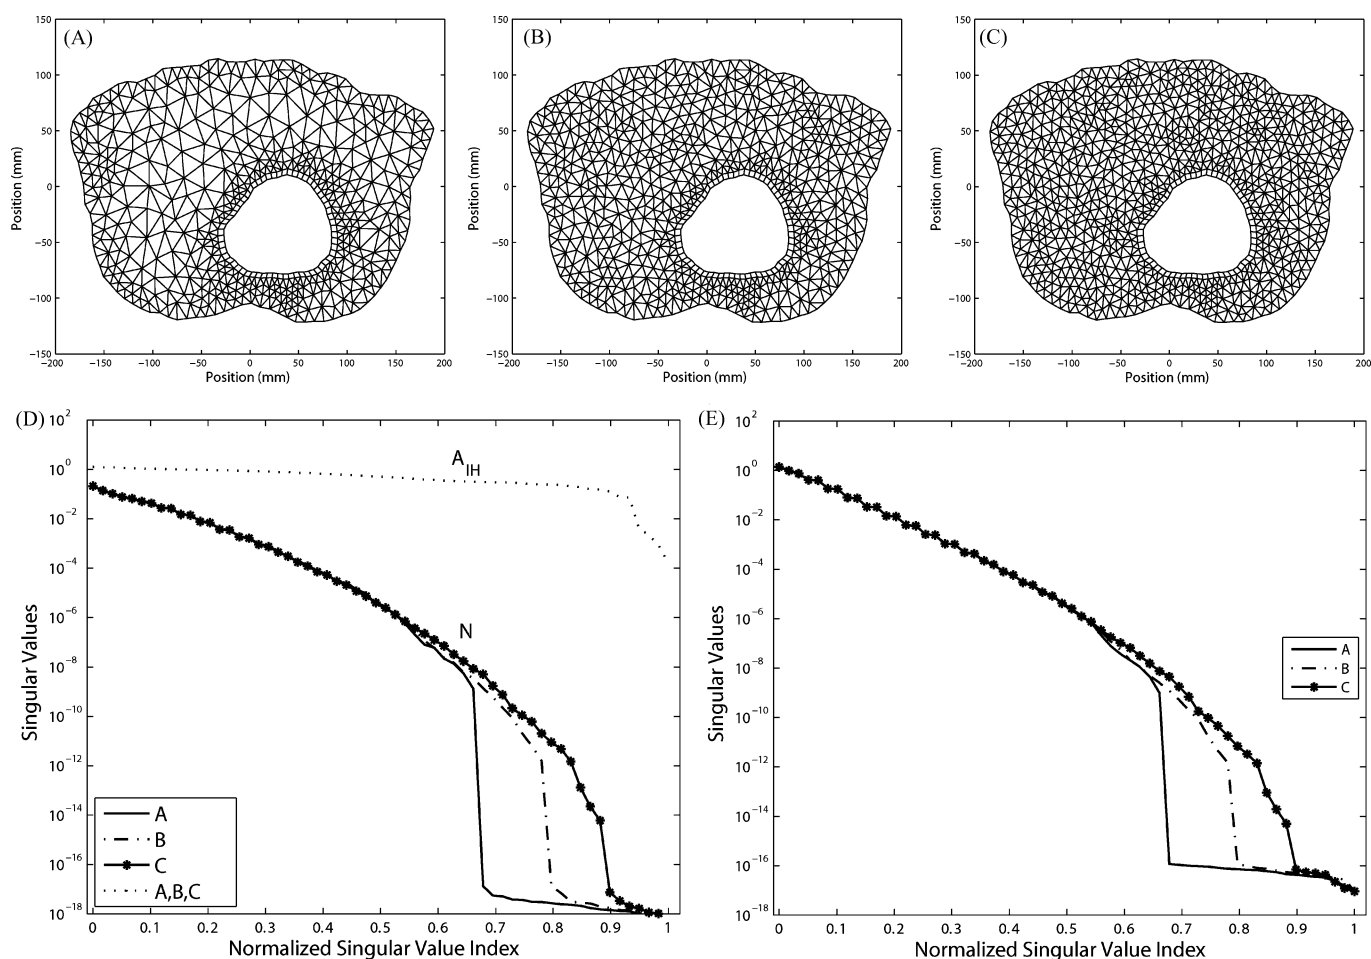


Fig. 8. Refining the volume conductor of a 2-D homogeneous torso slice. Three meshes share the same boundary resolution: 105 nodes on the torso and 60 nodes on the heart. (A) Mesh containing 943 elements, 584 nodes. (B) Mesh containing 1297 elements, 761 nodes. (C) Mesh containing 1899 elements, 1062 nodes. (D) Singular values of \mathbf{N} and \mathbf{A}_{IH} . (E) Singular values of \mathbf{K} .

TABLE II
EVALUATION OF INVERSE SOLUTIONS OF THE HOMOGENEOUS TORSO
SIMULATION SHOWN IN FIG. 8

	Noise Level	Mesh A	Mesh B	Mesh C
Mean RE	Noise Free	8.81%	5.85%	4.19%
	30 dB	9.02%	7.90%	7.17%
	20 dB	14.33%	15.45%	15.26%
λ	Noise Free	0.0077	0.0016	0.0005
	30 dB	0.0080	0.0042	0.0035
	20 dB	0.0205	0.0207	0.0215
MGL Error	Noise Free	1	1	1
	30 dB	1	1	0.98
	20 dB	1	0.82	0.7

The data represent the arithmetic average of 50 repeated simulations. White noise of 30 and 20 dB SNR is added to the torso-surface measurements before solving the inverse problem. MGL: maximum gradient location in the vector of the inverse solution.

to 0.0005, indicating that an improved \mathbf{K} requires less regularization; consequently, the error is reduced from 8.81% to 4.19%. Similar improvement is also observed in the case of 30 dB noise, but is not evident when the noise level increases to 20 dB. Although the quantitative metrics do not indicate evident improvement in the case of 20 dB input noise, a visual inspection of the inverse solution shows that the solution calculated from mesh C in Fig. 8 indeed better captures the features of the exact solution than does the solution from mesh A, as shown in




Fig. 9. Plots of real epicardial potential signals (inverse solutions) corresponding to Fig. 8 and Table II. Mesh C yields a better solution than mesh A.

Fig. 9. In summary, all results are consistent with those from the corresponding experiment conducted on the annulus model.

2) *Resolution in the Normal Direction:* We explore the impact of the resolution in the normal direction by refining a region

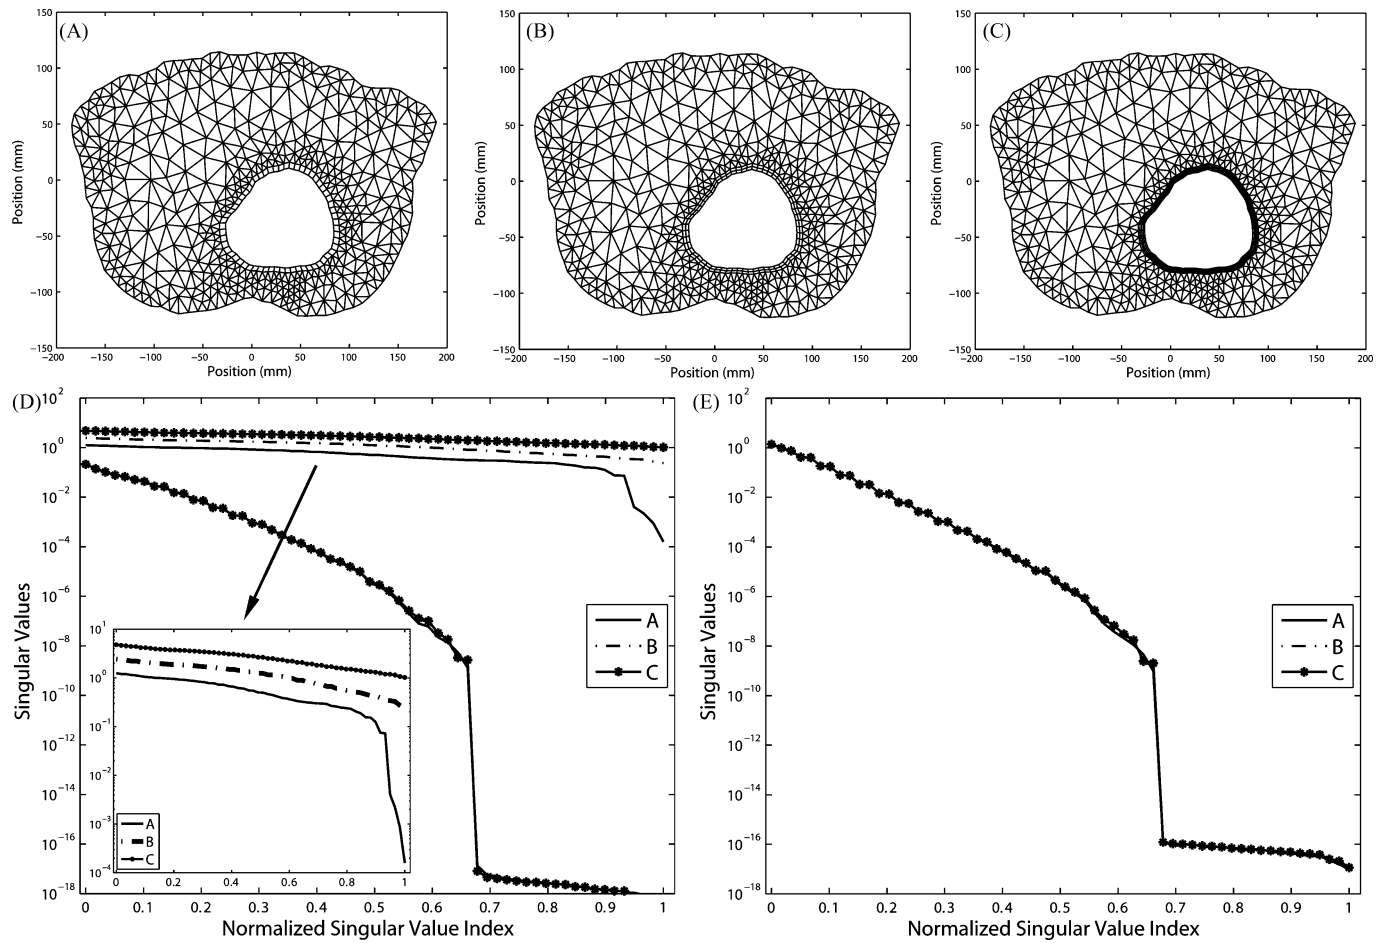


Fig. 10. Refining the resolution normal to the heart under a coarse volume discretization. (A) One layer of quadrilaterals around the heart. (B) Two layers of quadrilaterals. (C) Four layers of quadrilaterals. (D) Singular values of \mathbf{N} and \mathbf{A}_{IH} . (E) Singular values of \mathbf{K} .

of quadrilateral layers near the heart while fixing the rest of the volume mesh. Fig. 10 demonstrates the setup of three meshes, and the resulting \mathbf{A}_{IH} , \mathbf{N} , and \mathbf{K} in terms of singular values. The same conclusions are drawn as those from the annulus test in Section IV-B3. Increasing the resolution in the normal direction improves \mathbf{A}_{IH} , the boundary-to-volume lifting operator. The abrupt drop in the singular values implies that the azimuthal resolution in the volume still dominates the basic quality of the numerical system. This test uses an underrefined volume mesh. We also conducted the test with a well-refined homogeneous torso model, and drew the same results as in the annulus test in Section IV-B3. To avoid duplication, the latter test is not presented in this paper.

D. 2-D Heterogeneous Realistic Geometry

Simulations are performed using a heterogeneous 2-D torso mesh, which conforms to the interfaces between different physiological tissues. The same epicardial data are taken as with the homogeneous torso model. We repeat the refinement tests in the volume and in the normal direction around the heart. This model differs from the previous homogeneous model in that refinements must respect boundary interfaces between organs and tissues. This study intends to demonstrate that the conclusions

drawn from homogeneous meshes also hold in heterogeneous meshes in general. Fig. 11 presents results of volume refinement and torso refinement. Panel (A) shows the original mesh generated by tissue segmentation. One can discern the epicardium, lungs, skeletal muscle, and torso surface. To simplify the problem, we make two types of refinement: refining the lungs, as shown in panel (B), and refining the tissue outside the lungs, as shown in panel (c). Panel (D) displays the combination of both refinements. By inspecting \mathbf{N} in panel (E) and \mathbf{K} in panel (E), one can see that the lung refinement extends the singular value spectrum, reducing the proportion of trivial singular values, whereas the refinement on the torso surface extends the spectrum slightly but meanwhile “lifts” it as well. Compared to mesh A, mesh D combines the improvement brought by meshes B and C.

Fig. 12 shows the test of increasing the resolution normal to the heart boundary. The results are consistent with the corresponding observation in the homogeneous torso case: The boundary-to-interior lifting matrix \mathbf{A}_{IH} is improved while the improvement of \mathbf{N} and \mathbf{K} is limited.

V. DISCUSSION

Our primary concern is how a finite-element discretization of the ECG model influences the numerical conditioning of the

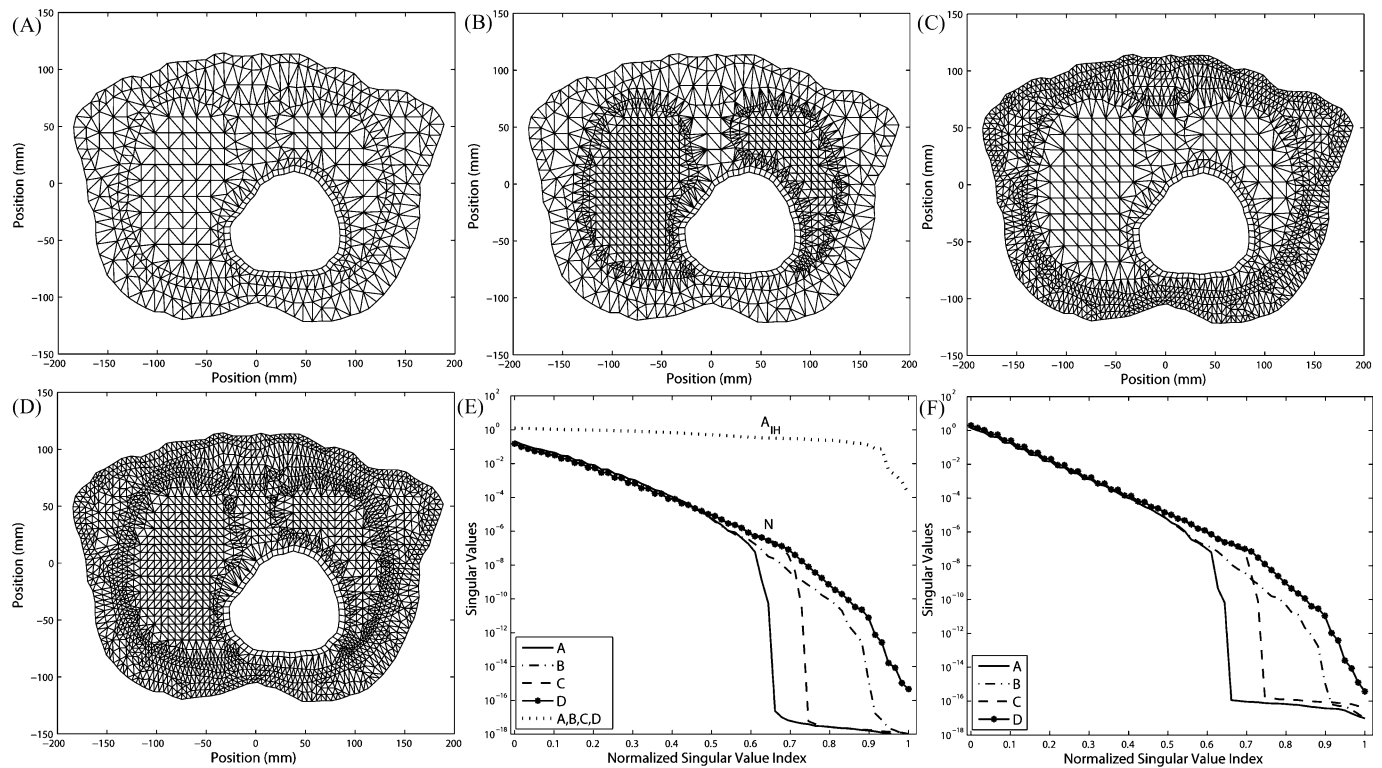


Fig. 11. Effects of refining different regions in the heterogeneous torso mesh while keeping the epicardium intact. (A) Original mesh. (B) Refining the lungs. (C) Refining skeletal muscles and torso surface. (D) Combining of B and C. (E) Singular values of \mathbf{N} and \mathbf{A}_{IH} . (F) Singular values of \mathbf{K} .

inverse problem, and consequently, the quality of the inverse solution. The impact of discretization choices is first reflected in the formulation of the forward transfer matrix \mathbf{K} . Since the inverse solution is obtained by conceptually “inverting” \mathbf{K} , the singular values of \mathbf{K} (and its components such as \mathbf{N} and \mathbf{A}_{IH}) provide a quantitative measure of the numerical quality of the discrete inverse problem. A well-formulated \mathbf{K} should be characterized by a slowly descending singular value spectrum and a small proportion of trivial singular values corresponding to the null space of \mathbf{K} (theoretically, source information that falls into the null space will be completely filtered out by \mathbf{K} in the forward problem, and therefore, can never be recovered in the inverse problem, thus worsening its ill-posedness).

Although singular values describe general properties of the system \mathbf{K} to be solved, they are not directly associated with the accuracy of the inverse solution, which relies heavily on regularization of the still ill-conditioned system. In the classic Tikhonov regularization, the parameter λ controls the amount of regularization, and therefore, indicates the ill-conditioning of \mathbf{K} reflected in the problem-solving stage. These two measures, along with the solution error evaluation, constitute the metrics to assess the impact of discretization in this study.

When the volume conductor is insufficiently discretized, Fourier analysis implies that the actually recoverable spatial-frequency band limit allowed by the discrete system is less than the frequency band limit assumed to be recovered on the heart boundary. This implication is reflected in the transfer matrix \mathbf{K} by the truncation of its nontrivial singular values and by the widening of its null space. This phenomenon is referred to as

artificial ill-conditioning because it is not due to the ill-posed nature of the continuum problem. Volume refinement effectively eliminates this type of ill-conditioning by improving the singular value spectrum and reducing the null space. Such improvement on the transfer matrix also shows its benefits in the problem-solving stage by reducing the necessary regularization amount, which is an improvement that ultimately leads to less error in the inverse solution. This sequence of improvements is evident especially when there is no measurement noise in the input torso potential vector, because in the noise-free case, the only error in (28) is the discretization error that lies in \mathbf{K} only. Discretization refinement promotes the accuracy of \mathbf{K} , and hence, directly improves the desired solution \mathbf{u}_H .

However, such improvement in the inverse solution becomes less evident as input noise increases, until the noise reaches a level that may obscure the improvement. This occurs because any discretization refinement cannot change the underlying ill-posed nature of the physical problem. The well-refined inverse problem, though improved, is still ill-conditioned and very sensitive to the perturbation of input noise (as can be seen, the most improved \mathbf{K} still has a condition number of 10^{14}). This observation suggests that volume refinement is necessary before reaching a certain level, beyond which it however becomes unnecessary for computational efficiency.

All results presented in our study corroborate the following set of “guidelines” for the placement of resolution with the finite-element-based discretization of the inverse ECG problem.

- 1) Increasing the resolution on the heart surface leads to a corresponding increase in the ill-conditioned nature of the

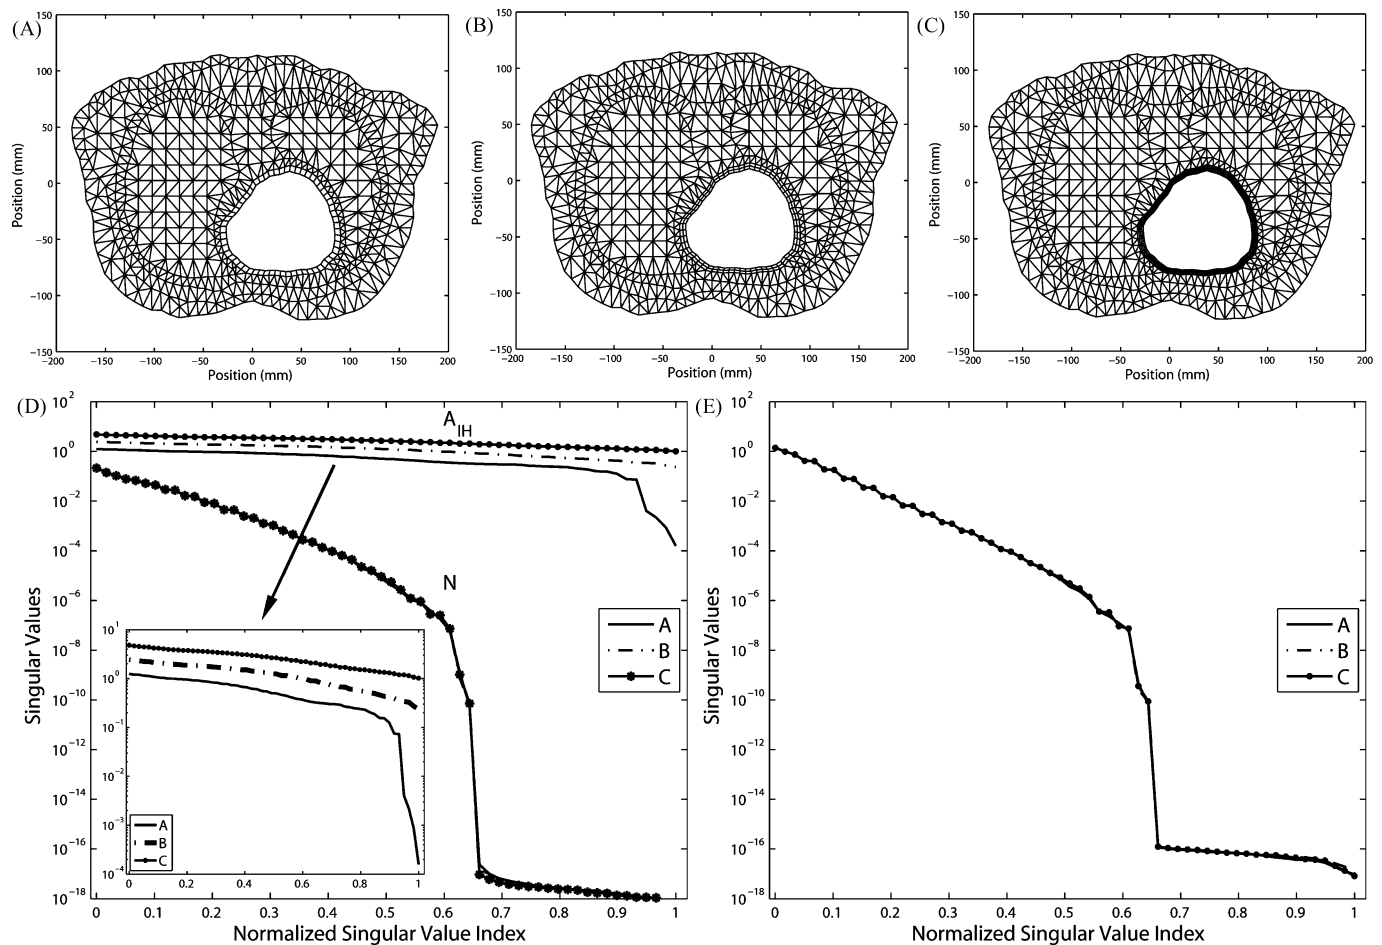


Fig. 12. Refining the resolution normal to the heart. (A) One layer of quadrilaterals around the heart. (B) Two layers of quadrilaterals. (C) Four layers of quadrilaterals. (D) Singular values of \mathbf{N} and \mathbf{A}_{IH} . (E) Singular values of \mathbf{K} .

discretized inverse system. One should realistically assess the resolution needed on the interior boundary to satisfy the problems of interest and be careful not to solve beyond the minimum resolution needed.

- 2) Some benefit can be gained from increasing the resolution normal to the heart boundary. This corresponds to a refinement of the right-hand side operator within the finite-element forward problem—the boundary-to-interior lifting operator. The challenge historically with unstructured discretizations has been to effectively increase normal resolution while maintaining tangential (or azimuthal in our test cases) resolution. With element types such as triangles or tetrahedra, the effort to strike a balance between the two normally leads to poorly shaped elements with their own conditioning issues. We advocate the use of hybrid discretizations—in particular, the use of quadrilateral elements for the heart surface in 2-D. The quadrilateral elements can be connected to triangular elements in the volume. We suggest the use of prismatic elements on the heart's surface in 3-D. The prism elements can be connected to tetrahedral or hexahedral elements within the volume. This allows for fixing the heart surface resolution while increasing the resolution normal to the heart.

- 3) Once the other two items are in place, one should increase the resolution of the volume conductor. Although, theoretically, one can argue that monotonic increases in the resolution of the volume conductor leads to continual improvement of the inverse problem, there is a fundamental practical limitation in terms of the amount of information provided on the torso boundary. The resolution within the volume conductor needs to be sufficient to capture both the features of the torso data and the features implied by the discretization of the heart boundary, but for computational efficiency, it should not be much greater.
- 4) Increasing the resolution on the exterior boundary can positively impact the conditioning of the inverse system, but only when this increase comes as a consequence of increasing the measured data available on the exterior boundary. Merely embedding the data in a higher dimensional (resolved) function space on the boundary does not necessarily lead to better inverse problem solution.

These guidelines mainly advocate refining the volume conductor while judiciously deciding the fidelity of the discretization on the heart surface. One might use these guidelines to advise under appropriate scenarios the use of the BEM, which achieves an exact solution in the volume, which is a property

sought for by volume refinement in the FEM. Given that the BEM only requires discretizing interfaces between regions of different conductivities, it is worth studying the impact of resolution on each interface.

VI. CONCLUSION

This study explored the impact of different discretizations on the inverse ECG problem and proposed refinement strategies oriented to specific considerations in the inverse problem. We showed that refinement strategies developed solely for optimizing the forward problem may be inappropriate for the corresponding inverse problem. We then proposed refinement strategies for the inverse problem that were based on the analysis of its ill-posedness. Such strategies include determining the resolution on the heart surface, determining the resolution in the volume conductor, and decomposing the resolution in tangential and normal direction near the heart by employing quadrilateral elements. Impacts of refinement strategies were assessed by inspecting the singular values, the regularization behaviors, and the error measurements of the inverse solution. The results obtained both from the annulus model and the realistic torso model consistently indicated that discretization refinement by itself is one form of regularization, and can be combined with other classic regularization techniques to further improve the inverse solution.

We currently considered the problem in two dimensions. Future work includes extending this study to realistic ECG problems in three dimensions. Our preliminary results indicate that the guidelines presented in the previous section carry over in a straightforward manner. Issues still to be resolved include how to assess the epicardial potential information on the heart surface and how to properly define its spatial frequency spectrum on the 3-D surface. It is also worth addressing how the anisotropy, as well as discontinuities in the conductivity fields, influences discretization considerations in the inverse problem. Given that the size of the problem in three dimensions is significantly larger than that in two dimensions, computational resource will pose another major constraint in practical simulations. Thus, another probable future research direction is to optimize resource distributions given limited computation resources.

ACKNOWLEDGMENT

The authors would like to thank Dr. R. S. MacLeod and Dr. D. H. Brooks for their helpful discussions and comments. The authors also gratefully appreciate the assistance of Dr. Stinstra in simulations.

REFERENCES


- [1] I. Babuska and J. T. Oden, "Verification and validation in computational engineering and science: Basic concepts," *Comput. Methods Appl. Mech. Eng.*, vol. 193, no. 36–38, pp. 4057–4066, 2004.
- [2] J. A. Schmidt, C. R. Johnson, J. C. Eason, and R. S. MacLeod, "Applications of automatic mesh generation and adaptive methods in computational medicine," in *Modeling, Mesh Generation, and Adaptive Methods for Partial Differential Equations*. Berlin, Germany: Springer-Verlag, 1995, pp. 367–390.
- [3] C. R. Johnson, "Computational and numerical methods for bioelectric field problems," *Crit. Rev. Biomed. Eng.*, vol. 25, no. 1, pp. 1–81, 1997.
- [4] A. N. Tikhonov and V. Y. Arsenin, *Solutions of Ill-Posed Problems*. Washington, DC: V. H. Winston & Sons, 1977.
- [5] A. N. Tikhonov, A. V. Goncharsky, V. V. Stepanov, and A. G. Yagola, *Numerical Methods for the Solution of Ill-Posed Problems*. Dordrecht, The Netherlands: Kluwer, 1995.
- [6] V. Isakov, *Inverse Problems for Partial Differential Equations*, 2nd ed. New York: Springer-Verlag, 2006.
- [7] H. W. Engl, M. Hanke, and A. Neubauer, *Regularization of Inverse Problems*. Norwell, MA: Kluwer, 2000.
- [8] A. Kirsch, *An Introduction to the Mathematical Theory of Inverse Problems*. New York: Springer-Verlag, 1996.
- [9] J. Kaipio and E. Somersalo, *Statistical and Computational Inverse Problems*. New York: Springer-Verlag, 2004.
- [10] F. Natterer, "Error bounds for Tikhonov regularization in Hilbert scales," *Appl. Anal.*, vol. 18, pp. 29–37, 1984.
- [11] P. Mathé and S. V. Pereverzev, "Optimal discretization of inverse problems in Hilbert scales, regularization and self-regularization of projection methods," *SIAM J. Numer. Anal.*, vol. 38, no. 6, pp. 1999–2021, 2001.
- [12] H. Engl and A. Neubauer, "Convergence rates for Tikhonov regularization in finite-dimensional subspaces of Hilbert scales," *Proc. Amer. Math. Soc.*, vol. 102, pp. 587–592, 1988.
- [13] J. Kaipio and E. Somersalo, "Statistical inverse problems: Discretization, model reduction and inverse crimes," *J. Comput. Appl. Math.*, vol. 198, no. 2, pp. 493–504, 2007.
- [14] S. Lasanen and L. Roininen, "Statistical inversion with Green's priors," presented at the 5th Int. Conf. Inverse Problems Eng.: Theory Pract., Cambridge, U.K., 2005.
- [15] M. Seger, G. Fischer, R. Modre, B. Messnarz, F. Hanser, and B. Tilg, "Lead field computation for the electrocardiographic inverse problem finite elements versus boundary elements," *Comput. Methods Programs Biomed.*, vol. 77, pp. 241–252, 2005.
- [16] D. M. Weinstein, C. R. Johnson, and J. A. Schmidt, "Effects of adaptive refinement on the inverse EEG solution," *Proc. SPIE*, vol. 2570, pp. 2–11, 1995.
- [17] C. R. Johnson and R. S. MacLeod, "Nonuniform spatial mesh adaptation using a posteriori error estimates: Applications to forward and inverse problems," *Appl. Numer. Math.*, vol. 14, pp. 311–326, 1994.
- [18] G. Shou, L. Xia, M. Jiang, F. Liu, and S. Crozier, "Forward and inverse solutions of electrocardiography problem using an adaptive BEM method," in *Functional Imaging and Modeling of the Heart* (Lecture Notes in Computer Science), vol. 4466. Berlin, Germany: Springer-Verlag, 2007, pp. 290–299.
- [19] C. R. Johnson, "Adaptive finite element local regularization methods for the inverse ECG problem," in *Inverse Problems in Electrocardiology* (Advances in Computational Biomedicine). Southampton, U.K.: WIT Press, 2001.
- [20] A. J. Pullan and C. P. Bradley, "A coupled cubic Hermite finite element/boundary element procedure for electrocardiographic problems," *Comput. Mech.*, vol. 18, pp. 356–368, 1996.
- [21] G. Fischera, B. Tilga, P. Wacha, R. Modrea, U. Lederb, and H. Nowak, "Application of high-order boundary elements to the electrocardiographic inverse problem," *Comput. Methods Programs Biomed.*, vol. 58, pp. 119–131, 1999.
- [22] B. Kaltenbacher, "V-cycle convergence of some multigrid methods for ill-posed problems," *Math. Comput.*, vol. 72, no. 244, pp. 1711–1730, 2003.
- [23] M. Mohr and U. Rude, "Multilevel techniques for the solution of the inverse problem of electrocardiography," presented at the Multigrid Methods VI, 6th Eur. Multigrid Conf., Gent, Belgium, vol. 14, 1999.
- [24] C. R. Johnson, M. Mohr, U. Rude, A. Samsonov, and K. Zyp, "Multilevel methods for inverse bioelectric field problems," in *Lecture Notes in Computational Science and Engineering—Multiscale and Multiresolution Methods: Theory and Applications*. Heidelberg, Germany: Springer-Verlag, 2001.
- [25] P. R. Johnston, S. J. Walker, J. A. K. Hyttinen, and D. Kilpatrick, "Inverse electrocardiographic transformations: Dependence on the number of epicardial regions and body surface data points," *Math. Biosci.*, vol. 120, pp. 165–187, 1994.
- [26] R. M. Gulrajani, F. A. Roberge, and G. E. Mailloux, "The forward problem of electrocardiography," in *Comprehensive Electrocardiology*, P. W. Macfarlane and T. D. V. Lawrie, Eds. Oxford, U.K.: Pergamon, 1989, pp. 197–236.

- [27] T. J. R. Hughes, *The Finite Element Method*. Englewood Cliffs, NJ: Prentice-Hall, 1987.
- [28] O. Axelsson, *Iterative Solution Methods*. New York: Cambridge Univ. Press, 1994.
- [29] G. E. Karniadakis and S. J. Sherwin, *Spectral/hp Element Methods for CFD*. New York: Oxford Univ. Press, 1999.
- [30] Y. Yamashita, "Theoretical studies on the inverse problem in electrocardiography and the uniqueness of the solution," *IEEE Trans. Biomed. Eng.*, vol. BME-29, no. 11, pp. 719–725, Nov. 1982.
- [31] T. Ohe and K. Ohnaka, "A precise estimation method for locations in an inverse logarithmic potential problem for point mass models," *Appl. Math. Model.*, vol. 18, pp. 446–452, 1994.
- [32] T. Ohe and K. Ohnaka, "An estimation method for the number of point masses in an inverse logarithmic potential problem using discrete Fourier transform," *Appl. Math. Model.*, vol. 19, pp. 429–436, 1995.
- [33] T. Nara and S. Ando, "A projective method for an inverse source problem of the Poisson equation," *Inverse Problems*, vol. 19, pp. 355–369, 2003.
- [34] B. Messnarz, M. Seger, R. Modre, G. Fischer, F. Hanser, and B. Tilg, "A comparison of noninvasive reconstruction of epicardial versus transmbrane potentials in consideration of the null space," *IEEE Trans. Biomed. Eng.*, vol. 51, no. 9, pp. 1609–1618, Sep. 2004.
- [35] P. C. Hansen, "Analysis of discrete ill-posed problems by means of the l-curve," *SIAM Rev.*, vol. 34, no. 4, pp. 561–580, 1992.
- [36] B. Punske, Q. Ni, R. L. Lux, R. S. MacLeod, P. R. Ershler, T. J. Dustman, M. J. Allison, and B. Taccardi, "Spatial methods of epicardial activation time determination in normal hearts," *Ann. Biomed. Eng.*, vol. 31, pp. 781–792, 2003.
- [37] A. Ghodrati, D. Brooks, G. Tadmor, and R. Macleod, "Wavefront-based models for inverse electrocardiography," *IEEE Trans. Biomed. Eng.*, vol. 53, no. 9, pp. 1821–1831, Sep. 2006.
- [38] I. Babuška, T. Strouboulis, C. S. Upadhyay, and S. K. Gangaraj, "A model study of the quality of a posteriori estimators for linear elliptic problems error estimation in the interior of patchwise uniform grids of triangles," *Comput. Methods Appl. Mech. Eng.*, vol. 114, pp. 307–378, 1994.
- [39] R. S. MacLeod, C. R. Johnson, and P. R. Ershler, "Construction of an inhomogeneous model of the human torso for use in computational electrocardiography," in *Proc. c-EMBS91*. Piscataway, NJ: IEEE Press, 1991, pp. 688–689.
- [40] C. R. Johnson, R. S. MacLeod, and P. R. Ershler, "A computer model for the study of electrical current flow in the human thorax," *Comput. Biol. Med.*, vol. 22, no. 3, pp. 305–323, 1992.



Dafang Wang (S'09) was born in Hangzhou, China, in 1984. He received the B.S. degree in computer science from Zhejiang University, Hangzhou, in 2005. He is currently working toward the Ph.D. degree in computer science at the School of Computing, University of Utah, Salt Lake City.

He is also a Research Assistant at the Scientific Computing and Imaging Institute, University of Utah. His current research interests include inverse problems, scientific computing, and computational electrocardiography.



current research interests

Robert M. Kirby (M'04) received the M.S. degree in applied mathematics, the M.S. degree in computer science, and the Ph.D. degree in applied mathematics from Brown University, Providence, RI, in 1999, 2001, and 2002, respectively.

He is currently an Associate Professor of computer science with the School of Computing, University of Utah, Salt Lake City, where he is also an Adjunct Associate Professor in the Departments of Bioengineering and Mathematics and a member of the Scientific Computing and Imaging Institute. His



Chris R. Johnson (SM'04) received the Ph.D. degree from the University of Utah, Salt Lake City, in 1989.

He founded the Scientific Computing and Imaging (SCI) Research Group (now SCI Institute), University of Utah, Salt Lake City, in 1992, where he is currently a Distinguished Professor of computer science and holds faculty appointments in the Departments of Physics and Bioengineering. His current research interests include scientific computing and scientific visualization. He is a member of editorial boards of several international journals. He is also a member of the advisory boards to several national research centers.

Prof. Johnson has received several awards, including the National Science Foundation Presidential Faculty Fellow Award from President Clinton in 1995 and the Governor's Medal for Science and Technology from Governor Michael Leavitt in 1999. He is a Fellow of the American Institute for Medical and Biological Engineering, the American Association for the Advancement of Science, and the Society of Industrial and Applied Mathematics.