Bayesian Solutions and Performance Analysis in Bioelectric Inverse Problems

Yeşim Serinağaoğlu*, *Member, IEEE*, Dana H. Brooks, *Member, IEEE*, and Robert S. MacLeod, *Member, IEEE*

***Abstract*—In bioelectric inverse problems, one seeks to recover bioelectric sources from remote measurements using a mathematical model that relates the sources to the measurements. Due to attenuation and spatial smoothing in the medium between the sources and the measurements, bioelectric inverse problems are generally ill-posed. Bayesian methodology has received increasing attention recently to combat this ill-posedness, since it offers a general formulation of regularization constraints and additionally provides statistical performance analysis tools. These tools include the estimation error covariance and the marginal probability density of the measurements (known as the "evidence") that allow one to predictively quantify and compare experimental designs. These performance analysis tools have been previously applied in inverse electroencephalography and magnetoencephalography, but only in relatively simple scenarios. The main motivation here was to extend the utility of Bayesian estimation techniques and performance analysis tools in bioelectric inverse problems, with a particular focus on electrocardiography. In a simulation study we first investigated whether Bayesian error covariance, computed without knowledge of the true sources and based on instead statistical assumptions, accurately predicted the actual reconstruction error. Our study showed that error variance was a reasonably reliable qualitative and quantitative predictor of estimation performance even when there was error in the prior model. We also examined whether the evidence statistic accurately predicted relative estimation performance when distinct priors were used. In a simple scenario our results support the hypothesis that the prior model that maximizes the evidence is a good choice for inverse reconstructions.**

***Index Terms*—Bayesian estimation, electrocardiography, inverse solutions.**

I. Introduction

THE PRIMARY GOAL in bioelectric inverse problems is to recover a bioelectric source distribution given a set of remote measurements and a mathematical model that relates the desired spatial source distribution to the measurement data distribution. The two best-known examples are inverse electrocardiography (ECG) or magnetocardiography (MCG) problems, in which one seeks to reconstruct cardiac electrical source distributions from the distributions of electrical potentials [1], [2] or the magnetic field on or outside the body surface, and inverse electroencephalography (EEG) or magnetoencephalography (MEG) problems, in which one seeks to reconstruct the distribution of sources in the brain given electrical or magnetic measurements on or outside the head [3], [4].

It is well known that due to attenuation and spatial smoothing inside the medium between the desired sources and the measurement sites, these bioelectric inverse problems are generally ill-posed; small disturbances in the measured data may yield large variations in the reconstructed source distribution. The most common approach to overcome this ill-posedness is to employ deterministic "regularization," in which the solution is a trade-off between a good fit to the measured data and an *a priori* constraint imposed on the solution [5], with a regularization parameter determining the trade-off point. One alternate approach to regularization which has received increasing attention recently in the literature is the Bayesian methodology [6], [7].

A key requirement of this Bayesian methodology is that one has prior information about the otherwise unknown sources with which to define prior probability density functions. In this study, we used a database of prior measurements from experiments, which served as training data from which to estimate models for these prior densities. The Bayesian formulation allows one to include such measurements and prior information in a "natural" way, offering a general formulation of the regularization constraints; but perhaps more importantly, it provides statistical performance analysis tools, such as estimation error covariance and the marginal probability density function of the measurements, known as the "evidence," that allow one to predictively quantify and compare a variety of experimental designs. For instance, one can use the Bayesian error covariance to predict whether obtaining extra measurements would lead to a significant improvement in one's confidence in the inverse estimates, or apply the Bayesian "evidence" statistic to compare different databases used to create a prior model. As we describe below, Bayesian approaches to the inverse electrocardiography problem have concentrated on the formulation of statistically-defined regularization methods. In the EEG/MEG literature, application of the performance analysis tools has been reported, but generally with relatively simple scenarios such as independent sources or the restriction of probabilistic parameters to the parameters of the medium [7], [8].

Manuscript received February 24, 2004; revised October 31, 2004. This work was supported in part by NIH National Center for Research Resources under Grant P41-RR12553, the Center for Bioelectric Field Modeling, Simulation, and Visualization, the Whitaker Foundation, the Nora Eccles Treadwell Foundation, and the Richard A. and Nora Eccles Harrison Fund for Cardiovascular Research. *Asterisk indicates corresponding author.*

*Y. Serinağaoğlu is with the Department of Electrical and Electronics Engineering, Middle East Technical University, Ankara, Turkey (e-mail: yserin@metu.edu.tr).

D. H. Brooks is with the Communications and Digital Signal Processing (CDSP) Center, Electrical, and Computer Engineering Department, Northeastern University, Boston, MA 02115 USA.

R. S. MacLeod is with the Bioengineering Department, The Scientific Computing and Imaging Institute, and Nora Eccles Harrison Cardiovascular Research and Training Institute, University of Utah, Salt Lake City, UT 84101 USA.

Digital Object Identifier 10.1109/TBME.2005.846725

In this paper, we extend the application of these statistical performance analysis tools to inverse ECG, a case in which the sources (in our formulation, the time-varying potentials at a mesh of nodes on the epicardium) are clearly correlated with each other across the surface. We expect that this extension will, in turn, be useful for EEG/MEG problems, although we do not directly treat them here. Also, for simplicity, we ignore the *temporal* correlation of the epicardial potentials, rather treating them as uncorrelated; the purpose of this study was to attempt to incorporate *spatially correlated* sources into a practical Bayesian framework.

To help further motivate the use of Bayesian tools, we consider cases in which there are two types of measurements, both resulting from the same underlying sources but each with its own biophysical or spatial relationship to those sources. We have two specific types of scenarios in mind. In the first, we assume one has the usual remote measurements from the body surface (e.g., ECG or EEG) but also access to a small number of measurement locations very near the sources; in this case, one would like to use these high quality but spatially sparse measurements to improve the inverse solution. Such a scenario could arise with venous-catheter-based endocardial [9] or epicardial [10] measurements, or perhaps in brain surgery in which a small number of cortical measurements might be available. In the second scenario, which is really a generalization of the first, one might have two distinct types of measurements, such as electrical and magnetic measurements. Although our formulation allows for both scenarios (with the restriction that we only deal with scalar measurements, although magnetic field measurements might be multicomponent), we concentrate on the former scenario applied to electrocardiology, specifically on the possibility that one has sparse epicardial potential measurements available, via multielectrode venous catheters, to use in an epicardial potential formulation of the inverse electrocardiography problem.

In the inverse ECG literature, Bayesian methods have generally assumed that the measurement noise and the epicardial potentials are Gaussian distributed and mutually (spatially) uncorrelated. Thus, the prior density of epicardial potentials could be defined in terms of a mean and a covariance matrix. This approach was introduced early in the development of the field by Martin *et al.* [11], who examined the feasibility of applying statistical constraints to inverse solutions. The mean vector and the covariance matrix were estimated through temporal averaging and Monte Carlo sampling approaches using potentials simulated from an activation sequence. The measurement noise model not only included the instrumentation noise, but also modeling errors. Results were encouraging enough to warrant further pursuit of statistical regularization methods. Barr *et al.* in [12] used a simplified version of this idea; zero-mean independent and identically-distributed (i.i.d.) epicardial potentials and noise. This type of Bayesian maximum *a posteriori* (MAP) formulation is equivalent to zero-order Tikhonov regularization with the regularization parameter equal to the ratio of the standard deviation of the noise in the torso measurements to the standard deviation of the epicardial potentials. Much more recently, van Oosterom [6], [13] improved this idea and used a more realistic volume conductor to obtain the covariance matrix used in the MAP equation. All of these studies ignored temporal correlation. Greensite, in [14] and [15], has proposed a statistical approach that combined temporal and spatial constraints, presenting ways to estimate temporal and spatial covariances under certain mathematical assumptions, based on structures inherent in the space-time correlation matrix.

In the EEG/MEG literature, there have been several reports using Bayesian tools for performance analysis as well as for inverse reconstruction [16]–[18]. In the Bayesian case, in which the unknowns are modeled as stochastic, the prior distribution function plays a role in determining error bounds [19]. Russell *et al.* [7] applied this idea to the inverse EEG/MEG problem using a discrete distributed model of cortical dipole current sources in a four-concentric-sphere head model. The stochastic source parameters were the dipole moments, which were assumed to be linearly related to the measurements. Both the sources and measurement noise were assumed to be i.i.d. and independent from each other. The only unknown statistical model parameters were the variances of the dipole moments and the measurement noise. In addition to using a Bayesian estimation error metric, the authors also used the evidence to estimate these variances, and calculated the information content of the measurements using differential entropy and mutual information. Radich *et al.* [8] derived error bounds for unbiased dipole location estimation under a general head model parameterized by both deterministic and stochastic parameters. In their derivation, the only stochastic terms were the head model parameters, such as conductivities of various layers within a four-shell spherical head. They did not study the effects of uncertainties in the source parameters. One non-Bayesian approach used by several investigators reported performance analysis by modeling the unknowns as deterministic; in this case, the investigators applied the well know Cramer-Rao bound (CRB) on the variance of unbiased estimates of deterministic parameters [20]–[22].

In this paper, we report a simulation study designed to address the following questions.

1) How reliable is the theoretical Bayesian estimation error covariance when the prior density distribution is not exactly known, but instead is estimated from previously recorded similar source distributions, i.e., training data? Can we use it as a substitute for the actual error variability, which can only be obtained from the actual estimation error, which in turn requires knowledge of the true solution?
2) How does the reliability of the estimation error covariance change when an "ideal" estimated prior model obtained by estimating it from the true sources is used?
3) If the source parameters are not i.i.d. but have a more complicated prior model, how can one adapt the approach of Russell *et al.* [7] of using the evidence for statistical parameter selection to accommodate this new model?

In what follows, we concentrate on the linear inverse epicardial potential problem, using Gaussian assumptions for both the unknown sources and any measurement noise. In a simulation study we investigate whether the Bayesian error covariance,

which can be computed based simply on the problem formulation without knowledge of the true sources, and is based on simplifying statistical assumptions and estimated prior statistics, accurately models the actual error of the methods, both with and without the use of sparse epicardial potentials. We also seek to establish whether the Bayesian evidence statistic accurately predicts estimation performance when different sets of training data are employed.

In Section II, we formulate a mathematical relationship between the epicardial potentials, and the measurements used for estimation of these potentials. In Section III, we first present Bayesian estimation and evaluation methods, then we explain how we obtained simulated measurements and our candidate training sets used in this study. In Section IV, we report on simulation results that examine the use of two Bayesian performance analysis tools: the error covariance and the evidence. Finally, in Section V, we discuss our results and present our conclusions.

II. Problem Definition

We denote the epicardial potentials as a vector $\mathbf{x}(k)$, of size $N \times 1$, obtained by stacking the values at each node on the heart surface into a column vector at time instant $k, k = 1, \ldots, K$, the corresponding torso potentials as $\mathbf{y}(k)$, of size $M \times 1$, and a forward solution matrix as $\mathbf{A}$, of size $M \times N$, which represents the mathematical model of the conducting medium between the sources and measurement sites. Then

$$\mathbf{y}(k) = \mathbf{A}\,\mathbf{x}(k) + \mathbf{n}(k) \tag{1}$$

where $\mathbf{n}(k)$ is the measurement noise vector. We assumed in this paper that the noise is additive, and we did not take into account possible errors or noise in the forward solution matrix.

In addition to $\mathbf{y}(k)$, we assume that we have available a second set of measurements, denoted as $\mathbf{z}(k)$, of size $L \times 1$, related to the desired source distribution via an equation

$$\mathbf{z}(k) = \mathbf{B}\mathbf{x}(k) + \mathbf{e}(k) \tag{2}$$

where $\mathbf{e}(k)$ is the measurement noise vector of appropriate size, and $\mathbf{B}$ is the mathematical model that relates $\mathbf{x}(k)$ to $\mathbf{z}(k)$. As discussed in Section I, although other scenarios can fit the same formulation, here we will call this second set of measurements, $\mathbf{z}(k)$, the "secondary measurements" and denote $\mathbf{y}(k)$ as the "primary measurements."

In the specific application we will test, $\mathbf{z}(k)$ contains direct measurements of a small subset of the elements of $\mathbf{x}(k)$. If we arrange our matrices and vectors so that the measured locations in $\mathbf{z}(k)$ correspond to the first L elements of $\mathbf{x}(k)\,(L \leq N), \mathbf{B}$ is equal to $[\mathbf{I_L}\ \ \mathbf{0}]$. We note that for other scenarios, such as combining electrical and magnetic measurements, $\mathbf{B}$ would be the forward matrix corresponding to the second measurement modality. Since we do not consider temporal correlation, and thus approach the problem independently at each time instant, for simplicity we drop the time index from our equations in the sequel.

A. Augmented Problem

There are three ways to choose the measurements used in an inverse solution: include only the primary measurements, include only the secondary measurements, or include both types of measurements.

If only the primary measurements are available, one solves the problem defined by (1) for $\mathbf{x}$. We will refer to this approach as "classical reconstruction." Using only the secondary measurements is not of interest here. To incorporate both types of measurements, we can reformulate the problem by combining (1) and (2) into an augmented form

$$\mathbf{v} = \mathbf{D}\,\mathbf{x} + \tilde{\mathbf{n}}, \quad \mathbf{v} = \begin{bmatrix} \mathbf{y} \\ \mathbf{z} \end{bmatrix}, \quad \tilde{\mathbf{n}} = \begin{bmatrix} \mathbf{n} \\ \mathbf{e} \end{bmatrix}, \quad \mathbf{D} = \begin{bmatrix} \mathbf{A} \\ \mathbf{B} \end{bmatrix} \tag{3}$$

and solve this equation for $\mathbf{x}$. We will refer to this approach as "augmented reconstruction."

All of the equations that use one or both types of measurements have a simple linear relation to the source distribution, with only the variable names changed. Therefore, we will present the following mathematical framework for only the classical reconstruction notation.

III. Methods

A. Bayesian Estimation

In the Bayesian approach, one estimates the value of the unknown parameters, which are assumed to be random with some known prior probability density function (pdf), so that the estimator is optimal "on the average." The following provides a brief summary of Bayesian methods used in this study. Refer to [23] for more details.

Both $\mathbf{x}$ and $\mathbf{y}$ are assumed to be random with a joint probability density function (pdf), $p(\mathbf{y}, \mathbf{x})$, and $\mathbf{x}$ has a prior pdf denoted by $p(\mathbf{x})$. Then the posterior pdf of $\mathbf{x}$ given the data $\mathbf{y}$ is

$$p(\mathbf{x}\backslash\mathbf{y}) = \frac{p(\mathbf{y}\backslash\mathbf{x})\,p(\mathbf{x})}{\int_X p(\mathbf{y}\backslash\mathbf{x})\,p(\mathbf{x})\,d\mathbf{x}} \tag{4}$$

where $p(\mathbf{y}\backslash\mathbf{x})$ is the conditional pdf of the measurement vector, $\mathbf{y}$, conditioned upon the parameter vector, $\mathbf{x}$, (also known as the likelihood function), and X is the parameter space.

One well-known estimation method is the Bayesian maximum *a posteriori* (MAP) estimation, in which the estimator is picked to maximize the posterior pdf. Another Bayesian estimation approach is the Bayesian minimum mean square error (MMSE) approach in which the estimator, chosen to minimize the square of the estimation error on the average, is the posterior mean. Since here we will assume that $\mathbf{x}$ and $\mathbf{y}$ are jointly normal, the Bayesian MMSE and MAP estimates coincide.

Specifically, we assume that $\mathbf{x}$ has a normal prior distribution with mean $\bar{\mathbf{x}}$ and covariance $\mathbf{C}_x$, the noise in the primary measurements is distributed as $\mathbf{n} \sim \mathrm{N}(\mathbf{0}, \mathbf{C}_n)$, with $\mathbf{C}_n = \sigma_n^2\mathbf{I}$, and the noise in the secondary measurements is $\mathbf{e} \sim \mathrm{N}(\mathbf{0}, \mathbf{C}_e)$, with $\mathbf{C}_e = \sigma_e^2\mathbf{I}$. Both noise terms are uncorrelated with $\mathbf{x}$, and

with each other. Under these conditions, the Bayesian MAP (or MMSE) solution equals

$$\hat{\mathbf{x}} = \left(\mathbf{A}^T\mathbf{C}_n^{-1}\mathbf{A} + \mathbf{C}_x^{-1}\right)^{-1}\left(\mathbf{A}^T\mathbf{C}_n^{-1}\mathbf{y} + \mathbf{C}_x^{-1}\bar{\mathbf{x}}\right) \quad (5)$$

where ^{T} and $^{-1}$ denote transpose and inverse of a matrix, respectively.

1) Estimation Error: One of the properties that makes Bayesian estimation attractive is that one can estimate the statistical characteristics of the estimation error, and then construct confidence intervals from these characteristics.

The Bayesian MAP reconstruction $\hat{\mathbf{x}}$, under the assumptions listed in the previous section (i.e., $\mathbf{x}$ and $\mathbf{y}$ are jointly normal, etc.), has mean equal to $\bar{\mathbf{x}}$. Then, the estimation error has zero mean (i.e., the estimate is unbiased) and covariance matrix equal to

$$\mathbf{C}_\epsilon = \left(\mathbf{A}^T\mathbf{C}_n^{-1}\mathbf{A} + \mathbf{C}_x^{-1}\right)^{-1}. \quad (6)$$

The diagonal elements of this error covariance matrix give the variances of the error in the estimate of the epicardial potential at the corresponding leads, i.e.,

$$\mathrm{var}(x_j - \hat{x}_j) = \mathbf{C}_\epsilon(j,j) \quad (7)$$

where x_j and $\hat{x}_j$ are the jth elements of $\mathbf{x}$ and $\hat{\mathbf{x}}$, respectively, and $\mathbf{C}_\epsilon(j,j)$ is the jth diagonal element of $\mathbf{C}_\epsilon$.

Using this variance and the Bayesian MAP solution, one can compute confidence intervals for the estimate. For example, with 95% probability, the true solution at lead j lies approximately within the range

$$\hat{\mathbf{x}}_j - 2\sqrt{\mathbf{C}_\epsilon(j,j)} \le \mathbf{x}_j \le \hat{\mathbf{x}}_j + 2\sqrt{\mathbf{C}_\epsilon(j,j)}. \quad (8)$$

If we map the error covariance or confidence intervals back onto the epicardium, we have a theoretical quantitative prediction of where on the epicardium we expect our method to give us more or less reliable results, and indeed an indication, in the appropriate physical unit, of this uncertainty.

B. Prior Selection

In the previous sections, we defined the prior density model of $\mathbf{x}$ in terms of the mean vector $\bar{\mathbf{x}}$ and the spatial covariance matrix $\mathbf{C}_x$ but did not discuss the specifics of how to obtain this model. In this section, we discuss a method to select the prior density from various candidate models using the evidence (i.e., the marginal probability distribution function of the measurement vector). This discussion extends the method reported by Russell *et al.* [7], in which only two independent parameters needed to be chosen, and these parameters were selected by plotting the evidence with respect to these variance parameters and choosing the values that yielded the maximum value. Here we assume a more complicated covariance matrix for which it is not straightforward to plot the evidence.

We start by assuming that we have available p different prior models, $H_1, H_2, \ldots, H_p$, where each model, H_i, can be represented by the mean vector $\bar{\mathbf{x}}^{(\mathbf{i})}$ and covariance matrix $\mathbf{C}_\mathbf{x}^{(\mathbf{i})}$. The hypothesis is that the prior model that maximizes the evidence is also the model that best explains the measurements. Under the Gaussian-distribution assumptions stated earlier, the marginal probability density function of the measurement vector, $\mathbf{y}$, given a prior model H_i, is

$$p(\mathbf{y}\backslash H_i) = \frac{1}{(2\pi)^{M/2}\left|\mathbf{C}_y^{(i)}\right|^{1/2}} \cdot \exp\left(-\frac{1}{2}\left(\mathbf{y} - \mathbf{A}\bar{\mathbf{x}}^{(i)}\right)^T\left(\mathbf{C}_y^{(i)}\right)^{-1}\left(\mathbf{y} - \mathbf{A}\bar{\mathbf{x}}^{(i)}\right)\right) \quad (9)$$

where $\mathbf{C}_y^{(i)} = \mathbf{A}^T\mathbf{C}_x^{(i)}\mathbf{A} + \mathbf{C}_n$ [23]. We pick the model that maximizes the evidence as our "best model" among the candidates.

We can equally well use the logarithm of the marginal density. Then, the best model becomes

$$H_{\mathrm{best}} = \arg\max_{H_i} \log[p(\mathbf{y}\backslash H_i)] \quad (10)$$

$$= \arg\max_{H_i}\left\{-\frac{1}{2}\left[M\log(2\pi) + \log\left(\left|\mathbf{C}_y^{(i)}\right|\right) + \left(\mathbf{y} - \mathbf{A}\bar{\mathbf{x}}^{(i)}\right)^T\left(\mathbf{C}_y^{(i)}\right)^{-1}\left(\mathbf{y} - \mathbf{A}\bar{\mathbf{x}}^{(i)}\right)\right]\right\}. \quad (11)$$

It is also possible to treat $\mathbf{C}_n$ as a variable parameter and try to choose it to maximize the evidence. In this work, we limit ourselves to choosing the prior model only, because there are often other ways to estimate the noise variance, for example by taking the sample variance of a region in the measurements in which we have only noise.

If there are a limited number of prior models available (e.g., from different training sets), one can pick the prior model that yields the largest evidence value. But it is also possible to solve the optimization problem introduced in (10) directly to find the best mean vector and covariance matrix.

One can either find a different "best model" for each vector $\mathbf{y}(k)$, where k denotes time instant, or find one model that will be used at all time instants and is the "best model" on the average.

C. Validation Data

In this paper, we report on simulation studies applying the Bayesian techniques described in previous sections to inverse electrocardiography. The source distribution in our formulation is the potential distribution on the heart surface. The primary measurements are the corresponding potential values on the body surface, and the secondary measurements are epicardial measurements measured at sites that lie over coronary veins. Such epicardial signals are equivalent to catheter measurements from the coronary veins [10] so that we refer to them as surrogate catheter measurements. These measurements are spatially very sparse. However, they are direct in the sense that they do not suffer from attenuation or spatial smoothing nearly as much as the body surface measurements.

1) Experiments: All data for this research came from experiments performed at the Nora Eccles Harrison Cardiovascular Research and Training Institute of the University of Utah. Two different canine heart preparations were used; an *in situ* preparation and an isolated heart preparation. In the *in situ* cases, the

heart remained in place in the animal. It was exposed through a midsternal opening and suspended in a pericardial cradle. In the isolated heart cases, a preparation described in [24], [25], consisting of an isolated dog heart suspended in a torso-shaped electrolytic tank, was used.

In each of preparations described above, we recorded epicardial potentials simultaneously at a sampling rate of 1000 Hz. from 490-electrode epicardial sock electrodes contained in a flexible array attached to a nylon stocking fitted over the ventricles [24], [25].

2) Simulated Measurements: We formed surrogate venous catheter signals by selecting 42 evenly spaced leads corresponding to the locations of the major cardiac veins in a 490-electrode sock applied to either an *in situ* or an isolated canine heart and added normally distributed zero mean i.i.d. noise at 30 dB SNR, approximately what we observe in the true catheter recordings. We simulated torso potentials using a boundary element solution to Laplace's equation for a human shaped torso tank in which the heart was suspended. We then added noise at 25 dB SNR to the simulated torso surface potentials. Again, this SNR value is approximately what we observe in the true body surface recordings. In (1) and (2), $N = 490$ and $M = 771$, and we set $L = 42$ wherever applicable.

D. A Priori Information From Training Sets

We obtained *a priori* information needed for Bayesian methods from training sets which consisted of epicardial potential maps by estimating the sample mean and covariance of epicardial potentials by averaging over time. We used two different protocols to create training and test data sets: "*Leave-one-beat-out*" and "*Leave-one-experiment-out*" protocols.

1) "Leave-One-Beat-Out" Protocol: In the "leave-one-beat-out" protocol, the beat we used for simulating the measurements (i.e., the test beat) and the training set beats came from the same experiment (same dog heart), but excluded the test beat. Clearly such a training set is unrealistic, since we do not normally expect to have access to recordings from the same heart surface. However, the prior distribution obtained from this kind of training set is a good approximation of the true prior, and it allows us to test the reliability of the theoretical error variances in a best case scenario.

2) "Leave-One-Experiment-Out" Protocol: In this second protocol, the test beat came from a different experiment (i.e., a different dog heart) than the experiments from which we obtained the training set beats. This scenario is more realistic than that of the "leave-one-beat-out" protocol and it allows us to study how prior-model mismatch affects the results.

IV. Results

In this section, we present the simulation results in two parts. Given that the main objective of this paper is to report on the use of Bayesian performance analysis tools, we include the results that examine the use of both the error covariance and the evidence. First, we report on the feasibility of using the theoretical Bayesian error covariance matrix, defined by (6), when we did not have the actual error values between the true epicardial potentials and their estimates. Next, we report on the feasibility of using the "evidence" as a prior-model selection criterion.

A. Theoretical Versus Computed Error Maps

We compared two types of error covariance matrices: a "theoretical error covariance" given in (6) which was predicted by the Bayesian model, and a "computed error covariance" calculated using the actual estimation error between the original potentials and the corresponding reconstructions. The goal was to determine if the theoretical error covariance could be a useful substitute for the computed one when we do not have the actual error values that are required to obtain the computed error covariance.

We carried out simulations using both the "leave-one-beat-out" and "leave-one-experiment-out" protocols, applied Bayesian MAP estimation, and solved both the classical inverse problem [**MAP** solution—(1)] and the augmented inverse problem [**MAP-ED** solution—(3)].

We used a simple simulation scenario: for K_p different test beats, we first found the **MAP** and **MAP-ED** solutions. If we define $\mathbf{X}_t, t = 1, \ldots, K_p$, as the true solution matrix from the test beat, where each row represents a node across time and each column a time instant across all nodes, and $\hat{\mathbf{X}}_t^{\mathrm{MAP}}$ and $\hat{\mathbf{X}}_t^{\mathrm{MAP-ED}}$ as the **MAP** and **MAP-ED** solution matrices, respectively, organized similarly to $\mathbf{X}_t$, then the computed error matrices corresponding to **MAP** and **MAP-ED** solutions are: $E_t^{\mathrm{MAP}} = \mathbf{X}_t - \hat{\mathbf{X}}_t^{\mathrm{MAP}}$ and $E_t^{\mathrm{MAP-ED}} = \mathbf{X}_t - \hat{\mathbf{X}}_t^{\mathrm{MAP-ED}}$, respectively. Using the error matrices obtained for every test beat, we obtained the augmented error matrices

$$E^{\mathrm{MAP}} = \left[E_1^{\mathrm{MAP}} \quad E_2^{\mathrm{MAP}} \quad \ldots \quad E_{K_p}^{\mathrm{MAP}} \right] \tag{12}$$

and

$$E^{\mathrm{MAP-ED}} = \left[E_1^{\mathrm{MAP-ED}} \quad E_2^{\mathrm{MAP-ED}} \quad \ldots \quad E_{K_p}^{\mathrm{MAP-ED}} \right]. \tag{13}$$

We estimated the "computed error covariance" matrices for **MAP** and **MAP-ED** solutions from these augmented error matrices by replacing the expectations in the definition of covariance matrix with averages over time and beat

$$\mathbf{C}_\epsilon^{\mathrm{MAP}} = \frac{1}{Q} E^{\mathrm{MAP}} \cdot (E^{\mathrm{MAP}})^T \tag{14}$$

and

$$\mathbf{C}_\epsilon^{\mathrm{MAP-ED}} = \frac{1}{Q} E^{\mathrm{MAP-ED}} \cdot (E^{\mathrm{MAP-ED}})^T \tag{15}$$

where Q is the number of columns in the augmented error matrices. If the number of time instances is the same for all test beats (i.e., K), then $Q = K_p \cdot K$.

To examine how the error variance was distributed over the surface of the heart, we obtained "theoretical" and "computed" error standard deviation maps by taking the square root of the diagonal entries of the "theoretical" and "computed" error covariance matrices, respectively, and mapping these standard deviation (std) values to their corresponding node locations on the heart surface. We used the visualization program *map3d* [26] to display these maps.

Next, we used confidence intervals to test the similarity of theoretical and computed error maps and their numerical accuracy as an error evaluation technique. The question posed was whether the percentage of actual error values of the computed solutions that fell within the interval obtained using the theoretical error variances in fact reflected the predicted confidence percentage. According to the theory presented in Section III-A1, the estimation error, $\epsilon_j(i)$, at lead j and time instant i, satisfies $-2\cdot\sqrt{\mathbf{C}_\epsilon(j,j)} \le \epsilon_j(i) \le 2\cdot\sqrt{\mathbf{C}_\epsilon(j,j)}$, with 95% probability, and $-3\cdot\sqrt{\mathbf{C}_\epsilon(j,j)} \le \epsilon_j(i) \le 3\sqrt{\mathbf{C}_\epsilon(j,j)}$, with 99% probability. We used the multibeat error matrices defined by (12) and (13) for this test. For each of the augmented error matrices, we divided the number of elements of the matrix that had values within the "confidence interval" by the total number of elements in the matrix. We compared "theoretical confidence interval," using the "theoretical" error standard deviations, and "computed confidence interval," using the "computed" error standard deviations. We performed this experiment for 95% and 99% confidence intervals. We then generated scatter plots of the theoretical error standard deviation values versus the computed ones, and fitted to them linear regression lines.

1) "Leave-One-Beat-Out" Protocol: For this protocol, we simulated 10 different torso and epicardial measurements using the same left ventricularly paced (LV-paced) test beat, the same noise variance, but different noise realizations. The training set included beats paced from both left ventricular (LV) and right ventricular (RV) surfaces.

Fig. 1 shows the "theoretical" and "computed" error standard deviation maps corresponding to **MAP** and **MAP-ED** solutions for the "leave-one-beat-out" protocol. This figure shows panels containing four different views of the heart and each panel consists of four error maps in which the top two maps are the theoretical error maps, and the bottom two are the computed error maps. The maps on the left correspond to **MAP** reconstruction and the ones on the right correspond to **MAP-ED** reconstruction. In all of these error maps, darker regions correspond to higher error standard deviation values. The range is fixed globally for all of the maps, across all of the views.

Based on these results (and others like them) we make the following observations.

- In the **MAP-ED** reconstruction, both the theoretical and the computed error maps displayed distinctly lower error standard deviation values around the epicardial measurement sites. In the anterior view, these small values lay along the 1 o'clock–8 o'clock line, which is along the left anterior descending artery (LAD), and in the posterior view, the small values appeared as continuous long and thin vertical areas that divide the posterior surface of the heart into three regions of larger error standard deviation values.

  The significance of this observation is not that the locations were correctly found, because we already start with the assumption that we know the measurement locations. Of greater interest is the fact that the error got smaller around the measurement sites, and that this reduction was visible not only in the computed error maps, but also in the theoretical error maps.

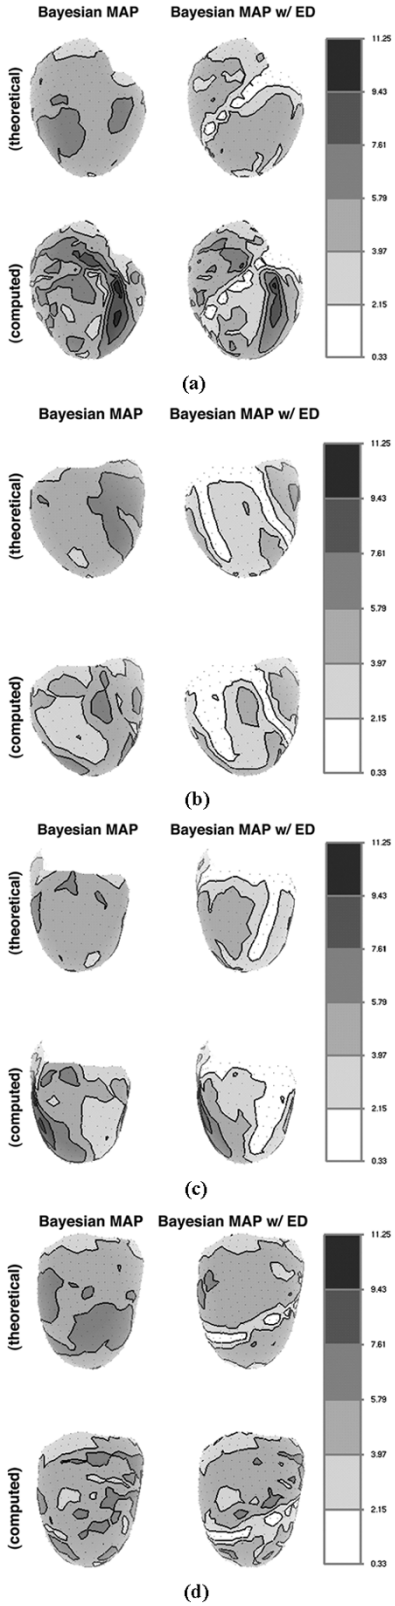

Fig. 1. Error standard deviation maps, observed from 4 different views of the heart using the "leave one beat out" protocol. In each view, the error plots on the left are based on **MAP** reconstruction, and the ones on the right are based on **MAP-ED** reconstruction with $K = 42$. The top two maps are the "theoretical error maps" and the two bottom maps are the "computed error maps." See text for details. (a) Anterior view, (b) posterior view, (c) LV view, and (d) RV view.

- Comparing the theoretical and computed error maps showed that both types of error maps show a decrease in standard deviation values in **MAP-ED** relative to **MAP**. This is consistent with a previous study [27], which reported that the **MAP-ED** detected small details missed by **MAP**.
- In the theoretical and computed error maps, we did not observe a detailed agreement, only a gross correspondence at some regions of the heart (e.g., along the coronary veins, in **MAP-ED**).

Tables I and II contain summaries of the percentage of error values that fell within the theoretical and computed confidence intervals for 95% and 99% predicted confidence intervals, respectively. We observe the following.

- The percentages of error values that fell within the confidence intervals, using both the computed and the theoretical standard deviation values, were consistent with the predicted percentage. The discrepancy was not more than 2% in both **MAP** and **MAP-ED** estimates.
- **MAP** and **MAP-ED** estimates showed very similar values in both 95% and 99% cases.

2) "Leave-One-Experiment-Out" Protocol: In this protocol, we simulated torso and epicardial measurements using 9 different test beats, each of which was paced at a different node on the LV surface. The noise variance was the same. The training set included beats paced from both LV and RV surfaces.

The results in Fig. 2 show the "theoretical" and "computed" error standard deviation maps for **MAP** and **MAP-ED** reconstructions. These results support the first two observations we made on Fig. 1. In addition, we make the following observations.

- Similar to the leave-one-beat-out case, there were some differences between the theoretical and computed error maps. On the average, computed error maps had higher error standard deviation values than the theoretical error maps.
- The distance between the measurement electrodes along the LAD was larger than the distance between the electrodes that lay on the posterior cardiac veins. This is reflected in both types of error maps, i.e., we observed small circular areas with small error standard deviation values along the LAD, and by looking at the centers of these circular areas, we could determine where we made the measurements. This was not as obvious in Fig. 1. The electrodes were placed closer to each other on the posterior veins than they were near the LAD, hence the small error standard deviation values appeared to be continuous, and we did not see gaps between the electrode locations.
- In the first test scenario, the theoretical maps "underpredicted" the regions of direct influence of the sparse measurements, in the sense that the regions of very low error variance were wider in the computed maps than in the theoretical ones. In the present scenario, the reverse is the case; the theoretical maps overpredicted the direct influence regions of the sparse measurements. This seems to indicate that mismatch between the actual test case and the prior model leads to a somewhat optimistic assessment of the error variance, especially close to the catheter measurement sites.

TABLE I
PERCENTAGE OF ERROR VALUES THAT FELL WITHIN THE CONFIDENCE INTERVAL, USING BOTH THE "COMPUTED" AND THE "THEORETICAL" STANDARD DEVIATION VALUES IN THE *LEAVE-ONE-BEAT-OUT* PROTOCOL. FOR THESE RESULTS, WE ASSUMED A 95% CONFIDENCE INTERVAL

	% error values within theor. interval	% error values within computed interval
MAP	0.931	0.932
MAP-ED	0.931	0.937

TABLE II
PERCENTAGE OF ERROR VALUES THAT FELL WITHIN THE CONFIDENCE INTERVAL, USING BOTH THE "COMPUTED" AND THE "THEORETICAL" STANDARD DEVIATION VALUES IN THE *LEAVE-ONE-BEAT-OUT* PROTOCOL. FOR THESE RESULTS, WE ASSUMED A 99% CONFIDENCE INTERVAL

	% error values within theor. interval	% error values within computed interval
MAP	0.976	0.984
MAP-ED	0.975	0.984

Comparisons of the percentage of error values that fell within the theoretical and computed confidence intervals for 95% and 99% predicted confidence intervals summarized in Tables III and IV suggested the following.

- The percentage of error values that fell within the confidence interval, according to the theoretical standard deviation values is smaller than that predicted by the computed values, in both **MAP** and **MAP-ED** estimates (the gap was generally around 5%). This is not surprising, since we observed previously that the computed standard deviations are larger than the theoretical ones, yielding a wider interval.
- The percentage of error values that fell within the confidence interval were almost the same with **MAP** and **MAP-ED**, for both 95% and 99% cases. This shows that even though the percentage of error values that fell within the confidence interval is not the same as the proposed percentage, the relative performance of the error covariance matrix is independent of the two approaches employed.

3) Scatter Plots and Line Fitting: In the previous two sections, we showed error covariance maps and calculated confidence intervals from these maps. The maps showed only gross correspondence, but the confidence interval results suggested a higher level of agreement. In order to better understand the relation between the theoretical and computed error standard deviation values, we obtained scatter plots of the two, and fitted linear regression lines. These scatter plots, and the corresponding regression lines are shown in Fig. 3 for leave-one-beat-out and leave-one-experiment-out protocols,

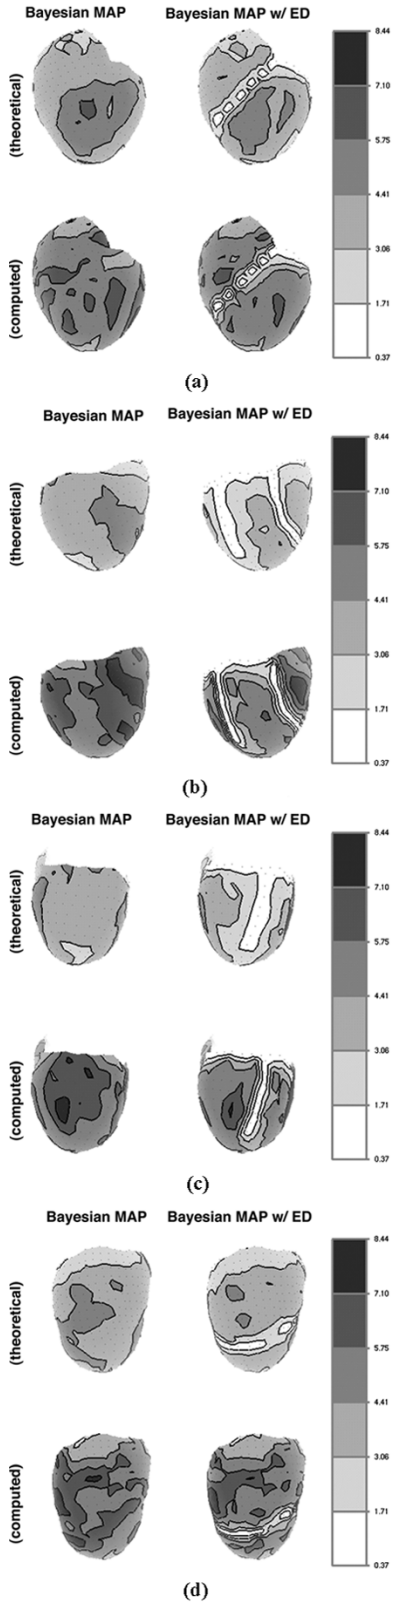

Fig. 2. Error standard deviation maps, observed from 4 different views of the heart using the "leave one experiment out" protocol. The layout and description of each panel is the same as Fig. 1. (a) Anterior view, (b) posterior view, (c) LV view, and (d) RV view.

using methods **MAP** and **MAP-ED**. The regression line, in each method, had slope smaller than one, suggesting a consistent underprediction of theoretical values.

TABLE III
PERCENTAGE OF ERROR VALUES THAT FELL WITHIN THE CONFIDENCE INTERVAL, USING BOTH THE "COMPUTED" AND THE "THEORETICAL" STANDARD DEVIATION VALUES IN THE *LEAVE-ONE-EXPERIMENT-OUT* PROTOCOL. FOR THESE RESULTS, WE ASSUMED A 95% CONFIDENCE INTERVAL

	% error values within theor. interval	% error values within computed interval
MAP	0.87	0.93
MAP-ED	0.88	0.93

TABLE IV
PERCENTAGE OF ERROR VALUES THAT FELL WITHIN THE CONFIDENCE INTERVAL, USING BOTH THE "COMPUTED" AND THE "THEORETICAL" STANDARD DEVIATION VALUES IN THE *LEAVE-ONE-EXPERIMENT-OUT* PROTOCOL. FOR THESE RESULTS, WE ASSUMED A 99% CONFIDENCE INTERVAL

	% error values within theor. interval	% error values within computed interval
MAP	0.93	0.97
MAP-ED	0.94	0.97

B. Comparison of Prior Models

The results in the previous section showed that the reliability of the error variance maps increases if a "good" prior distribution is used. Moreover, earlier studies have shown that the Bayesian MAP reconstructions using a "good" prior density produced reconstructions with good fidelity to the true potential distributions [6]. Therefore, it is important to create a training set that best represents the unknown source distribution.

According to the theory in Section III-B, one can obtain better reconstructions if one chooses the prior model that maximizes the evidence. We studied the feasibility of using the evidence to compare different models based on a very simple scenario, in which there are only three prior models, obtained using the leave-one-experiment-out protocol:

1) use a training set composed of left ventricularly beats only (LV-paced);
2) use a training set composed of right ventricularly beats only (RV-paced);
3) use a training set composed of both LV and RV paced beats (LVRV-paced).

The goal was to study the relation between the evidence value corresponding to each of these prior models, and the reliability of the reconstructions. The test beat used for this study was paced from the LV surface of the heart. The prior parameters (i.e., the mean vector and the covariance matrix) obtained from the LV-paced training set were coarsely similar to the 'true' mean and covariance (i.e., parameters estimated from the test beat), but those obtained from the RV-paced and LVRV-paced training sets were very different from the true mean and covariance.

Fig. 4 contains plots of evidence values found using logarithm of (9) at all time instances in the QRS interval (the interval that corresponds to the depolarization phase of the ventricular cells), for the three training sets. We note that the evidence stays approximately the same for all three priors until around 25 ms, at which a region of the epicardium around the pacing site has

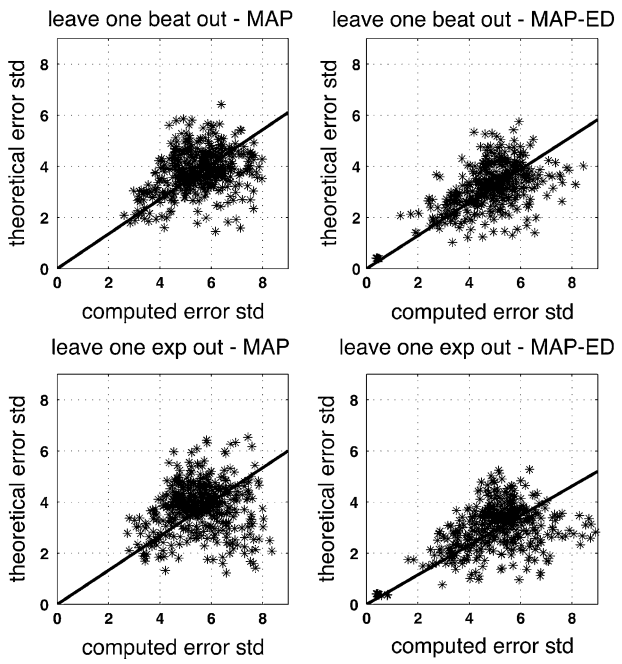

Fig. 3. Scatter plots and the corresponding regression lines. Slopes of the regression lines are: in the leave-one-beat-out protocol, 0.68 for **MAP** and 0.65 for **MAP-ED**, and in the leave-one-experiment-out protocol, 0.67 for **MAP** and 0.58 for **MAP-ED**.

depolarized. After that time, the evidence value corresponding to the RV-paced prior becomes smaller than the other two. This gap widens even more between time instances 70 and 100 ms. The evidence value corresponding to LV-paced prior is indistinguishable from that of LVRV-paced prior until around 75 ms, at which most of the epicardial surface has depolarized, and after that becomes smaller.

To compare the evidence statistics with inverse reconstruction results using different priors, we show reconstructed epicardial potential maps in the QRS interval for all three priors using the **MAP** approach. The two panels in Fig. 5 show these reconstructions as well as the original isopotential maps at 42 (top) and 80 (bottom) ms after pacing. At 42 ms, the evidence value for RV-paced prior is smaller than the other two, so we would expect the reconstruction with the RV-paced prior to be worse than the others. At 80 ms, the evidence values from largest to smallest correspond to the LVRV-paced prior, the LV-paced prior and the RV-paced prior. According to our hypothesis, we expect to obtain the best reconstruction among the three priors using the LVRV-paced prior, and the worst using the RV-paced prior. Each panel contains the original isopotential map at the top, and the three reconstructions below. To the right of the original maps is a single time signal from one lead with the current time instant marked for reference. In each panel, the range is fixed for all isopotential maps. In all of these isopotential maps, darker regions represent negative potentials, lighter regions represent positive potentials, and the wavefront lies at the transition from darker to lighter regions.

At 42 ms, reconstruction using the RV-paced prior was indeed significantly worse than for the other two reconstructions. It was noisy and the shape of the wavefront was not accurate.

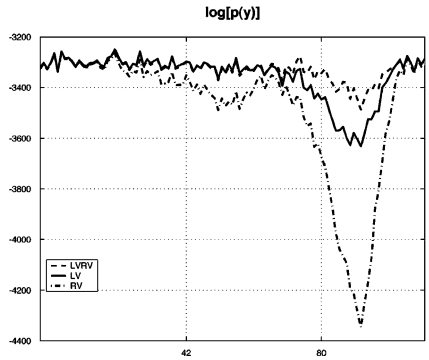

Fig. 4. Evidence values across time for three different prior models.

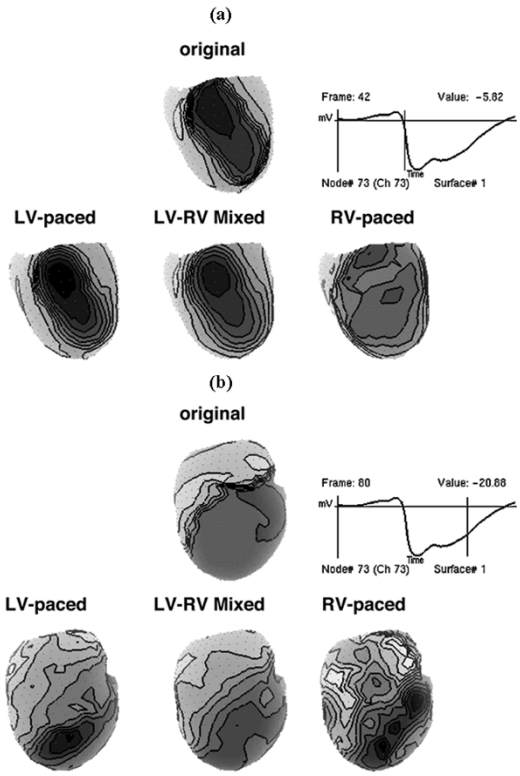

Fig. 5. Isopotential maps at two different time instances, using three different prior densities, compared to the original isopotential maps. (a) At 42 ms and (b) at 80 ms.

On the other hand, reconstructions using the other two priors had better fidelity to the original potential distribution. Between the reconstructions that use LV-paced and LVRV-paced priors, the one using LV-paced prior was slightly better.

At 80 ms, the reconstruction using the RV-paced prior produced many noisy contours, and it was hard to differentiate one single dominant wavefront. The best reconstruction was with the LVRV-paced prior; the wavefront that lay along the 2 o'clock–8 o'clock line in the original isopotential map was better reconstructed than the other two. This finding was also consistent with the evidence values: at this time instant, LVRV-paced prior produced the highest evidence value, and RV-paced produced the lowest evidence value.

V. Discussion and Conclusions

The main motivation of this paper was to advance the utility of Bayesian estimation techniques and performance analysis tools in bioelectric inverse problems, with a specific focus on inverse electrocardiography. We are interested in a Bayesian estimation framework for three main reasons.

1) Previous simulation studies showed that when a "good" prior density was available, its performance was much better than that of standard Tikhonov regularization.
2) The performance analysis tools available to Bayesian estimation methods (e.g., the error covariance matrix) hold promise to be a good tool to evaluate the methods, in advance of actually carrying out any measurements, and even when true solutions are not available.
3) Once one has measurements, one can use statistical information derived from the prior and posterior probability density distribution to design the solution methodology and the experiment. For example, we investigated using the Bayesian evidence to choose a good prior model for the epicardial potentials.

Our study showed that the error covariance was a reasonably reliable qualitative and quantitative predictor of estimation performance despite error in the estimated prior statistical model. When the training data used for the prior estimate closely matched the test beat, confidence intervals were quite accurate. When the training data matched less closely (as here, test and training data came from different animals, and from both *in situ* and tank experiments lumped together), numerical predictions were somewhat less accurate but still reasonable and qualitative results were predictive. Moreover, confidence intervals were equally predictive both with and without the use of sparse measurements (comparing **MAP** and **MAP-ED** methods.) These findings suggest that the use of the sparse measurements does not introduce any further inaccuracy in the statistical modeling, despite the expectation of a more pronounced mismatch between training and test specifications when the blurring effect of the volume conductor is absent, lending impetus to the development of better methods to construct a training data set and to estimate a prior covariance.

The question of appropriate error measures is one which still challenges inverse electrocardiography: even when one has validation data available in experimental scenarios, the standard quantitative measures such as relative error and correlation coefficient are only somewhat representative of the evaluation of results. The results here suggest that the error covariance, although it quantifies *reliability* rather than *accuracy* directly, corresponds well to reconstruction accuracy. In a companion study [28] we illustrate that the error covariance indeed seems to better represent comparative accuracy of significant features in the reconstruction than does relative error.

Moreover, because the error covariance does not depend on the measurements, but rather quantifies the reliability inherent in the experimental design itself, one can quantify expected reliability of an inverse solution, for example in terms of confidence intervals, before taking any measurements; the calculation depends only on the geometry of the problem and the prior information. However, the validity of the error covariance may be suspect due to errors in the prior statistics (and/or the forward model). The statistical model that we used in this paper is approximate since it only accounts for additive noise. The question of whether this approximation is close enough to be useful or whether a more sophisticated statistical model is required is one for future investigation.

In Section IV, we compared the theoretical Bayesian error covariance with the computed error covariance. In the error covariance maps, we did not observe a detailed agreement, only a gross correspondence at some regions of the heart. On the other hand, the confidence intervals were quite accurate, implying that the theoretical values, even when we used a prior model that fits the test data only in a general sense, could be used instead of the computed error covariance values to obtain meaningful results. In order to quantify the correspondence between the two error standard deviation values, we obtained scatter plots. The slopes of the associated regression lines were smaller than one, implying that there seems to be a systematic bias toward under-predicting the error variance (or over-predicting the reliability of the method). This is understandable because of the simplified assumptions we have made on the statistical characteristics of the data. The leave-one-experiment-out method is more biased toward under-prediction than the leave-one-beat-out, and the **MAP-ED** method is more biased than **MAP**.

In addition to quantifying reliability, one could use the error covariance maps to design a measurement scheme. For example, one could choose the torso measurement electrode locations to reduce the overall error variance values; one set of measurement leads might yield lower error variances than another set. Similarly, one could try to locate the secondary measurement sites, if possible, near regions where "poor" estimation were expected.

The observations on the error covariance metric clearly showed the importance of choosing a "good" prior density if we would like to use the theoretical error maps as a reliable evaluation technique, and to increase the reliability of the Bayesian reconstructions. In this paper, we studied a second Bayesian metric, the "evidence," and designed a very simple prior selection scenario, with only three prior models. Even by using this simple simulation, we obtained results that support the hypothesis that the prior model that maximizes the evidence is a good choice of prior. Here, we did not really maximize the evidence, but just labeled the prior among the proposed candidates that yields the largest evidence value as the "best" prior. This idea can easily be extended to choose from a larger number of candidate priors. An alternative use is to design an adaptive algorithm in which one can test whether to accept a training beat into a training set by testing if including that beat increases or decreases the evidence. More research is necessary to evaluate the evidence as a prior selection criterion for inverse bioelectric problems.

When we compared the inverse solutions using LV-paced and LVRV-paced training sets, we observed that the LVRV-paced training set (whose prior density parameters were different from the parameters obtained from the test beat) works better than the LV-paced set (whose prior density parameters were coarsely similar to the parameters obtained from the test beat) even though the test beat was an LV-paced sample. This result was somewhat surprising and should be a topic of further

investigation; we are not at this point certain whether this is an artifact of the simulation (we had data from only two experiments in the training sets) or a more general phenomenon that comes from having a more varied dataset to estimate the prior.

A third Bayesian metric that could be considered as a performance analysis tool in a future study is the differential entropy (or mutual information derived from the differential entropy). The measure of information gain in the system results in a reduction in entropy (or uncertainty), or an increase in the mutual information: the higher the level of the information gained from data, the lower the uncertainty [29]. This idea could be used to evaluate the information content introduced by primary and extra measurements. For example, one could use mutual information to decide whether including a certain type of measurement in addition to primary measurements would add enough new information to justify the effort required to obtain it, or to select the number and/or location of measurement leadsets.

As we have described in Section I, Bayesian techniques were applied previously to inverse EEG/MEG problems with simplified statistical assumptions. In this paper, we attempted to advance that work to include spatially correlated sources. We expect that the same approach can be useful in several related inverse problems. Some suggestions of possible applications might include the following.

- The idea of incorporating various information sources in a Bayesian framework can be extended to other types of inverse problems. For example, endocardial mapping data via noncontact electrodes could be combined with sparse endocardial measurements via contact electrodes. Rao *et al.* presented an inverse solution in [9], in which both endocardial mapping data via noncontact electrodes and sparse endocardial measurements via contact electrodes were available, but the latter were used just for validation. One can imagine combining both types of measurements into a single inverse solution.
- Estimation error covariance could be used to evaluate other types of inverse problems. Russell *et al.* showed the feasibility of this idea when both the sources and the measurement noise are i.i.d. This idea could easily be extended in a fashion similar to our derivation to handle full prior covariance matrices, such as the one used in the inverse EEG/MEG problem of [16].

Future research will include formally treating error in the forward model and the covariance estimate, refining the prior model selection algorithm using evidence to create a training set that is in some sense optimized for the problem, and incorporating temporal correlations.

ACKNOWLEDGMENT

The authors wish to express their gratitude to Dr. B. Taccardi for his assistance in obtaining the experimental data, and B. Yılmaz for his assistance in preparing data for analysis. In addition, Y. S. would like to thank Turkish Higher Education Council for their support of her Graduate Studies.

REFERENCES

[1] R. M. Gulrajani, P. Savard, and F. A. Roberge, "The inverse problem in electrocardiography: Solutions in terms of equivalent sources," in *CRC Crit. Rev. Biomed. Eng.*, vol. 16, 1988, pp. 171–214.

[2] R. S. MacLeod and D. H. Brooks, "Recent progress in inverse problems in electrocardiology," *IEEE Eng. Med. Biol. Mag.*, vol. 17, no. 1, pp. 73–83, Jan./Feb. 1998.

[3] R. L. S. Baillet and J. C. Mosher, "Electromagnetic brain mapping," *IEEE Signal Processing Mag.*, vol. 18, no. 6, pp. 14–30, Nov. 2001.

[4] J. C. Mosher, R. M. Leahy, and P. S. Lewis, "EEG and MEG: Forward solutions for inverse methods," *IEEE Trans. Biomed. Eng.*, vol. 46, no. 3, pp. 245–259, Mar. 1999.

[5] A. N. Tikhonov and V. Y. Arsenin, *Solutions of Ill-Posed Problems*. New York: Halsted, 1977.

[6] A. van Oosterom, "The use of the spatial covariance in computing pericardial potentials," *IEEE Trans. Biomed. Eng.*, vol. 46, no. 7, pp. 778–787, Jul. 1999.

[7] G. S. Russell, R. Srinivasan, and D. M. Tucker, "Bayesian estimates of error bounds for EEG source imaging," *IEEE Trans. Med. Imag.*, vol. 17, no. 6, pp. 1084–1089, Dec. 1998.

[8] B. M. Radich and K. M. Buckley, "EEG dipole localization bounds and MAP algorithms for head models with parameter uncertainties," *IEEE Trans. Biomed. Eng.*, vol. 42, no. 3, pp. 233–241, Mar. 1995.

[9] L. Rao and D. S. Khoury, "System and methods for electrical-anatomical imaging of the heart," in *Proc. IEEE 2nd Joint EMBS-BMES Conf.*, Houston, TX, 2002.

[10] R. O. Kuenzler, R. S. MacLeod, B. Taccardi, Q. Ni, and R. L. Lux, "Estimation of epicardial activation maps from intravascular recordings," *J. Electrocardiol.*, vol. 32, no. 2, pp. 77–92, 1999.

[11] R. O. Martin, T. C. Pilkington, and M. N. Morrow, "Statistically constrained inverse electrocardiography," *IEEE Trans. Biomed. Eng.*, vol. BME-22, no. 11, pp. 487–492, Nov. 1975.

[12] R. C. Barr and M. S. Spach, "Inverse calculation of QRS-T epicardial potentials from body surface potential distributions for normal and ectopic beats in the intact dog," *Circ. Res.*, vol. 42, pp. 661–675, 1978.

[13] A. van Oosterom, "The spatial covariance used in computing the pericardial potential distribution," in *Computational Inverse Problems in Electrocardiography*, P. R. Johnston, Ed. Southampton, U.K.: WIT, 2001, pp. 1–50.

[14] F. Greensite, "A new treatment of the inverse problem of multivariate analysis," *Inv. Probl.*, vol. 18, pp. 363–379, 2002.

[15] ——, "The temporal prior in bioelectromagnetic source imaging problems," *IEEE Trans. Biomed. Eng.*, vol. 50, no. 9, pp. 1152–1159, Oct. 2003.

[16] J. W. Philips, R. M. Leahy, and J. C. Mosher, "MEG-based imaging of focal neuronal current sources," *IEEE Trans. Med. Imag.*, vol. 16, no. 3, pp. 338–348, Jun. 1997.

[17] S. Baillet and L. Garnero, "A Bayesian approach to introducing anatomo-functional priors in the EEG/MEG inverse problem," *IEEE Trans. Biomed. Eng.*, vol. 44, no. 5, pp. 374–385, May 1997.

[18] D. M. Schmidt, J. S. George, and C. C. Wood, "Bayesian inference applied to the electromagnetic inverse problem," *Hum. Brain Map.*, vol. 7, pp. 195–212, 1999.

[19] H. L. V. Trees, *Detection, Estimation, and Modulation Theory, Part I Detection Estimation and Linear Modulation Theory*. New York: Wiley, 1968.

[20] J. C. Mosher, M. E. Spencer, R. M. Leahy, and P. S. Lewis, "Error bounds for EEG and MEG dipole source localization," *Electroensephal. Clin. Neurophysiol.*, vol. 86, pp. 303–321, 1993.

[21] J. Malmivuo, V. Suihko, and H. Eskola, "Sensitivity distributions of EEG and MEG measurements," *IEEE Trans. Biomed. Eng.*, vol. 44, no. 3, pp. 196–208, Mar. 1997.

[22] C. H. Muravchik and A. Nehorai, "EEG/MEG error bounds for a static dipole source with a realistic head model," *IEEE Trans. Signal Processing*, vol. 49, no. 3, pp. 470–484, Mar. 2001.

[23] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Englewood Cliffs, NJ: Prentice-Hall, 1993, ch. 10–12.

[24] R. MacLeod, Q. Ni, B. Punske, P. Ershler, B. Yilmaz, and B. Taccardi, "Effects of heart position on the body-surface ECG," *J. Electrocardiol.*, vol. 33, pp. 229–238, 2000.

[25] R. MacLeod, R. Lux, and B. Taccardi, "A possible mechanism for electrocardiographically silent changes in cardiac repolarization," *J. Electrocardiol.*, vol. 30 (suppl.), pp. 114–121, 1998.

[26] R. S. MacLeod and C. R. Johnson, "Map3d: Interactive scientific visualization for bioengineering data," in *Proc. 15th Annu. Int. Conf. IEEE Engineering Medicine and Biology Soc.*, 1993, pp. 30–31.
[27] Y. Serinagaoglu, R. S. MacLeod, B. Yilmaz, and D. H. Brooks, "Multielectrode venous catheter mapping as a high quality constraint for electrocardiographic inverse solution," *J. Electrocardiol.*, vol. 35 (suppl.), pp. 65–73, 2002.
[28] Y. Serinagaoglu, D. H. Brooks, and R. S. MacLeod, "Improved performance of Bayesian solutions for inverse electrocardiography using multiple information sources," , to be published.
[29] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: Wiley, 1991, ch. 9.

Yeşim Serinağaoğlu (S'93–M'03) received the B.S. degree in 1995 and the M.S. degree in 1997 in electrical and electronics engineering from Middle East Technical University, Ankara, Turkey, and Ph.D. degree in 2003 from Northeastern University, Boston, MA.

She is currently an Assistant Professor of electrical and electronics engineering at Middle East Technical University. She is also a Member of the Institute of Applied Mathematics at the same university. Her research interests are in statistical and digital signal processing, with particular application to biomedical problems. Specific current interests include forward and inverse problems in electrocardiography, modeling the electrical activity of the heart, and electrocardiography signal processing.

Dana H. Brooks (S'86-M'91) received the B.A. degree in English in 1972 from Temple University, Philadelphia, PA, and the B.S.E.E., M.S.E.E., and Ph.D. degrees in electrical engineering in 1986, 1988, and 1991, respectively, from Northeastern University, Boston, MA.

He is an Associate Professor of electrical and computer engineering, Associate Director of the Center for Communications and Digital Signal Processing, a Member of the Institute for Complex Software Systems and the Center for Subsurface Sensing and Imaging Systems, and PI of the BioMedical Imaging and Signal Processing Laboratory, all at Northeastern University, and a Member of the Center for Bioelectric Field Modeling, Simulation, and Visualization headquartered at the University of Utah, Salt Lake City. He was a Visiting Professor during 1999–2000 at the Universitat Polit'ecnica de Catalunya, Barcelona, Spain. His research interests lie in application of statistical and digital signal and image processing to biomedical signal processing and medical and biological imaging, and in open-source software systems for these applications.

Robert S. MacLeod (S'87-M'87) received the B.Sc. degree in 1979 in engineering physics and the Ph.D. degree in 1990 in physiology and biophysics from Dalhousie University, Halifax, NS, Canada. He received the M.Sc. degree in electrical engineering from the Technische Universität, Graz, Austria, in 1985.

He is an Associate Professor in the Bioengineering Department and the Department of Internal Medicine (Division of Cardiology) at the University of Utah, Salt Lake City, and an Associate Director of the Scientific Computing and Imaging Institute and the Nora Eccles Harrison Cardiovascular Research and Training Institute. His research interests include computational electrocardiography (forward and inverse problems), experimental investigation and clinical detection of cardiac ischemia and repolarization abnormalities, and scientific computing and visualization.