A Bayesian Approach to Quantifying Uncertainty in Tikhonov Solutions for the Inverse Problem of Electrocardiography

Jessie J France, Yaniv Gur, Robert M Kirby and Chris R Johnson

Abstract

This study aimed to quantify differences in uncertainty in Tikhonov solutions arising from mesh discretization, conductivity, and zeroth, first, and second order Tikhonov (ZOT, FOT, and SOT) solutions for the inverse problem of electrocardiography. We indirectly analyzed levels of uncertainty in Tikhonov solutions through deriving their equivalent Bayesian maximum a posteriori (MAP) estimates, and then performing regularized sampling from the Bayesian posterior distributions to form credible intervals (CIs). We calculated the percentage of the true heart voltages that fell between the 95% CIs. For all noise levels, the 95% mean CIs for FOT and SOT always captured 11% to 42% more of the true heart voltages than ZOT, suggesting that regularization with FOT and SOT may provide a greater level of certainty in reconstructing heart voltages. In summary, we provide a methodology for quantifying uncertainty in Tikhonov solutions, and use it to study different regularization techniques.

Keywords—**electrocardiography; inverse problem; Tikhonov regularization, Bayesian uncertainty quantification, sampling**

1. Introduction

The inverse problem of electrocardiography aims to reconstruct heart voltages from torso surface recordings, and holds the potential to transition the traditional electrocardiogram (ECG) into a new, patient-specific imaging modality [1]. Given torso surface ECG voltage recordings $\mathbf{t}$, and a *transfer* or *lead field* matrix K, the epicardial potential-based inverse problem of electrocardiography seeks to reconstruct the heart surface voltages $\mathbf{h}$ from the linear system $K\mathbf{h} = \mathbf{t}$, where $K \in \mathbb{R}^{m \times n}$, $\mathbf{h} \in \mathbb{R}^n$, and $\mathbf{t} \in \mathbb{R}^m$. In practice $\mathbf{t}$ contains the true torso voltages $\mathbf{t}^*$ plus noise $\mathbf{e}$, or $\mathbf{t} = \mathbf{t}^* + \mathbf{e}$.

Despite its great potential, the inverse ECG problem is severely ill-posed, meaning small changes in torso surface ECG recordings lead to correspondingly large changes in simulated heart surface voltages. One popular method for stabilizing the reconstruction involves Tikhonov regularization, where the solution to the original equation $K\mathbf{h} = \mathbf{t}$, is found by solving a regularized least squares optimization problem

$$\mathbf{h} = \operatorname{argmin} \left\{ \left|\left| K\mathbf{h} - \mathbf{t} \right|\right|_2^2 + \lambda^2 \left|\left| L\mathbf{h} \right|\right|_2^2 \right\}. \qquad (1)$$

In the above minimization problem, additional information is introduced into the original inverse problem in the form of a positive parameter λ and a matrix L, and this new modified problem is solved for an optimal Tikhonov solution [2]. For a fixed λ, the Tikhonov solution solves the system $A\mathbf{h} = \mathbf{b}$, where $A = \begin{bmatrix} K \\ \lambda L \end{bmatrix}$, and $\mathbf{b} = \begin{bmatrix} \mathbf{t} \\ \mathbf{0} \end{bmatrix}$. In zeroth-order Tikhonov (ZOT), the matrix L is the identity matrix, while in first and second order Tikhonov (FOT and SOT), L approximates the first and second derivative, respectively [2]. However, while a Tikhonov solution provides us with a single deterministic solution, a Bayesian approach can not only provide us with an equivalent maximum a *posteriori* (MAP) estimate, but also an entire probability distribution that we can evaluate to study sources of uncertainty. The level of noise in the ECG recordings, the choice of L, and parameters altering the formation of the lead field matrix K (such as the model's governing equations, mesh discretization [3], and conductivity), all present sources of uncertainty for the inverse problem of electrocardiography.

In this study, we indirectly analyze the level of credibility in Tikhonov solutions by deriving equivalent probabilistic Bayesian MAP estimates, and then sample from the Bayesian posterior distributions to form credible intervals (CIs). Using this approach, we use the concept of functional band depth [4] to compare differences in how much of the true solution the 95% CIs capture for different choices of L regularizers as well as hybrid and non-hybrid meshes for homogeneous and inhomogeneous conductivity. In our hierarchical Bayesian formulation, we show how the noise variance in torso surface voltage recordings can be estimated (if it is unknown) for a given Tikhonov solution, leading to the efficient implementation of a block-wise Gibbs sampler. Furthermore, for severely ill-conditioned transfer matrices, we introduce the concept of regularized sampling. Our preliminary results show that

France, Gur, Kirby, and Johnson are with the Scientific Computing and Imaging Institute, University of Utah. Email: {jfrance,yanivg,kirby,crj}@sci.utah.edu. This project was supported by grants from the National Institute of General Medical Sciences (8 P41 GM103545-15) from the National Institutes of Health.

our algorithm accurately predicted the true noise variance 82% of the time to within 8% absolute error, giving some initial verification for our method. Additionally, we found that the 95% CIs for FOT and SOT captured between 11% and 42% more of the true heart voltages than ZOT, suggesting FOT and SOT may help reduce uncertainty for the inverse ECG problem. In summary, we propose a new method that quantifies the level of uncertainty via CIs in Tikhonov solutions, and use it to compare different regularizers and meshes as well.

2. Methods

We describe how we obtained a Tikhonov solution, our Bayesian approach to uncertainty quantification for this Tikhonov solution, and our numerical simulations.

2.1. Obtaining a Tikhonov Solution

Based on earlier inverse ECG studies [3], we used a quasi-static approximation of Maxwell's equations to model the voltage potential u in the torso using Laplace's equation, $\nabla \cdot (\sigma(\mathbf{x})\nabla u(\mathbf{x})) = 0$, where σ denotes a conductivity tensor. We assume no current leaves the torso surface into the air, and consider the heart surface as a Dirichlet boundary. We discretized this model using finite elements to form the equation $K\mathbf{h} = \mathbf{t}$ [3]. We then solved the inverse problem via Tikhonov regularization as in Equation (1), using the L-curve method [5] to find the optimal λ.

2.2. Bayesian Approach

In the Bayesian setting, all unknown parameters are modeled as random variables [2,6]. Because the noise $\mathbf{e}$ in the torso surface ECG recordings is unknown, we model the noise as being identical and independently distributed (IID) with zero mean and variance v_e, $\mathbf{E} \sim \mathcal{N}(\mathbf{0}, v_e I)$, to form the *likelihood* density

$$\pi(\mathbf{t}|\mathbf{h}, v_e) \propto v_e^{-m/2} exp\left(-\frac{1}{2v_e}||K\mathbf{h} - \mathbf{t}||_2^2\right). \quad (2)$$

In the Bayesian formulation, the prior acts as the regularizer, and introduces outside information in the same way that λ and L introduce outside information in Tikhonov regularization. For example, in first order Tikhonov (FOT) regularization, L approximates the first spatial derivative, and choosing this L assumes that the epicardial voltage distribution should be relatively "flat" [2]. An analogous probabilistic version of this assumption may state that the difference in voltage between two neighboring positions on the epicardial surface in two dimensions (2D), H_i and H_{i-1}, is a random variable Q_i, $H_i - H_{i-1} = Q_i$, where $Q_i \sim \mathcal{N}(0, v_h)$ for all i [6]. Likewise, if our prior belief is that epicardial heart voltages should be "smooth," we would use an L in Tikhonov regularization that approximates the second spatial derivative [2], resulting in second order Tikhonov (SOT). In the Bayesian sense, we can model this as a probabilistic 2D Laplacian stencil as $H_{i+1} - 2H_i + H_{i-1} = Q_i$.

For both of the aforementioned 2D stencils, and for any arbitrary choice of L, we may generalize the results as a matrix system, where $\mathbf{Z} = L\mathbf{H} = \mathbf{Q}$. Consequently, $\mathbf{Q}$ and $\mathbf{Z}$ are both IID with similar multivariate Gaussian distributions [6], and we can therefore write

$$\pi(\mathbf{z} = L\mathbf{h}|v_h) \propto v_h^{-n/2} exp\left(-\frac{1}{2v_h}||L\mathbf{h}||_2^2\right). \quad (3)$$

Using Bayes' rule, $\pi(\mathbf{z}|\mathbf{t}) \propto \pi(\mathbf{t}|\mathbf{z})\pi(\mathbf{z})$, we can easily combine Equations (2) and (3) to get

$$\pi(\mathbf{z}|\mathbf{t}) \propto exp\left(-\frac{1}{2v_e}||K\mathbf{h} - \mathbf{t}||_2^2 - \frac{1}{2v_h}||L\mathbf{h}||_2^2\right). \quad (4)$$

Thus, to obtain a MAP estimate for $\mathbf{h}$ equal to the deterministic Tikhonov solution, we can see this will only occur in Equation (4) when we factor out a v_e term, and observe that $\lambda^2 = \frac{v_e}{v_h}$. We thus rewrite Equation (4) as

$$\pi(\mathbf{z}|\mathbf{t}) \propto exp\left(-\frac{1}{2v_e}\left(||K\mathbf{h} - \mathbf{t}||_2^2 + \lambda^2||L\mathbf{h}||_2^2\right)\right). \quad (5)$$

At this point, if the variance v_e in the noise is known, we can sample from Equation (5) for $\mathbf{h}$ as described in section 2.2.1. However, if the variance v_e is unknown for a multivariate normal distribution, then it can be modeled as an inverse gamma (IG) distribution with a shape parameter α and scale parameter v_0, as $v_e \sim IG(\alpha, v_0)$ [6], or

$$\pi(v_e; \alpha, v_0) \propto v_e^{-\alpha-1} exp\left(-\frac{v_0}{v_e}\right). \quad (6)$$

In this case, the joint PDF $\pi(\mathbf{z}, \mathbf{t}, v_e)$ can be written as the product of Equations (2), (3), and (6), with the relationship that $v_h = \frac{v_e}{\lambda^2}$. This gives

$$\pi(\mathbf{h}, \mathbf{t}, v_e) \propto \quad (7)$$
$$exp\left(-\frac{1}{2v_e}\left|\left|A\mathbf{h} - \mathbf{b}\right|\right|_2^2 - (\tilde{\alpha} + 1)\ln(v_e) - \frac{v_0}{v_e}\right)$$

where we define $\tilde{\alpha} = \frac{m+n}{2} + \alpha$, and with A and $\mathbf{b}$ defined previously in reference to Tikhonov regularization in section 1. In Equation (7), we make the key assumption that λ has already been found in obtaining an optimal Tikhonov solution, and is therefore a known variable.

From here, if we are able to "observe" the torso potentials $\mathbf{t}$, and the variance v_e, then the conditional density for $\mathbf{h}$ is given as in Equation (5). Likewise, if we are then able to "observe" $\mathbf{t}$ and $\mathbf{h}$, then the posterior variance $v_e|\mathbf{t}, \mathbf{h}$ can be written as an IG distribution with new shape and scale parameters $\tilde{\alpha}$ and $\tilde{v_0}$ as $v_e|\mathbf{h}, \mathbf{t} \sim IG(\tilde{\alpha}, \tilde{v_0})$:

$$\pi(v_e|\mathbf{h}, \mathbf{t}) \propto v_e^{-\tilde{\alpha}-1} exp\left(-\frac{\tilde{v_0}}{v_e}\right), \quad (8)$$

where $\tilde{v_0} = \frac{1}{2}||A\mathbf{h} - \mathbf{b}||_2^2 + v_0$.

Although we can determine a MAP estimate for v_e based on the properties of the IG distribution, this MAP estimate depends heavily on the original hyperparameters α and v_0 in Equation (6), and we therefore describe a novel method for setting these parameters when a good estimate of the variance v_e is unknown. First, let $\hat{v_e}$ and $\hat{\mu}$ represent the unbiased estimators for the variance and mean of the noise $\mathbf{e}$. The unbiased variance can be written as

$$\hat{v_e} = \left(\frac{1}{m-1}\right)\sum_{i=1}^{m}(e_i - \hat{\mu})^2. \quad (9)$$

Since we previously assumed that the noise had zero mean, the unbiased mean estimator should also be approximately zero, $\hat{\mu} \approx 0$. Therefore we can approximate the unbiased variance as

$$\hat{v_e} \approx \left(\frac{1}{m-1}\right)||\mathbf{e}||_2^2. \quad (10)$$

In choosing the optimal Tikhonov solution $\mathbf{h_\lambda}$, the Morosov discrepancy principle selects this solution such that $||K\mathbf{h_\lambda} - \mathbf{t}||_2 \approx ||\mathbf{e}||_2$ [7]. If we are confident in our algorithm for choosing the optimal Tikhonov solution via the L-curve or some other method (besides the discrepancy principle), we may apply Morosov's discrepancy principle in reverse here and Equation (10) becomes

$$\hat{v_e} \approx \left(\frac{1}{m-1}\right)||K\mathbf{h_\lambda} - \mathbf{t}||_2^2. \quad (11)$$

Next, we can incorporate this result as prior information by setting the mode of the conditional variance, $v_e|\mathbf{t}, \mathbf{h}$ equal to the result in Equation (11). With the mode of the conditional variance given as $\frac{\tilde{v_0}}{\tilde{\alpha}+1}$, we set this equal to the unbiased estimate $\hat{v_e}$ in our model:

$$\frac{\tilde{v_0}}{\tilde{\alpha}+1} = \left(\frac{1}{m-1}\right)\left|\left|K\mathbf{h_\lambda} - \mathbf{t}\right|\right|_2^2, \quad (12)$$

With $\tilde{v_0}$ and $\tilde{\alpha}$ previously defined, we can solve Equation (12) for the original scale parameter v_0, and choose the original shape parameter α such that v_0 is greater than zero. In our simulations, we chose $\alpha = 20$.

2.2.1. Regularized Sampling

After finding a Tikhonov solution, Figure 1 illustrates the Gibbs sampling algorithm to generate an ensemble of heart voltages $\mathbf{h}$ and noise variance v_e (denoted as v in Figure 1) realizations which we used to calculate credible intervals (CIs). First, since the Bayesian MAP estimate equals the Tikhonov solution in our derivation, we consider the Tikhonov solution as the first sample $\mathbf{h^{(0)}}$. Next, if the variance is unknown, we determine a MAP estimate for the noise variance by setting the hyperparameters v_0 and α as in Equation (12), and set this MAP estimate equal

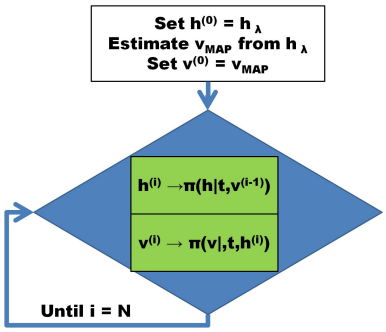

Figure 1. Overview of block-wise Gibbs sampler.

to our zeroth variance sample $v_e^{(0)}$. Then, we form the heart of the block-wise Gibbs sampling algorithm by iteratively sampling from the conditional densities given in Equations (4) and (8) until we have collected N total samples. Once the algorithm is complete, we sorted the samples into percentiles to determine the 95% CIs. Note that in our algorithm, we dealt with the case where we assumed the noise variance is unknown. If it is known, we could easily forgo the block-wise Gibbs sampling algorithm, and instead sample iteratively (or in parallel) from Equation (4) for heart voltage realizations as described next.

As described in Kapio [7], we may sample from the multivariate Gaussian distribution in Equation (4) as:

1. Sample $\boldsymbol{\xi} \sim N(\mathbf{0}, \sqrt{v}I)$,
2. Set $\mathbf{h} = A^\dagger(\mathbf{b} + \boldsymbol{\xi})$,

where, in the above two steps, we simplify our notation for the i^{th} sample to write $\mathbf{h}$ for $h^{(i)}$ and v for $v^{(i-1)}$, and $\dagger$ denotes the pseudoinverse.

However, when the K that comprises the matrix A is severely ill-conditioned (K had a condition number of 10^{17} or more in our simulations), Step 2 above may require special care, as the λ in Tikhonov regularization may not be enough to produce stable results. Step 2 is essentially a new inverse problem, similar to our original inverse problem $K\mathbf{h} = \mathbf{t}$, except that K, and $\mathbf{t}$ are now replaced with A and the value $\mathbf{b} + \boldsymbol{\xi}$. Furthermore, if $\boldsymbol{\xi} = \mathbf{0}$, Step 2 will produce a result equal to the Tikhonov solution. Thus while the goal in the original inverse problem was to find the value $\mathbf{h}$ closest to the *true solution*, an analogous goal that we propose for the sampling in Step 2 is to find the value $\mathbf{h}$ closest to the Tikhonov solution $\mathbf{h_\lambda}$, which is presumably known. To achieve this, one could easily perform singular value decomposition (SVD) on A, and determine the sample from the k largest singular values that minimize the difference between $\mathbf{h_k}$ and $\mathbf{h_\lambda}$ as $k = min\{||\mathbf{h_\lambda} - \sum_{j=1}^{k}\frac{\mathbf{u_j^T b}}{\sigma_j}\mathbf{v_j}||_2\}$, and $\mathbf{h} = \mathbf{h_k}$, with $k \leq n$. However, performing the SVD may be infeasible for large scale problems, and in this paper we therefore used the conjugate gradient-least squares (CGLS) algorithm, and used the properties of semi-convergence to terminate at the iteration k that minimized the difference between $\mathbf{h_k}$ and $\mathbf{h_\lambda}$. While we used the CGLS algorithm

in this paper, we note that other Krylov subspace algorithms, such as GMRES, might also be feasible.

2.3. Numerical Simulations

We ran 50 simulations at three different noise levels (at $\frac{||\mathbf{e}||_2}{||\mathbf{t}^*||_2}$ = 1%, 3%, and 5%). For each individual simulation, we generated a total of 4,000 heart voltage and noise variance samples, and calculated the 95% credible intervals (CIs) for heart voltage. Using the concept of functional band depth [4], we calculated the percentage of the true heart voltage that fell within the 95% CIs for different regularizers (ZOT, FOT, and SOT), hybrid and non-hybrid meshes (Figure 2), and conductivity combinations to give a total of $3 \times 2 \times 2 = 12$ different experimental combinations for a single noise level, $\frac{||\mathbf{e}||_2}{||\mathbf{t}^*||_2}$.

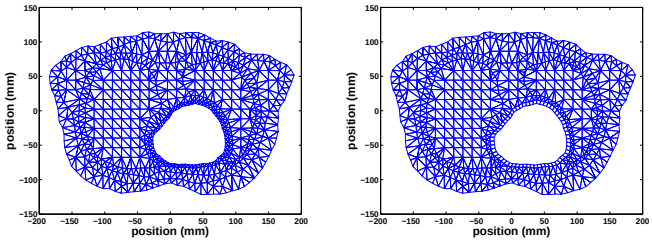

Figure 2. Non-hybrid (triangle only) and hybrid (triangle-quadrilateral) meshes.

3. Results and Conclusion

Figure 3 shows the results of individual simulations for zeroth and second order Tikhonov (ZOT and SOT) (top left and top right), and the mean percentage of the true heart voltages that fell within the 95% CIs at $\frac{||\mathbf{e}||_2}{||\mathbf{t}^*||_2} = 3\%$ with homogeneous (bottom left) and inhomogeneous (bottom right) conductivity. While mesh discretization and conductivity had some effect on the results, the choice of L in Tikhonov regularization had the greatest impact on how well the CIs captured the true solution at all noise levels. Across all noise levels, conductivities, and mesh discretizations, the mean 95% CIs from first order Tikhonov (FOT) and SOT always captured 11% to 42% more of the true heart voltages than ZOT. Out of all simulations combined, SOT on a hybrid mesh captured more than the other regularizer and mesh combination 57% of the time. Furthermore, our method for estimating the noise variance in block-wise Gibbs sampling accurately predicted the true noise variance 82% of the time to within 8% absolute error, demonstrating our method produces meaningful results. In summary, our results suggest a method for quantifying uncertainty in Tikhonov solutions, and also suggest that FOT and SOT may provide more accurate credible intervals than ZOT.

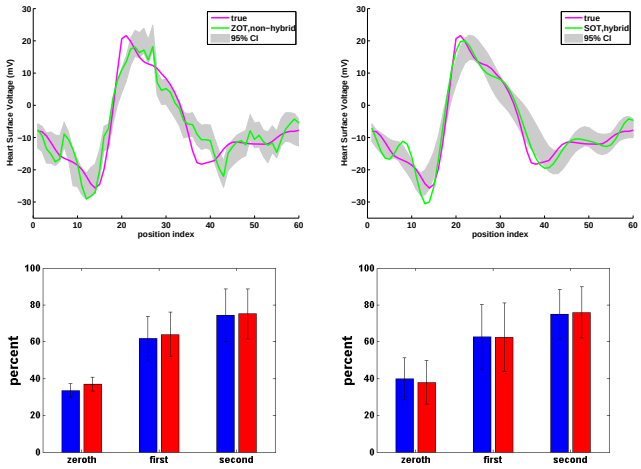

Figure 3. Example and summary results. Top left and top right figures compare the zeroth order Tikhonov (ZOT) reconstruction on a non-hybrid mesh with the second order Tikhonov (SOT) reconstruction on a hybrid mesh in reference to the true voltages. Grey areas show the 95% credible intervals (CIs) in both figures. The bottom left and bottom right figures compare the mean percentage (of 50 simulations) of the true solution that falls within the 95% CIs for zeroth, first, and second order Tikhonov on non-hybrid (blue) and hybrid (red) meshes. Simulations in the bottom left figure were ran with homogeneous conductivity, while the bottom right figure used inhomogeneous conductivity values.

References

[1] van der Graaf AW, Bhagirath P, Ramanna H, van Driel VJ, de Hooge J, de Groot NM, Götte MJ. Noninvasive imaging of cardiac excitation: current status and future perspective. Ann Noninvasive Electrocardiol 2014;19(2):105–13.

[2] Aster RC, Borchers B, Thurber CH. Parameter Estimation and Inverse Problems. New York: Elsevier, 2005.

[3] Wang D, Kirby RM, Johnson CR. Resolution strategies for the finite-element-based solution of the ECG inverse problem. IEEE Trans Biomed Eng 2010;57(2):220–37.

[4] López-Pintado S, Jornsten R. Functional analysis via extensions of the band depth, volume 54 of Lecture Notes–Monograph Series. Beachwood, Ohio, USA: Institute of Mathematical Statistics, 2007; 103–120.

[5] Hansen P, Jensen T, Rodriguez G. An adaptive pruning algorithm for the discrete L-curve criterion. Journal of Computational and Applied Mathematics 2007;198(2):483 – 492.

[6] Calvetti D, Somersalo E. Hypermodels in the Bayesian imaging framework. Inverse Problems 2008;24(3):034013.

[7] Kaipio J, Somersalo E. Statistical and Computational Inverse Problems. New York,N.Y.: Springer, 2004.