
Sequential Estimation of Discretization Errors in Inverse Problems

Oliver Ernst



Cleveland, 2014.

Sequential Estimation of Discretization Errors in Inverse Problems

Oliver K. Ernst

Submitted in partial fulfillment of the requirements of a
Bachelor of Science in Physics with a Mathematical Physics Concentration

Department of Physics
Case Western Reserve University
2076 Adelbert Road, Cleveland, OH 44106

2 May 2014.

This thesis was carried out under the supervision of
Daniela Calvetti and Erkki Somersalo
Department of Mathematics, Applied Mathematics and Statistics
Case Western Reserve University

Committee Members:
Kenneth Kowalski
Harsh Mathur

Abstract

Inverse problems are by nature computationally intensive, and a key challenge in practical applications is to reduce the computing time without sacrificing the accuracy of the solution. When using Finite Element (FE) or Finite Difference (FD) methods, the computational burden of using a fine discretization is often unrealizable in practical applications. Conversely, coarse discretizations introduce a modeling error, which may become the predominant part of the noise, particularly when the data is collected with high accuracy. In the Bayesian framework for solving inverse problems, it is possible to attain a super-resolution in a coarse discretization by treating the modeling error as an unknown that is estimated as part of the inverse problem. It has been previously proposed to estimate the probability density of the modeling error in an off-line calculation that is performed prior to solving the inverse problem. In this thesis, a dynamic method to obtain these estimates is proposed, derived from the idea of Ensemble Kalman Filtering (EnKF). The modeling error statistics are updated sequentially based on the current ensemble estimate of the unknown quantity, and concomitantly these estimates update the likelihood function, reducing the variance of the posterior distribution. A small ensemble size allows for rapid convergence of the estimates, and the need for any prior calculations is eliminated. The viability of the method is demonstrated in application to Electrical Impedance Tomography (EIT), an imaging method that measures spatial variations in the conductivity distribution inside a body.

“Thinking is more interesting than knowing, but less interesting than looking.”
“Denken ist interessanter als Wissen, aber nicht als Anschauen.”
-Johann Wolfgang von Goethe

Acknowledgements

I wish to thank my advisors Prof. Erkki Somersalo and Prof. Daniela Calvetti for their guidance and for countless enlightening discussions. Gauß wrote: “I have had my results for a long time: but I do not yet know how I am to arrive at them” - I am grateful that they always took the time to explain what I was looking at, and why.

I also wish to express my gratitude to Prof. Kenneth Kowalsi and Prof. Kenneth Singer for organizing the senior project class, and furthermore for the countless hours spent lecturing, viewing presentations, and reading reports, without which this work would not have been possible.

Last but not least, I wish to thank my parents, for everything.

Contents

1	Introduction	1
2	Inverse Problems in the Bayesian Framework	3
2.1	Tikhonov Regularization	3
2.2	Bayesian Inverse Problems	6
2.3	Discretization Error	7
3	Estimation of the Modeling Error	10
3.1	Off-Line Modeling Error Approximation	10
3.2	Dynamic Modeling Error Approximation	11
4	Application to EIT	14
4.1	Forward Model	14
4.1.1	Boundary Conditions	16
4.2	Discretized Model	17
5	Computed Examples in EIT	22
5.1	Prior Density	22
5.2	Conductivity Reconstructions	22
5.3	Convergence and Discussion	24
6	Conclusions	30
	Appendices	32
A	Derivations of EIT Equation from Maxwells Equations	33
B	Gauß-Newton Algorithm	37
B.1	Wiener Filtering	39
C	Implementation details in EIDORS	40

Chapter 1

Introduction

The tradeoff between computational efficiency and the accuracy of the solution is a key challenge in solving inverse problems. In practical applications, the discretization of the inverse problem, for example using FE or FD methods, leads to a difference between the model and the reality. Using a finer discretization decreases this discrepancy, but may increase the computational burden beyond what is feasible. In particular, when the forward map is non-linear, a coarse discretization is used that differs significantly from the model used to produce the data, synthetic or experimental.

The insufficiency of a model to describe reality is referred to as model discrepancy, and accounting for it is a topic of active study [14, 18]. The model discrepancy due to numerical discretization error has been addressed in the literature in the context of Bayesian inverse problems [1, 17], together with other types of model mismatches such as poorly known geometry [21, 22], boundary clutter in the data [5], or insufficient modeling of the underlying physics [12, 13, 27]. Particularly when experimental measurements are precise, the modeling error may predominate the noise, and will significantly decrease the quality of the solution if not compensated for.

This thesis addresses the problem of accounting for the discretization error in the Bayesian framework for solving inverse problems. The error is treated as a realization of a random variable with some unknown probability density. Since the actual value of the modeling error depends on the unknown parameter of interest, the estimates of the error depend on the information available about the unknown. Approximating the probability density of the modeling error therefore becomes part of the inverse problem.

Previously, the error's probability density has been estimated in an off-line calculation that is performed prior to solving the inverse problem. A sample of draws from the prior density is generated, and the modeling error is estimated by computing the model predictions from two meshes, one representing an accurate approximation of the continuous model, the other representing the discretized mesh used in the inverse solver. Using the same mesh to both generate the data and solve the inverse problem is known as an "inverse crime", and leads to results that are overly optimistic.

In this thesis, we investigate the possibility of estimating the modeling error in a dynamic fashion, by updating it while we iteratively improve the solution of the inverse problem. In particular, unlike in the earlier works, this estimate is not based on draws from the prior but rather on forward simulations using the current ensemble of estimates, analogous to the classical estimation of the error of numerical solvers for differential equations. The proposed

algorithm is a modification of an EnKF algorithm, with an artificial time parameter that refers to the updating round. A similar idea of interpreting iterations as an artificial time step in the EnKF framework for time-invariant inverse problems was proposed recently in [15].

The method is applied to the discrete ill-posed inverse problem in EIT, an imaging method which measures spatial variations in the electric conductivity distribution inside a body. Currents are applied through electrodes on the surface of the body, and resulting relative voltages are measured. Benefits of the method include its non-invasive nature and rapid measurement time, with applications ranging from biomedical imaging to industrial process tomography. The partial differential equation model that describes the forward model in EIT is assumed to provide a good approximation of the physical measurement setting, while the discretization of the equation generates an approximation error.

The thesis is divided into six chapters. Following the Introduction in Chapter 1, Chapter 2 presents an introduction to solving inverse problems in the Bayesian framework, and the connection of this method to Tikhonov regularization. The origin of the discretization error follows and its connection to the inverse problem solution is made explicit. Chapter 3 reviews the previously proposed off-line algorithm for estimating the modeling errors probability density, and presents the dynamic method that is the subject of this thesis. Chapter 4 derives the partial differential equation model that describes EIT, and discretizes the model using the FE method. Chapter 5 presents a number of computed examples that demonstrate the convergence and efficiency of the method in comparison to the off-line calculation. Finally, Chapter 6 discusses the significance of the work presented in this thesis, and suggestions for future research directions are given.

Chapter 2

Inverse Problems in the Bayesian Framework

In this chapter, regularization methods for ill-posed inverse problems are reviewed [11], followed by regularization in the Bayesian framework [16]. Finally, the origin of discretization errors is discussed [1, 16, 17], arriving at a formulation of the forward model.

2.1 Tikhonov Regularization

We begin by considering the discrete inverse problem of recovering an unknown x from a linear measurement model

$$b = Ax + e, \quad x \in \mathbb{R}^n, \quad b, e \in \mathbb{R}^m, \quad A \in \mathbb{R}^{m \times n}, \quad (2.1)$$

where b is a given data vector, and the additive measurement noise e is some zero mean Gaussian.

We consider the general case where $m \neq n$, and furthermore where A is not invertible. A first approach to solve (2.1) in this case is to find the solution that minimizes the residual,

$$x = \arg \min_{x'} \{ \|b - Ax'\|_2^2 \}, \quad (2.2)$$

where $\|\cdot\|_2$ denotes the two norm.

However, the solution with the minimum residual is in general not the solution that most closely resembles the true unknown. To see this, consider the singular value decomposition (SVD) of the matrix A ,

$$A = U\Sigma V^T, \quad U \in \mathbb{R}^{m \times m}, \quad \Sigma \in \mathbb{R}^{m \times n}, \quad V \in \mathbb{R}^{n \times n}, \quad (2.3)$$

where the matrices satisfy the following properties

$$\begin{aligned} UU^T &= U^T U = I_m, \\ VV^T &= V^T V = I_n, \\ \Sigma &= \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_r), \end{aligned}$$

where \mathbf{I}_m is the unit matrix of dimension $m \times m$, r is the rank of A , and $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r \geq 0$ are the *singular values* of A .

Using the SVD, the solution to (2.1) is

$$x_{\text{SVD}} = \mathbf{A}^{-1}b = \mathbf{V}\Sigma^{-1}\mathbf{U}^T b = \sum_{i=1}^r \frac{(u_i^T b)}{\sigma_i} v_i, \quad (2.4)$$

where u_i and v_i are the i th columns of U and V , respectively. In practice, the noiseless data vector $b_0 = b - e$ is not known. However, if it were, the effect of the noise on the exact solution can be written as

$$x_{\text{SVD}} = \sum_{i=1}^r \frac{(u_i^T (b - e))}{\sigma_i} v_i. \quad (2.5)$$

In this form, it is apparent that despite minimizing the residual, the solution will tend to infinity if the singular values either

1. Span several orders of magnitude, or
2. Are close to zero.

Inverse problems of this type are called linear *discrete ill-posed problems*. If the singular values span several orders of magnitude, the unknown error components e in the data dominate the behavior of the solution. Even when there is no measurement error present in the data but the singular values are close to zero, problems of this type do not admit a numerically stable solution.

A simple modification of (2.4) to overcome this difficulty is the Truncated SVD (TSVD) solution, in which the sum is truncated at some index $1 \leq k \leq r$. The choice of k is determined both by the noise level in the data, and by examining the singular values specific to the problem. While effective at estimating the general behavior of the solution, this approach is not particularly elegant, as there is only a single parameter to adjust. The truncation discards a portion of the information about the solution, limiting further refinement.

These challenges motivate the introduction of some additional parameters into the problem to finely control the behavior of the solution, known as *regularization*. Certainly the most widespread form of regularization is Tikhonov regularization, where the minimization problem (2.2) is replaced by

$$x_{\lambda, \mathbf{L}} = \arg \min_x \left\{ \|b - \mathbf{A}x\|_2^2 + \lambda^2 \|\mathbf{L}x\|_2^2 \right\}, \quad (2.6)$$

where $\lambda \in \mathbb{R}$ is referred to as the *regularization parameter* and $\mathbf{L} \in \mathbb{R}^{p \times n}$ is the *regularization matrix*.

Tikhonov regularization alters the original problem statement (2.1). The minimization problem's equivalent augmented matrix is

$$\begin{bmatrix} \mathbf{A} \\ \lambda \mathbf{L} \end{bmatrix} \begin{bmatrix} x \end{bmatrix} = \begin{bmatrix} b \\ \mathbf{0} \end{bmatrix},$$

where $\mathbf{0}_p \in \mathbb{R}^p$ is a column vector with all zero entries. The effect of this change can be seen by manipulating the analogous form of the normal equations $\mathbf{A}^\top \mathbf{A} x = \mathbf{A}^\top b$ to attain

$$x_{\lambda, \mathbf{L}} = (\mathbf{A}^\top \mathbf{A} + \lambda^2 \mathbf{L}^\top \mathbf{L})^{-1} \mathbf{A}^\top b. \quad (2.7)$$

It is worth mentioning that this Tikhonov regularized solution closely resembles the form of *Wiener Filtering*. In general, it can be shown by singular value decomposition that the Tikhonov and Wiener solutions are equivalent, given as

$$x_{\text{Tikhonov}} = (\mathbf{A}^\top \mathbf{C}^{-1} \mathbf{A} + \lambda^2 \mathbf{\Gamma}^{-1})^{-1} \mathbf{A}^\top \mathbf{C}^{-1} \tilde{b} \quad (2.8)$$

$$x_{\text{Wiener}} = \mathbf{\Gamma} \mathbf{A}^\top (\mathbf{A} \mathbf{\Gamma} \mathbf{A}^\top + \lambda^2 \mathbf{C})^{-1} \tilde{b}. \quad (2.9)$$

From the normal equations (2.7), the effect of the regularization parameter and matrix can be elucidated by considering the general singular value decomposition (GSVD) of the matrices \mathbf{A} , \mathbf{L} , defined as

$$\mathbf{A} = \mathbf{U} \mathbf{C} \mathbf{X}^\top, \quad \mathbf{U} \in \mathbb{R}^{m \times m}, \quad \mathbf{C} \in \mathbb{R}^{m \times n}, \quad \mathbf{X} \in \mathbb{R}^{n \times n} \quad (2.10)$$

$$\mathbf{L} = \mathbf{V} \mathbf{S} \mathbf{X}^\top, \quad \mathbf{V} \in \mathbb{R}^{p \times p}, \quad \mathbf{S} \in \mathbb{R}^{p \times n}, \quad (2.11)$$

where the following conditions are satisfied

$$\mathbf{U}^\top \mathbf{U} = \mathbf{U} \mathbf{U}^\top = \mathbf{I}_m$$

$$\mathbf{V}^\top \mathbf{V} = \mathbf{V} \mathbf{V}^\top = \mathbf{I}_n$$

$$\mathbf{X}^\top \mathbf{X} = \mathbf{X} \mathbf{X}^\top = \mathbf{I}_p$$

$$\mathbf{C}^\top \mathbf{C} + \mathbf{S}^\top \mathbf{S} = \mathbf{I}_n$$

$$\mathbf{C} = \begin{pmatrix} \mathbf{\Sigma} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{n-p} \end{pmatrix}$$

$$\mathbf{S} = \begin{pmatrix} \mathbf{M} & \mathbf{0} \end{pmatrix}$$

$$\mathbf{\Sigma} = \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_p), \quad 1 \geq \sigma_p \geq \dots \geq \sigma_1 \geq 0$$

$$\mathbf{M} = \text{diag}(\mu_1, \mu_2, \dots, \mu_p), \quad 1 \geq \mu_1 \geq \dots \geq \mu_p \geq 0,$$

where we have assumed for simplicity that $p \leq n \leq m$.

The solution to the normal equations is

$$x_{\lambda, \mathbf{L}} = \sum_{i=1}^n \frac{\sigma_i^2}{\sigma_i^2 + \lambda^2 \mu_i^2} \frac{(u_i^\top b)}{\sigma_i} x_i, \quad (2.12)$$

where x_i is the i th column of \mathbf{X} . The coefficients $f_i = \sigma_i^2 / (\sigma_i^2 + \lambda^2 \mu_i^2)$ are known as *filter factors*. For a common choice of the regularization matrix $\mathbf{L} = \mathbf{I}_p$, the singular values are simply $\mu_i = 1$.

The solution (2.12) shows that through a careful choice of the regularization parameter and matrix, it is possible to compensate for the behavior of the singular values σ_i of \mathbf{A} . This makes it possible to balance minimizing the residual and minimizing the norm of the solution. A small regularization parameter λ recovers the previous minimization problem (2.2). A large regularization parameter effectively minimizes the norm of the solution, $\|\mathbf{L}x\|_2$.

Clearly, the choice of the regularization parameters has a significant effect on the quality of the solution. In the next section, we discuss how to choose the parameters based on a statistical interpretation.

2.2 Bayesian Inverse Problems

Regularization is inherently an *ad hoc* method for obtaining solutions to ill-posed inverse problems - there is no single correct choice of regularization parameters for a given problem. Likewise, there is no single method for all problems that yields the optimal regularization parameters. Despite this, several common methods exist that can be used in practice to obtain reasonable parameters. L-curve analysis plots in a logarithmic scale the norm of the solution against the norm of the residual, and attempts to find a balance between the two. However, it requires solving the inverse problem at several values of λ , and for a fixed matrix \mathbf{L} . A similar approach is to find a value of λ that yields a solution whose residual is approximately equal to the noise level in the data, if this information is known.

Another interpretation of the parameters in Tikhonov regularization is found in the Bayesian framework. A closer examination of these ideas is enlightening both in the choice of regularization parameters and in the physical interpretation of regularization, and follows below.

We begin by restating our intention, namely, to determine the unknown $x \in \mathbb{R}^n$. At the heart of the Bayesian method lies the idea that random variables are used to describe our uncertainty about some unknown, leading to the term *subjective* probability. Let $\pi(x)$ be the probability density of some random variable \mathbf{X} , of which x is a realization. We express our prior belief that any given solution x is the true solution through the *prior probability density* $\pi_{\text{pr}}(x)$.

The *likelihood density* is defined as the conditional probability density $\pi(b|x)$ of measuring a data sample b , given a fixed solution x . Furthermore, let the measurement noise be a realization e of a random variable \mathbf{E} , with probability density $\pi_{\text{noise}}(e)$. For the measurement model (2.1), these two densities are equal,

$$\pi(b|x) = \pi_{\text{noise}}(e) = \pi_{\text{noise}}(b - \mathbf{A}x). \quad (2.13)$$

The desired information we wish to attain is the *posterior density* $\pi(x|b)$, or the conditional probability that x is the true solution given a fixed data sample b . Bayes formula states that the posterior density is given by

$$\pi(x|b) = \frac{\pi(b|x)\pi_{\text{pr}}(x)}{\pi(b)}. \quad (2.14)$$

The normalization $\pi(b)$ is uninteresting for determining a solution, and since b is the measured data, it is always possible to choose a normalization such that $\pi(b) = 1$. Bayes formula (2.14) is therefore more commonly written as

$$\pi(x|b) \propto \pi(b|x)\pi_{\text{pr}}(x). \quad (2.15)$$

A particularly revealing example is the case where X and E are independent Gaussian random variables

$$\mathbf{X} \sim \mathcal{N}(0, \Gamma_X), \quad \mathbf{E} \sim \mathcal{N}(0, \sigma^2 \mathbf{I}), \quad (2.16)$$

where Γ_X is the covariance matrix of \mathbf{X} , and σ is the standard deviation of \mathbf{E} . In this case, the prior density for \mathbf{X} is given by

$$\pi_{\text{pr}}(x) \propto \exp\left(-\frac{1}{2}x^T \Gamma_X^{-1} x\right). \quad (2.17)$$

Since Γ_X is a covariance matrix, it must be both symmetric and positive definite, and so Γ_X^{-1} admits a Cholesky factorization $\mathbf{L}_X^\top \mathbf{L}_X = \Gamma_X^{-1}$. The prior can therefore equivalently be written as

$$\pi_{\text{pr}}(x) \propto \exp\left(-\frac{1}{2} \|\mathbf{L}_X x\|_2^2\right). \quad (2.18)$$

The likelihood density (2.13) is

$$\pi(b|x) \propto \exp\left(-\frac{1}{2}(b - \mathbf{A}x)^\top (\sigma^2 \mathbf{I})^{-1} (b - \mathbf{A}x)\right) = \exp\left(-\frac{1}{2} \left\| \frac{1}{\sigma} (b - \mathbf{A}x) \right\|_2^2\right). \quad (2.19)$$

The posterior (2.15) in this case is given by

$$\pi(x|b) \propto \exp\left\{-\frac{1}{2} \left(\left\| \frac{1}{\sigma} (b - \mathbf{A}x) \right\|_2^2 + \|\mathbf{L}_X x\|_2^2 \right)\right\}. \quad (2.20)$$

The posterior is a probability distribution, rather than a point estimate such as in Tikhonov regularization. If normalized, the distribution describes at each point the probability that the solution is the true unknown. It is possible to construct point estimates from the posterior, such as the maximum probability or *maximum a posteriori* (MAP) estimate, denoted by x_{MAP} . For (2.20), this estimate can be found by solving the minimization problem

$$x_{MAP} = \arg \min_x \left\{ \|b - \mathbf{A}x\|_2^2 + \sigma^2 \|\mathbf{L}_X x\|_2^2 \right\}. \quad (2.21)$$

Comparing to the minimization problem in Tikhonov regularization (2.6) offers a statistical interpretation for the regularization parameter and matrix. The regularization matrix \mathbf{L} is the Cholesky factorization of the inverse of the covariance matrix Γ_X , and therefore expresses the prior belief about the sample. The regularization parameter is related to the measurement noise, specifically, for it is equal to the standard deviation for Gaussian white noise.

2.3 Discretization Error

We next turn to the origin of discretization errors. In particular, we consider the effect of discretizing a partial differential equation with some finite number of unknown coefficients.

Let the unknown be described by a parameter vector θ . For an input source term y , and finite difference (FD) or finite elements (FE) matrix \mathbf{A}_θ , the discretized solution x is given by

$$y = \mathbf{A}_\theta x. \quad (2.22)$$

Here, \mathbf{A}_θ is both square and invertible.

The linear observation model (2.1) with additive noise is

$$b = \mathbf{B}x + e = \mathbf{B}\mathbf{A}_\theta^{-1}y + e = f(\theta) + e, \quad (2.23)$$

where f represents the non-linear mapping from the unknown θ to the noiseless data.

Assume that the measurement noise is independently distributed Gaussian with covariance matrix \mathbf{C} ,

$$e \sim \mathcal{N}(0, \mathbf{C}), \quad (2.24)$$

and that the prior is also Gaussian with mean θ_0 and covariance Γ ,

$$\pi_{\text{prior}}(\theta) = \mathcal{N}(\theta_0, \Gamma). \quad (2.25)$$

The MAP estimate (2.21) in the Bayesian framework in this case is

$$\theta_{\text{MAP}} = \arg \min_{\theta} \left\{ \|b - f(\theta)\|_{\mathbf{C}}^2 + \|\theta - \theta_0\|_{\Gamma}^2 \right\}, \quad (2.26)$$

where

$$\|z\|_{\mathbf{M}} = z^{\top} \mathbf{M}^{-1} z. \quad (2.27)$$

The discretization (2.22) from the continuous model introduces a modeling error. Ideally, it would be possible to increase the level of discretization to obtain an arbitrarily accurate solution. However, a fine discretization leads to an increased computation time, further burdened by the non-linearity of the forward model f . Thus, in many practical applications, computational efficiency limits the number of elements used in the FE or FD model when solving the inverse problem. Since the significance of the modeling error grows with the coarseness of the discretization, it is necessary to estimate and account for this modeling error to attain viable solutions (2.26).

To understand the effect of the discretization error on the data b , we consider a fine discretization N , and a coarse discretization n ,

$$\mathbf{A}_{\theta}^N x = y^N, \quad \mathbf{A}_{\theta}^n x = y^n, \quad (2.28)$$

where $n < N$ is the number of nodes in the discretization. The corresponding noiseless observation models are

$$b^N = f^N(\theta), \quad b^n = f^n(\theta). \quad (2.29)$$

The fine discretization in this case is used to represent the true θ , while the coarse discretization is used in practice to solve the inverse problem. In other words, it is assumed that the data is a product of the fine mesh, that is

$$b = f^N(\theta) + e = f^n(\theta) + \{f^N(\theta) - f^n(\theta)\} + e. \quad (2.30)$$

Evidently, the modeling error vector introduced by the discretization from the fine to the coarse mesh can be written as

$$m = F^{N,n}(\theta) = f^N(\theta) - f^n(\theta). \quad (2.31)$$

In the Bayesian framework, this modeling error can be approximated as a random variable, with some mean and covariance. Assume that this error is Gaussian distributed

$$m \sim \mathcal{N}(\bar{m}, \Sigma), \quad (2.32)$$

where \bar{m} is the mean and Σ is the covariance. The complete measurement model in the coarse mesh may thus be written more compactly as

$$b = f^n(\theta) + E, \quad E \sim \mathcal{N}(\bar{m}, \mathbf{C} + \Sigma). \quad (2.33)$$

The MAP solution for this error enhanced model is given by

$$\theta_{\text{MAP}} = \arg \min_{\theta} \left\{ \|b - f(\theta) - \bar{m}\|_{\mathbf{C} + \Sigma}^2 + \|\theta - \theta_0\|_{\Gamma}^2 \right\}. \quad (2.34)$$

While the modeling error is treated as a random variable, it is not independent of the unknown θ . If the probability density of θ is denoted as π_{θ} , then the density of m can be written implicitly as the push-forward density

$$\pi_m = F_*^{N,n} \pi_{\theta}, \quad (2.35)$$

that is,

$$\mathbb{P}\{m \in A\} = \int_A F_*^{N,n} \pi_{\theta}(m) dm. \quad (2.36)$$

Chapter 3

Estimation of the Modeling Error

To solve the error enhanced model (2.34), it is necessary to estimate the probability density of the modeling error. The prior is taken to be a Gaussian distribution, and likewise the modeling error is assumed to be approximately Gaussian distributed. A previously proposed method to obtain the mean and covariance of the error is reviewed below, in which the prior density is sampled in an off-line calculation that is carried out before solving the inverse problem. The motivation for a dynamic estimation method is discussed, and an alternative algorithm is introduced which concomitantly updates the modeling error while solving the inverse problem.

3.1 Off-Line Modeling Error Approximation

In [16, 17], it was proposed to use the prior to estimate the modeling error probability density. In practice, a Gaussian approximation of the modeling error distribution was proposed, leading to Algorithm 1.

Algorithm 1: Off-Line Modeling Error Approximation

1. Draw a sample of size K of conductivity distributions θ in the coarse mesh from the prior π_{prior} ,

$$\mathcal{S} = \{\theta_1, \theta_2, \dots, \theta_K\}.$$

2. Compute the modeling error for each draw θ_l as

$$m_l = F^{N,n}(\theta_l) = f^N(\theta_l) - f^n(\theta_l).$$

Estimate the sample mean and covariance of this sample as

$$\bar{m} = \frac{1}{K} \sum_{l=1}^K m_l, \quad \Sigma = \frac{1}{K} \sum_{l=1}^K (m_l - \bar{m})(m_l - \bar{m})^\top$$

3. Compute the MAP estimate for θ from the error enhanced model (2.34).
-

The algorithm only needs to be run once, and the estimates for the mean and covariance can be used for repeatedly solving the inverse problem thereafter. The practical viability of the super-resolution attained has been demonstrated in numerous applications [1, 4, 20]. However, the algorithm above suffers from two primary drawbacks. First, for a large sample size K , computing the forward estimates in step two is computationally expensive due to the non-linear nature of the problem. Second, the algorithm neglects the correspondence between the modeling error m and the unknown θ [5, 17], a question that is further explored in this work.

3.2 Dynamic Modeling Error Approximation

The idea to dynamically update the modeling error is inspired by the observation that the information about the modeling error increases as the information about the true unknown increases. Before, sampling the prior density yields a preliminary estimate of the error probability distribution. However, the true dependence of the modeling error on the unknown is captured by the likelihood function. This motivates a scheme that iteratively samples the posterior density as solutions are obtained and the density theoretically narrows, yielding increasingly accurate estimates of the modeling error.

Algorithm 2, outlined below, is based on the idea of Ensemble Kalman Filtering (EnKF), where an ensemble of conductivity distributions providing estimates of the modeling error and the unknown. Over a series of time-steps, the likelihood function is updated at each time-step to take the latest information about the modeling error into account, and the subsequent solutions are used to update the estimate for the modeling error.

The ability to have a small ensemble size in EnKF algorithms allows the approach to be computationally efficient. In particular, the effect of varying ensemble size is explored in detail later on. A further benefit is the ability to monitor the convergence of the modeling error and solution, and to terminate the iterations prematurely while retaining these estimates.

To sample the posterior, Algorithm 2 uses a randomized optimization scheme [23, 24], which bears similarity to the updating steps in EnKF algorithms [10], and also to parametric bootstrapping [8, 9].

Let $\bar{m}^{(t)}$ denote the most recent, cumulative estimate for the mean at time-step t , and $\Sigma^{(t)}$ the estimate for the covariance. Initially, $\bar{m}^{(0)} = 0$ and $\Sigma^{(0)} = 0$, and a sample \mathcal{S}_k^0 of size k of realizations is drawn independently from the prior π_{prior} in the coarse mesh,

$$\mathcal{S}_k^0 = \{\theta_1^0, \theta_2^0, \dots, \theta_k^0\}. \quad (3.1)$$

At every time-step t , an artificial data sample is constructed as

$$b_l = b + w_l, \quad w_l \sim \mathcal{N}(0, \mathbf{C} + \Sigma^{(t)}), \quad 1 \leq l \leq k. \quad (3.2)$$

From this, a new estimate for each element of the ensemble is obtained as

$$\mathcal{S}_k^{t+1} = \{\theta_1^{t+1}, \theta_2^{t+1}, \dots, \theta_k^{t+1}\}, \quad (3.3)$$

where θ_l^{t+1} is given as the solution to the minimization problem

$$\theta_l^{t+1} = \arg \min_{\theta} \left\{ \|b_l - f(\theta) - \bar{m}^{(t)}\|_{\mathbf{C} + \Sigma^{(t)}}^2 + \|\theta - \theta_l^t\|_{\Gamma}^2 \right\}. \quad (3.4)$$

The new sample \mathcal{S}_k^{t+1} represents a model for the posterior density, which narrows as the time-step grows. The ensemble therefore effectively samples the posterior for $t > 0$, and provides increasingly accurate estimates of the modeling error as the algorithm continues.

The modeling error of the sample at the current time-step $t + 1$ is calculated as before

$$m_l^{t+1} = F^{N,n}(\theta_l^{t+1}) = f^N(\theta_l^{t+1}) - f^n(\theta_l^{t+1}), \quad 1 \leq l \leq k. \quad (3.5)$$

The cumulative mean and covariance are updated as a moving average and covariance using the recursive relations

$$\begin{aligned} \bar{m}^{(t+1)} &= \frac{t}{t+1} \bar{m}^{(t)} + \frac{1}{t+1} \bar{m}^{t+1} \\ \Sigma^{(t+1)} &= \frac{t}{t+1} \Sigma^{(t)} + \frac{1}{t+1} \Sigma^{t+1} + \frac{t}{(t+1)^2} (\bar{m}^{(t)} - \bar{m}^{t+1})(\bar{m}^{(t)} - \bar{m}^{t+1})^{\top}. \end{aligned}$$

where $\bar{m}^{(t+1)}, \Sigma^{(t+1)}$ denote the new cumulative estimates, $\bar{m}^{(t)}, \Sigma^{(t)}$ denote the estimates up to time-step t , and $\bar{m}^{t+1}, \Sigma^{t+1}$ denote the estimates at time-step $t + 1$.

Algorithm 2: Dynamic Modeling Error Approximation

1. Draw an initial ensemble \mathcal{S}_k^0 of size k of conductivity distributions θ^0 in the coarse mesh from the prior π_{prior} ,

$$\mathcal{S}_k^0 = \{\theta_1^0, \theta_2^0, \dots, \theta_k^0\}.$$

2. Initialize the modeling error mean $\bar{m} = 0$ and covariance $\Sigma = 0$. Set $t = 0$.
3. Begin iteration for time-step t :

- (a) Generate an artificial data sample $\{b_l\}$ for all $1 \leq l \leq k$, where

$$b_l = b + w_l, \quad w_l \sim \mathcal{N}(0, \mathbf{C} + \Sigma).$$

- (b) Compute the solutions to the inverse problem using the error enhanced model (2.34), where for each member of the ensemble the data sample is b_l and the mean is θ_l , that is

$$\theta_l^{t+1} = \arg \min_{\theta} \left\{ \|b_l - f(\theta) - \bar{m}\|_{\mathbf{C} + \Sigma}^2 + \|\theta - \theta_l^t\|_{\Gamma}^2 \right\},$$

leading to a new ensemble of conductivity distributions \mathcal{S}_k^{t+1} given by

$$\mathcal{S}_k^{t+1} = \{\theta_1^{t+1}, \theta_2^{t+1}, \dots, \theta_k^{t+1}\}.$$

- (c) Compute the modeling error for each element θ_l^{t+1} of the ensemble as

$$m_l^{t+1} = F^{N,n}(\theta_l^{t+1}) = f^N(\theta_l^{t+1}) - f^n(\theta_l^{t+1}).$$

Estimate the sample mean and covariance of this sample as

$$\bar{m}^{t+1} = \frac{1}{K} \sum_{l=1}^K m_l^{t+1}, \quad \Sigma^{t+1} = \frac{1}{K} \sum_{l=1}^K (m_l^{t+1} - \bar{m}^{t+1})(m_l^{t+1} - \bar{m}^{t+1})^{\top}.$$

- (d) Updated the cumulative modeling error mean \bar{m} and covariance Σ with the estimates at $t + 1$ as a moving average and covariance,

$$\begin{aligned} \bar{m}^* &= \frac{t}{t+1} \bar{m} + \frac{1}{t+1} \bar{m}^{t+1} \\ \Sigma^* &= \frac{t}{t+1} \Sigma + \frac{1}{t+1} \Sigma^{t+1} + \frac{t}{(t+1)^2} (\bar{m} - \bar{m}^{t+1})(\bar{m} - \bar{m}^{t+1})^{\top}. \end{aligned}$$

Set $\bar{m} = \bar{m}^*$ and $\Sigma = \Sigma^*$.

- (e) If the convergence criteria for modeling error mean and covariance are satisfied, stop. Else, set $t = t + 1$ and continue from (a) for the next iteration.
-

Chapter 4

Application to EIT

Electrical impedance tomography (EIT) is an imaging method which measures spatial variations in the electric admittivity distribution inside a sample. Currents are applied through electrodes attached to the surface of the sample, and resulting relative voltages are measured. Benefits of the method include its non-invasive nature and rapid measurement time, with applications ranging from biomedical imaging to industrial process tomography [6].

4.1 Forward Model

To examine the forward model for EIT, consider applying currents with angular frequency w to a body Ω through L electrodes. For typical values in medical applications, for example, applied current frequencies are on the order of hundreds of kHz, and magnitudes on the order of tens of mA. Define the admittivity distribution $\gamma(x, w)$ of a point x inside the body as

$$\gamma(x, w) = \sigma(x, w) + iw\epsilon(x, w), \quad (4.1)$$

where $\sigma(x, w)$ is the electric conductivity distribution and $\epsilon(x, w)$ is the electric permittivity distribution. The conductivity is the inverse of the resistivity, and the two are often used interchangeably.

Let the matrix $I = (I_1, \dots, I_L)^\top \in \mathbb{C}^{L \times N}$ represent a set of N current patterns applied to the L electrodes. Specifically, each column of I represents a current pattern, with the i^{th} element of the column representing the magnitude of the current applied at the i^{th} electrode. The set of noiseless voltage measurements made at the electrodes can be written by Ohm's law as

$$U(\gamma) = R(\gamma)I, \quad (4.2)$$

where $U(\gamma) = (U_1, \dots, U_L)^\top \in \mathbb{C}^{L \times N}$ is the noiseless resulting voltage pattern and $R(\gamma) \in \mathbb{C}^{L \times L}$ is the impedance matrix of the body Ω .

While the voltage patterns are a linear function of the applied current patterns, the impedance matrix $R(\gamma)$ is a non-linear function of the admittivity distribution. For the discretized version of the EIT model, the exact form of the matrix R is derived in Section 4.2




Figure 4.1: Illustrations of EIT setup. (a) Sample 2D conductivity setup, with 8 electrodes e_l applied to the surface of a body Ω . (b) Experimental EIT setup in biomedical imaging [6].

below. In general, it is assumed that the body consists of linear and isotropic media, that is,

$$\vec{D} = \epsilon \vec{E} \quad (4.3)$$

$$\vec{B} = \mu \vec{H} \quad (4.4)$$

$$\vec{j} = \sigma \vec{E}, \quad (4.5)$$

where $\mu(x, w)$ is the magnetic permeability, $\vec{E}(x, w)$ is the electric field, $\vec{B}(x, w)$ is the magnetic field, $\vec{D}(x, w)$ is the electric displacement, and $\vec{H}(x, w)$ is the magnetic induction.

These assumptions lead to a specific set of equations that govern the behavior of the electromagnetic fields inside the body. Using Maxwell equations, it can be shown that the electric potential $u(x, w)$ inside the body is described by

$$\vec{\nabla} \cdot (\gamma \vec{\nabla} u) = 0 \text{ for } x \in \Omega. \quad (4.6)$$

The derivation of this result from the Maxwell equations is provided in Appendix A.

While the admittivity distribution is complex, in general the approximation is invoked that the imaginary term is negligible, so that $\gamma \approx \sigma$. This approximation is valid in most industrial applications where the angular frequency w is very low, typically on the order of kHz. The approximation also holds in medical applications, for example, in the diagnosis of breast cancer. A primary objective is to distinguish between cancerous and normal breast tissue, and the large difference in conductivity values between the two tissues makes EIT a promising method for accurate diagnosis. It has been shown that, depending on the frequency, for cancerous tissue the maximum ratio is $w\epsilon/\sigma \approx 10^{-1}$, and for normal tissue a lower ratio is encountered [19].

From here on, it is assumed that $\gamma = \sigma$. The behavior of the electric potential inside the body in this approximation is governed by

$$\vec{\nabla} \cdot (\sigma \vec{\nabla} u) = 0 \text{ for } x \in \Omega. \quad (4.7)$$

4.1.1 Boundary Conditions

The electric potential is subject to boundary conditions on the surface of the body $\partial\Omega$, governed by the injection of currents through a set of L discrete electrodes with area e_l , $1 \leq l \leq L$. The derivation of the boundary conditions is presented below [6], leading to a discussion of the most accurate electrode model [26, 7].

The currents injected into the electrodes must satisfy the condition corresponding to conservation of charge,

$$\sum_{l=1}^L I_l = 0. \quad (4.8)$$

Similarly, the ground voltage is chosen such that

$$\sum_{l=1}^L U_l = 0. \quad (4.9)$$

Define a unit normal vector $\hat{\nu}$, pointing inward at every point on the surface of the body. It is tempting to define a current density over the surface of the body, whose normal component j is given as

$$j = \sigma \frac{\partial u}{\partial \nu}. \quad (4.10)$$

This is known as the *continuum* electrode model. In reality, it is necessary to account for two effects: first, the electrodes are discrete. The integral of the current density over the surface of the electrode must be equal to the total current injected into the electrode, that is,

$$\int_{e_l} \gamma \frac{\partial u}{\partial \nu} dS = I_l. \quad (4.11)$$

Between the electrodes, the current density must be zero

$$\sigma \frac{\partial u}{\partial \nu} = 0 \text{ between electrodes.} \quad (4.12)$$

The second effect that needs to be accounted for is the extra conductive material added by the electrodes themselves. In practice, this leads to a *shunting* effect, in which the current utilizes the lower conductivity of the electrodes to traverse the path of minimum energy. A first *ansatz* to describe the voltage measurements on the electrodes is to assume that the potential is constant on each electrode, $U_l = u$ on e_l .

It has been experimentally established that this purely shunting model fails to account for an electrochemical effect that occurs at the area of contact, forming a region of low conductance around the electrode. This region is characterized by the contact impedance of each electrode z_l . A simple model for estimating this effect is the *lumped-z* model

$$u + z_l I_l = U_l \text{ on } e_l. \quad (4.13)$$

This model, too, is insufficient to accurately capture the change in conductivity caused by the electrodes. The electrodes are not point objects but rather have finite area. Applying a current at the center of the electrode from an external source leads to the well known effect that the current density is greater at the corners than in the center of the electrode. Accounting for this, the measured voltages U_l actually satisfy

$$u + z\sigma \frac{\partial u}{\partial v} = U_l \text{ on } e_l. \quad (4.14)$$

The most accurate electrode model is the *complete* model, which consists of the description of the potential inside the body (4.7), the boundary conditions (4.11), (4.12), (4.14), and the conservation conditions (4.8) and (4.9) [6].

In Figure 4.2, the behavior of the currents in a FEM mesh due to the discreteness, shunting and lumping effects in the complete electrode model are illustrated.

Figure 4.3 is a copy of a previous result that demonstrates the accuracy of the various electrode models in comparison to experimental data [26]. The characteristic resistance ρ corresponding to a current pattern I_l is defined as the eigenvalues

$$R(\gamma)I_l = \rho I_l. \quad (4.15)$$

Data is generated from an experimental phantom tank containing a homogenous saline solution of conductivity $1/248 \text{ Sm}^{-1}$ with 16 electrodes. Various current patterns I_l are applied at varying spatial frequencies, and the product of ρ and spatial frequency is plotted as a function of spatial frequency. This procedure is repeated through simulations for the various electrode models.

The parabolic shape of the data is a product of the shunting effect, where the number of unstimulated electrodes, which is determined by the spatial frequency, affects the characteristic resistivity measured. The results show that when the complete electrode model is employed, the forward model described accurately reproduces the data.

4.2 Discretized Model

To attain a forward map of the form (4.2), the complete EIT model described above can be discretized to a finite mesh using the Finite Element method [3]. Define two 2D meshes: a fine mesh with N nodes and corresponding N_e elements, and a coarse mesh with n nodes and corresponding n_e elements. The meshes used are shown in Figure 4.4, and are chosen so that the fine mesh is refined at the boundary near the electrodes, where the dominant effect of the modeling error is observed. The meshes were generated in the public domain research tool EIDORS [28] using DistMesh, a MATLAB code distributed under GNU GPL [25].

To both meshes, $L = 16$, electrodes are applied on the boundary, where each 1D electrode has the same size and equal spacing. The size of the electrodes is chosen such that $\approx 70\%$ of the total perimeter is covered by electrodes, with a slight error stemming from the number of nodes available on the boundary to represent each electrode.

The conductivity is represented in the discretized form as a constant over each mesh element. In the coarse mesh, this can be expressed as

$$\sigma(x) = \sigma_0 + \sum_{j=1}^{n_e} \theta_j \chi_j^n(x), \quad (4.16)$$

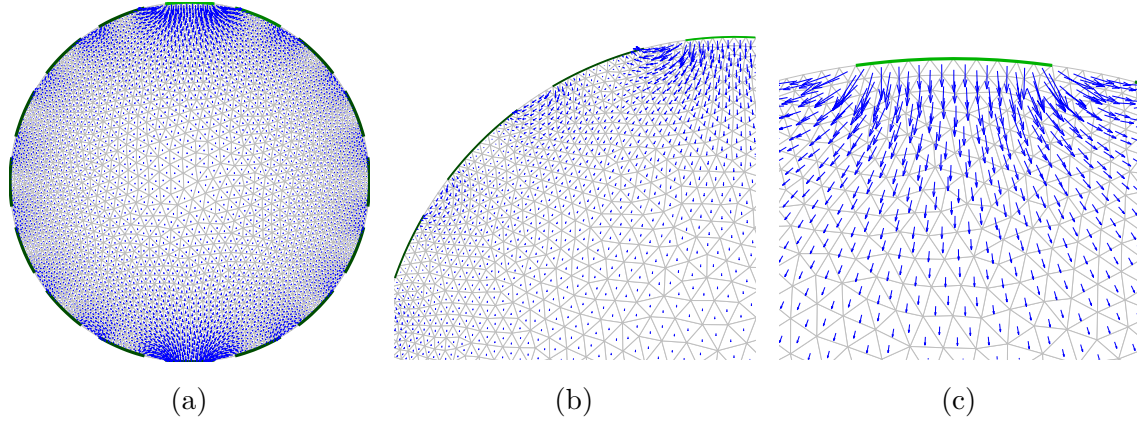


Figure 4.2: Illustrations of the behavior of the current in the complete electrode model for a homogenous sample conductivity distribution. A current of magnitude I is applied at the top electrode and of $-I$ at the bottom electrode, and no currents are applied elsewhere. The vector attached to each element in the FEM mesh represents the magnitude and direction of the current in that element. The panels show: (a) the entire distribution, (b) the shunting effect, where the current utilizes the lower conductivity of the electrodes, and (c) the lumping effect, where the current density on an electrode is greater at the edge than at the center.

where x is a coordinate point in the mesh, χ_j^n is the characteristic function of the j th simplex, θ_j is the constant conductivity value in the j th mesh, and σ_0 is the background conductivity. The background conductivity is chosen as a constant value that is the same for all elements in the mesh.

To project from the coarse mesh into the fine mesh, define the projection operator $P : \mathbb{R}^n \rightarrow \mathbb{R}^N$, which computes the value of the conductivity in an element in the fine mesh as the fractional overlap area with an element in the coarse mesh times the value of that element.

Consider first the discretization of the EIT model in the coarse mesh. The complete observation model consists of applying one frame of current patterns, and measuring the resulting frame of voltage patterns. Here, a frame is defined to consist of the maximum number of linearly independent current patterns. Recall that due to the conservation of charge condition (4.8), there exist $L - 1$ independent current patterns $[I^{(1)}, I^{(2)}, \dots, I^{(L-1)}] \in \mathbb{R}^{L(L-1)}$ in a single frame, where each $I^{(l)}$ is a current pattern that consists of the currents applied to the L electrodes $[I_1^{(l)}, I_2^{(l)}, \dots, I_L^{(l)}]^\top \in \mathbb{R}^L$. A resulting frame of voltage patterns $[U^{(1)}, U^{(2)}, \dots, U^{(L-1)}] \in \mathbb{R}^{L(L-1)}$ is measured, with each voltage pattern consisting of the measurements made on the L electrodes $[U_1^{(l)}, U_2^{(l)}, \dots, U_L^{(l)}] \in \mathbb{R}^L$. For $L = 16$, this corresponds to 240 measurements.

To derive the discretized model, first consider applying a single current pattern $I = I^{(l)}$ and measuring a single resulting voltage pattern $U = U^{(l)}$.

The electric potential is approximated by the values α_i at each node $1 \leq i \leq n$ in the mesh,

$$u(x) = \sum_{i=1}^n \alpha_i \psi_i(x), \quad (4.17)$$

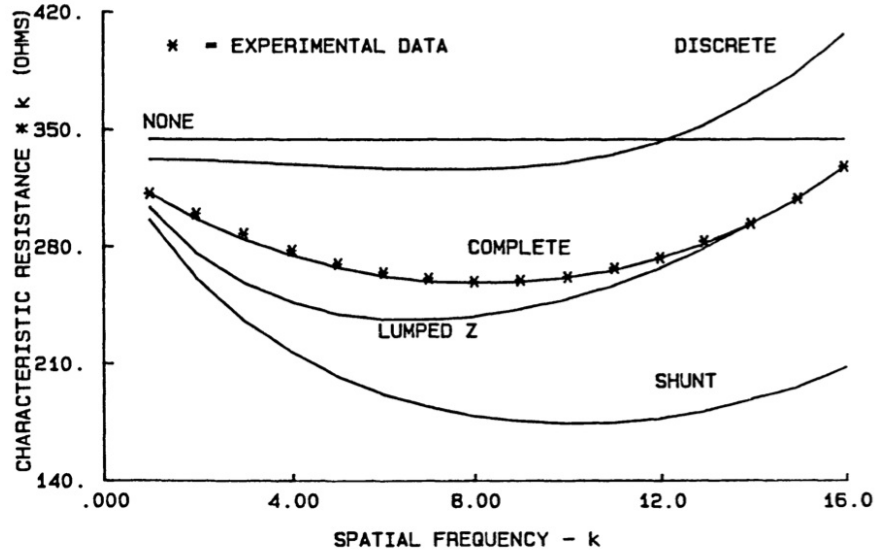


Figure 4.3: Characteristic resistance times spatial frequency as a function of spatial frequency for various electrode models, compared to experimental data. The data is produced from a phantom tank containing a homogenous saline solution of conductivity $1/248 \text{ Sm}^{-1}$ with 16 electrodes [26].

where $\psi(x)$ are the basis vectors. Also define a set of $L-1$ basis vectors $\{U^1, U^2, \dots, U^{L-1}\}$, where $U^i \in \mathbb{R}^L$, and is not to be confused with the measurements made at the electrodes U_l which are the elements of U . For some constants β_i , expand the current pattern U as

$$U = \sum_{i=1}^{L-1} \beta_i U^i, \quad (4.18)$$

where one degree of freedom has been eliminated due to the choice of ground (4.9).

It has been shown that for any potential and voltage combinations u, U and v, V that satisfy the complete electrode model, the following variational form is satisfied,

$$\mathcal{B}((v, V), (u, U)) = \sum_{l=1}^L V_l I_l, \quad (4.19)$$

where $\mathcal{B}((v, V), (u, U))$ is the *sesquilinear* form, defined as

$$\mathcal{B}((v, V), (u, U)) = \int_{\Omega} dx \sigma \vec{\nabla} v \cdot \vec{\nabla} u + \sum_{l=1}^L \frac{1}{z} \int_{e_l} dS (v - V_l)(u - U_l). \quad (4.20)$$

Expand the left side of (4.19) as

$$\sum_{i=0}^n \alpha_i \mathcal{B}((v, V), (\psi_i, 0)) + \sum_{i=1}^{L-1} \beta_i \mathcal{B}((v, V), (0, U^i)) = \sum_{l=1}^L V_l I_l. \quad (4.21)$$

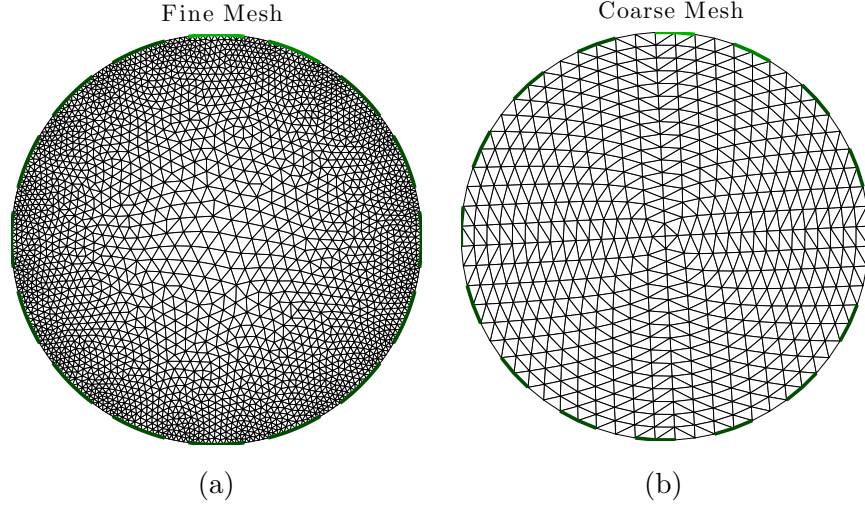


Figure 4.4: (a) The fine grid, consisting of $N_e = 4579$ elements, and (b) the coarse grid (right) of $n_e = 1024$ elements. The number of nodes are $N = 2675$ and $n = 545$, respectively. The green arcs on the boundary represent the $L = 16$ electrodes. Figures were generated in EIDORS [28] using DistMesh [25].

This can be written in the matrix equation form

$$\mathbf{A}_\theta^n x = y^n, \quad (4.22)$$

where the superscript n makes explicit that the forms for the fine and coarse mesh are different, and

$$x = \begin{bmatrix} \alpha \\ \beta \end{bmatrix} \in \mathbb{R}^{n+L-1}, \quad y = \begin{bmatrix} \mathbf{0}_n \\ [(U^i)^\top I]_{i=1:L-1} \end{bmatrix} \in \mathbb{R}^{n+L-1}, \quad (4.23)$$

where $\alpha = [\alpha_1, \alpha_2, \dots, \alpha_n]^\top \in \mathbb{R}^n$, $\beta = [\beta_1, \beta_2, \dots, \beta_{L-1}]^\top \in \mathbb{R}^{L-1}$, $[(U^i)^\top I]_{i=1:L-1} = [(U^1)^\top I, (U^2)^\top I, \dots, (U^{L-1})^\top I] \in \mathbb{R}^{L-1}$, and $\mathbf{0}_n \in \mathbb{R}^n$ is a column vector with all zero entries. The matrix $\mathbf{A}_\theta^n \in \mathbb{R}^{(n+L-1) \times (n+L-1)}$ is given in block form as

$$\mathbf{A}_\theta^n = \begin{bmatrix} [\mathcal{B}((\psi_i, 0), (\psi_j, 0))]_{i,j=1:n} & [\mathcal{B}((\psi_i, 0), (0, U^j))]_{i=1:n, j=1:L-1} \\ [\mathcal{B}((0, U^i), (\psi_j, 0))]_{i=1:L-1, j=1:n} & [\mathcal{B}((0, U^i), (0, U^j))]_{i,j=1:L-1} \end{bmatrix}. \quad (4.24)$$

To derive the equations describing the full frame of measurements, generalize this discretization to any current pattern $I^{(l)} = I = [I_1, I_2, \dots, I_L]^\top$. Explicitly, let $y^{(l)}$ denote the vector y corresponding the l th current pattern,

$$y = \begin{bmatrix} \mathbf{0}_n \\ [(U^i)^\top I^{(l)}]_{i=1:L-1} \end{bmatrix} \in \mathbb{R}^{n+L-1},$$

and write

$$\begin{bmatrix} \alpha^{(l)} \\ \beta^{(l)} \end{bmatrix} = (\mathbf{A}_\theta^n)^{-1} y^{(l)}. \quad (4.25)$$

To each current pattern $I^{(l)}$, the corresponding voltage pattern measured $U^{(l)}$ is given by (4.18)

$$U^{(l)} = \mathbf{U}\beta^{(l)} = \mathbf{U} \left[\mathbf{0}_{(L-1) \times n} \mathbf{I}_{L-1} \right] (\mathbf{A}_\theta^n)^{-1} y^{(l)} = \mathbf{B}^n(\mathbf{A}_\theta^n) y^{(l)}, \quad (4.26)$$

where \mathbf{I}_{L-1} is the identity matrix, and \mathbf{U} is defined as

$$\mathbf{U} = [U^1, U^2, \dots, U^{L-1}] \in \mathbb{R}^{L \times (L-1)}. \quad (4.27)$$

The complete observation model is described by stacking the vectors $I^{(l)}$ into a frame of current patterns, and likewise stacking the vectors $U^{(l)}$ into a frame of voltage measurements.

The discretization to a fine mesh is equivalent to procedure for a coarse mesh, with the appropriate change of number of nodes to N , and with the conductivity σ being projected into the fine mesh using the projection operator.

The voltage patterns in (4.26) represent the noiseless data. As before, it is assumed that the observation model has additive Gaussian measurement noise,

$$V = U(\theta) + e, \quad e \sim \mathcal{N}(0, \mathbf{C}), \quad (4.28)$$

where Γ_η is the covariance matrix of the noise and $U(\theta) = \mathbf{B}^n(\mathbf{A}_\theta^n) y^n$ are the noiseless measurements. It is left as implicit that this forward map refers to the full frame of measurements, so that $V \in \mathbb{R}^{L(L-1)}$.

Chapter 5

Computed Examples in EIT

The algorithm for dynamically estimating the modeling error is examined in the context of EIT and compared to the performance of Algorithm 1 [3]. The computations are performed using a 2D model with finite discretization. The circular discs are assumed to be of radius 0.2 m. It has been shown that the optimal current patterns for circular conductivity distributions are trigonometric patterns, where the amplitude of the current is chosen as 10 mA, and the contact impedance is $z = 3 \times 10^{-3} \Omega \text{m}^2$. These values consistent with what is typically encountered in biomedical applications [6, 26, 19].

5.1 Prior Density

The prior is assumed to be Gaussian distributed with mean $\sigma_0 \in \mathbb{R}^{n_e}$ and covariance $\Gamma \in \mathbb{R}^{n_e \times n_e}$,

$$\pi_{\text{prior}} \sim \mathcal{N}(\sigma_0, \Gamma). \quad (5.1)$$

Let x_j denote the Cartesian coordinates of the center of the j th coarse grid element, $1 \leq j \leq n_e$. Define the covariance by the correlation matrix

$$\Gamma_{i,j} = \gamma \exp\left(-\frac{\|x_j - x_i\|_2}{\lambda}\right), \quad (5.2)$$

where λ is the correlation length and γ is the prior marginal variance of each pixel, not to be confused with the admittivity distribution referenced previously. In principle, it is possible for the conductivity distribution to attain negative values. However, these low probability occurrences are rejected in practice.

In Figure 5.1, six draws from the prior distribution are shown for $\lambda = 0.2 \text{ m}$ and $\gamma = 0.05 \text{ S}^2 \text{ m}^2$. These values are chosen *ad hoc* such such that the realizations approximately resemble the true distribution. The background conductivity is chosen for all elements in the mesh as

$$\sigma_0 = 1.3 \text{ Sm}^{-1}. \quad (5.3)$$

5.2 Conductivity Reconstructions

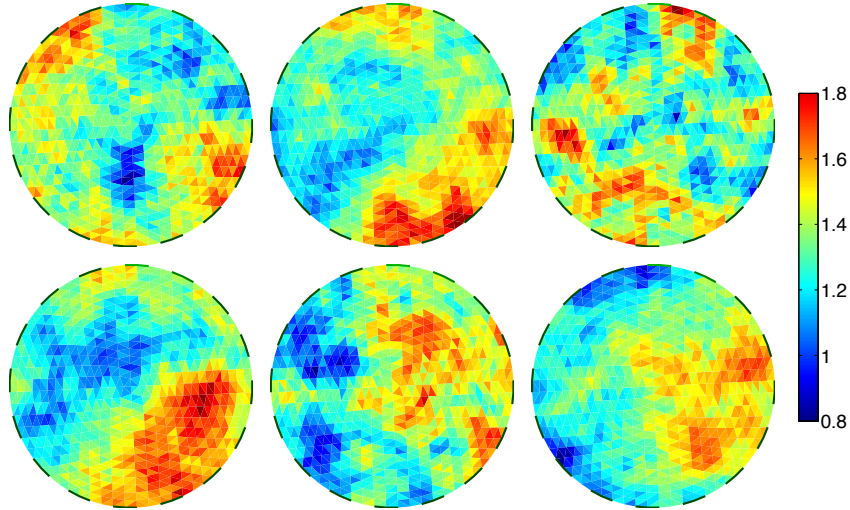


Figure 5.1: Six random realizations of the conductivity model (4.16) with the coefficient vector θ drawn from the prior distribution (5.1). The radius of the disc is 0.2 m, and the correlation length in (5.2) is $\lambda = 0.2$ m. The background conductivity is $\sigma_0 = 1.3 \text{ Sm}^{-1}$, and the prior variance is $\gamma = 0.05 \text{ S}^2 \text{ m}^2$.

To ease the implementation of the proposed algorithms, the public domain EIT research EIDORS is used [28]. In particular, a MATLAB implementation is used for mesh generation, solving of the forward problem and the display of the images. The minimization problems in the two algorithms require non-linear optimization, and are solved using a Gauß-Newton algorithm with backtracking.

For testing the algorithms, a sample conductivity distribution is generated in the fine mesh to represent the true physical conductivities, shown in the top-left panel of Figure 5.2. Test data b is generated using this distribution, representing the noiseless voltage vector in the measurement model.

To demonstrate the effect of the modeling error, the level of the measurement error in the data is taken to be low in comparison, corresponding to precise measurements. The variance η^2 of the Gaussian white noise is determined as

$$\eta^2 = \delta \left\| U^N(\theta) - U^N(0) \right\|_2, \quad (5.4)$$

where the proportionality constant is chosen as $\delta = 7 \times 10^{-3}$. Here, $U^N(0)$ refers to the forward calculation corresponding to a homogenous conductivity distribution at the background value σ_0 .

The reconstructions produced as a result of estimating the modeling error dynamically in Algorithm 2 are shown in Figure 5.2. For an ensemble size of $k = 20$, panels show the average of the ensemble \mathcal{S}_k^t at time-steps $t = 1, 2, 3, 10,$ and 20 .

At the first time step, reconstructions are based off the initial zero values for the modeling error mean and covariance. This leads to particularly strong artifacts at the boundary that corrupt the quality of the image. After the first estimate of the modeling error is obtained, the reconstructions at $t = 2$ have already corrected for a majority of the artifacts. Beyond this, the images quickly converge, and the quality of the reconstructions is unaffected

by the further updating at time-steps beyond $t \approx 10$. The persistent fluctuations in the background conductivity are due to the sample size, and may be corrected by increasing k , with corresponding a tradeoff in computational efficiency.

5.3 Convergence and Discussion

The convergence of the modeling error is of interest as both a measure of the accuracy and the computational efficiency of the dynamic estimation algorithm. Since the algorithm can be terminated at any time-step, once the modeling error has converged, the estimate for its mean and covariance can be used without further updating. A quick convergence therefore can significantly reduce computation time in the dynamic algorithm, particularly if the error converges to the same values as ion the prior sampling algorithm.

Figure 5.3 shows the convergence of the mean and covariance to the values attained from Algorithm 1. The horizontal axis plots kt , corresponding to the cumulative number of computationally expensive inverse problems that have been solved at time-step t , and is a measure of the efficiency of the method. The vertical axes plot the relative differences $\|\bar{m}^{(T)} - \bar{m}\|_2 / \|\bar{m}\|_2$ and $\|\Sigma - \Sigma^{(T)}\|_F / \|\Sigma\|_F$, where $\|\cdot\|_F$ denotes the Frobenius norms of the matrices, and \bar{m}, Σ are the estimates obtained from Algorithm 1 using $K = 2500$.

The mean converges rapidly to the prior sampled values, stabilizing to a difference level of $\approx 5\%$. After the first time-step, the difference level is already less than 20% for all ensemble sizes, suggesting that only few time steps are necessary to obtain accurate estimates. The convergence in the number of time-steps is distinctly independent of the ensemble size k , indicating that even a small ensemble size $k = 2$ may be used to estimate the mean.

The covariance matrix also shows convergent behavior, but it surprisingly does *not* converge to Σ from the prior-sampling. In particular, the covariance is lower in the dynamic estimation method, expressing a higher certainty about the modeling error. This is a consequence of the posterior having a narrower distribution than the prior.

The effect of discretization is captured by the likelihood function in the Bayes formula, rather than the prior. Sampling the prior in Algorithm 1 therefore cannot accurately capture the dependency of the modeling error on the true θ . In the sequential method, the ensemble of conductivity distributions is also initially drawn from the prior. However, at each time step, the samples are updated as the most recent solutions to the inverse problem. After several time steps, the ensemble is therefore effectively sampling the posterior, rather than the prior distribution, incorporating the most recent information about θ at each time step into the estimate for the errors mean and covariance. Since the posterior is guaranteed to have a narrower distribution than the prior, this in turn leads to a lower variance in the distribution for the modeling error.

To study the uncertainty in the posterior between the two methods, consider making draws from the posterior using randomized sampling. For Algorithm 1, use the latest information about the modeling error from step (2) to construct an artificial data sample

$$b_l = b - w_l, \quad w_l \sim \mathcal{N}(0, C + \Sigma), \quad 1 \leq l \leq K^*, \quad (5.5)$$




Figure 5.2: Reconstructions of the true conductivity distribution used to generate the data, shown in top row, left column, for $k = 20$ at $t = 1, 2, 3, 10, 20$ as indicated. The background conductivity is $\sigma_0 = 1.3 \text{ Sm}^{-1}$, and the maximum conductivity (red) and minimum conductivity (blue) are 1.6 Sm^{-1} and 0.9 Sm^{-1} , respectively. The contact impedance value is fixed and assumed to be known, with $z = 3 \times 10^{-3} \Omega \text{ m}^2$.

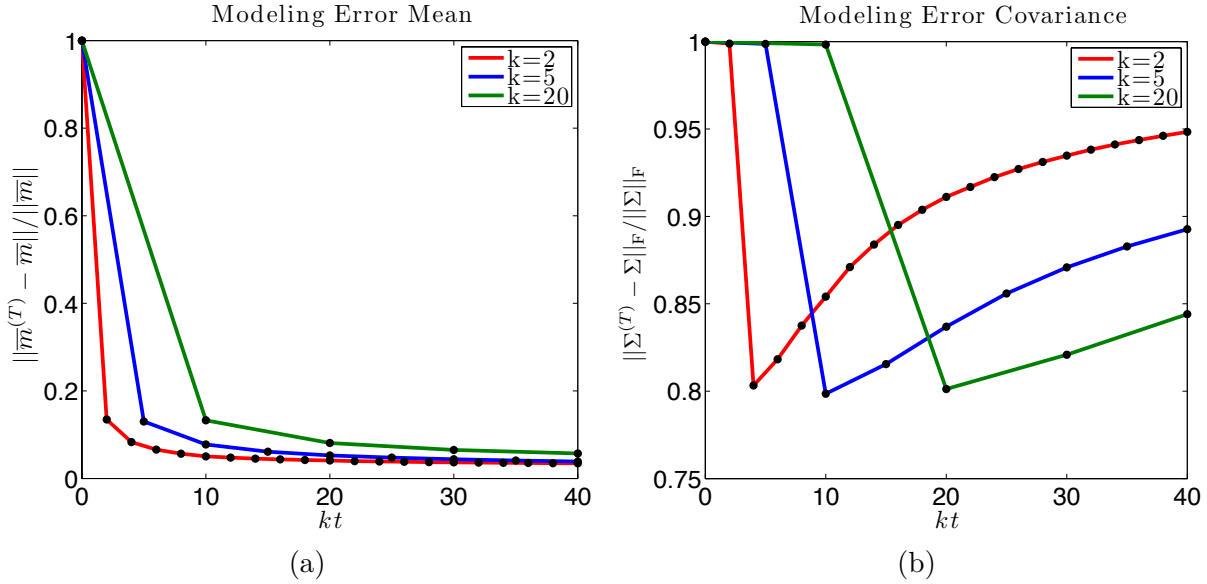


Figure 5.3: The convergence of the relative difference between the dynamic and prior-based modeling error (a) mean and (b) covariance for $k = 2, 5$ and 20 . The computation of the prior-based quantities is done with a fixed number $K = 2500$ of random draws from the prior. The horizontal axis indicates the cumulative number of inverse problems solved at a given time-step.

where $K^* = 1000$ is the sample size. Using the latest estimate θ^+ for the unknown from step (3) as the mean, the solution to the inverse problem is given as

$$\theta_l^* = \arg \min_{\theta^*} \left\{ \|b_l - f(\theta^*) - \bar{m}\|_{C+\Sigma}^2 + \|\theta^* - \theta^+\|_{\Gamma}^2 \right\}. \quad (5.6)$$

Define the covariance matrix Ξ of the sample as

$$\Xi = \frac{1}{K^*} \sum_{l=1}^{K^*} (\theta_l^* - \bar{\theta})(\theta_l^* - \bar{\theta})^\top, \quad (5.7)$$

then a measure of the standard deviation in each element in the mesh is given by the square root of the diagonal of Ξ .

For Algorithm 2, the procedure is similar, with the estimates for \bar{m}, Σ being replaced by the estimates at the latest time-step $\bar{m}^{(T)}, \Sigma^{(T)}$, and the mean θ^+ being replaced by the average of the ensemble S_k^T .

Figure 5.4 shows the standard deviation in each element for both algorithms. The deviation in posterior of Algorithm 1 is larger than in Algorithm 2 by about a factor of ten, with the minimum $\approx 0.1 \text{ Sm}^{-1}$ being almost twice the maximum of Algorithm 2 $\approx 0.06 \text{ Sm}^{-1}$. The plots also agree with the expectation that the uncertainty is larger in the center of the domain than near the boundary. The ratio of the deviation at the boundary to the center is smaller for Algorithm 2, indicating that the posterior is more relatively certain about the boundary than in Algorithm 1.

The curious plateau in the shape of the convergence curves for the covariance in Figure 5.3 is a product of the initialization of the mean and covariance to zero in the dynamic estimation




Figure 5.4: Standard deviation in each pixel in the posterior distributions from (a) Algorithm 1 and (b) Algorithm 2. A sample size of $K^* = 1000$ draws was made from the prior by randomized sampling using the estimates \bar{m}, Σ for Algorithm 1 and $\bar{m}^{(20)}, \Sigma^{(20)}$ for Algorithm 2. The plots are of the square root of the diagonal entries of the covariance matrices of θ for the two samples.

method. The solutions at the first time-step $t = 1$ are obtained using $\bar{m} = 0$ and $\Sigma = 0$, and consequently the ensemble \mathcal{S}_k^1 consists of conductivity distributions that have no information about the modeling error. These almost identical solutions are dominated by a strong modeling error artifact, as illustrated in Figure 5.2, and thus $\Sigma - \Sigma^{(1)} \approx \Sigma$. In Figure 5.3, this is reflected by a relative error that is close to 100% regardless of the sample size k .

The estimates for the modeling error mean obtained from \mathcal{S}_k^1 already have a relative error of less than 20%. At the second step then, the algorithm has started to learn about the modeling error, leading to a covariance more similar Σ . As the iterations proceed, the samples \mathcal{S}_k^t sample the posterior, which is of decreasing variance as more information about the modeling error incorporated into the likelihood function. Consequently the covariance Σ converges to a smaller value than in Algorithm 2, leading to a tendency in the relative error away toward 1.

This behavior is further illustrated in Figure 5.5. For sample size $k = 20$ and time-steps $t = 1, 2, 3, 10, 20$, and 50, the standard deviation in each pixel is shown, calculated from the current ensemble \mathcal{S}_k^t . At $t = 1$, the similarity of the solutions due to the lack of information about the modeling error is reflected in a low variance in the ensemble. Similarly at $t = 2$, the solutions were obtained using a the low estimate $\Sigma^{(1)} \approx 0$, and thus have a similarly low variance. Only for $t \geq 3$ does the method incorporate reasonable estimates for both the mean and covariance of the modeling error. As expected, the uncertainty is greater in the center than at the boundary. Eventually, the posterior converges, decreasing the variance in the ensemble and leading to a gradual decrease in the uncertainty from $t = 20$ to $t = 50$.

In summary, the iterative updating of the modeling error estimate produces qualitatively




Figure 5.5: The standard deviation in each pixel, based on the ensemble of size $k = 20$ in Algorithm 2, at various time-steps as indicated. The plots are of the square root of the diagonal entries of the covariance matrices of θ for the ensemble \mathcal{S}_k^t at each time-step.

similar results as the off-line algorithm based on random draws from the prior. However, the modeling error variance with dynamic updating is significantly smaller, which leads to a posterior for the parameter θ that is smaller by about a factor of ten. This variance is reflected in the evolution of the variance in the ensemble, which smoothes gradually after the first several time-steps as the algorithm progresses. Convergence of the modeling error mean can be achieved with extremely small sample sizes ($k = 2$, $k = 5$), and can be terminated when convergence criteria are satisfied. Thereafter, Algorithm 2 can make use of the estimate of the modeling error mean and covariance for all further inverse problem solutions, similar to Algorithm 2. This leads to a fast algorithm with no need of off-line computations.

The very small sample size, however, may result in noisier sample mean estimates. In Figure 5.6, the sample means after $T = 20$ iterations with different sample sizes are shown, together with the MAP estimate calculated by Algorithm 1 using the off-line modeling error statistics estimate with $K = 2500$ random draws from the prior. As expected, small ensemble sizes $k = 2, 5$ lead to a mean which is less smooth than for large sample sizes $k = 20, 50$. The background the solution from Algorithm 2 is less noisy than in Algorithm 1 - however, the inclusions are slightly more diffuse. These observations are consistent with the larger variance in the posterior of Algorithm 1, shown in Figure 5.4.

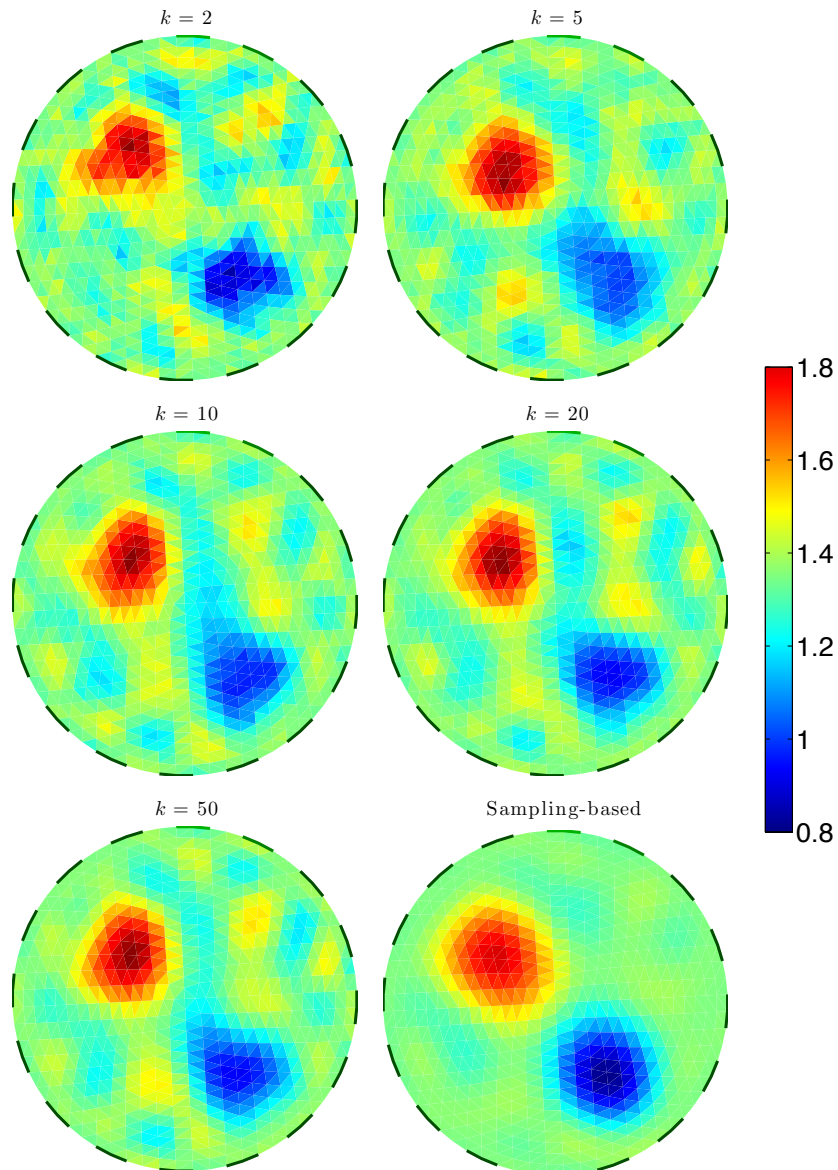


Figure 5.6: Final sample mean estimates from Algorithm 2 after $T = 20$ iteration rounds with different sample sizes, $k = 2, 5, 10, 20$ and 50 . For comparison, the conductivity estimate at the bottom of the right column is computed from the modeling error estimates from Algorithm 1, where $K = 2500$.

Chapter 6

Conclusions

This thesis describes a novel sampling-based algorithm for inverse problems with discretization-based model discrepancies, based on the idea of dynamically updating in an alternating fashion the parameter sample and the modeling error statistics. The viability of this algorithm is demonstrated with an application to the inverse problem of electrical impedance tomography.

Unlike the previous approach of prior-sampling, discretization error estimates from the dynamic algorithm do not require off-line computations, yielding immediate solutions. Furthermore, the adequacy of the sample size can be assessed as the iterations proceed, as well as the convergence of the estimates and solutions. The speed of the algorithm can be adjusted by controlling the sample size, although there exists a corresponding tradeoff in the noise in the posterior. Computed examples demonstrate that even with extremely low sample size, the estimates of the modeling error probability are of reasonably good quality, and the associated computational efficiency makes the method ideal for time sensitive applications. The algorithm can be adapted to time-varying inverse problems such as real time imaging, where the forward model needs not be stationary, thus extending the error estimate methodology to genuinely dynamic models.

Appendices

Appendix A

Derivations of EIT Equation from Maxwells Equations

The derivation of the EIT equations governing the behavior of electromagnetic fields inside the body is reproduced following [6]. We start with a general approach from Maxwell's equations from electrodynamics in the time domain, that is,

$$\vec{\nabla} \times \vec{E}(x, t) = -\frac{\partial \vec{B}(x, t)}{\partial t} \quad (\text{A.1})$$

$$\vec{\nabla} \times \vec{B}(x, t) = \epsilon\mu \frac{\partial \vec{E}(x, t)}{\partial t} + \mu \vec{j}(x, t), \quad (\text{A.2})$$

where x is a point in the domain Ω , $\mu(x, t)$ is the magnetic permeability, $\epsilon(x, t)$ is the electric permittivity, and $\vec{j}(x, t)$ is the current density.

Using the inverse Fourier transform for \vec{E} ,

$$\vec{E}(x, t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} dw \vec{E}(x, w) \exp(iwt), \quad (\text{A.3})$$

and similarly for \vec{B} , we can derive the fixed frequency versions of the Maxwell equations, where w is the angular frequency of the applied current. From (A.1), we have

$$\frac{1}{2\pi} \int_{-\infty}^{\infty} dw (\vec{\nabla} \times \vec{E}(x, w)) \exp(iwt) = \frac{1}{2\pi} \int_{-\infty}^{\infty} dw (-\vec{B}(x, w)) \frac{\partial}{\partial t} \exp(iwt) \quad (\text{A.4})$$

$$= \frac{1}{2\pi} \int_{-\infty}^{\infty} dw (-iw\vec{B}(x, w)) \exp(-iwt) \quad (\text{A.5})$$

$$\therefore \vec{\nabla} \times \vec{E}(x, w) = -iw\vec{B}(x, w), \quad (\text{A.6})$$

and from (A.2),

$$\frac{1}{2\pi} \int_{-\infty}^{\infty} dw \left(\vec{\nabla} \times \vec{B}(x, w) \right) \exp(iwt) = \frac{1}{2\pi} \int_{-\infty}^{\infty} dw \left(\epsilon\mu \vec{E}(x, w) \frac{\partial}{\partial t} \exp(iwt) + \mu \vec{j}(x, w) \exp(iwt) \right) \quad (\text{A.7})$$

$$= \frac{1}{2\pi} \int_{-\infty}^{\infty} dw \left(iw\epsilon\mu \vec{E}(x, w) + \mu \vec{j}(x, w) \right) \exp(-iwt) \quad (\text{A.8})$$

$$\therefore \vec{\nabla} \times \vec{B}(x, w) = iw\epsilon\mu \vec{E}(x, w) + \sigma\mu \vec{E}(x, w). \quad (\text{A.9})$$

Here, in the last line, we used the approximation that the current density in the body Ω is proportional to the electric field

$$\vec{j}(x, w) \approx \sigma(x, w) \vec{E}(x, w), \quad (\text{A.10})$$

where $\sigma(x, w)$ is the electric conductivity. From here on, we implicitly work in the frequency domain.

We make the *ansatz* from electrostatics that \vec{E} can be written as the gradient of the electric potential

$$\vec{E} = -\vec{\nabla}u, \quad (\text{A.11})$$

where $u(x, w)$ is the electric potential. Using this and taking the divergence of (A.9), we can obtain the desired result

$$\vec{\nabla} \cdot \vec{\nabla} \times \vec{B} = \vec{\nabla} \cdot (-\sigma\mu \vec{\nabla}u - iw\epsilon\mu \vec{\nabla}u) \quad (\text{A.12})$$

$$0 = \vec{\nabla} \cdot ((\sigma + iw\epsilon)\vec{\nabla}u) \quad (\text{A.13})$$

$$\therefore \vec{\nabla} \cdot (\gamma \vec{\nabla}u) = 0 \quad (\text{A.14})$$

where $\gamma(x, w) = \sigma(x, w) + iw\epsilon(x, w)$ is the electric admittivity.

To verify this *ansatz* (A.11), consider writing (A.6) in terms of unit-less vectors, that is, let $\vec{E} = [E]\tilde{E}$, $\vec{B} = [B]\tilde{B}$, and $\vec{x} = [x]\tilde{x}$, where $[x]$ is the unit of distance. Here, tilde denotes a vector without dimensions, and the quantity in brackets is a scalar with dimensions as appropriate. (A.6) therefore reads

$$[x]^{-1}[E]\vec{\nabla} \times \tilde{E} = -iw[B]\tilde{B}, \quad (\text{A.15})$$

where we have written the nabla operator in terms of a unit-less vector, $\vec{\nabla} = [x]^{-1}\tilde{\nabla}$, as before. Next, we make a convenience choice for our units for \tilde{E} and \tilde{B} such that

$$\sigma = \frac{[B]}{[E][x]}. \quad (\text{A.16})$$

Therefore, (A.15) can be written as

$$\tilde{\nabla} \times \tilde{E} = -iw\sigma[x]^2\tilde{B}. \quad (\text{A.17})$$




Figure A.1: Integration pillbox Ω_δ of width δ on the boundary $\partial\Omega$. $\hat{\nu}$ is the unit normal vector pointing out of the body.

The general application of EIT is to medical imaging, where $w\sigma[x]^2$ is negligible. For example, applied currents typically operate on the order of 10 kHz, conductivities are less than $1\ \Omega^{-1}\text{m}^{-1}$, and distances are less than 1 m. The curl of the electric field is therefore approximately zero, which by Stoke's theorem is equivalent to (A.11), and so (A.14) holds true as the desired result inside the body Ω .

Next, we derive the equation for the potential at points y on the surface $\partial\Omega$ of the body. It is necessary to account for the applied currents $\vec{j}^{appl}(y, w)$ on the surface by replacing the approximation (A.10) with

$$\vec{j}(y, w) \approx \sigma(y, w)\vec{E}(y, w) + \vec{j}^{appl}(y, w). \quad (\text{A.18})$$

For the fixed frequency Maxwell equation (A.9), this leads to

$$\vec{\nabla} \times \vec{B}(y, w) = iw\epsilon\mu\vec{E}(y, w) + \sigma\mu\vec{E}(y, w) + \mu\vec{j}^{appl}(y, w). \quad (\text{A.19})$$

Taking the divergence of both sides and using (A.11), we have

$$\vec{\nabla} \cdot (\gamma\vec{\nabla}u) = \vec{\nabla} \cdot \vec{j}^{appl}. \quad (\text{A.20})$$

We now integrate both sides over a pillbox of thickness δ on the boundary $\partial\Omega$, as shown in Figure A.1. Using the divergence theorem, we have

$$\iiint_{\Omega_\delta} \vec{\nabla} \cdot (\gamma\vec{\nabla}u) = \iiint_{\Omega_\delta} \vec{\nabla} \cdot \vec{j}^{appl} \quad (\text{A.21})$$

$$\iint_{\partial\Omega_\delta} \gamma\vec{\nabla}u \cdot \hat{\nu} = \iint_{\partial\Omega_\delta} \vec{j}^{appl} \cdot \hat{\nu}, \quad (\text{A.22})$$

where $\hat{\nu}$ is the unit normal vector pointing out of the body Ω .

In the limit where $\delta \rightarrow 0$, only the outside (of Ω) and inside ends of the pillbox contribute, and we have

$$\gamma_{out}\partial_\nu u - \gamma_{in}\partial_\nu u = \vec{j}_{out}^{appl} \cdot \hat{\nu} - \vec{j}_{in}^{appl} \cdot \hat{\nu}. \quad (\text{A.23})$$

Here, ∂_ν denotes a partial derivative along $\hat{\nu}$, and *out* and *in* demote the respective values on the outside and inside of Ω . Since we are assuming that currents are applied only to the outside of the body, we have $\vec{j}_{in}^{appl} = 0$. Furthermore, in most imaging applications, γ outside of the body is negligible, so that we may approximate $\gamma_{out} \approx 0$ and write $\gamma = \gamma_{in}$. We are left with the desired result,

$$\gamma \frac{\partial u}{\partial \nu} = \vec{j}, \quad (\text{A.24})$$

where we have defined $\vec{j} \equiv -\vec{j}_{out}^{appl} \cdot \hat{\nu}$ as the current entering the body.

Appendix B

Gauß-Newton Algorithm

The minimization problem is given as

$$\arg \min_{\theta} \left\{ \|b - f(\theta)\|_{\mathbf{C} + \Sigma}^2 + \lambda^2 \|\theta - \theta_0\|_{\Gamma}^2 \right\} \quad (\text{B.1})$$

$$= \arg \min_{\theta} \left\{ (b - f(\theta))^T (\mathbf{C} + \Sigma)^{-1} (b - f(\theta)) + \lambda^2 (\theta - \theta_0)^T \Gamma^{-1} (\theta - \theta_0) \right\}, \quad (\text{B.2})$$

where λ is a regularization parameter.

Using the Cholesky factorizations:

$$(\mathbf{C} + \Sigma)^{-1} = R^T R \quad (\text{B.3})$$

$$\Gamma^{-1} = M^T M, \quad (\text{B.4})$$

we can rewrite the problem as

$$\arg \min_{\theta} \left\{ \|R(b - f(\theta))\|_2^2 + \lambda^2 \|M(\theta - \theta_0)\|_2^2 \right\} \quad (\text{B.5})$$

$$= \arg \min_{\theta} \left\{ \left\| \begin{bmatrix} Rf(\theta) \\ \lambda M\theta \end{bmatrix} - \begin{bmatrix} Rb \\ \lambda M\theta_0 \end{bmatrix} \right\|_2^2 \right\} \quad (\text{B.6})$$

Define the following quantities:

$$F(\theta) = \begin{bmatrix} Rf(\theta) \\ \lambda M\theta \end{bmatrix} \quad (\text{B.7})$$

$$r = \begin{bmatrix} Rb \\ \lambda M\theta_0 \end{bmatrix} \quad (\text{B.8})$$

$$g(\theta) = \|F(\theta) - r\|_2^2, \quad (\text{B.9})$$

and so the minimization problem is given by $\arg \min_{\theta} \{g(\theta)\}$.

Linearize $F(\theta) = F(\theta_c + \vartheta) \approx F(\theta_c) + \partial F(\theta_c)\vartheta$, where $\theta = \theta_c + \vartheta$. Then

$$g(\theta) \approx \|F(\theta_c) + \partial F(\theta_c)\vartheta - r\|_2^2 = \|J\vartheta - y\|_2^2 \quad (\text{B.10})$$

where we have defined $y = r - F(\theta_c)$ and $J = \partial F(\theta_c)$. Expanding the 2-norm, we have:

$$g(\theta) \approx (J\vartheta - y)^T(J\vartheta - y) \quad (\text{B.11})$$

$$= (\vartheta^T J^T - y^T)(J\vartheta - y) \quad (\text{B.12})$$

$$= \vartheta^T J^T J\vartheta - \vartheta^T J^T y - y^T J\vartheta + \|y\|_2^2 \quad (\text{B.13})$$

Since y and ϑ are column vectors and J is a matrix, then $y^T J\vartheta = (y^T J\vartheta)^T = \vartheta^T J^T y$ and so:

$$g(\theta) \approx \vartheta^T J^T J\vartheta - 2\vartheta^T J^T y + \|y\|_2^2 \quad (\text{B.14})$$

To minimize this function, differentia with respect to ϑ and set to zero:

$$\nabla_{\vartheta} g(\theta) = 0 \quad (\text{B.15})$$

$$\Rightarrow \nabla_{\vartheta} (\vartheta^T J^T J\vartheta - 2\vartheta^T J^T y + \|y\|_2^2) = 0 \quad (\text{B.16})$$

$$\Rightarrow J^T J\vartheta + \vartheta^T J^T J - 2J^T y = 0 \quad (\text{B.17})$$

$$\Rightarrow J^T J\vartheta - J^T y = 0 \quad (\text{B.18})$$

$$\therefore \vartheta = (J^T J)^{-1} J^T y \quad (\text{B.19})$$

where we used the fact that $J^T J\vartheta = (J^T J\vartheta)^T = \vartheta^T J^T J$.

To evaluate these components, recall that $J = \partial F(\theta_c)$, and so we can write:

$$J = \begin{bmatrix} \partial(Rf(\theta_c)) \\ \partial(\lambda M\theta_c) \end{bmatrix} = \begin{bmatrix} R\partial f(\theta_c) \\ \lambda M \end{bmatrix} \quad (\text{B.20})$$

$$J^T J = \begin{bmatrix} \partial f(\theta_c)^T R^T & \lambda M^T \end{bmatrix} \begin{bmatrix} R\partial f(\theta_c) \\ \lambda M \end{bmatrix} \quad (\text{B.21})$$

$$= \partial f(\theta_c)^T R^T R\partial f(\theta_c) + \lambda^2 M^T M \quad (\text{B.22})$$

$$= \partial f(\theta_c)^T (C + \Sigma)^{-1} \partial f(\theta_c) + \lambda^2 \Gamma^{-1} \quad (\text{B.23})$$

This leads to the complete solution at the next step

$$\theta = \theta_c + \vartheta. \quad (\text{B.24})$$

For computational efficiency, it is desirable to ensure that every iteration of the algorithm contributes to the minimization of g , that is, $g(\theta) < g(\theta_c)$. To ensure that the step ϑ does not overstep the minimum, a technique known as *backtracking* may be employed, where (B.24) is replaced by

$$\theta = \theta_c + s\vartheta, \quad (\text{B.25})$$

where $s \in (0, 1]$ is chosen by interpolation such that $g(\theta) < g(\theta_c)$. This ensures that the next step in the iteration is guaranteed to be lower than the previous, and so the algorithm always works to minimize g .

B.1 Wiener Filtering

Tikhonov regularization is closely related to Wiener filtering, and it may be useful in terms of computational efficiency for some applications to reformulate the solution in this way.

Starting by expanding the solution (B.24) without backtracking:

$$\theta = \theta_c + \vartheta \quad (\text{B.26})$$

$$= \theta_c + \left(\partial f(\theta_c)^T (\mathbf{C} + \Sigma)^{-1} \partial f(\theta_c) + \lambda^2 \Gamma^{-1} \right)^{-1} \begin{bmatrix} \partial f(\theta_c)^T R^T & \lambda M^T \end{bmatrix} \begin{bmatrix} Rb - Rf(\theta_c) \\ \lambda M\theta_0 - \lambda M\theta_c \end{bmatrix} \quad (\text{B.27})$$

$$= \theta_c + \left(\partial f(\theta_c)^T (\mathbf{C} + \Sigma)^{-1} \partial f(\theta_c) + \lambda^2 \Gamma^{-1} \right)^{-1} \partial f(\theta_c)^T (\mathbf{C} + \Sigma)^{-1} (b - f(\theta_c)) \quad (\text{B.28})$$

$$+ \left(\partial f(\theta_c)^T (\mathbf{C} + \Sigma)^{-1} \partial f(\theta_c) + \lambda^2 \Gamma^{-1} \right)^{-1} \lambda^2 \Gamma^{-1} (\theta_0 - \theta_c)$$

$$= \theta_c + \left(\mathcal{J}^T (\mathbf{C} + \Sigma)^{-1} \mathcal{J} + \lambda^2 \Gamma^{-1} \right)^{-1} \mathcal{J}^T (\mathbf{C} + \Sigma)^{-1} (b - f(\theta_c)) \quad (\text{B.29})$$

$$+ \left(\mathcal{J}^T (\mathbf{C} + \Sigma)^{-1} \mathcal{J} + \lambda^2 \Gamma^{-1} \right)^{-1} \lambda^2 \Gamma^{-1} (\theta_0 - \theta_c)$$

where we defined $\mathcal{J} = \partial f(\theta_c)$.

The first term in (B.29) can be rewritten in the form

$$\left(A^T \mathcal{C}^{-1} A + \lambda^2 \Gamma^{-1} \right)^{-1} A^T \mathcal{C}^{-1} z = \Gamma A^T (A \Gamma A^T + \lambda^2 \mathcal{C})^{-1} z, \quad (\text{B.30})$$

where the left side is referred to as the *Tikhonov regularized* solution and the term on the right is referred to as the *Wiener filtered* solution.

It is straightforward although involved to prove this identity using the singular value decomposition [2], and that the equivalent form of the solution follows

$$\theta = \theta_c + \Gamma \mathcal{J}^T \left(\mathcal{J} \Gamma \mathcal{J}^T + \lambda^2 (\mathbf{C} + \Sigma) \right)^{-1} (b - f(\theta_c - \theta_0)). \quad (\text{B.31})$$

Appendix C

Implementation details in EIDORS

EIDORS simplifies the construction of meshes, and the calculation of the Jacobians. In this work, its MATLAB implementation is primarily used to solve the forward problem, and produce the figures shown. Nonetheless, EIDORS introduces several points of difficulty. In particular, forward calculations are performed based on units that are assumed to be scaled to the unit disk. By assuming that the given values correspond to a circular distribution of radius $r = 0.2$ m, it is necessary to scale the values appropriately in the calculation. For clarity, this procedure is outlined here.

Define a scaling parameter $s = 5$, such that scaling to the unit disk of radius $r' = 1$ is defined as

$$r \rightarrow r' = sr = 1.$$

In solving the forward problem, the appropriate quantities are scaled as follows

$$\begin{aligned}\theta &\rightarrow \theta' = \theta/s \\ \sigma_0 &\rightarrow \sigma'_0 = \sigma_0/s \\ z &\rightarrow z' = s^2z.\end{aligned}$$

All other parameters remain unscaled. After the algorithms have run to completion, the results are scaled back appropriately to the $r = 0.2$ m disk for analysis.

References

- [1] Arridge, S.R., Kaipio, J.P., Kolehmainen, V., Schweiger, M., Somersalo, E., Tarvainen, T., Vauhkonen, M.: Approximation errors and model reduction with an application in optical diffusion tomography. *Inverse Problems* **22**(1), 175 (2006). URL <http://stacks.iop.org/0266-5611/22/i=1/a=010>
- [2] Calvetti, D.: Introduction to Bayesian scientific computing ten lectures on subjective computing. Springer Science+Business Media, New York (2007)
- [3] Calvetti, D., Ernst, O., Somersalo, E.: Dynamic updating of numerical model discrepancy using sequential sampling. Submitted (2014)
- [4] Calvetti, D., McGivney, D., Somersalo, E.: Left and right preconditioning for electrical impedance tomography with structural information. *Inverse Problems* **28**(5), 055,015 (2012). URL <http://stacks.iop.org/0266-5611/28/i=5/a=055015>
- [5] Calvetti, D., Somersalo, E.: Statistical elimination of boundary artefacts in image deblurring. *Inverse Problems* **21**(5), 1697 (2005). URL <http://stacks.iop.org/0266-5611/21/i=5/a=012>
- [6] Cheney, M., Isaacson, D., Newell, J.: Electrical impedance tomography. *SIAM Review* **41**(1), 85–101 (1999). DOI 10.1137/S0036144598333613. URL <http://epubs.siam.org/doi/abs/10.1137/S0036144598333613>
- [7] Cheng, K.S., Isaacson, D., Newell, J.C., Gisser, D.: Electrode models for electric current computed tomography. *Biomedical Engineering, IEEE Transactions on* **36**(9), 918–924 (1989). DOI 10.1109/10.35300
- [8] Efron, B.: Bootstrap methods: Another look at the jackknife. *The Annals of Statistics* **7**(1), 1–26 (1979). DOI 10.1214/aos/1176344552. URL <http://dx.doi.org/10.1214/aos/1176344552>
- [9] Efron, B.: An introduction to the bootstrap. Chapman & Hall, New York (1994)
- [10] Evensen, G.: Sequential data assimilation with a nonlinear quasi-geostrophic model using monte carlo methods to forecast error statistics. *Journal of Geophysical Research: Oceans* **99**(C5), 10,143–10,162 (1994). DOI 10.1029/94JC00572. URL <http://dx.doi.org/10.1029/94JC00572>

- [11] Hansen, P.C.: Rank-Deficient and Discrete Ill-Posed Problems. Society for Industrial and Applied Mathematics (1998). DOI 10.1137/1.9780898719697. URL <http://epubs.siam.org/doi/abs/10.1137/1.9780898719697>
- [12] Heino, J., Somersalo, E.: A modelling error approach for the estimation of optical absorption in the presence of anisotropies. *Physics in Medicine and Biology* **49**(20), 4785 (2004). URL <http://stacks.iop.org/0031-9155/49/i=20/a=009>
- [13] Heino, J., Somersalo, E., Kaipio, J.P.: Compensation for geometric mismodelling by anisotropies in optical tomography. *Opt. Express* **13**(1), 296–308 (2005). DOI 10.1364/OPEX.13.000296. URL <http://www.opticsexpress.org/abstract.cfm?URI=oe-13-1-296>
- [14] Higdon, D., Gattiker, J., Williams, B., Rightley, M.: Computer model calibration using high-dimensional output. *Journal of the American Statistical Association* **103**(482), 570–583 (2008). DOI 10.1198/016214507000000888. URL <http://dx.doi.org/10.1198/016214507000000888>
- [15] Iglesias, M.A., Law, K.J.H., Stuart, A.M.: Ensemble kalman methods for inverse problems. *Inverse Problems* **29**(4), 045,001 (2013). URL <http://stacks.iop.org/0266-5611/29/i=4/a=045001>
- [16] Kaipio, J.: *Statistical and computational inverse problems*. Springer, New York (2005)
- [17] Kaipio, J., Somersalo, E.: Statistical inverse problems: Discretization, model reduction and inverse crimes. *Journal of Computational and Applied Mathematics* **198**(2), 493 – 504 (2007). DOI <http://dx.doi.org/10.1016/j.cam.2005.09.027>. URL <http://www.sciencedirect.com/science/article/pii/S0377042705007296>. Special Issue: Applied Computational Inverse Problems
- [18] Kennedy, M.C., O’Hagan, A.: Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **63**(3), 425–464 (2001). DOI 10.1111/1467-9868.00294. URL <http://dx.doi.org/10.1111/1467-9868.00294>
- [19] Lee, E., Ts, M.E., Seo, J.K., Woo, E.J.: Breast eit using a new projected image reconstruction method with multi-frequency measurements. *Physiological Measurement* **33**(5), 751 (2012). URL <http://stacks.iop.org/0967-3334/33/i=5/a=751>
- [20] Lipponen, A., Seppänen, A., Kaipio, J.P.: Nonstationary approximation error approach to imaging of three-dimensional pipe flow: experimental evaluation. *Measurement Science and Technology* **22**(10), 104,013 (2011). URL <http://stacks.iop.org/0957-0233/22/i=10/a=104013>
- [21] Nissinen, A., Kolehmainen, V., Kaipio, J.: Compensation of modelling errors due to unknown domain boundary in electrical impedance tomography. *Medical Imaging, IEEE Transactions on* **30**(2), 231–242 (2011). DOI 10.1109/TMI.2010.2073716

- [22] Nissinen, A., Kolehmainen, V., Kaipio, J.P.: Reconstruction of domain boundary and conductivity in electrical impedance tomography using the approximation error approach. *International Journal for Uncertainty Quantification* **1**(3), 203–222 (2011)
- [23] Oliver, D.: *Inverse theory for petroleum reservoir characterization and history matching*. Cambridge University Press, Cambridge New York (2008)
- [24] Oliver, D., Chen, Y.: Recent progress on reservoir history matching: a review. *Computational Geosciences* **15**(1), 185–221 (2011). DOI 10.1007/s10596-010-9194-2. URL <http://dx.doi.org/10.1007/s10596-010-9194-2>
- [25] Persson, P., Strang, G.: A simple mesh generator in matlab. *SIAM Review* **46**, 2004 (2004)
- [26] Somersalo, E., Cheney, M., Isaacson, D.: Existence and uniqueness for electrode models for electric current computed tomography. *SIAM Journal on Applied Mathematics* **52**(4), 1023–1040 (1992). DOI 10.1137/0152060. URL <http://epubs.siam.org/doi/abs/10.1137/0152060>
- [27] Tarvainen, T., Kolehmainen, V., Pulkkinen, A., Vauhkonen, M., Schweiger, M., Arridge, S.R., Kaipio, J.P.: An approximation error approach for compensating for modelling errors between the radiative transfer equation and the diffusion approximation in diffuse optical tomography. *Inverse Problems* **26**(1), 015,005 (2010). URL <http://stacks.iop.org/0266-5611/26/i=1/a=015005>
- [28] Vauhkonen, M., Lionheart, W.R.B., Heikkinen, L.M., Vauhkonen, P.J., Kaipio, J.P.: A matlab package for the eidors project to reconstruct two-dimensional eit images. *Physiological Measurement* **22**(1), 107 (2001). URL <http://stacks.iop.org/0967-3334/22/i=1/a=314>