

Statistical inversion and Monte Carlo sampling methods in electrical impedance tomography

Jari P Kaipio†§, Ville Kolehmainen†, Erkki Somersalo‡ and Marko Vauhkonen†

† Department of Applied Physics, University of Kuopio, PO Box 1627, 70211, Kuopio, Finland
‡ Helsinki University of Technology, Institute of Mathematics, PO Box 1000, FIN-02015, Finland

E-mail: kaipio@venda.uku.fi

Received 19 April 2000, in final form 15 June 2000

Abstract. This paper discusses the electrical impedance tomography (EIT) problem: electric currents are injected into a body with unknown electromagnetic properties through a set of contact electrodes. The corresponding voltages that are needed to maintain these currents are measured. The objective is to estimate the unknown resistivity, or more generally the impedivity distribution of the body based on this information. The most commonly used method to tackle this problem in practice is to use gradient-based local linearizations. We give a proof for the differentiability of the electrode boundary data with respect to the resistivity distribution and the contact impedances. Due to the ill-posedness of the problem, regularization has to be employed. In this paper, we consider the EIT problem in the framework of Bayesian statistics, where the inverse problem is recast into a form of statistical inference. The problem is to estimate the posterior distribution of the unknown parameters conditioned on measurement data. From the posterior density, various estimates for the resistivity distribution can be calculated as well as *a posteriori* uncertainties. The search of the maximum *a posteriori* estimate is typically an optimization problem, while the conditional expectation is computed by integrating the variable with respect to the posterior probability distribution. In practice, especially when the dimension of the parameter space is large, this integration must be done by Monte Carlo methods such as the Markov chain Monte Carlo (MCMC) integration. These methods can also be used for calculation of *a posteriori* uncertainties for the estimators. In this paper, we concentrate on MCMC integration methods. In particular, we demonstrate by numerical examples the statistical approach when the prior densities are non-differentiable, such as the prior penalizing the total variation or the L^1 norm of the resistivity.

1. Introduction

Consider the electrical impedance tomography (EIT) problem: given a body with unknown electromagnetic properties, the goal is to estimate the unknown resistivity distribution in the body from a finite number of current/voltage measurements on the surface of the body. In this work, the measurement configuration can be described as follows: a set of contact electrodes are attached on the body surface and different electric currents are injected into the body. The corresponding voltages needed to maintain these currents are then recorded. The current/voltage pairs constitute the data.

The potential applications of this imaging modality are numerous. The medical applications of EIT vary from the detection of cancerous tumours from breast tissue [55, 59] to monitoring of pulmonary or gastric functions [5, 20, 42]. For a review on the medical

§ Corresponding author.

applications see [13]. In industry EIT has applications such as monitoring of industrial processes [12, 37, 49, 62, 78] and non-destructive testing and evaluation of materials [2, 21, 60]. For a recent review on EIT see also [8].

This review concerns statistical inversion methods applied to the EIT problem. By statistical inversion methods, we refer in particular to the Bayesian approach, recasting the traditional inverse problem in the form of statistical inference from the distribution of the unknown parameters.

The paper is organized as follows. In section 2, we consider the mathematical model of both the forward problem with a particular emphasis on the computational aspects and the conventional formulation of the inverse problem. We discuss briefly the linearized inverse problem, iterative gradient-based solvers and differentiability issues. A new result, to the knowledge of the authors, concerning the Fréchet differentiability of the forward electrode model is proved. Section 3 is devoted to the description of the general idea of statistical inversion methods, i.e. Bayesian formulation of the problem, and results concerning the Markov chain Monte Carlo (MCMC) iteration schemes for sampling from probability distributions. In section 4, we formulate the EIT problem as a Bayesian problem of statistical inference and apply the MCMC algorithm with various prior distributions. In particular, we demonstrate the use of the statistical approach with non-differentiable prior distributions that lead to computational difficulties when traditional optimization-based iterative inversion methods are employed. Finally, in section 5 we discuss briefly some related questions and problems for further study. Two technical appendices are included at the end of the paper.

2. EIT problem

We start by discussing the mathematical model of the forward problem of determining the electrode voltages on the surface of the body when the injected currents and the electromagnetic properties of the body are known. The implementation of the forward problem with a finite element approximation is also discussed. Further, the inverse problem and its linearized version is then discussed in detail.

2.1. Forward problem

Let $\Omega \subset \mathbb{R}^n$, $n = 2, 3$, be a bounded domain with a connected complement. We assume that Ω has a smooth boundary. Here, Ω represents the body with known electromagnetic properties. We consider time-harmonic electromagnetic fields in Ω with low frequencies. In the quasi-static approximation, the fields can be described in terms of a scalar voltage potential u satisfying the equation

$$\nabla \cdot \sigma \nabla u = 0 \quad (2.1)$$

in Ω . Within this approximation, the function σ is complex valued and describes the admittivity, the inverse of the impedivity, of the body. In this paper, we restrict ourselves to the case where the admittivity is real and positive, describing the conductivity of the body, i.e. $\sigma : \Omega \rightarrow \mathbb{R}_+$. Physically, this corresponds to the static measurement. An extension to complex admittivities is straightforward.

The following definition fixes the admissible class of conductivities considered in this paper.

Definition 2.1. A conductivity distribution $\sigma : \Omega \rightarrow \mathbb{R}_+$ is in the admissible class of conductivities, denoted by $\mathcal{A} = \mathcal{A}(\Omega)$, if the following conditions are satisfied:

(1) For some $N \geq 1$, there is a family $\{\Omega_j\}_{j=1}^N$ of open disjoint sets, $\Omega_j \subset \Omega$, having piecewise smooth boundaries and for which

$$\bar{\Omega} = \bigcup_{j=1}^N \bar{\Omega}_j.$$

Furthermore, we require that $\sigma|_{\Omega_j} \in C(\bar{\Omega}_j)$, $1 \leq j \leq N$, i.e., σ restricted to each subset Ω_j allows a continuous extension up to the boundary of the subset.

(2) For some constants c and C ,

$$0 < c \leq \sigma(x) \leq C < \infty.$$

In medical applications, which is our major interest, the subsets Ω_j in the forward problem may represent the organs. In the inverse problem, the set of admissible conductivities provides a natural discretization basis.

Due to the possible discontinuities of $\sigma \in \mathcal{A}$, the equation (2.1) must be interpreted in the weak sense, discussed in detail below.

To describe the current injection and voltage measurements on the surface of the body, we define a set of surface patches $e_\ell \subset \partial\Omega$, $1 \leq \ell \leq L$, as a mathematical model of the contact electrodes. The electrodes are strictly disjoint, i.e. $\bar{e}_\ell \cap \bar{e}_k = \emptyset$ for $\ell \neq k$. If $\Omega \in \mathbb{R}^2$, the electrodes are disjoint intervals of the boundary, and in the case $\Omega \in \mathbb{R}^3$, they are sets with a piecewise smooth simple boundary curve on $\partial\Omega$. Let I_ℓ be the electric current injected through the electrode e_ℓ . We call the vector $I = (I_1, \dots, I_L)^T \in \mathbb{R}^L$ a current pattern if it satisfies the charge conservation condition

$$\sum_{\ell=1}^L I_\ell = 0. \tag{2.2}$$

Let U_ℓ denote the voltage on the ℓ th electrode, the ground voltage being chosen so that

$$\sum_{\ell=1}^L U_\ell = 0. \tag{2.3}$$

The vector $U = (U_1, \dots, U_L)^T \in \mathbb{R}^L$ is called a voltage vector. In terms of the current patterns and voltages, the appropriate boundary condition for the electric potential is given as

$$\int_{e_\ell} \sigma \frac{\partial u}{\partial n} dS = I_\ell, \quad 1 \leq \ell \leq L, \tag{2.4}$$

$$\sigma \frac{\partial u}{\partial n} \Big|_{\partial\Omega \setminus \cup e_\ell} = 0, \tag{2.5}$$

$$\left(u + z_\ell \sigma \frac{\partial u}{\partial n} \right) \Big|_{e_\ell} = U_\ell, \quad 1 \leq \ell \leq L. \tag{2.6}$$

Here, the numbers z_ℓ are the presumably known contact impedances between the electrodes and the body. We use the notation $z = (z_1, \dots, z_L)^T$ in what follows. For simplicity, we assume that the contact impedances are real. Note that in the forward problem, only the current patterns on the boundary are specified. However, conditions (2.4) and (2.5) alone are not sufficient to uniquely determine the potential u , but one needs to require $u + z_\ell \partial u / \partial n$ to be constant on e_ℓ , the constant being U_ℓ . Finding these voltages is part of the forward problem.

The electrode model described above was originally suggested in [9]. The well-posedness and unique solvability of the direct problem was discussed in the articles [61] and [68]. The

following proposition was proved in [68]. It serves also as a tool for further discussion in this section. In the following, we use the notation

$$\mathbb{H} = H^1(\Omega) \oplus \mathbb{R}^L, \tag{2.7}$$

where $H^1(\Omega)$ is the L^2 -based Sobolev space. Further, we denote

$$\dot{\mathbb{H}} = \mathbb{H}/\mathbb{R} \tag{2.8}$$

equipped with the quotient norm,

$$\|(u, U)\|_{\dot{\mathbb{H}}} = \inf_{c \in \mathbb{R}} \|(u - c, U - c)\|_{\mathbb{H}}.$$

Thus, $(u, U) \in \mathbb{H}$ and $(v, V) \in \mathbb{H}$ are equivalent in $\dot{\mathbb{H}}$ if

$$u - v = U_1 - V_1 = \dots = U_L - V_L = \text{constant}. \tag{2.9}$$

With these notations, the following proposition fixes the notion of the weak solution of the electrode model.

Proposition 2.2. *Let $\sigma \in \mathcal{A}(\Omega)$. The problem (2.1), (2.4)–(2.6) has a unique weak solution $(u, U) \in \dot{\mathbb{H}}$ in the following sense. There is a unique $(u, U) \in \dot{\mathbb{H}}$ satisfying the equation*

$$\mathcal{B}_{\sigma,z}((u, U), (v, V)) = \sum_{\ell=1}^L I_\ell V_\ell \tag{2.10}$$

for all $(v, V) \in \dot{\mathbb{H}}$, where the quadratic form $\mathcal{B}_{\sigma,z}$ is given as

$$\mathcal{B}_{\sigma,z}((u, U), (v, V)) = \int_{\Omega} \sigma \nabla u \cdot \nabla v \, dx + \sum_{\ell=1}^L \frac{1}{z_\ell} \int_{e_\ell} (u - U_\ell)(v - V_\ell) \, dS. \tag{2.11}$$

Furthermore, the quadratic form is coercive in $\dot{\mathbb{H}}$, i.e., we have the inequalities

$$\alpha_0 \|(u, U)\|_{\dot{\mathbb{H}}}^2 \leq \mathcal{B}_{\sigma,z}((u, U), (u, U)) \leq \alpha_1 \|(u, U)\|_{\dot{\mathbb{H}}}^2 \tag{2.12}$$

for some constants $0 < \alpha_0 \leq \alpha_1 < \infty$.

The proof is based on the observation that the space $\dot{\mathbb{H}}$ can be equipped with an equivalent norm defined as

$$\|(u, U)\|_*^2 = \int_{\Omega} |\nabla u|^2 \, dx + \sum_{\ell=1}^L \int_{e_\ell} |u - U_\ell|^2 \, dS,$$

and the quadratic form $\mathcal{B}_{\sigma,z}$ satisfies the estimate

$$c \|(u, U)\|_*^2 \leq |\mathcal{B}_{\sigma,z}((u, U), (u, U))| \leq C \|(u, U)\|_*^2 \tag{2.13}$$

with some constants $0 < c \leq C \leq \infty$ depending on σ and z . With this estimate, the result is a direct consequence of the Lax–Milgram lemma. The details are not repeated here.

We define the *resistance matrix* of the complete model as follows. If $\sigma \in \mathcal{A}(\Omega)$ and $z \in \mathbb{R}^L, z_\ell > 0$ are given, the resistance matrix $R(\sigma, z)$ is the $L \times L$ matrix with the property

$$U = R(\sigma, z)I,$$

where I is any current pattern that satisfies (2.2).

2.2. Numerical implementation of the forward problem

The numerical implementation of the forward problem based on the electrode model of the previous subsection has been discussed in detail, for example in [75] and [35]. A brief summary of the details is given here.

Let u^h denote a finite-dimensional approximation of the solution u satisfying the equations (2.1), (2.4)–(2.6) in the weak sense specified by proposition 2.2, corresponding to a given current pattern I . The approximation u^h is in a finite-dimensional subspace $Q^h = \text{span}\{\varphi_i | 1 \leq i \leq N_n\}$ of dimension N_n ,

$$u^h = \sum_{i=1}^{N_n} \alpha_i \varphi_i, \tag{2.14}$$

where the functions φ_i are basis functions related to a finite element mesh, the number of nodes being N_n . The finite element meshing of Ω is in practice chosen to conform with the underlying partition defined by $\sigma \in \mathcal{A}$. As usual, the super-index h indicates the mesh size. In practice, we use linear or quadratic basis functions. In order that the condition (2.3) be satisfied, the voltage vector is represented as

$$U^h = \sum_{j=1}^{L-1} \beta_j n_j, \tag{2.15}$$

where the vectors $n_j \in \mathbb{R}^L$ are chosen as $n_1 = (1, -1, 0, \dots, 0)^T$, $n_2 = (1, 0, -1, \dots, 0)^T, \dots, n_{L-1} = (1, 0, \dots, -1)^T$.

Using the theory of finite elements [4], a substitution of the approximations (2.14) and (2.15) to the weak form (2.10) yields a matrix equation

$$Ab = f, \tag{2.16}$$

where $b = (\alpha, \beta)^T \in \mathbb{R}^{N_n+L-1}$ and the data vector f is

$$f = \left(\begin{array}{c} \mathbf{0} \\ \sum_{\ell=1}^L I_\ell (n_j)_\ell \end{array} \right) = \left(\begin{array}{c} \mathbf{0} \\ C^T I \end{array} \right), \tag{2.17}$$

where $\mathbf{0} = (0, \dots, 0)^T \in \mathbb{R}^{N_n}$ and $C \in \mathbb{R}^{L \times (L-1)}$ is the sparse matrix given as

$$C = \begin{pmatrix} 1 & 1 & 1 & \dots & 1 \\ -1 & 0 & \dots & & 0 \\ 0 & -1 & 0 & \dots & \vdots \\ \vdots & & \ddots & & \\ 0 & \dots & & \ddots & -1 \end{pmatrix}. \tag{2.18}$$

The stiffness matrix $A \in \mathbb{R}^{(N_n+L-1) \times (N_n+L-1)}$ is the sparse block matrix of the form

$$A = \begin{pmatrix} B & C \\ C^T & G \end{pmatrix} \tag{2.19}$$

with

$$B_{i,j} = \int_{\Omega} \sigma \nabla \varphi_i \cdot \nabla \varphi_j \, dx + \sum_{\ell=1}^L \frac{1}{z_\ell} \int_{e_\ell} \varphi_i \varphi_j \, dS, \quad 1 \leq i, \quad j \leq N_n,$$

$$C_{i,j} = -\left(\frac{1}{z_1} \int_{e_1} \varphi_i \, dS - \frac{1}{z_{j+1}} \int_{e_{j+1}} \varphi_i \, dS \right), \quad 1 \leq i \leq N_n, \quad 1 \leq j \leq L-1$$

$$G_{i,j} = \sum_{\ell=1}^L \frac{1}{z_\ell} \int_{e_\ell} (n_i)_\ell (n_j)_\ell \, dS$$

$$= \begin{cases} \frac{|e_1|}{z_1}, & i \neq j \\ \frac{|e_1|}{z_1} + \frac{|e_{j+1}|}{z_{j+1}}, & i = j, \end{cases} \quad 1 \leq i, \quad j \leq L - 1.$$

By solving equation (2.16) as $b = A^{-1}f$, an approximate solution for the forward problem is obtained. The N_n first coefficients in b give the solution u^h in the nodes and the last $L - 1$ coefficients give the referenced voltages $\beta = (\beta_1, \dots, \beta_{L-1})^T$ on the electrodes. The potentials U_ℓ on the electrodes are calculated with the aid of (2.15) to yield

$$U^h = C\beta. \tag{2.20}$$

The relation between the injected currents and the computed voltages on the electrodes can be written in the form

$$U^h = C\beta = C\tilde{R}^h(\sigma, z)\mathcal{C}^T I = R^h(\sigma, z)I, \tag{2.21}$$

where $\tilde{R}^h(\sigma, z) \in \mathbb{R}^{(L-1) \times (L-1)}$ is a block $(A^{-1})_{i,j}$, $N_n + 1 \leq i, j \leq N_n + L - 1$ of the inverse of the matrix A . Thus the computed approximation of the resistance matrix in the finite element model is given as $R^h(\sigma, z) = C\tilde{R}(\sigma, z)\mathcal{C}^T \in \mathbb{R}^{L \times L}$.

2.3. Inverse problem

Within the framework of the electrode model of the previous section, the corresponding inverse problem is to estimate the conductivity distribution from a set of measured current/voltage pattern pairs. The estimation of the contact impedances z_ℓ of the electrodes can be included in the inverse problem. In the following discussion, we assume that the contact impedances are known.

Let $\{I^{(k)}\}_{k=1}^K \subset \mathbb{R}^L$, $K \leq L - 1$, be a set of linearly independent current patterns. These current patterns are applied through the electrodes into the body whose conductivity distribution is unknown. Let us denote by $\{V^{(k)}\}_{k=1}^K \subset \mathbb{R}^L$ the set of measured voltage patterns. The traditional formulation of the EIT inverse problem is to find an estimate $\hat{\sigma} \in \mathcal{A}(\Omega)$ of the true conductivity distribution $\sigma \in \mathcal{A}(\Omega)$ of the body such that the set $\{U^{(k)}\}_{k=1}^K \subset \mathbb{R}^L$ of computed voltage patterns are in some sense close to the measured voltages.

In the EIT problem as formulated here, the input data are the known current patterns and the measured output data are the corresponding voltage vectors. Thus, the information of the internal structure of the body is in the resistance matrix. Since the resistance matrix is more closely related to the resistivity distribution of the body, rather than to its inverse, the conductivity distribution, it is more natural to seek to estimate the resistivity. In the case of complex admittivity, the resistivity is replaced by complex impedivity. By the assumptions of the admissible class of conductivities, we may assume that the resistivity is in the same class, i.e.

$$\rho = \sigma^{-1} \in \mathcal{A}(\Omega).$$

We introduce here the following notations. If K linearly independent current patterns are used in the EIT measurement, we stack the measured voltages to a single vector and denote

$$V = (V_1^{(1)}, \dots, V_L^{(1)}, \dots, V_1^{(K)}, \dots, V_L^{(K)})^T \in \mathbb{R}^{KL}.$$

Similarly, we denote

$$U(\rho) = (U_1^{(1)}(\rho), \dots, U_L^{(1)}(\rho), \dots, U_1^{(K)}(\rho), \dots, U_L^{(K)}(\rho))^T \in \mathbb{R}^{KL},$$

where $U^{(k)}(\rho) = R(\rho^{-1}, z)I^{(k)}$.

2.4. Least squares methods, regularization and linearization

The most straightforward approach to attack the inverse problem is to seek to minimize the weighted least squares error, i.e., seeking ρ such that the functional

$$F(\rho) = \|U(\rho) - V\|_W^2 = \sum_{k=1}^K \sum_{\ell=1}^L w_{k,\ell} (U_\ell^{(k)}(\rho) - V_\ell^{(k)})^2 \tag{2.22}$$

attains its minimum. Here, $W = (w_{k,\ell})$ is a symmetric positive definite weight matrix. Due to the ill-posedness, the problem necessitates regularization. The classical way of implementing a regularizing side constraint is to use a Tikhonov-type regularization: instead of minimizing the least squares functional (2.22), one minimizes the regularized modification,

$$F_\alpha(\rho) = \|U(\rho) - V\|_W^2 + \alpha A(\rho), \tag{2.23}$$

where A is a properly chosen regularizing functional and $\alpha > 0$ is a regularization parameter.

The choice of the regularization functional $A(\rho)$ is often done *ad hoc*. We shall discuss specific choices in more detail in the next section. The goal of this paper is not to concentrate on least squares methods *per se*. However, few comments concerning the minimization process of the functional F_α are in order, since this issue touches also on the statistical inversion algorithms. What is more, the standard least squares algorithms provide a good reference for the statistical inversion methods discussed later.

The most commonly used algorithms for minimizing the regularized least squares functional (2.23) are iterative gradient-based optimization algorithms. We discuss these algorithms briefly.

Let $H_M \subset \mathcal{A}(\Omega)$ be an M -dimensional subspace of $\mathcal{A}(\Omega)$. If $H_M = \text{span}\{\eta_m | 1 \leq m \leq M\}$, we write $\rho = \sum_{m=1}^M \rho_m \eta_m \in H_M$ and identify ρ with the vector $(\rho_1, \dots, \rho_M)^T \in \mathbb{R}^M$.

Consider the restriction of the functional F_α in (2.23) to H_M . Starting from an initial guess $\rho^{(0)} \in \mathbb{R}^M$, the iterative gradient-based methods produce a sequence $\{\rho^{(j)}\} \subset \mathbb{R}^M$ of resistivity distributions that hopefully converge towards the (global) minimum point of the functional $F_\alpha(\rho)$. We consider first the choice of the initial guess.

Often, when no prior information of the internal structure of the body is available, the initial guess of the resistivity is chosen to be constant. This constant may be chosen so that the corresponding computed voltages are fitted to the measured voltages. In [36] and [69], the fitting was based on the observation that the quadratic form $\mathcal{B}_{\sigma,z}$ in (2.11) has the property

$$\mathcal{B}_{\sigma,z}(\lambda(u, U), (v, V)) = \mathcal{B}_{\lambda\sigma, \lambda^{-1}z}((u, U), (v, V)), \quad \lambda > 0,$$

implying that the resistance matrix $R(\sigma, z)$ satisfies

$$R(\lambda\sigma, \lambda^{-1}z) = \lambda^{-1}R(\sigma, z).$$

We consider now a model where both the resistivity and contact impedances are constants, $\rho = \rho_0(1, \dots, 1)^T$, $z = z_0(1, \dots, 1)^T$. We have

$$U^{(k)}(\rho_0) = R(\rho_0^{-1}, z_0)I^{(k)} = \rho_0 R(1, z_0/\rho_0)I^{(k)}.$$

We assume, for a moment, that the ratio $\xi = z_0/\rho_0$ is known. We can then compute the resistance matrix $R(1, \xi)$ and define the reference voltage vectors $\tilde{U}^{(k)} = R(1, \xi)I^{(k)}$, and further $\tilde{U} = (\tilde{U}_1^{(1)}, \dots, \tilde{U}_L^{(1)}, \dots, \tilde{U}_1^{(K)}, \dots, \tilde{U}_L^{(K)})^T$. The optimal background resistance in the weighted least squares sense is now found by minimizing the deviation

$$\|U(\rho_0) - V\|_W^2 = \|\rho_0 \tilde{U} - V\|_W^2,$$

yielding

$$\rho_0 = \frac{\tilde{U}^T W V}{\|\tilde{U}\|_W^2}. \tag{2.24}$$

The above procedure was based on the assumption that the ratio ξ is known. In general, this is not the case. However, as pointed out in [9], in medical applications there is an empirical value that gives a good approximation for this ratio. In general, the optimization of both ρ_0 and z_0 would require iterative methods.

In the following discussion the parenthesized super-index refers to the iteration number. Having the initial value $\rho^{(0)}$ fixed, the conventional Gauss–Newton-type iteration step $\rho^{(j)} \rightarrow \rho^{(j+1)}$, $j \geq 0$ is computed now as

$$\rho^{(j+1)} = \rho^{(j)} - \lambda_s^{(j)} (H_\alpha^{(j)})^{-1} g^{(j)}, \tag{2.25}$$

where the Hessian $H_\alpha^{(j)} \in \mathbb{R}^{M \times M}$ is given by

$$H_\alpha^{(j)} = (DU(\rho^{(j)}))^T W (DU(\rho^{(j)})) + \frac{1}{2} \alpha D^2 A(\rho^{(j)}), \tag{2.26}$$

and the gradient $g^{(j)}$ by

$$g^{(j)} = (DU(\rho^{(j)}))^T W (U(\rho^{(j)}) - V) + \frac{1}{2} \alpha DA(\rho^{(j)}), \tag{2.27}$$

and $\lambda_s > 0$ is a relaxation parameter controlling the step size. Here, DU and DA denote the differentials of the maps $\rho \mapsto U(\rho)$ and $\rho \mapsto A(\rho)$, respectively, and $D^2 A = (\partial^2 A / \partial \rho_j \partial \rho_k)$.

This algorithm (and other gradient-based methods) rely on the differentiability of the map $\rho \mapsto R(\rho^{-1}, z)$ as well as on the differentiability of the regularizing functional A . Since to our knowledge the differentiability proof of the resistance matrix has not been presented in the literature (although it is stated without proof in [36]), we give the proof here. For a reference to results concerning the Fréchet differentiability and related inverse problems, see [10].

The differentiability assumption of the regularizing functional excludes many useful and interesting choices of the functional. This issue is discussed in later sections with examples. In the following theorem, we use the notation $z = 1/y = (1/y_1, \dots, 1/y_L) \in \mathbb{R}^L$, $y \in \mathbb{R}^L$ being the electrode contact admittance.

Theorem 2.3. *The mapping*

$$\mathcal{M} : \mathcal{A}(\Omega) \oplus \mathbb{R}^L \rightarrow \mathring{\mathbb{H}}, \quad (\sigma, y) \mapsto (u, U),$$

is Fréchet differentiable. The derivative $\mathcal{M}'(\sigma, y)$ satisfies the following equation. Let $h = (s, \eta) \in \mathcal{A} \oplus \mathbb{R}^L$, with $\sigma + s > 0$ and $y_\ell + \eta_\ell > 0$. Denoting $(w, W) = \mathcal{M}'(\sigma, y)h$, we have

$$B_{\sigma, \frac{1}{y}}((w, W), (v, V)) = - \int_D s \nabla u^0 \cdot \nabla v \, dx - \sum_{\ell=1}^L \eta_\ell \int_{e_\ell} (u^0 - U_\ell^0)(v - V_\ell) \, dS \tag{2.28}$$

for all $(v, V) \in \mathring{\mathbb{H}}$.

Above, the space $\mathcal{A} \oplus \mathbb{R}^L$ is equipped with the norm $\|(\sigma, y)\| = \|\sigma\|_\infty + \|y\|_\infty$. Especially, theorem 2.3 implies that the resistance matrix $(\sigma, y) \mapsto R(\sigma, 1/y)$ is Fréchet differentiable.

Proof. Let us denote $\mathcal{M}(\sigma + s, y + \eta) = (u, U)$ and $\mathcal{M}(\sigma, y) = (u^0, U^0)$. We need to prove the estimate

$$\|(u, U) - (u^0, U^0) - (w, W)\|_* \leq C(\|s\|_\infty + \|\eta\|_\infty)^2 c$$

with some constant C independent of s and η , where (w, W) is the solution of the equation (2.28). We have

$$B_{\sigma, 1/y}((u^0, U^0), (v, V)) = \sum_{\ell=1}^L I_\ell V_\ell = B_{\sigma_0+h, 1/(y+\eta)}((u, U), (v, V)),$$

or

$$\int_{\Omega} (\sigma + s) \nabla u \cdot \nabla v \, dx - \int_{\Omega} \sigma \nabla u^0 \cdot \nabla v \, dx + \sum_{\ell=1}^L (y_{\ell} + \eta) \int_{e_{\ell}} (u - U_{\ell})(v - V_{\ell}) \, dS - \sum_{\ell=1}^L y_{\ell} \int_{e_{\ell}} (u^0 - U_{\ell}^0)(v - V_{\ell}) \, dS = 0.$$

In terms of the quadratic form $B_{\sigma,1/y}$, we have

$$B_{\sigma,1/y}((u - u^0, U - U^0), (v, V)) = - \int_{\Omega} s \nabla u \cdot \nabla v \, dx - \sum_{\ell=1}^L \eta_{\ell} \int_{e_{\ell}} (u - U)(v - V) \, dS. \tag{2.29}$$

By subtracting the equation (2.28), we obtain

$$B_{\sigma,1/y}((u - u^0 - w, U - U^0 - W), (v, V)) = - \int_{\Omega} s \nabla (u - u^0) \cdot \nabla v \, dx - \sum_{\ell=1}^L \eta_{\ell} \int_{e_{\ell}} ((u - u^0) - (U_{\ell} - U_{\ell}^0))(v - V_{\ell}) \, dS.$$

Therefore, by the estimate (2.13), we have

$$\begin{aligned} & \| (U, u) - (U^0, u^0) - (w, W) \|_*^2 \\ & \leq C |B_{\sigma,1/y}((u - u^0 - w, U - U^0 - W), (u - u^0 - w, U - U^0 - W))| \\ & \leq C \left(\int_{\Omega} |s| |\nabla (u - u^0)| |\nabla (u - u^0 - w)| \, dx + \sum_{\ell=1}^L |\eta_{\ell}| \int_{e_{\ell}} |u - u^0 - (U_{\ell} - U_{\ell}^0)| |u - u^0 - w - (U_{\ell} - U_{\ell}^0 - W_{\ell})| \, dS \right) \\ & \leq C (\|s\|_{\infty} + \|\eta\|_{\infty}) \| (u, U) - (u^0, U^0) \|_* \| (u, U) - (u^0, U^0) - (w, W) \|_*. \end{aligned}$$

The claim thus follows if we show the estimate

$$\| (u, U) - (u^0, U^0) \|_* \leq C (\|s\|_{\infty} + \|\eta\|_{\infty}).$$

This inequality is a consequence of the equation (2.29), since

$$\begin{aligned} \| (u, U) - (u^0, U^0) \|_*^2 & \leq C |B_{\sigma,y}((u, U) - (u^0, U^0), (u, U) - (u^0, U^0))| \\ & \leq \int_{\Omega} |s| |\nabla u| |\nabla (u - u^0)| \, dx + \sum_{\ell=1}^L |\eta_{\ell}| \int_{e_{\ell}} |u - U_{\ell}| |u - u^0 - (U_{\ell} - U_{\ell}^0)| \, dS \\ & \leq C (\|s\|_{\infty} + \|\eta\|_{\infty}) \| (u, U) \|_* \| (u, U) - (u^0, U^0) \|_*, \end{aligned}$$

giving the desired estimate. The proof is complete. □

In practice, the full iterative Gauss–Newton algorithm is seldom used. Instead, one usually performs just one updating step $\rho^{(0)} \rightarrow \rho^{(1)}$ and then stops. It is our experience that in order to get a converging iterative algorithm, one usually has to decrease the relaxation parameter λ_s in formula (2.25) as the iteration marches on. Often, the reconstruction after several iterations does not differ significantly from the reconstruction after the first iteration step, especially when the initial estimate has been chosen to be the best matching constant resistivity. The numerical computation of the differential $DU(\rho)$ is explained in appendix B.

3. Statistical inversion

The philosophy behind the statistical inversion methods is to recast the inverse problem in the form of statistical *quest for information*. We have directly observable quantities and those that cannot be observed. These quantities depend on each other through a more or less well known model. The objective is to extract all the possible information of some of the variables based on all knowledge of the measurements, model and information obtained prior to the measurement. In the statistical inversion approach, all the variables included in the model are modelled as random variables. It can be thought that the randomness is related to our uncertainty of their true values, and this uncertainty is expressed in terms of probability distributions. In contrast to the traditional formulation of the inverse problem discussed above, the solution of the statistical inverse problem is the posterior probability distribution of the sought parameter conditioned on the measurement. The identification of inverse problems with statistical inference has been promoted, especially in the geophysical literature, see [46, 52, 54, 71, 72] and [53]. Below, a summary of the basic ideas is given with the assumption that the variables are finite-dimensional random vectors, although an extension to infinite-dimensional spaces is possible; see [45].

3.1. Bayesian model

We start by formulating the statistical inverse problem in the framework of Bayesian statistics. As a general reference on probability theory, see, e.g., [56] or [66]. As is customary in probability theory, random variables are denoted by capital letters and their values by lower case letters. Let (S, \mathcal{B}, P) denote a probability space, \mathcal{B} being the σ -algebra of measurable subsets of S and $P : \mathcal{B} \rightarrow [0, 1]$ a probability measure. Let

$$(X, N) : S \rightarrow \mathbb{R}^{n+k}, \quad V : S \rightarrow \mathbb{R}^m \quad (3.1)$$

be random vectors. In our formulation, the vector (X, N) represents all those quantities that cannot be directly measured while V represents a vector of observable quantities. We assume that $X \in \mathbb{R}^n$ represents those variables that we are primarily interested in while $N \in \mathbb{R}^k$ contains unknown but uninteresting variables such as the measurement noise or model parameters of which we have incomplete knowledge. These variables are assumed to be tied together through a model

$$V = F(X, N), \quad (3.2)$$

i.e., we assume that the model is complete in the sense that the vectors X and N determine the observable V uniquely. The function $F : \mathbb{R}^{n+k} \rightarrow \mathbb{R}^m$ is assumed here to be a known deterministic function representing a model of the measurement. We assume that the variables X and N take the values $X = x \in \mathbb{R}^n$ and $N = n \in \mathbb{R}^k$. Then, the probability distribution of the random variable V conditioned on $X = x$ and $N = n$ is formally given by

$$\pi(v|x, n) = \delta(v - F(x, n)),$$

where δ is the Dirac delta in \mathbb{R}^m . Let $\pi_{\text{pr}}(x, n)$ denote the prior probability density of the unknown vector (X, N) . Then the joint probability density of (X, N) and V can be written as

$$\pi(x, n, v) = \pi(v|x, n)\pi_{\text{pr}}(x, n) = \delta(v - F(x, n))\pi_{\text{pr}}(x, n).$$

Since we have arranged the variables so that N represents all the variables whose values are not of primary interest, we may integrate the variable n out and define the joint probability density of the variables X and V as a marginal distribution,

$$\pi(x, v) = \int_{\mathbb{R}^k} \pi(x, n, v) dn = \int_{\mathbb{R}^k} \delta(v - F(x, n))\pi_{\text{pr}}(x, n) dn. \quad (3.3)$$

Often in the applications, one considers a simple model where the variables X and N are independent. In terms of the probability densities, this can be written as

$$\pi_{\text{pr}}(x, n) = \pi_{\text{pr}}(x)\pi_{\text{noise}}(n),$$

where we have identified the variable N as noise. Furthermore, the posterior distribution attains a particularly simple form in the case where the noise is additive. Then the model equation (3.2) is written as

$$V = f(X) + N,$$

and the integration with respect to v in (3.3) can be performed. We have

$$\pi(x, v) = \int_{\mathbb{R}^k} \delta(v - f(x) - n)\pi_{\text{pr}}(x)\pi_{\text{noise}}(n) \, dn = \pi_{\text{pr}}(x)\pi_{\text{noise}}(v - f(x)). \tag{3.4}$$

The solution of the statistical inverse problem is defined as the *posterior distribution of X* , i.e. the conditional probability density $\pi_{\text{post}}(x) = \pi(x|v)$, given by the well known *Bayes formula*

$$\pi(x|v) = \frac{\pi(x, v)}{\int \pi(x, v) \, dx}. \tag{3.5}$$

Thus, the solution of the statistical inverse problem, in contrast to traditionally formulated ones, is not just an estimate of the unknown vector X but a probability distribution. We can also write $\pi(x, v) = \pi_{\text{pr}}(x)\pi_{\text{noise}}(v - f(x)) = \pi_{\text{pr}}(x)\pi(v|x)$. The term $\pi(v|x)$ is called *the likelihood density*.

Let us mention that the formulation above covers implicitly also the *hierarchical Bayes models*. More precisely, we assume that the prior density $\pi_{\text{pr}}(x)$ depends on model parameters $\alpha \in \mathbb{R}^\ell$ that are not exactly known, $\pi_{\text{pr}}(x) = \pi_{\text{pr}}(x, \alpha)$. The uncertainty concerning the value of α is expressed in hierarchical Bayes models by saying that α is distributed according to a presumably known *hyper-prior distribution* $\pi_{\text{h}}(\alpha)$. Merging the vector α into the model vector n renders the hierarchical model formally, similar to the discussed one.

3.2. Estimates from posterior distribution

We assume that we have solved an inverse problem in the statistical sense, i.e., we have determined the posterior probability distribution of the random variable X . This distribution can be used to calculate various estimates for X as well as *a posteriori* uncertainties for these estimates. Assuming that the probability distribution is given in terms of a probability density as above, a commonly used estimate based on this density is the (possibly non-unique) *maximum a posteriori (MAP) estimate*

$$x_{\text{MAP}} = \arg \max_x \pi(x|v). \tag{3.6}$$

The computation of the MAP estimate leads typically to an optimization problem such as the regularized least squares problem discussed in section 2.4. It is worth noting that the well known *maximum likelihood estimate* which amounts to the determination of the maximum of the likelihood density only, that is

$$x_{\text{ML}} = \operatorname{argmax}_x \pi(v|x),$$

corresponds to the solution of the non-regularized inverse problem. Thus in the case of ill-posed inverse problems, maximum likelihood estimates are unstable.

Another commonly used estimate that will be discussed in more detail here is the *conditional expectation*

$$x|_v = \int_{\mathbb{R}^n} x \pi(x|v) \, dx. \tag{3.7}$$

In contrast to the search for the maximum posterior estimate, the search for the conditional expectation is an integration problem.

We assume that $\tilde{x} \in \mathbb{R}^n$ is an estimate of the variable X . To get an idea of the reliability of this estimate, it is also desirable to obtain credibility limits for this estimate. Let us consider the j th component of the vector x . In the following, we use the notation

$$x_{-j} = (x_1, \dots, x_{j-1}, x_{j+1}, \dots, x_n)^T \in \mathbb{R}^{n-1},$$

i.e., x_{-j} is the vector x with the j th component deleted. The *conditional posterior probability density* of the single component X_j conditioned on \tilde{x}_{-j} is given by

$$\pi(x_j | \tilde{x}_{-j}, v) \sim \pi((\tilde{x}_1, \dots, \tilde{x}_{j-1}, x_j, \tilde{x}_{j+1}, \dots, \tilde{x}_n) | v). \quad (3.8)$$

The *marginal posterior probability density* of X_j is given as

$$\pi(x_j | v) \sim \int_{\mathbb{R}^{n-1}} \pi(x | v) dx_{-j}, \quad (3.9)$$

i.e., we integrate out the $n - 1$ variables from the posterior distribution. Both of these one-dimensional distributions are informative in the evaluation of the credibility intervals of the estimate \tilde{x}_j . However, there is a remarkable difference: if the components of the random variable $X - \tilde{x}$ are strongly correlated with respect to the posterior measure, it is clear that the conditional probability density of X_j given by (3.8) depends heavily on the values \tilde{x}_{-j} , while the marginal probability density is independent of them. Therefore, it is advisable to calculate the *conditional correlation matrix*

$$\text{corr}(x - \tilde{x} | v) = \int_{\mathbb{R}^n} (x - \tilde{x})(x - \tilde{x})^T \pi(x | v) dx \quad (3.10)$$

before jumping to conclusions based on the conditional distribution. Large off-diagonal components compared with the diagonal ones indicate large correlation and thus heavy dependence of the conditional posterior density on the corresponding components.

To summarize, the estimation of the random variable based on the posterior density as well as the estimation of the *a posteriori* uncertainties require integration of the posterior probability density. In the following sections, we discuss sampling-based integration techniques.

3.3. General idea of MCMC

The statistical formulation of the inverse problem gives a posterior probability density (or more generally, a posterior probability distribution) as a solution rather than single estimates. To obtain single estimates and information as to their credibility, one often faces the problem of integrating numerically with respect to the posterior probability distribution. It is clear that if the dimension of the parameter space \mathbb{R}^n is large, the use of numerical quadrature methods is out of the question. In this section we discuss the MCMC methods. The MCMC methods are, at least on the conceptual level, relatively simple algorithms to generate sample ensembles for Monte Carlo integration. The literature on MCMC methods is extensive. Our main references in this section are [65, 67, 74] and [24].

Let ν denote a probability measure over a state space that is chosen here to be \mathbb{R}^n and let f be a scalar or vector-valued measurable function on \mathbb{R}^n , $f \in L^1(\nu(dx))$. The objective is to calculate the integral of f by Monte Carlo integration. Hence, if $\{x^{(1)}, x^{(2)}, \dots, x^{(N)}\} \subset \mathbb{R}^n$ is a representative ensemble of samples distributed according to the distribution ν , we seek to approximate the integral of f by an ergodic average

$$\int_{\mathbb{R}^n} f(x) \nu(dx) \approx \frac{1}{N} \sum_{j=1}^N f(x^{(j)}). \quad (3.11)$$

The MCMC methods are just a systematic way of generating a sample ensemble such that (3.11) holds.

Let \mathcal{B} denote the Borel sets over \mathbb{R}^n . A mapping $P : \mathbb{R}^n \times \mathcal{B} \rightarrow [0, 1]$ is called a *probability transition kernel*, if

- (i) for each $B \in \mathcal{B}$, the mapping $\mathbb{R}^n \rightarrow [0, 1], x \mapsto P(x, B)$ is a measurable function;
- (ii) for each $x \in \mathbb{R}^n$, the mapping $\mathcal{B} \rightarrow [0, 1], B \mapsto P(x, B)$ is a probability distribution.

A *time-homogenous Markov chain* with the transition kernel P is a stochastic process $\{X^{(j)}\}_{j=1}^\infty$ with the properties

$$P(X^{(j+1)} \in A | X^{(1)}, \dots, X^{(j)}) = P(X^{(j+1)} \in A | X^{(j)}), \tag{3.12}$$

$$P(x, A) = P(X^{(j+1)} \in A | X^{(j)} = x). \tag{3.13}$$

More generally, we define

$$\begin{aligned} P^{(k)}(x, A) &= P(x^{(j+k)} \in A | x^{(j)} = x) \\ &= \int_{\mathbb{R}^n} P(y, A) P^{(k-1)}(x, dy), \quad k \geq 2 \end{aligned}$$

where $P^{(1)}(x, A) = P(x, A)$. In particular, if ν is the probability distribution of $X^{(j)}$, the distribution of $X^{(j+1)}$ is given by

$$\nu P(A) = \int_{\mathbb{R}^n} P(x, A) \nu(dx). \tag{3.14}$$

The measure ν is an *invariant measure* of $P(x, A)$ if

$$\nu P = \nu, \tag{3.15}$$

i.e., the distribution of the random variable after one transition step is the same as before the step.

Given a probability measure ν , the transition kernel P is called *irreducible* (with respect to ν) if for each $x \in \mathbb{R}^n$ and $A \in \mathcal{B}$ with $\nu(A) > 0$ there exists an integer k such that $P^{(k)}(x, A) > 0$. Thus, regardless of the starting point, the Markov chain enters with a positive probability any set of positive measures.

Let P be an irreducible kernel. We call P *periodic* if for some integer $m \geq 2$ there is a set of disjoint non-empty sets $\{E_1, \dots, E_m\} \subset \mathbb{R}^n$ such that for all $j = 1, \dots, m$ and all $x \in E_j$, $P(x, E_{j+1(\text{mod } m)}) = 1$. Otherwise, P is *aperiodic*.

The following result is of crucial importance for MCMC methods.

Proposition 3.1. *Let ν be an invariant measure of the transition kernel P , and we assume that P is irreducible and aperiodic. Then for all $x \in \mathbb{R}^n$,*

$$\lim_{N \rightarrow \infty} P^{(N)}(x, A) = \nu(A) \quad \text{for all } A \in \mathcal{B}, \tag{3.16}$$

and for $f \in L^1(\nu(dx))$,

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{j=1}^N f(X^{(j)}) = \int_{\mathbb{R}^n} f(x) \nu(dx) \tag{3.17}$$

almost certainly.

See [67], proofs in [74] based on [58].

The property (3.17) is the important ergodicity property that is used in Monte Carlo integration. The convergence (3.16), stating that ν is a limit distribution for the transition kernel P can be stated also in a slightly stronger form. We skip the detailed discussion here and refer to article [74].

3.4. Construction of the transition kernel

To avoid measure theoretic notations, we assume that the probability distribution ν on \mathbb{R}^n is defined by a probability density, $\nu(dx) = \pi(x) dx$. We wish to determine a transition kernel $P(x, A)$ such that ν is its invariant measure. To this end, we write P as a sum of absolutely continuous and singular part

$$\begin{aligned} P(x, A) &= \int_A K(x, y) dy + r(x)\chi_A(x) \\ &= \int_A (K(x, y) + r(x)\delta(x - y)) dy, \end{aligned}$$

where χ_A is the characteristic function of the set A and δ denotes the Dirac delta. Although $K(x, y) \geq 0$ is actually a density in the setting of continuous random variables, it can loosely be interpreted as the probability of the move from x to y while $r(x) \geq 0$ is the probability of x remaining inert. This would be correct in the setting of discrete random variables. The condition $P(x, \mathbb{R}^n) = 1$ implies that

$$r(x) = 1 - \int_{\mathbb{R}^n} K(x, y) dy. \quad (3.18)$$

In order for $\pi(x) dx$ to be an invariant measure of P , we must have the identity

$$\begin{aligned} \nu P(A) &= \int_A \left(\int_{\mathbb{R}^n} \pi(x) K(x, y) dx + r(y)\pi(y) \right) dy \\ &= \int_A \pi(y) dy \end{aligned}$$

for all $A \in \mathcal{B}$, or

$$\pi(y)(1 - r(y)) = \int_{\mathbb{R}^n} \pi(x) K(x, y) dx. \quad (3.19)$$

By formula (3.18), this is tantamount to

$$\int_{\mathbb{R}^n} \pi(y) K(y, x) dx = \int_{\mathbb{R}^n} \pi(x) K(x, y) dx, \quad (3.20)$$

referred to as the *balance equation*. In particular, if K satisfies the *detailed balance equation*

$$\pi(y) K(y, x) = \pi(x) K(x, y) \quad (3.21)$$

for all pairs $x, y \in \mathbb{R}^n$, then the balance equation holds *a fortiori*. The conditions (3.20) and (3.21) constitute the starting point in constructing the Markov chain transition kernels used for stochastic sampling.

We start by discussing an MCMC scheme known as the *Metropolis–Hastings* algorithm; see [33, 51]. The aim is to construct a transition kernel K that satisfies the detailed balance equation (3.21). One starts with some candidate kernel. Thus, let $q : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}_+$ be a given function with the property $\int q(x, y) dy = 1$. The function q defines a transition kernel

$$Q(x, A) = \int_A q(x, y) dy.$$

If q happens to satisfy the detailed balance equation, we set simply $K(x, y) = q(x, y)$, $r(x) = 0$ and we are done. Otherwise, we define

$$K(x, y) = \alpha(x, y)q(x, y), \quad (3.22)$$

where α is a correction term to be determined. We assume that for some $x, y \in \mathbb{R}^n$, instead of the detailed balance, we have

$$\pi(y)q(y, x) < \pi(x)q(x, y). \quad (3.23)$$

Our aim is to choose α such that

$$\pi(y)\alpha(y, x)q(y, x) = \pi(x)\alpha(x, y)q(x, y). \tag{3.24}$$

This is achieved if we set

$$\alpha(y, x) = 1, \quad \alpha(x, y) = \frac{\pi(y)q(y, x)}{\pi(x)q(x, y)} < 1. \tag{3.25}$$

By reversing x and y , we see that the kernel K defined through (3.22) satisfies (3.21) if we choose

$$\alpha(x, y) = \min \left(1, \frac{\pi(y)q(y, x)}{\pi(x)q(x, y)} \right). \tag{3.26}$$

This transition kernel is called the Metropolis–Hastings kernel. The kernel q is called *the candidate generating kernel* for reasons that become evident when the implementation of the method is discussed.

A slightly different algorithm is obtained if the candidate generating kernel is defined by using the density π directly and a block segmentation of the vectors in \mathbb{R}^n . Let us write $\mathbb{R}^n = \mathbb{R}^{k_1} \times \dots \times \mathbb{R}^{k_m}$, where $k_1 + \dots + k_m = n$, and denote correspondingly

$$x = (x_{(1)}, \dots, x_{(m)})^T \in \mathbb{R}^n, \quad x_{(j)} \in \mathbb{R}^{k_j}. \tag{3.27}$$

If X is an n -variate random variable with the probability density π , the probability density of the i th block $X_{(i)}$ conditioned on the variables $X_{(j)} = x_{(j)}$, $i \neq j$, is given by the density

$$\pi(x_{(i)} | x_{(1)}, \dots, x_{(i-1)}, x_{(i+1)}, \dots, x_{(m)}) = C_i \pi(x_{(1)}, \dots, x_{(i-1)}, x_{(i)}, x_{(i+1)}, \dots, x_{(m)}),$$

where C_i is a normalization constant. With these notations, we can define a transition kernel K by the formula

$$K(x, y) = \prod_{i=1}^m \pi(y_{(i)} | y_{(1)}, \dots, y_{(i-1)}, x_{(i+1)}, \dots, x_{(m)}), \tag{3.28}$$

and

$$r(x) = 0.$$

This transition kernel does not in general satisfy the detailed balance equation (3.21). However, it satisfies the weaker but sufficient condition (3.20). To prove this claim, we observe first that

$$\int_{\mathbb{R}^{k_i}} \pi(y_{(i)} | y_{(1)}, \dots, y_{(i-1)}, x_{(i+1)}, \dots, x_{(m)}) \, dy_{(i)} = 1$$

for all $i = 1, \dots, m$, implying that

$$\int_{\mathbb{R}^n} \pi(y) K(y, x) \, dx = \pi(y) \int_{\mathbb{R}^n} K(y, x) \, dx = \pi(y).$$

Furthermore, we observe that

$$\begin{aligned} & \int_{\mathbb{R}^{k_i}} \pi(y_{(1)}, \dots, y_{(i-1)}, x_{(i)}, \dots, x_{(m)}) \pi(y_{(i)} | y_{(1)}, \dots, y_{(i-1)}, x_{(i+1)}, \dots, x_{(m)}) \, dx_{(i)} \\ &= \pi(y_{(1)}, \dots, y_{(i-1)}, x_{(i+1)}, \dots, x_{(m)}) \pi(y_{(i)} | y_{(1)}, \dots, y_{(i-1)}, x_{(i+1)}, \dots, x_{(m)}) \\ &= \pi(y_{(1)}, \dots, y_{(i-1)}, y_{(i)}, x_{(i+1)}, \dots, x_{(m)}). \end{aligned}$$

Applying this identity recursively we obtain

$$\int_{\mathbb{R}^n} \pi(x) K(x, y) \, dx = \pi(y),$$

showing that (3.20) holds. The algorithm thus obtained is known as the *Gibbs sampler*. For the special case when $m = n$ and $k_1 = \dots = k_m = 1$, the algorithm is referred to as the

single component Gibbs sampler. The algorithm was originally proposed in article [23] for image restoration purposes, and later generalized and developed to its present form in [22]. Although the construction of the Gibbs sampler, as opposed to the Metropolis–Hastings kernel, is not based on the requirement of the detailed balance of the transition kernel, it is often seen as a special case of the Metropolis–Hastings algorithm. As we shall see in the next section, algorithmically they yield essentially the same method.

The success of the Metropolis–Hastings and Gibbs sampler algorithms depend largely on whether they satisfy the assumptions of proposition 3.1. There are known sufficient conditions concerning the density π that guarantee the ergodicity of these methods. In the following proposition, we give such conditions, which are relatively general. The proposition summarizes lemmas 1 and 2 in [67]. The proofs can be found in article [64].

Proposition 3.2. (i) Let $\pi : \mathbb{R}^n \rightarrow \mathbb{R}_+$ be a probability density, and $E_+ = \{x \in \mathbb{R}^n | \pi(x) > 0\}$, and let $q : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}_+$ be a candidate generating kernel and $Q(x, A)$ be the corresponding transition probability function. If Q is aperiodic, then the Metropolis–Hastings chain is also aperiodic. Further, if Q is irreducible and $\alpha(x, y) > 0$ for all $(x, y) \in E_+ \times E_+$, then the Metropolis–Hastings chain is irreducible.

(ii) Let π be a lower semicontinuous density and $E_+ = \{x \in \mathbb{R}^n | \pi(x) > 0\}$. The Gibbs sampler defines an irreducible and aperiodic transition kernel if E_+ is connected and each $(n - 1)$ -dimensional marginal $\pi(x_{-j}) = \int_{\mathbb{R}} \pi(x) dx_j$ is locally bounded.

When implemented in practice, the question of convergence is a difficult issue, since it is not a straightforward matter to decide when the sample set is large enough to represent good coverage of the probability distribution. Another practical issue that needs to be taken into account is related to the initiation of the sampling. If the initial value $x^{(1)}$ is far away from the maxima of the distribution, it may take a long sequence of updates before the algorithm starts to sample the distribution near the maxima where most of the contribution should come from. Therefore, in practice, it is advisable to retain the samples only after a period of time that is usually referred to as the *burn-in phase*. One can monitor the posterior probabilities of the updates as the sampling proceeds and retain the samples only after this random function stabilizes around some value. This will be demonstrated in section 4. It is also proposed in the literature to restart the MCMC sampling with different initial values. This procedure guarantees on one hand that the samples are less correlated and it may also lead to a better coverage of the probability distribution. More details of the convergence issues may be found, e.g., in [24] and [11].

3.5. Implementing MCMC

One of the reasons why the MCMC methods are becoming increasingly popular in various fields of application is that they are conceptually simple and the algorithms are easy to describe on a formal level. We start by giving the implementation steps of the single-component Gibbs sampler. A generalization to a block version is straightforward.

- (1) Pick the initial value $x^{(1)} \in \mathbb{R}^n$ and set $k = 1$.
- (2) For $1 \leq j \leq n$, draw $y_j^{(k)} \in \mathbb{R}$ from the one-dimensional distribution $\pi(y_j | y_1^{(k)}, \dots, y_{j-1}^{(k)}, x_{j+1}^{(k)}, \dots, x_n^{(k)})$.
- (3) Set $x^{(k+1)} = y^{(k)}$. When $k = K$, the desired sample size, stop. Else, increase $k \rightarrow k + 1$ and repeat from (2).

Similarly, the Metropolis–Hastings algorithm consists of a few simple steps.

- (1) Pick the initial value $x^{(1)} \in \mathbb{R}^n$ and set $k = 1$.

- (2) Draw $y^{(k)} \in \mathbb{R}^n$ from the proposal distribution $q(x^{(k)}, y)$ and calculate the acceptance ratio

$$\alpha(x^{(k)}, y) = \min \left(1, \frac{\pi(y)q(y, x^{(k)})}{\pi(x^{(k)})q(x^{(k)}, y)} \right).$$

- (3) Draw $t \in [0, 1]$ from uniform probability density.
 (4) If $\alpha(x^{(k)}, y) \geq t$, set $x^{(k+1)} = y$, else $x^{(k+1)} = x^{(k)}$. When $k = K$, the desired sample size, stop, else increase $k \rightarrow k + 1$ and repeat from (2).

The implementation of both of the above methods requires care. In implementing the Gibbs sampler, the posterior distribution is typically given in a non-parametric way. Therefore, the one-dimensional updatings are usually done by approximating numerically the probability densities. More precisely, to perform the second step with a given j , one needs to calculate the function

$$\Phi(t) = C \int_{-\infty}^t \pi(y_j | y_1^{(k)}, \dots, y_{j-1}^{(k)}, x_{j+1}^{(k)}, \dots, x_n^{(k)}) dy_j, \quad (3.29)$$

where we have included a normalization constant C to guarantee that $\Phi(t) \rightarrow 1$ as $t \rightarrow \infty$. The value $y_j^{(k)}$ is then found by setting $y_j^{(k)} = \Phi^{-1}(t)$, where t is a drawn value of a uniformly distributed random variable in the interval $[0, 1]$. To calculate the function Φ above, one needs a relatively good understanding of the interval where the distribution function differs clearly from zero. If the direct model is complicated, the numerical search as well as the integration itself may make the algorithm prohibitively slow.

In the Metropolis–Hastings algorithm, the proposal algorithm is usually chosen to be a parametric density (e.g. Gaussian) to make the drawing simple. The problem here is to choose the proposal distribution so that the acceptance ratio has a reasonable value. If q is chosen too wide, the drawn y is practically never accepted. On the other hand, if q is so narrow that the proposed y is always accepted, the step $x^{(k)} \rightarrow x^{(k+1)}$ gets too short and a proper sampling of the distribution requires a prohibitively large sample set. There are results concerning an optimal acceptance ratio (see, e.g., [24, 64]). A good rule of thumb is that of all y , roughly 20–30% should be accepted.

Let us mention here that in the literature adaptive Metropolis–Hastings schemes have been suggested; see, e.g., [18, 24, 25, 29, 30, 57, 70]. The idea in some of the adaptive methods is to change the proposed distribution based on recent draws as the algorithm progresses.

4. EIT as a statistical inverse problem

Consider now the impedance tomography problem as a statistical inverse problem. As in section 2.4 where the gradient-based methods were discussed, we consider the discretized problem, i.e. we identify the resistivity with the coefficient vector, so that we have $\rho \in \mathbb{R}^M$. The discrete model is

$$V = U(\rho) + N,$$

where the additive noise N is assumed to be independent of ρ . What is more, we assume that the contact impedances are known. An extension to the case of unknown contact impedances is quite straightforward.

We assume that the basis functions $\eta_k \in H_M$ spanning the subspace H_M are positive. Typically, the functions are the characteristic functions of the pixels (or voxels when $\Omega \subset \mathbb{R}^3$) of the discretized body. To implement the positivity constraint in this situation, we consider prior densities of the form

$$\pi_{\text{pr}}(\rho) = \pi_+(\rho) \tilde{\pi}_{\text{pr}}(\rho), \quad (4.1)$$

where $\tilde{\pi}_{\text{pr}}$ is to be discussed later and π_+ is the positivity prior

$$\pi_+(\rho) = \begin{cases} 1 & \text{if } 0 \leq \rho_{\min} \leq \rho_j \leq \rho_{\max} \leq \infty \quad \text{for all } j, \quad 1 \leq j \leq M \\ 0 & \text{otherwise.} \end{cases}$$

In practice, one has to choose $\rho_{\min} > 0$ in order that the forward problem be solvable in the admissible class of H_M where $\pi_{\text{pr}}(\rho) > 0$.

Since the EIT problem with merely the positivity constraint is in practice still too sensitive to noise, we may call the density $\tilde{\pi}_{\text{pr}}$ the *regularizing prior density*. We write the regularizing prior densities in the form

$$\tilde{\pi}_{\text{pr}}(\rho) \sim e^{-\alpha A(\rho)}, \quad (4.2)$$

where $\alpha > 0$ is a parameter that is related to the confidence on the regularizing prior. Let us assume that the noise vector N is a zero mean Gaussian random vector with positive definite covariance matrix C . With this choice, the posterior distribution given by formulae (3.3) and (3.5) assumes the form

$$\pi(\rho|V) \sim \pi_+(\rho) \exp\left(-\frac{1}{2}(U(\rho) - V)^T C^{-1}(U(\rho) - V) - \alpha A(\rho)\right).$$

Hence, without the positivity constraint π_+ , the weighted least squares solution discussed in section 2.4 corresponds to the MAP solution when $W = \frac{1}{2}C^{-1}$. This rather obvious observation is our starting point in the discussion of the EIT prior densities. For earlier studies on the statistical approach to EIT inverse problems, see [18, 39, 40, 43, 44, 50, 57, 69].

4.1. EIT priors

In impedance tomography, the fundamental prior information of the resistivity distribution is that it is a positive function. Additional information may be related to the overall size of the resistivity [6, 7], smoothness properties of the function, see e.g. [1, 34, 80], or to the presumably known internal structure of the body. For example, in magnetic resonance imaging (MRI)-guided medical impedance tomography, such information is the approximately known anatomical structure of the body; see, e.g., [38, 39, 75, 77].

In the discussion below, we use as a starting point the most commonly used regularizing functionals for generating different regularizing priors. We start by specifying the discretization subspace H_M . We assume that the body Ω is divided into M pixels Ω_j , $1 \leq j \leq M$. We define the subspace $H_M \subset \mathcal{A}$ as consisting of resistivities of the simple form

$$\rho(x) = \sum_{j=1}^M \rho_j \chi_j(x), \quad (4.3)$$

where χ_j is the characteristic function of Ω_j and identify $\rho \in H_M$ with the vector $(\rho_1, \dots, \rho_M) \in \mathbb{R}^M$.

We consider first the simplest Tikhonov-type regularization, where the least squares functional to be minimized is

$$F_\alpha(\rho) = \|U(\rho) - V\|_W^2 + \alpha \|\rho - \rho_*\|_{L^2(\Omega)}^2.$$

More generally and especially when gradient-based methods are employed, the regularizing functional A is often chosen to be a quadratic functional

$$A(\rho) = \|\mathcal{L}(\rho - \rho_*)\|_{L^2(\Omega)}^2.$$

Here, ρ_* is a ‘good guess’ for the resistivity distribution and \mathcal{L} is a properly chosen linear regularization operator.

Consider the discrete case where L is a matrix approximation of the operator \mathcal{L} , and

$$\tilde{\pi}_{\text{pr}}(\rho) \sim e^{-\alpha \|L(\rho - \rho_*)\|^2}. \tag{4.4}$$

When L is a discrete approximation of a differential operator, we arrive at a class of regularizing priors that can be called *smoothness priors*. For a short discussion on the utilization of conventional smoothness priors in EIT inverse problems, see [34]. A typical difficulty in using the smoothness-type priors is that the functions in H_M are not differentiable. However, reasonable discrete approximations can be defined. One such construction is given in appendix A.

A different meaning of the quadratic regularization functional is obtained by observing that the probability mass of the density (4.4) is concentrated near the affine space

$$\rho_* + \text{Ker}(L) = \{\rho \in H_M \mid \rho - \rho_* \in \text{Ker}(L)\}.$$

Thus, the prior density tends to pull the resistivity distributions towards this space without strictly constraining the solutions to this space. In article [77] this point of view was adopted and the matrix L was chosen as

$$L = I - P,$$

the matrix P being an orthogonal projection $P : H_M \rightarrow H_M^0$. The space H_M^0 was chosen such that the vectors $\rho \in H_M^0$ correspond to (anatomically) realistic resistivity distributions. A basis for H_M^0 was constructed by using principal components computed from an acceptable ensemble of resistivity distributions. This class of priors is called *subspace priors*.

Anatomical prior information can also be considered, loosely speaking, as knowledge about the spatial and directional variations in the size of the first-order derivatives of the resistivity distribution. In article [39] structural prior information was incorporated into the regularization operator penalizing the first-order smoothness of ρ with the aid of a properly constructed matrix-valued field. Also, the statistical interpretation of the method was explained. Especially, the assessment of the implicit assumptions that are carried out when the prior model is used, was explained via the investigation of certain conditional covariances of the prior distribution.

Obviously, if we have strict confidence on the anatomical prior information, it can be incorporated as a *structural constraint* to the EIT inverse problem via the construction of the space H_M . These approaches lead typically to a very small-dimensional inverse problem that is no longer ill-posed. Thus, a maximum likelihood approach can yield reasonable estimates. These methods are not discussed further here; for details see [16, 26, 27, 76, 79].

The statistical inversion method allows us to analyse also priors that lead to non-differentiable regularization functionals but which conform well with the assumption $\rho \in \mathcal{A}$, i.e. priors that allow moderate discontinuities in the resistivities. In article [14], the authors considered the linearized EIT problem with continuous-type boundary data and used an image enhancement technique based on the minimization of the *total variation* of the image. Modifying this idea, the total variation can also be considered as a regularizing functional.

Let $f : \Omega \rightarrow \mathbb{R}$ be a function in $L^1(\Omega)$. We define the total variation of f , denoted by $\text{TV}(f)$, as

$$\text{TV}(f) = \int_{\Omega} |Df| = \sup \left\{ \int_{\Omega} f \operatorname{div} g \, dx \mid g = (g_1, \dots, g_n) \in C_0^1(\Omega; \mathbb{R}^n), |g(x)| \leq 1 \right\}.$$

A function is said to have bounded variation if $\text{TV}(f) < \infty$.

To find a discrete analogue, let us consider the two-dimensional problem. We assume that ρ is of the form (4.3). Let J be the number of edges of positive length between the pixels. In

this case, the total variation of ρ can be calculated as

$$\text{TV}(\rho) = \sum_{j=1}^J d_j |\Delta_j^T \rho|, \tag{4.5}$$

where d_j is the length of the j th edge between the adjacent pixels $\Omega_{i_1^j}$ and $\Omega_{i_2^j}$ and $\Delta_j \in \mathbb{R}^N$ is the vector

$$\Delta_j = (0, \dots, \overset{(i_1^j)}{1}, 0, \dots, 0, \overset{(i_2^j)}{-1}, 0, \dots, 0)^T.$$

The *total variation prior* is defined as

$$\tilde{\pi}_{\text{pr}}(\rho) \sim e^{-\alpha \text{TV}(\rho)}, \tag{4.6}$$

where the total variation is calculated as in (4.5). The statistical approach to EIT with the total variation prior was adopted in [40,69]. Estimates from the posteriors were computed by using the sampling-based approach.

Another non-quadratic prior is obtained by considering the regularization functional

$$A(\rho) = \|\rho - \rho_*\|_{L^1(\Omega)}.$$

A piecewise constant discretization leads us to consider the prior

$$\tilde{\pi}_{\text{pr}}(\rho) \sim \exp\left(-\alpha \sum_{j=1}^M |\Omega_j| |\rho_j - \rho_*|\right). \tag{4.7}$$

Hence, the corresponding regularizing functional is also non-differentiable in this case and the direct utilization of gradient-based optimization methods for the search of the MAP estimate is not possible. An approximation for the L^1 -prior has been used in connection with EIT and statistical inversion in [50].

References [18, 57] give very interesting ideas for the statistical inversion of EIT data. In [18] the authors consider a situation in which the *possible* material types, say materials denoted by $\{1, \dots, N\}$, making up the object Ω are known. In their model the resistivities of the materials $\{1, \dots, N\}$ are assumed to be known exactly *a priori* but the proportion and the spatial distribution of the different materials is unknown. In their model, they use a regularizing prior of the form

$$\pi_{\text{pr}}(\rho) \propto \exp\left(\alpha \sum_{m=1}^M H_m(\rho)\right), \quad H_m(\rho) = \sum_{m' \sim m} \delta(\rho_m - \rho_{m'}), \tag{4.8}$$

where δ is the Dirac delta and the notation $m \sim m'$ means that the summation is carried over the immediate neighbours m' of the pixel m on the lattice. The exact knowledge about the resistivities of the (possible) materials in Ω is taken into account via the construction of the transition kernel for the Metropolis–Hastings algorithm. In a uniform square lattice this prior can be basically thought of as an analogue of the total variation prior (4.6) for processes with discrete states. Priors of the form (4.8) are often referred to as Markov random field priors.

In [57] the authors extend the model by assuming that the resistivities of the materials $\{1, \dots, N\}$ are not known *a priori* but instead they may vary from pixel to pixel. Let $\tau = (\tau_1, \dots, \tau_M)^T$ denote some estimate for the material-type distribution in Ω with states $\tau_i \in \{1, \dots, N\}$. By assuming that the resistivity of each material j obeys Gaussian density with mean $\eta(j)$ and standard deviation $\xi(j)$, they defined a conditional prior of the form

$$\pi_{\text{pr}}(\rho|\tau) \propto \pi_+(\rho) \exp\left(-\alpha \sum_{m=1}^M G_m(\rho)\right) \prod_{m=1}^M \exp\left(-\frac{1}{2\xi(\tau_m)^2}(\rho_m - \eta(\tau_m))^2\right), \tag{4.9}$$




Figure 1. Left: the discretization of the domain Ω into 260 quadrilateral pixels. Right: the i th and $(i + 1)$ th edge between the adjacent pixels. The dotted lines denote the boundaries of the triangular elements that are used in the forward computations.

where $G_m(\rho) = \sum_{m' \sim m} (\rho_m - \rho_{m'})^2$ and $m' \sim m$ mean the same as in (4.8). For the material-type distribution τ the authors used a prior $\pi_{\text{pr}}(\tau)$ of the form (4.8) and the overall prior was obtained as $\pi_{\text{pr}}(\rho, \tau) = \pi_{\text{pr}}(\rho|\tau)\pi_{\text{pr}}(\tau)$. These sophisticated prior models could be especially useful in certain industrial applications of EIT.

There are also some very interesting results concerning the use of Bayesian neural networks in connection with impedance imaging [43, 44]. In this approach the prior information as well as the nonlinearity of the inverse problem are coded in a multilayer perceptron (MLP) network. This approach was shown to be very successful in a process tomography application of EIT in which the task was to detect and estimate non-conducting bubbles in a conducting liquid. Although it is evident from these results that Bayesian networks have great potential in statistical inversion of the EIT problem—especially in relation to constructing very nonlinear prior densities with non-connected support—the treatise of MLP networks is outside the scope of this review.

To obtain further insight into the significance of the selection of the prior, we generate a set of spatial resistivity distributions based on different prior distributions. We will use the term resistivity distribution to refer to the spatial interpretation of the associated random vector. In the numerical calculations, we use a two-dimensional circular body divided into quadrilateral pixels shown in figure 1. The number of pixels is 260. This discretization is essentially identical to the ‘Joshua tree’ of [7].

We consider four different prior distributions of the form (4.1). We refer to the first one as the *white noise distribution*, obtained by choosing $L = I$, the $M \times M$ unit matrix in (4.4). Thus, the pixel resistivities are assumed to be uncorrelated Gaussian random variables of equal variance centred around ρ_* . Here we use $\rho_* = \rho_0$. The second prior distribution in our selection is a *second-order smoothness prior*. Here, the form of the *regularizing* prior is again that of (4.4), but the matrix L is a discrete approximation of the Laplacian. The construction of the matrix L is given in detail in appendix A. In addition to the (truncated) Gaussian priors above, we consider the two non-differentiable priors discussed above, where the regularizing prior $\tilde{\pi}_{\text{pr}}(\rho)$ is either the total variation prior (4.6) or the L^1 -prior given in (4.7). The latter one could also be called the *impulse prior* since it is concentrated around spatial resistivity distributions which are almost constant but may have large deviations with small support. Thus, these resistivity distributions look like impulse noise on a relatively flat background.

For an extensive discussion of this feature, see article [15].

In figure 2, we show sets of five resistivity distributions from a large randomly drawn collection from each of the above-mentioned distributions. The random drawing is produced from the *conditional* priors $\pi_{\text{pr}}(\rho_{33}, \dots, \rho_{260} | \rho_1, \dots, \rho_{32})$ using the Gibbs sampler, and the displayed resistivity distributions are picked so that their probabilities are relatively large. The drawing is produced from the conditional priors due to the fact that the smoothness prior and the total variation prior are partially uninformative since Δ and $\text{TV}(\rho)$ have non-trivial kernels[†]. The fixed pixels $1, \dots, 32$ form the outermost layer of the mesh. The value of resistivity in the fixed pixels was 1 and the parameter α controlling the width of the distribution is chosen in each case so that the dynamical ranges of the resistivity distributions are roughly of the same magnitude.

4.2. MCMC sampling of resistivities

In the following simulations, we assume that there are 16 identical electrodes with equal spacing attached to the boundary of the circular domain Ω . The width of the electrodes is chosen so that they cover 50% of the boundary $\partial\Omega$. We use 15 linearly independent trigonometric current patterns; see, e.g., [41]. Taking the symmetry of the resistance matrix into account, the data consist of 120 independent real numbers. However, we use all the 240 measured voltages despite this redundancy in the data. For the computation of the simulated measurements \mathbf{V} the domain Ω was divided into 1432 triangular elements, the total number of vertex nodes being 749. To the computed voltages, Gaussian random errors with a standard deviation of 1% of the corresponding voltage values were added.

We consider the EIT problem where the resistivity is discretized using the pixel mesh of figure 1. For the computation of the forward solution $U(\rho)$ the mesh was divided into 1056 triangular elements, the number of nodes being 561; see figure 1. Although this is strictly speaking incorrect, we use the (false) assumption that the noise vector is Gaussian zero mean, all the channels being independent and having equal variance. The reason for this is that the noise covariance matrix is diagonal and the scale difference of the individual variances is not great. Thus, the noise covariance matrix is $C = \sigma^2 I$, and the posterior distribution is

$$\pi(\rho | \mathbf{V}) \sim \pi_+(\rho) \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{U}(\rho) - \mathbf{V}\|^2 - \alpha A(\rho)\right). \quad (4.10)$$

In experimental measurement systems an approximation for the noise covariance C can be obtained by using a set of repeated measurements or by analysing the measurement system. For further details, see e.g. [1, 16, 80].

Our primary interest in the numerical studies is in those prior distributions that correspond to the non-differentiable regularizing functionals, the total variation prior (4.6) and the L^1 -prior (4.7). To shorten the burn-in phase of the MCMC calculation, it is advisable to find a relatively good approximation for the MAP estimate. The first step is to compute estimates for the optimal background resistivity and the electrode contact impedance (assuming that the ratio $\xi = z_0/\rho_0$ is known). These values, denoted by (ρ_0, z_0) , are calculated by using the formula (2.24). Next, we apply the Gauss–Newton-type search described in section 2.4. Here, we face two difficulties. First, the total variation and L^1 -priors are non-differentiable, due to the presence of the absolute value function. To overcome this problem, we use the smooth approximation

$$|t| \approx h_\beta(t) = \frac{1}{\beta} \log(\cosh(\beta t)),$$

[†] In practice, however, the numerical kernel of the approximate Laplacian L can depend on the mesh structure.




Figure 2. Spatial resistivity distributions which are sampled from different prior distributions. From left to right: white noise prior, second-order smoothness prior, total variation prior and impulse prior, i.e. L^1 -prior.

where $\beta > 0$ is a small parameter adjusting the accuracy of the approximation. The second problem comes from the positivity constraint $\pi_+(\rho)$. To avoid negative values, we use an




Figure 3. Convergence of the MCMC iteration: the values of the exponent terms in the posterior densities for 400 states of the Markov chains corresponding to the total variation prior (top) and L^1 -prior (bottom). The end of the burn-in period is denoted with vertical lines.

interior point search and augment the least squares functional with a penalty term of the form

$$J_\gamma(\rho) = \gamma_j \sum_{k=1}^M \frac{1}{\rho_k},$$

where $\{\gamma_j\}$ is a sequence of decreasing positive coefficients (j denotes the iteration index) [17]. After finding a reasonable approximation $\hat{\rho}_\beta = \rho^{(j)}$ for the MAP estimate, we start the MCMC run with the posterior probability density of the form (4.10). Notice that for MCMC, there is no need to smooth the functional $A(\rho)$. However, to speed up the calculations, we linearize the mapping $\rho \mapsto U(\rho)$ around the initial guess $\hat{\rho}_\beta$

$$U(\rho) \approx U(\hat{\rho}_\beta) + DU(\hat{\rho}_\beta)(\rho - \hat{\rho}_\beta). \quad (4.11)$$

It is our experience that with noisy data, the effect of the linearization is usually not large on the outcome but simplifies the numerics considerably since we do not need to solve the finite-element problem after each updating of the resistivity. However, with some target distribution, such as large-scale deviations from the background near the boundary and very small noise levels, further iteration can prove to be worthwhile. The computation of the differential $DU(\rho)$ is explained in appendix B.

We generate a set of 10 000 samples using both the total variation and the L^1 -priors. The MCMC is based on a Gibbs sampling algorithm. Due to the relatively well chosen initial resistivity, we have reason to believe that the burn-in period is short. This belief is supported by the values of the exponent terms in the posterior distribution, depicted in figure 3.

We used the burn-in length of 200 samples. Figures 4 and 5 display the results. In both figures the top-left image shows the resistivity distribution that was used to generate the data. The mesh used in the inversion is depicted in the same image to facilitate the evaluation of the results. Also, we have marked four pixels for further discussion. The top-right image is the outcome of the positivity constrained Gauss–Newton iteration, i.e. the initial value $\hat{\rho}_\beta$. When the smooth approximation of the total variation was used, the Gauss–Newton iteration does not significantly improve the estimator, and the estimator is found by performing a single step.

When the approximate L^1 -functional was used, we performed six steps of the gradient-based iteration to find the estimate depicted in the top-right image of figure 5. In figures 4 and 5, the centre-left image shows the conditional mean of the resistivity based on the MCMC sampling, the centre-right image shows the square root of the diagonal of the posterior covariance matrix, i.e. the standard deviation of each pixel around the conditional mean. Notice that in the case of the total variation prior, the largest uncertainty is in the directions corresponding to the pixels located at the boundaries of the objects. Finally, the bottom rows show the 90% credibility limits: for each pixel, we have searched upper and lower bounds so that 5% of the sampled pixel values lie below or above these limits.

Finally, figures 6 and 7 show the marginal densities of single pixels marked in figures 4 and 5 when the prior distribution is the total variation prior or the L^1 -prior, respectively. In these figures, the solid line denotes the conditional expectations, the dotted lines the credibility limits of 90% probability and the dashed line indicates the initial value found as an approximative MAP estimate $\hat{\rho}_\beta$.

Figures 4–7 reveal clearly two things. First, the effect of the choice of the prior is remarkable. As was expected by the construction of the priors, the total variation prior prefers reconstruction with relatively flat details, while the L^1 -prior details tend to have a small support and a relatively good dynamical range. Also, it is clear from these examples that there can be a significant difference between the conditional mean and MAP estimates. This is evident both with respect to the total variation prior and the L^1 -prior.

Above, we have implemented the MCMC method by using the Gibbs sampler. It is our experience that the Gibbs sampler works more reliably in EIT than the Metropolis–Hastings method. In the Metropolis–Hastings method, we used the Gaussian approximation of the posterior distribution at $\hat{\rho}_\beta$ as the proposal distribution. One of the problems with the Metropolis–Hastings method is that in order to keep the positivity constraints intact, the proposal distribution needs to be quite narrow, leading to prohibitively slow convergence. It has proven to be very difficult to hit the gap between these two extreme cases. On the other hand, the reason for linearization of the mapping $\rho \mapsto U(\rho)$ is precisely the need for numerous re-evaluations of the voltage vector when running the Gibbs sampler.

In articles [18, 57] the authors proposed an adaptive Metropolis–Hastings scheme which is suitable when using priors of the form (4.8) and (4.9). Based on the results they obtained, the proposed adaptive Metropolis–Hastings scheme works well in such cases in which the values of the resistivities are (approximately) known.

4.3. Interpretation of results

The proper interpretation of the results obtained by Bayesian inversion methods is a subtle issue. It is very easy to jump into false conclusions and over-interpretation. This is true especially when applied to notoriously ill-posed inverse problems such as the EIT problem.

In the previous section we used the Monte Carlo sampling method to calculate the conditional expectation of the resistivity distribution as well as its *a posteriori* uncertainties using two different priors but the same data. The differences between the results are remarkable. In particular, when the total variation prior was used, the true values of the resistivities in the selected pictures in figure 6 lie outstandingly outside the 90% credibility limits. This raises the question about the meaning of these limits. One would feel tempted to say that if the measurements were repeated with independent noise, roughly 90% of the conditional expectations would have credibility limits including the true resistivity. This interpretation is *strictly incorrect*. The posterior probability distribution measures our uncertainty based on the measurement *and* the prior. The pitfall here is that the true resistivity distribution may

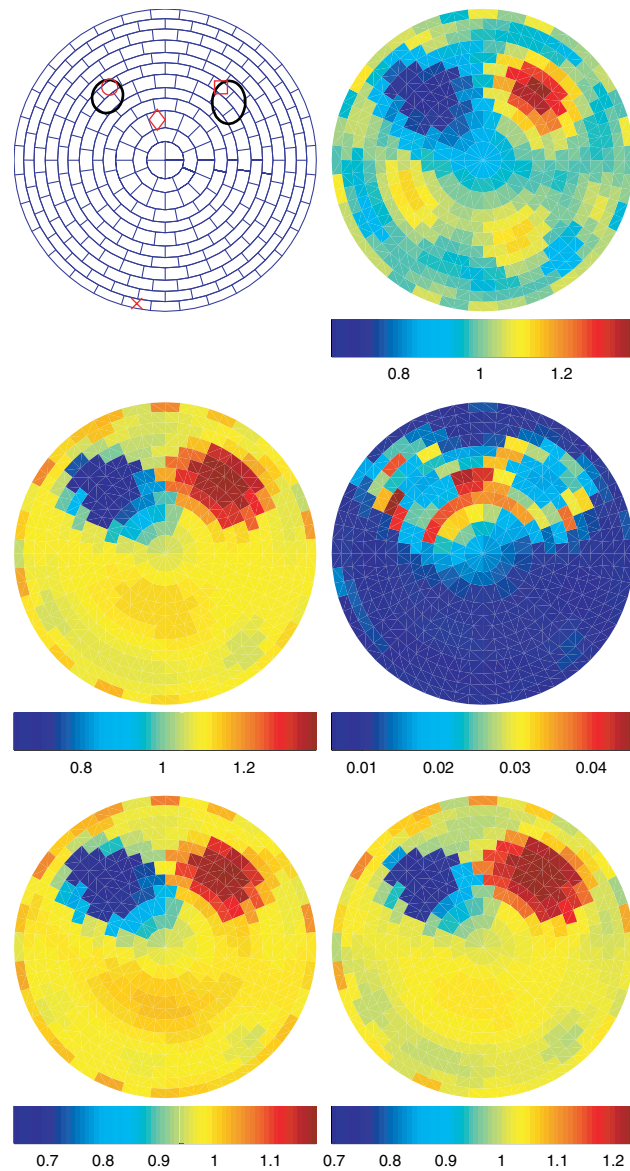


Figure 4. Total variation prior: in top-left image, the true resistivities are $\rho = 1$ for the background, $\rho = 0.1$ for the left inclusion and $\rho = 2$ for the right one. In the simulation of the data we used $z = 0.24$ for the contact impedances. Top right is the Gauss–Newton reconstruction, centre row shows the conditional expectation and the standard deviation, and the bottom row the upper and lower bounds for 90% credibility interval.

have a very small probability with respect to the postulated prior probability measure. On the other hand, the ill-posedness of the problem necessitates that useful priors are ‘informative’ with respect to certain subspaces from which the likelihood carries only little information. This problem is reflected in the fact that the results are usually sensitive to the selection of the prior distribution. Summarizing, the *a posteriori* uncertainties are reliable only in relation to our confidence in the postulated prior. If the prior is selected *ad hoc* and not based on actual

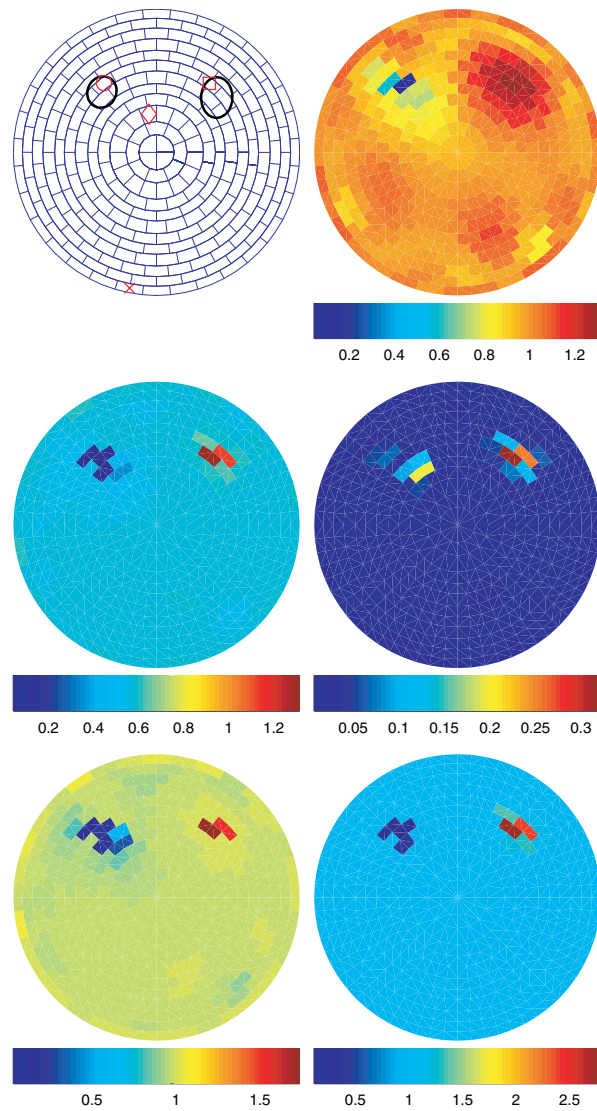


Figure 5. L^1 -prior, the images are as in figure 4.

prior information coming, e.g., from other measurement modalities, the estimates and their credibility limits may even be misleading.

5. Further topics

As we have seen in the previous sections, the choice of the prior density gives us the corresponding Tikhonov regularization functional if we identify the inverse problem with the search of the MAP estimator. Conversely, the statistical inversion approach can be used as a tool to analyse in general the underlying prior assumptions of various regularization schemes. As an example, consider the NOSER reconstruction algorithm described and discussed in articles [7, 8]. The algorithm performs one Newton iteration step starting from an optimally




Figure 6. Marginal distributions of single pixels with the total variation prior. The mid solid line is the conditional expectation, the left and right solid lines the 90% credibility limits and the dashed line the initial value found by the Gauss–Newton method.

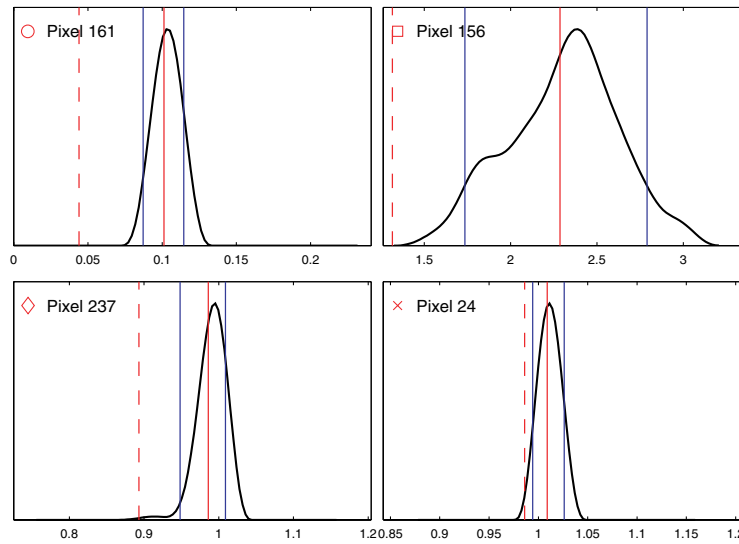


Figure 7. As in figure 6 when the prior is the L^1 -prior.

chosen constant background to minimize the least squares error. Since the Hessian is an ill-conditioned matrix, the inversion requires regularization which is done in the NOSER algorithm by adding a diagonal weight, i.e. the reconstruction is analogous to formula (2.25)

$$\rho = \rho_0 - (H + \alpha \text{diag}(H))^{-1} g,$$

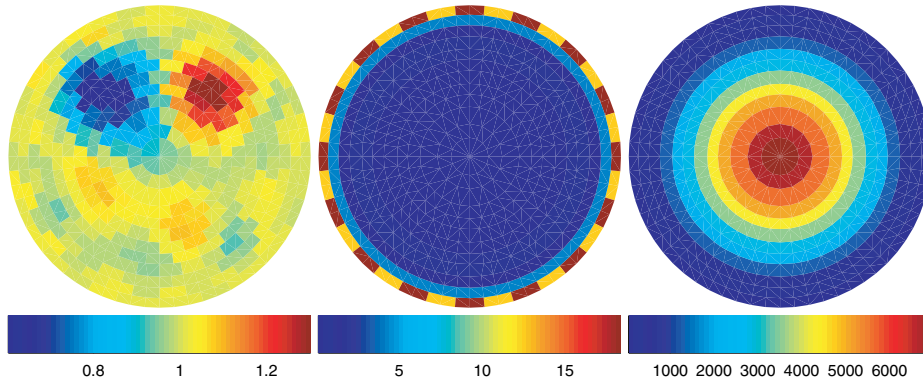


Figure 8. Left: NOSER reconstruction with $\alpha = 0.1$ (the data V are the same as those used in figures 4 and 6). Centre: the diagonal elements of the Hessian αH . Right: covariance of the pixels corresponding to the prior (5.1) in the NOSER algorithm.

where $\alpha > 0$ is a regularization parameter, and H and g are given similar formulae to (2.26) and (2.27)

$$H = (DU(\rho_0))^T (DU(\rho_0)), \quad g = (DU(\rho_0))^T (U(\rho_0) - V).$$

A NOSER reconstruction from the data V which was used in the examples of the previous section is shown in the left image of figure 8. The values of the regularization matrix $\alpha \text{diag}(H)$ are depicted in the centre image of figure 8. This image gives us a view of the implicit assumptions that the NOSER algorithm poses on the actual resistivity distribution. It can be seen that the measured voltages are much less sensitive to the values of ρ in the centre of Ω than to the values close to the boundary $\partial\Omega$. In other words, the sensitivity of the EIT reconstruction to the measurement noise increases when moving towards the centre of Ω .

From the point of view of statistical inversion, this reconstruction corresponds to (positivity constraints ignoring) the MAP estimate of the (linearized) posterior distribution, when the measurement noise consists of additive mutually independent Gaussian random variables with equal variances (white noise) and the prior probability density is

$$\pi_{\text{pr}}(\rho) \sim \exp(-\alpha(\rho - \rho_0)^T H_d(\rho - \rho_0)), \quad H_d = \text{diag}(H), \quad (5.1)$$

where ρ_0 is the optimal background resistivity distribution. This prior, again ignoring the positivity constraint, is a Gaussian density with covariance proportional to H_d^{-1} . Since H_d is diagonal, the pixels are assumed to be mutually independent. However, since H_d is clearly not of the form cI , the pixels are assumed to have different variances. The meaning of this prior is easy to display graphically; see figure 8.

From figure 8 we see that the prior is wider the deeper inside the body the pixel resides, i.e. the effect of the regularization gets weaker with increasing sensitivity to the measurement noise. One can justify this choice by saying that the prior density is chosen to be in balance with the likelihood density to avoid bias due to the prior. It must be noted here that in contrast to the ideas of Bayesian inference explained above, the implicit prior distribution of the NOSER algorithm is determined by the measurement setting. In the statistical inversion approach the prior should be independent of the measurements and carry such information that the measurement setting, at least partially, destroys. From this point of view, one could say that the NOSER algorithm does not belong to the class of statistical inversion algorithms although it has a statistical interpretation.

In section 3 we mentioned the hyper-priors. This topic is intimately related to the choice of the regularization parameter in the Tikhonov regularization. The literature concerning this topic is extensive; see, e.g., [31, 32, 47, 63, 73] and references therein. It would be tempting to include the parameter α appearing in the regularizing prior distributions in (4.2) as one unknown parameter that is eventually integrated out by MCMC sampling. More specifically, one could write the prior as

$$\tilde{\pi}_{\text{pr}}(\rho, \alpha) \sim \exp(-\alpha A(\rho))\pi_{\text{h}}(\alpha),$$

where $\pi_{\text{h}}(\alpha)$ is the hyper-prior of the parameter α and integrate α out to get

$$\pi(\rho|\mathbf{V}) \sim \int_0^\infty \exp(-\frac{1}{2}\|\mathbf{U}(\rho) - \mathbf{V}\|_{C^{-1}}^2 - \alpha A(\rho))\pi_{\text{h}}(\alpha) \, d\alpha,$$

the integration being done by MCMC. There are several problems with this procedure. First, if the hyper-prior is not very restrictive with relatively narrow support, the MCMC algorithm tends to favour small values of α and fits the resistivity distribution to satisfy the data only with the cost of the prior. On the other hand, restrictive hyper-priors are questionable since the goal was to do the opposite, not to fix the value of α . At the moment, the authors are not aware of a satisfactory solution to this problem.

Finally, we discuss briefly some computational aspects of the MCMC calculations. In section 4.2 we used the linearized model (4.11) to calculate approximations of the voltage. This simplification was done only to speed up the Gibbs sampling. When the full nonlinear version for computing the voltage vectors is used, one should take advantage of the special structure of the finite element system matrix. Consider the Gibbs sampling algorithm described in section 3.5. The updating of the resistivity requires the computation of the probability density (3.29), and to this end, one has to solve successively the forward problem (2.16). Let us write $A = A(\sigma)$ to indicate the dependence of the matrix (2.19) on the conductivity. Let σ denote the current value of the conductivity and assume that we need to update the value of the conductivity in the k th pixel Ω_k , $\sigma_k \rightarrow \sigma_k + t$, $t \in \mathbb{R}$. This updating yields a low rank perturbation of the system matrix $A(\sigma)$. Indeed, let Ω_k be the k th (open) pixel and

$$\ell \in I_k = \{k_1, \dots, k_j\} \quad \text{if and only if} \quad \text{supp } \varphi_\ell \cap \Omega_k \neq \emptyset.$$

By denoting with e_k the natural k th basis vector, we have

$$A(\sigma + te_k) = A(\sigma) + tV_k\Lambda_kV_k^T,$$

where $V_k = [e_{k_1}, \dots, e_{k_j}] \in \mathbb{R}^{(N_n+L-1) \times j}$, $\Lambda_k \in \mathbb{R}^{j \times j}$ is a symmetric positive definite matrix with entries of the form

$$(\Lambda_k)_{p,q} = \int_{\Omega_k} \nabla\varphi_{k_p} \cdot \nabla\varphi_{k_q} \, dx, \quad k_p, k_q \in I_k.$$

By denoting

$$b_0 = A(\sigma)^{-1}f, \quad B_k = A(\sigma)^{-1}V_k,$$

the solution of the equation $A(\sigma + te_k)b = f$ can be written by using the Sherman–Morrison–Woodbury formula [28] as

$$b = b_0 - tB_k\Lambda_k(1 + tV_k^TB_k\Lambda_k)^{-1}V_k^Tb_0,$$

i.e. the evaluation requires only successive solutions of a $j \times j$ problem. The matrices Λ_k can be pre-computed when the system matrix A is created. Let us mention here article [19], where the idea of using the Sherman–Morrison–Woodbury formula for solving partial differential equations repeatedly is also presented. Optimal updating ordering to minimize the computational work is an important issue that will be studied in the future.

6. Conclusions

The aim of this work was to demonstrate the application of a general approach, the statistical, or Bayesian, inversion method, to the EIT problem, and show how the MCMC sampling techniques work in this context with different priors. In contrast to the traditional way of looking at this inverse problem, the aim of the statistical approach is not to find a single estimate of the unknown resistivity but to estimate the whole *a posteriori* distribution. Since the visualization of the posterior distribution directly is impossible, the visualization is usually done by showing single estimates calculated on the basis of the *a posteriori* distribution, such as the conditional expectation or MAP conductivity distributions. This visual representation easily leads us astray, giving the impression that the method is simply another regularization method of producing single reconstruction images. The key features of the statistical approach, however, are (1) the possibility of incorporating prior information in a controlled way to the calculations; and, above all, (2) obtaining the full statistical description of the information that is immersed in the prior information and the measurements. The cost of this is, of course, the increased computational complexity, which is recognized as a major challenge for future research.

Acknowledgments

This work was supported by the Academy of Finland, Saastamoinen Foundation, Jenny and Antti Wihuri Foundation, Vilho, Yrjö and Kalle Väisälä Foundation, The Savo Foundation for Advanced Technology and TEKES.

Appendix A. Computation of the approximat Laplace operator

In this appendix, we describe the discrete approximation used in the smoothness prior construction when the pixel mesh is not a structured one. Let p_0 denote the midpoint of a given pixel and by p_j , $1 \leq j \leq n$ we denote the midpoints of the neighbouring pixels. We are looking for coefficients α_j , $0 \leq j \leq n$, such that $\Delta\rho(p_0) - \sum_{j=0}^n \alpha_j \rho(p_j)$ is small in some sense. Let us denote by H^s the L^2 -based Sobolev space in \mathbb{R}^2 with smoothness index $s \in \mathbb{R}$. We equip H^s with the norm

$$\|f\|_s = \left(\int_{\mathbb{R}^2} (1 + \lambda^2 |\xi|^2)^s |\hat{f}(\xi)|^2 d\xi \right)^{1/2}$$

and with the corresponding inner product $(\cdot, \cdot)_s$, where \hat{f} is the Fourier transform of f and $\lambda > 0$ is a scaling parameter. By the Sobolev imbedding theorem, $H^s \subset C^k$ if $k > s + 1$ in \mathbb{R}^2 , so if we choose $s > 3$, we have $\delta, \Delta\delta \in H^{-s}$, the dual of H^s . Denoting by $\langle \cdot, \cdot \rangle$ the distribution duality between H^s and H^{-s} , for $f \in H^s$, we obtain the estimate

$$\left| \sum_{j=0}^n \alpha_j f(p_j) - \Delta f(p_0) \right| = \left| \left\langle f, \sum_{j=0}^n \alpha_j \delta_{p_j} - \Delta \delta_{p_0} \right\rangle \right| \leq \|f\|_s \left\| \sum_{j=0}^n \alpha_j \delta_{p_j} - \Delta \delta_{p_0} \right\|_{-s}.$$

Thus, we seek to choose the coefficients α_j so that the norm attains its minimum. Since H^{-s} is a Hilbert space, the minimum is attained by orthogonal projection and the corresponding vector $\alpha = (\alpha_0, \dots, \alpha_n)$ satisfies $A\alpha = g$, where

$$A_{i,j} = (\delta_{p_i}, \delta_{p_j})_{-s}, \quad g_j = (\Delta \delta_{p_0}, \delta_{p_j})_{-s}.$$

In fact, explicit expressions for these coefficients can be found. By integration in polar coordinates of \mathbb{R}^2 , one finds that

$$A_{i,j} = 2\pi \int_0^\infty \frac{t}{(1 + \lambda^2 t^2)^s} J_0(t|p_i - p_j|) dt,$$

where J_0 is the Bessel function of order zero. From this expression we obtain

$$A_{j,j} = \frac{\pi}{\lambda^2(s-1)},$$

and for $i \neq j$, the Hankel–Nicholson formula [48, p 105] yields

$$A_{i,j} = \frac{\pi|p_i - p_j|^{s-1}}{\Gamma(s)2^{s-2}\lambda^{s+1}} K_{s-1}(\lambda^{-1}|p_i - p_j|),$$

where K_{s-1} denotes the modified Bessel function of order $s-1$. Similarly,

$$g_j = -2\pi \int_0^\infty \frac{t^3}{(1 + \lambda^2 t^2)^s} J_0(t|p_0 - p_j|) dt,$$

giving further

$$g_0 = -\frac{\pi}{\lambda^4(s-1)(s-2)},$$

and for $j > 0$,

$$g_j = -\frac{\pi|p_0 - p_j|^{s-2}}{\Gamma(s-1)2^{s-3}\lambda^{s+2}} \left(K_{s-2}(\lambda^{-1}|p_0 - p_j|) - \frac{|p_0 - p_j|}{2\lambda(s-1)} K_{s-1}(\lambda^{-1}|p_0 - p_j|) \right).$$

In practice, we have used the values $\lambda = 2$ and $s = 4$ in our calculations.

Appendix B. Computation of the Jacobian matrix $DU(\rho)$

The Jacobian matrix $DU(\rho)$ of the mapping $\rho \mapsto U(\rho)$ is computed as follows. First, let us denote

$$\tilde{\mathcal{C}} = (0\mathcal{C}) \tag{B.1}$$

$$\tilde{f} = [f^{(1)}, f^{(2)}, \dots, f^{(K)}] \tag{B.2}$$

$$\tilde{b} = [b^{(1)}, b^{(2)}, \dots, b^{(K)}], \tag{B.3}$$

where $\mathcal{C} \in \mathbb{R}^{L \times L-1}$ is defined by equation (2.18), $f^{(\cdot)} \in \mathbb{R}^{N_n+L-1}$ by equation (2.17), $0 \in \mathbb{R}^{L \times N_n}$ and $b^{(\cdot)} \in \mathbb{R}^{N_n+L-1}$ is the solution of the equation (2.16) corresponding to $f^{(\cdot)}$ and K is the number of current patterns. Furthermore, we assume that the resistivity distribution is of the form (4.3).

By equations (B.1)–(B.3) and (2.20), the set of computed voltages $\tilde{U}^h \in \mathbb{R}^{L \times K}$, corresponding to all K current patterns, is obtained as $\tilde{U}^h = \tilde{\mathcal{C}}\tilde{b}$. Since we are interested in the derivatives of the voltages \tilde{U}^h with respect to each pixel ρ_ℓ , $\ell = 1, \dots, M$, we need to determine

$$\frac{\partial(\tilde{\mathcal{C}}\tilde{b})}{\partial\rho_\ell}. \tag{B.4}$$

If the reference of the actual measurement is different from the choice (2.3), the actual measurements are obtained by multiplying \tilde{U}^h from the left with a properly chosen measurement operator (matrix) $M \in \mathbb{R}^{Q \times L}$ (Q is the number of actual measurements for the single current pattern). Thus, we obtain $M\tilde{\mathcal{C}}\tilde{b} = \tilde{M}\tilde{b}$. Now the derivative becomes

$$\frac{\partial(\tilde{M}\tilde{b})}{\partial\rho_\ell} = \tilde{M} \frac{\partial\tilde{b}}{\partial\rho_\ell}. \tag{B.5}$$

Consider the term $\partial \tilde{b} / \partial \rho_\ell$. From equation (2.16) we obtain

$$\frac{\partial \tilde{b}}{\partial \rho_\ell} = \frac{\partial(A^{-1} \tilde{f})}{\partial \rho_\ell}. \quad (\text{B.6})$$

The right-hand side of (B.6) can be expanded as

$$\frac{\partial(A^{-1} \tilde{f})}{\partial \rho_\ell} = -A^{-1} \frac{\partial A}{\partial \rho_\ell} A^{-1} \tilde{f} = -A^{-1} \frac{\partial A}{\partial \rho_\ell} \tilde{b}. \quad (\text{B.7})$$

Since only the matrix B in A (see equation (2.19)) depends on ρ , the derivative $\partial A / \partial \rho_\ell$ can be computed as

$$\frac{\partial A}{\partial \rho_\ell} = \begin{pmatrix} \frac{\partial B}{\partial \rho_\ell} & 0 \\ 0 & 0 \end{pmatrix} \quad (\text{B.8})$$

where

$$\frac{\partial B(i, j)}{\partial \rho_\ell} = -\frac{1}{\rho_\ell^2} \int_{\Omega_\ell} \nabla \varphi_i \cdot \nabla \varphi_j \, dx \, dy, \quad i, j = 1, 2, \dots, N_n. \quad (\text{B.9})$$

Now the derivatives with respect to the actual measurements become

$$\begin{aligned} \tilde{M} \frac{\partial \tilde{b}}{\partial \rho_\ell} &= -\tilde{M} A^{-1} \frac{\partial A}{\partial \rho_\ell} \tilde{b} = -((A^{-1})^T \tilde{M}^T)^T \frac{\partial A}{\partial \rho_\ell} \tilde{b} \\ &= -((A^{-1}) \tilde{M}^T)^T \frac{\partial A}{\partial \rho_\ell} \tilde{b} = -\Gamma^T \frac{\partial A}{\partial \rho_\ell} \tilde{b} = -\tilde{\Gamma}^T \frac{\partial B}{\partial \rho_\ell} \tilde{\alpha} \end{aligned} \quad (\text{B.10})$$

where $\Gamma = A^{-1} \tilde{M}^T$, $\tilde{\Gamma} \in \mathbb{R}^{Q \times N_n}$ and $\tilde{\alpha} \in \mathbb{R}^{N_n \times K}$. The last steps are based on the fact that the matrix A is symmetric. In (B.10) there are two separate potential distributions, b from the true injected currents and Γ that is due to the set of measurement patterns $\tilde{M}^T = (M\tilde{C})^T$. The matrix obtained from equation (B.10) is $\mathbb{R}^{Q \times K}$ where Q is the number of the actual measurements for the single current pattern and K the number of current patterns (right-hand sides). After reshaping we obtain the ℓ th column of the Jacobian. The same procedure is carried out for each ρ_ℓ to obtain the other columns of $DU(\rho)$. A similar procedure to compute the Jacobian has been used, for example, in optical tomography and is known as the *adjoint differentiation* [3].

References

- [1] Adler A and Guardo R 1996 Electrical impedance tomography: regularized imaging and contrast detection *IEEE Trans. Med. Imaging* **15** 170–79
- [2] Alessandrini G and Rondi L 1998 Stable determination of a crack in a planar inhomogeneous conductor *SIAM J. Math. Anal.* **30** 326–40
- [3] Arridge S 1999 Optical tomography in medical imaging *Inverse Problems* **15** R41–93
- [4] Brenner S and Scott L 1994 *The Mathematical Theory of Finite Element Methods* (Berlin: Springer)
- [5] Brown B, Barber D and Seagar A 1985 Applied potential tomography: possible clinical applications *Clin. Phys. Physiol. Meas.* **6** 109–21
- [6] Cohen-Bacrie Y G C and Guardo R 1997 Regularized reconstruction in electrical impedance tomography using a variance uniformization constraint *IEEE Trans. Med. Imaging* **16** 562–71
- [7] Cheney M, Isaacson D, Newell J, Simske S and Goble J 1990 NOSER: an algorithm for solving the inverse conductivity problem *Int. J. Imaging Syst. Technol.* **2** 66–75
- [8] Cheney M, Isaacson D and Newell J C 1999 Electrical impedance tomography *SIAM Rev.* **41** 85–101
- [9] Cheng K-S, Isaacson D, Newell J and Gisser D 1989 Electrode models for electric current computed tomography *IEEE Trans. Biomed. Eng.* **36** 918–24
- [10] Connolly T and Wall D 1990 On Fréchet differentiability of some nonlinear operators occurring in inverse problems: an implicit function theorem approach *Inverse Problems* **6** 949–66

- [11] Cowles M and Carlin B 1996 Markov chain Monte Carlo diagnostics: a comparative review *J. Am. Stat. Assoc.* **91** 883–904
- [12] Dickin F and Wang M 1996 Electrical resistance tomography for process tomography *Meas. Sci. Technol.* **7** 247–60
- [13] Dijkstra A, Brown B, Leathard A, Harris N, Barber D and Endbrooke D 1993 Review: clinical applications of electrical impedance tomography *J. Med. Eng. Technol.* **17** 89–98
- [14] Dobson D and Santosa F 1994 An image enhancement technique for electrical impedance tomography *Inverse Problems* **10** 317–34
- [15] Donoho D L, Johnstone I M, Hoch J C and Stern A S 1992 Maximum entropy and the nearly black object *J. R. Stat. Ser. B* **54** 41–81
- [16] Eyüboğlu B, Pilkington T and Wolf P 1994 Estimation of tissue resistivities from multiple-electrode impedance measurements *Phys. Med. Biol.* **39** 1–17
- [17] Fiacco A and McCormick G 1990 *Nonlinear Programming. Sequential Unconstrained Minimization Techniques* 2nd edn (*Classics in Applied Mathematics 4*) (Philadelphia, PA: SIAM)
- [18] Fox C and Nicholls G 1997 Sampling conductivity images via MCMC *The Art and Science of Bayesian Image Analysis: Proc. Leeds Ann. Statistics Research Workshop* ed K V Mardia, C A Gill and R G Aykroyd pp 91–100
- [19] Fox C, Nicholls G and Palm M 2000 Efficient solution of boundary value problems for image reconstruction via sampling *J. Electron. Imaging* at press
- [20] Frerichs I, Hahn G and Hellige G 2000 Thoracic electrical impedance tomographic measurements during volume controlled ventilation-effects of tidal volume and positive end-expiratory pressure *IEEE Trans. Med. Imaging* **18** 764–73
- [21] Friedman A and Vogelius M 1989 Determining cracks by boundary measurements *Indiana Univ. Math. J.* **38** 527–56
- [22] Gelfand A E and Smith A F M 1990 Sampling based approaches to calculating marginal densities *J. Am. Stat. Assoc.* **85** 398–409
- [23] Geman S and Geman D 1984 Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images *IEEE Trans. Pattern Anal. Mach. Intell.* **6** 721–41
- [24] Gilks W, Richardson S and Spiegelhalter D 1996 *Markov Chain Monte Carlo in Practice* (London: Chapman and Hall)
- [25] Gilks W R, Roberts G O and George E I 1994 Adaptive direction sampling *Stat. Sci.* **43** 179–89
- [26] Glidewell M and Ng K 1995 Anatomically constrained electrical impedance tomography for anisotropic bodies via a two-step approach *IEEE Trans. Med. Imaging* **14** 498–503
- [27] Glidewell M and Ng K 1997 Anatomically constrained electrical impedance tomography for three dimensional anisotropic bodies *IEEE Trans. Med. Imaging* **16** 572–80
- [28] Golub G and van Loan C 1989 *Matrix Computations* (Baltimore, MD: Johns Hopkins University Press)
- [29] Haario H, Saksman E and Tamminen J 1999 Adaptive proposal distribution for random walk Metropolis algorithm *Comput. Stat.* **14** 375–95
- [30] Haario H, Tamminen J and Saksman E 1998 An adaptive Metropolis algorithm *Reports of the Department of Mathematics* 187 University of Helsinki
- [31] Hanke M and Hansen P 1993 Regularization methods for large-scale problems *Surv. Math. Ind.* **3** 253–315
- [32] Hansen P 1992 Analysis of discrete ill-posed problems by means of the L-curve *SIAM Rev.* **34** 561–80
- [33] Hastings W K 1970 Monte Carlo sampling methods using Markov chains and their applications *Biometrika* **57** 97–109
- [34] Hua P, Woo E, Webster J and Tompkins W 1991 Iterative reconstruction methods using regularization and optimal current patterns in electrical impedance tomography *IEEE Trans. Med. Imaging* **10** 621–8
- [35] Järvenpää S 1996 A finite element model for the inverse conductivity problem *Phil. Lic. Thesis* University of Helsinki, Finland
- [36] Järvenpää S and Somersalo E 1997 Impedance imaging and electrode models inverse problems in medical imaging and nondestructive testing *Proc. Conf. in Oberwolfach (Federal Republic of Germany, Feb. 1996)* ed H W Engl, A K Louis and W Rundell (New York: Springer)
- [37] Jones O C, Lin J-T, Ovacik L and Shu H 1993 Impedance imaging relative to gas-liquid systems *Nucl. Eng. Design* **141** 159–76
- [38] Kaipio J, Kolehmainen V, Vauhkonen M and Somersalo E 1998 Construction of anatomy-based priors with anisotropic characteristics with application to electrical impedance tomography *Proc. 20th Ann. Int. Conf. IEEE Eng. Med. Biol. Soc. (Hong Kong, China, Oct. 1998)* pp 1032–5
- [39] Kaipio J, Kolehmainen V, Vauhkonen M and Somersalo E 1999 Inverse problems with structural prior information *Inverse Problems* **15** 713–29

- [40] Kolehmainen V, Somersalo E, Vauhkonen P, Vauhkonen M and Kaipio J 1998 A Bayesian approach and total variation priors in 3D electrical impedance tomography *Proc. 20th Ann. Int. Conf. IEEE Eng. Med. Biol. Soc. (Hong Kong, China, Oct. 1998)* pp 1028–31
- [41] Kolehmainen V, Vauhkonen M, Karjalainen P and Kaipio J 1997 Assessment of errors in static electrical impedance tomography with adjacent and trigonometric current patterns *Physiol. Meas.* **18** 289–303
- [42] Kunst P W A, Noordegraaf A V, Hoekstra O S, Postmus P E and de Vries P M J 1998 Ventilation and perfusion imaging by electrical impedance tomography: a comparison with radionuclide scanning *Physiol. Meas.* **19** 481–90
- [43] Lampinen J, Vehtari A and Leinonen K 1999 Application of Bayesian neural network in electrical impedance tomography *Proc. IJCNN'99 (Washington, DC, USA, July 1999)*
- [44] Lampinen J, Vehtari A and Leinonen K 1999 Using Bayesian neural network to solve the inverse problem in electrical impedance tomography *SCIA 99: Proc. 11th Scandinavian Conf. on Image Analysis (Gangerlussuaq, Greenland, June 1999)* ed B K Ersboll and P Johansen pp 87–93
- [45] Lehtinen M, Päivärinta L and Somersalo E 1989 Linear inverse problems for generalized random variables *Inverse Problems* **5** 599–612
- [46] Lehtinen M S 1988 On statistical inversion theory *Theory and Applications of Inverse Problems* ed H Haario (London: Longman)
- [47] Lukas M 1998 Comparisons of parameter choice methods for regularization with discrete noisy data *Inverse Problems* **14** 161–84
- [48] Magnus W, Oberhettiger F and Soni R P 1966 *Formulas and Theorems for The Special Functions of Mathematical Physics* (Berlin: Springer)
- [49] Mann R, Dickin F J, Wang M, Dyakowski T, Williams R A, Edwards R B, Forrest A E and Holden P J 1997 Application of electrical resistance tomography to interrogate mixing processes at plant scale *Chem. Eng. Sci.* **52** 2087–97
- [50] Martin T and Idier J 1997 A FEM based nonlinear MAP estimator in electrical impedance tomography *Proc. IEEE ICIP'97*
- [51] Metropolis N, Rosenbluth A W, Rosenbluth M N, Teller A H and Teller E 1953 Equations of state calculations by fast computing machine *J. Chem. Phys.* **21** 1087–91
- [52] Mosegaard K 1998 Resolution analysis of general inverse problems through inverse monte carlo sampling *Inverse Problems* **14** 405–26
- [53] Mosegaard K and Rygaard-Hjalsted C 1999 Probabilistic analysis of implicit inverse problems *Inverse Problems* **15** 573–83
- [54] Mosegaard K and Tarantola A 1995 Monte Carlo sampling of solutions to inverse problems *J. Geophys. Res. B* **14** 12 431–47
- [55] Mueller J, Isaacson D and Newell J 1999 A reconstruction algorithm for electrical impedance tomography data collected on rectangular electrode arrays *IEEE Trans. Biomed. Eng.* **46** 1379–86
- [56] Neveu J 1965 *Mathematical Foundations of the Calculus of Probability* (San Francisco: Holden-Day)
- [57] Nicholls G K and Fox C 1998 Prior modelling and posterior sampling in impedance imaging *Bayesian Inference for Inverse Problems (Proc. SPIE vol 3459)* ed A Mohammad-Djafari (Bellingham, WA: SPIE Optical Engineering Press) pp 116–27
- [58] Nummelin E 1984 *General Irreducible Markov Chains and Non-Negative Operators* (Cambridge: Cambridge University Press)
- [59] Osterman K S, Kerner T E, Williams D B, Hartov A, Poplack S P and Paulsen K D 2000 Multifrequency electrical impedance imaging: preliminary *in vivo* experience in breast *Physiol. Meas.* **21** 99–109
- [60] Kaup F S P G and Vogelius M 1996 Method of imaging corrosion damage in thin plates from electrostatic data *Inverse Problems* **12** 279–93
- [61] Paulson K, Breckon W and Pidcock M 1992 Electrode modelling in electrical impedance tomography *SIAM J. Appl. Math.* **52** 1012–22
- [62] Pinheiro P A T, Loh W W and Dickin F J 1997 Smoothness-constrained inversion for two-dimensional electrical resistance tomography *Meas. Sci. Technol.* **8** 293–302
- [63] Qi-nian J and Zong-yi H 1999 On a *posteriori* parameter choice strategy for Tikhonov regularization of nonlinear ill-posed problems *Numer. Math.* **83** 139–59
- [64] Roberts G O, Gelman A and Gilks W R 1994 Weak convergence and optimal scaling of random walk Metropolis algorithms *Preprint* URL <http://www.stats.bris.ac.uk/MCMC/>
- [65] Roberts G O and Smith A F M 1994 Simple conditions for the convergence of the Gibbs sampler and Metropolis–Hastings algorithms *Stoch. Processes. Appl.* **49** 207–16
- [66] Shiryaev A N 1994 *Probability* 2nd edn (Berlin: Springer)
- [67] Smith A F M and Roberts G O 1993 Bayesian computation via the Gibbs sampler and related Markov chain

- Monte Carlo methods *J. R. Stat. Soc. B* **55** 3–23
- [68] Somersalo E, Cheney M and Isaacson D 1992 Existence and uniqueness for electrode models for electric current computed tomography *SIAM J. Appl. Math.* **52** 1023–40
- [69] Somersalo E, Kaipio J, Vauhkonen M, Baroudi D and Järvenpää S 1997 Impedance imaging and Markov chain Monte Carlo methods *Computational Experimental and Numerical Methods for Solving Ill-Posed Inverse Imaging Problems: Medical and Nonmedical Applications: Proc. SPIE 42nd Ann. Meeting (San Diego, USA, June–Aug. 1997)* ed R Barbour, M Carvlin and M Fiddy pp 175–85
- [70] Tamminen J 1999 MCMC methods for inverse problems *Geophysical Publications vol 49* Finnish Meteorological Institute
- [71] Tarantola A 1987 *Inverse Problems* (Amsterdam: Elsevier)
- [72] Tarantola A and Valette B 1982 Inverse problems = quest for information *J. Geophys.* 159–70
- [73] Tautenhahn U and Hämarik U 1999 The use of monotonicity for choosing the regularization parameter in ill-posed problems *Inverse Problems* **15** 1487–505
- [74] Tierney L 1994 Markov chains for exploring posterior distributions *Ann. Stat.* **22** 1701–62
- [75] Vauhkonen M 1997 Electrical impedance tomography and prior information *PhD Thesis* University of Kuopio, Kuopio, Finland
- [76] Vauhkonen M, Kaipio J, Somersalo E and Karjalainen P 1997 Electrical impedance tomography with basis constraints *Inverse Problems* **13** 523–30
- [77] Vauhkonen M, Vadsász D, Karjalainen P, Somersalo E and Kaipio J 1998 Tikhonov regularization and prior information in electrical impedance tomography *IEEE Trans. Med. Imaging* **17** 285–93
- [78] Williams R and Beck M (ed) 1995 *Process Tomography, Principles, Techniques and Applications* (Oxford: Butterworth-Heinemann)
- [79] Woo E, Hua P, Webster J and Tompkins W 1992 Measuring lung resistivity using electrical impedance tomography *IEEE Trans. Biomed. Eng.* **39** 756–60
- [80] Woo E, P H J, Webster G and Tompkins W 1993 A robust image reconstruction algorithm and its parallel implementation in electrical impedance tomography *IEEE Trans. Med. Imaging* **12** 137–46