Sampling Conductivity Images via MCMC

Colin Fox and Geoff Nicholls
Mathematics Department, Auckland University, Private Bag 92019, Auckland, New Zealand

May 1997

Abstract

Electrical impedance tomography (EIT) is a technique for imaging the conductivity of material inside an object, using current/voltage measurements at its surface. We demonstrate Bayesian inference from EIT data. A prior probability distribution modeling the unknown conductivity distribution is given. A MCMC algorithm is specified which samples the posterior probability for the conductivity given the prior and the EIT data. In order to compute the likelihood of a conductivity distribution it is necessary to solve a second order linear partial differential equation (PDE). This would appear to make the sampling problem computationally intractable. However by using a Markov chain with Metropolis-Hastings dynamics, and treating the conductivity update as a small perturbation, we are able to avoid solving the PDE for those updates which are rejected by the MCMC sampling process. For real applications the likelihood will need to be sensitive to very small changes in the state, so that the posterior distribution may be sharply peaked as well as multi-modal. The details of the Metropolis-Hastings dynamics are chosen so that ergodic behavior is displayed on useful time scales. We show that the sampling problem is tractable, and illustrate inference from a simple synthetic data set.

1 Introduction

Conductivity imaging has been a promising new imaging technique for some time now. The equipment (essentially a current source and a voltmeter) is cheap, portable and non-invasive. Electrodes are attached to the surface of an object, and a small current is applied. The electrical potential is measured (in volts) at points over the surface. This surface potential depends on the unknown conductivity of materials inside the object. In EIT we attempt to reconstruct the conductivity at each point in the object using the known currents and the observed potentials.

Consider an open two-dimensional region, Ω with boundary $\partial\Omega$. The conductivity σ (units ohms^{-1}) is an unknown function of position, $\sigma(\mathbf{x}), \mathbf{x} \in \Omega$, in the region. Electrodes producing fixed current source distributions $J_{\partial\Omega} = J_{\partial\Omega}(\mathbf{x})$ for $\mathbf{x} \in \partial\Omega$ and $J_\Omega = J_\Omega(\mathbf{x})$ for $\mathbf{x} \in \Omega$ are applied at the boundary and in the interior of the region. The dimensions of $J_{\partial\Omega}$ and J_Ω are current per unit length and current per unit area respectively. The electrical potential, ϕ, is measured along the boundary.

When measurements are made at low frequency, Ohm's law (current is proportional to potential gradient), and Kirchoff's law (current is conserved) together imply that within the region the potential satisfies the partial differential equation

$$\nabla \cdot (\sigma \nabla \phi) = -J_\Omega. \tag{1}$$

Many materials (for example, animal muscle) conduct better along some axes than others. We restrict attention to isotropicly conducting materials, so that σ is a scalar. If the region is otherwise insulated, the current crossing into $\partial\Omega$ is just $J_{\partial\Omega}$. Hence if $\mathbf{n}$ is the unit outward normal on the boundary $\partial\Omega$, the boundary current density is

$$\sigma \frac{\partial \phi}{\partial \mathbf{n}} = J_{\partial\Omega}. \tag{2}$$

This is just Ohm's Law, and corresponds to a Neumann boundary condition at $\partial\Omega$. Note that, since we are assuming conservation of current, there is a consistency condition on the applied currents. Let $d\nu(\mathbf{x})$ be the measure of the element of area at $\mathbf{x} \in \Omega$ and let $d\ell(\mathbf{x})$ be the measure of the element of length at $\mathbf{x} \in \partial\Omega$. The applied currents also satisfy Kirchoff's law,

$$\int_\Omega J_\Omega(\mathbf{x})\, d\nu(\mathbf{x}) + \int_{\partial\Omega} J_{\partial\Omega}(\mathbf{x})\, d\ell(\mathbf{x}) = 0. \tag{3}$$

In practice one does not attach current sources in the interior, so J_Ω is often zero, and this is the case we treat. The "forward problem" is then to calculate the field ϕ at each $\mathbf{x} \in \partial\Omega$ given $J_{\partial\Omega}$ and σ. That is, given an applied current distribution $J_{\partial\Omega}$ satisfying Equation (3), and a conductivity distribution σ, find a solution to Equation (1) satisfying the boundary condition Equation (2). The "inverse problem" is, given $J_{\partial\Omega}$ and ϕ satisfying Equations (1) and (2) for $\mathbf{x} \in \partial\Omega$, to find σ. We wish to solve the inverse problem.

A review of classical methods for regularizing the ill-posed inverse mapping from the data to the conductivity is given in Barber (1989). We propose Bayesian inference based on MCMC sampling of the posterior. Martin and Idier (1997) have given a Bayesian regularization of the inverse problem, extending work by Hua *et al* (1991). However, these authors use optimization methods as alternatives to sampling. They aim to return just a single state. This cannot, by itself, give any indication of the range of plausible reconstructions. No MCMC has been attempted (including simulated annealing) for an EIT likelihood determined via a PDE. We wish to see to what extent the sampling problem is tractable.

2 The Observation process and Likelihood

Suppose K electrodes are attached at the boundary of a region Ω at points $\mathbf{x}_k \in \partial\Omega, k = 1, 2 \ldots K$. Let $\mathcal{E}_\Omega = \{\mathbf{x}_1, \mathbf{x}_2 \ldots \mathbf{x}_K\}$ be the set of electrode positions. A pair of electrodes are chosen as anode and cathode and a current of magnitude I_0 is applied. A third reference electrode is chosen, and a voltmeter is connected between the reference electrode and each of the other electrodes in turn, not including the anode and cathode, and a potential is measured. Let e_n^a, e_n^c and e_n^r denote respectively the label of the anode, cathode, and reference electrode used in each of the $n = 1, 2 \ldots N$ electrode arrangements. Let $e_n = \{e_n^a, e_n^c, e_n^r\}$. Since there are $N = K(K-1)/2$ distinct anode-cathode choices, we measure N, $K-3$ component vectors, $v_n, n = 1, 2 \ldots N$. We measure one vector v_n for each of the current distributions $J_1, \ldots J_N$ generated by the $n = 1, 2 \ldots N$ distinct current electrode pairs. Note that these measurements are not all independent. We have a total of $K(K-1)/2 \times (K-3)$ values in all, while only $K(K-1)/2$ are independent, as we will see. For notational convenience we let v_n be a K component vector, so that $v_n(k)$ is the potential measured at the kth electrode in the nth electrode configuration. The redundant components $v_n(k)\ k \in e_n$ are left undefined, since no measurement is made at these electrodes.

Let the components of the current due to the n'th electrode pair be $J_n = \{J_{\Omega,n}, J_{\partial\Omega,n}\}$, corresponding to the interior and boundary components of the applied current. We model the electrodes as point sources of magnitude I_0, located at the boundary, so $J_{\Omega,n} = 0$ and

$$J_{\partial\Omega,n}(\mathbf{x}) = I_0 \delta(\mathbf{x} - \mathbf{x}_{e_n^a}) - I_0 \delta(\mathbf{x} - \mathbf{x}_{e_n^c}). \tag{4}$$

We assume the electrode positions are measured with negligible uncertainty, and defer discussion of electrode position uncertainty and possible internal electrode structure to later work, in which we treat real data.

Let $s(\mathbf{x})$ be the true conductivity at each point $\mathbf{x} \in \Omega$, and let $\phi(\mathbf{x}|n, s)$ be the potential at $\mathbf{x}$ generated by current distribution J_n in the conductivity s. Refer to Figure 1b for an example. Let $\phi(s)$ denote the N component vector $(\phi(\mathbf{x}|1, s), \phi(\mathbf{x}|2, s) \ldots \phi(\mathbf{x}|N, s))$. We suppose the voltage measurements are made in noise that is additive and Gaussian, mean zero and standard deviation d. Thus if $\epsilon_n(k)$ is a Gaussian random variable, $\epsilon_n(k) \sim N(0, d^2)$, we observe

$$v_n(k) = \phi(\mathbf{x}_k|n, s) + \epsilon_n(k), \qquad \text{for each } n = 1..N, k = 1..K, k \notin e_n. \tag{5}$$

The potential at the reference electrode is $\phi(\mathbf{x}_k|n, s) = 0$. The solution to the PDE, Equation (1), is determined only up to an additive constant, and by setting $\phi(\mathbf{x}_k|n, s) = 0$ we fully determine the solution. Now, given applied currents J and data v, the likelihood of a conductivity distribution $\sigma(\mathbf{x})$ is proportional to $\exp(L(\sigma))$, where $L(\sigma)$ is equal to the log-likelihood,

$$L(\sigma) = -|v - \phi(\sigma)|/2d^2, \tag{6}$$

up to a term independent of σ, and

$$|v - \phi(\sigma)| \equiv \sum_{n=1}^{N} \sum_{\substack{k=1 \\ k \notin e_n}}^{K} (v_n(k) - \phi(\mathbf{x}_k|n, \sigma))^2.$$

In order to compute $\phi(\mathbf{x}_k|n,\sigma)$ we must solve the forward problem for the conductivity σ and the current J_n. Thus, to compute the likelihood, we must solve Equations (1) and (2) for $\phi(\mathbf{x}|n,\sigma)$, with $J_{\Omega,n}=0$, that is,

$$\nabla_x \cdot(\sigma\nabla_x\phi(\mathbf{x}|n,\sigma))=0 \tag{7}$$

with

$$\sigma\partial\phi(\mathbf{x}|n,\sigma)/\partial\mathbf{n}_x=J_{\partial\Omega,n}, \tag{8}$$

for each $n=1,2\ldots N$. One of the main features of EIT imaging is that information loss occurs in the mapping $\phi(s)$ from the unknown conductivity s to the observable electrode potentials. Even if the potentials are measured perfectly at the electrodes at the boundary, there may still be a great deal of uncertainty in the reconstructed conductivity in the interior.

3 Computing the Likelihood

Let current distributions $J=(J_1,J_2\ldots J_N)$ and a conductivity σ be given. In order to compute the likelihood of σ, we need the fields $\phi(\mathbf{x}|n,\sigma)$ for $\mathbf{x}\in\mathcal{E}_\Omega$. That is, we need the field at the electrodes only.

Let $g(\mathbf{x}|\mathbf{y},\sigma)$ be the Neumann Green's function for the potential at $\mathbf{x}\in\Omega$ given a point-like current source located at $\mathbf{y}\in\Omega$, with a uniform current sink along the boundary. That is, $g(\mathbf{x}|\mathbf{y},\sigma)$ satisfies the PDE

$$\nabla_x \cdot(\sigma\nabla_x g(\mathbf{x}|\mathbf{y},\sigma))=-\delta^{(2)}(\mathbf{x}-\mathbf{y}) \tag{9}$$

with

$$\sigma\partial g(\mathbf{x}|\mathbf{y},\sigma)/\partial\mathbf{n}_x=1/|\partial\Omega|, \tag{10}$$

where $|\partial\Omega|$ is the length of the boundary. When $\mathbf{y}\in\partial\Omega$ these relations become

$$\nabla_x \cdot(\sigma\nabla_x g(\mathbf{x}|\mathbf{y},\sigma))=0 \tag{11}$$

and

$$\sigma\partial g(\mathbf{x}|\mathbf{y},\sigma)/\partial\mathbf{n}_x=-\delta^{(1)}(\mathbf{x}-\mathbf{y})+1/|\partial\Omega|. \tag{12}$$

Thus when $\mathbf{y}\in\Omega$ (respectively, $\mathbf{y}\in\partial\Omega$) the Green's functions are defined by Equations (9) and (10) (respectively Equations (11) and (12)). In either case, the Green's function is determined up to an additive constant. We fix this "gauge" by referring our Green's functions to a mean zero voltage over the boundary,

$$\int_{\partial\Omega} g(\mathbf{x}|\mathbf{y},\sigma)\,dl(\mathbf{x})=0. \tag{13}$$

It is straightforward to show that the Green's functions we have defined are symmetric, ie, $g(\mathbf{y}|\mathbf{x},\sigma)=g(\mathbf{x}|\mathbf{y},\sigma)$, although this depends delicately on the boundary conditions, Equations (10) and (12), and on the choice of reference potential made in Equation (13).

Since the potential fields, $\phi(\mathbf{x}|n,\sigma)$, $n=1,2\ldots N$, satisfy Equations (7) and (8), they are given in terms of the Green's functions by

$$\begin{aligned}\phi(\mathbf{x}|n,\sigma) &= \int_\Omega g(\mathbf{x}|\mathbf{y},\sigma)J_{\Omega,n}(\mathbf{y})\,d\nu(\mathbf{y})+\int_{\partial\Omega} g(\mathbf{x}|\mathbf{y},\sigma)J_{\partial\Omega,n}(\mathbf{y})\,d\ell(\mathbf{y})-\phi_0(\mathbf{x}_{e_n^r}|n,\sigma) &(14)\\ &= I_0 g(\mathbf{x}|\mathbf{x}_{e_n^a},\sigma)-I_0 g(\mathbf{x}|\mathbf{x}_{e_n^c},\sigma)-\phi_0(\mathbf{x}_{e_n^r}|n,\sigma) &(15)\end{aligned}$$

using Equation (4), where

$$\phi_0(\mathbf{x}_{e_n^r}|n,\sigma)\equiv I_0 g(\mathbf{x}_{e_n^r}|\mathbf{x}_{e_n^a},\sigma)-I_0 g(\mathbf{x}_{e_n^r}|\mathbf{x}_{e_n^c},\sigma).$$

We subtract $\phi_0(\mathbf{x}_{e_n^r}|n,\sigma)$ so that the potential at the reference electrode located at $\mathbf{x}_{e_n^r}$ is zero. It follows that, in order to compute the $N=K(K-1)/2$ fields $\phi(\mathbf{x}|n,\sigma)$, we must compute K Green's functions, $g(\mathbf{x}|\mathbf{x}_k,\sigma)$, $k=1,2\ldots K$, one for each point, $\mathbf{x}_k$, at which we attach an anode or cathode. The potential fields are given as linear combinations of this set of "fundamental solutions" by Equation (15). Since the fields we measure depend on just

K Green's functions, no potential measurement can be made at the source electrode, and the Greens functions are symmetric, $K(K-1)/2$ is the number of genuinely independent measurements that can be made on K electrodes.

In the examples which follow, we solve for the Green's functions in a square region. We discretize the region in an M-element regular square array (mesh) of pixels labeled $m = 1, 2 \dots M$. All functions in Ω are approximated by a corresponding function piecewise-constant over pixels. Thus the conductivity $\sigma(\mathbf{x})$, the fields $\phi(\mathbf{x}|n, \sigma)$, and the Green's functions $g(\mathbf{x}|\mathbf{y}, \sigma)$ have discrete partners σ_m, $\phi_m(n, \sigma)$ and $g_m^{m'}(\sigma)$. For the results presented a simple finite difference scheme was implemented in C and called to solve for the Green's functions. We would do better to concentrate mesh nodes where the solution is a sensitive function of the conductivity and electrode position. We might also choose better basis functions to represent the solution within elements. However the discretization we have given seems to be adequate for simple prototype problems. Notice that the conductivity may be represented via a higher level description (as in Model C below) in which case the pixel conductivity σ_m will be determined by the local conductivity value $\sigma(\mathbf{x})$ in the higher level description.

Computing all K Green's functions is a brute force approach and imposes a big computational burden. One of the main tasks in a Bayes/sampling approach to EIT is to find ways to avoid this calculation at each evaluation of a likelihood ratio, and this is what we have done.

4 Modeling the Conductivity

We now specify a model for the conductivity distribution in Ω. We assume the true conductivity s is piecewise constant, and takes a finite number $\mathcal{C}$ of fixed values $C_1, C_2 \dots C_{\mathcal{C}}$ values in Ω. We use a simple clustering model for the conductivity: we represent the conductivity in Ω on the same square mesh we use to solve for $\phi(\sigma)$, using the same symbol σ for the vector of lattice values σ_m, $m \in 1, 2 \dots M$ and the distribution function $\sigma(\mathbf{x})$, $\mathbf{x} \in \Omega$. Let Σ_Ω denote the set of all possible pixel-based conductivity distributions σ, so that $\sigma \in \Sigma_\Omega$. We model σ on the lattice as an Ising or Potts binary or higher order Markov random field, with distribution

$$\Pr(\sigma) \propto \exp\left(\theta \sum_{m=1}^{M} H_m(\sigma)\right) \qquad \text{where} \quad H_m(\sigma) \equiv \sum_{m' \sim m} \delta_{\sigma_m, \sigma_{m'}}. \tag{16}$$

The sum over $m' \sim m$ is a sum over the four neighbors m' of m on the square lattice. The function $\delta_{a,b}$ is the indicator function for the event $a = b$. The posterior for a conductivity distribution σ, given observations v is then $\Pr(\sigma|v) \propto \Pr(\sigma)\exp(L(\sigma))$. We intend to sample the distribution $\Pr(\sigma|v)$ and use these samples to answer questions concerning the unknown true conductivity s.

The sampling framework which we describe below can be extended to cover more sophisticated priors. Indeed, there would be little point in setting up such a sampler without taking advantage of the flexibility it allows in the modeling stage of the analysis. We proceed now to simply list a few possible improvements which could be made to the prior model.

Model A The present model. The types of possible materials in the object are known but not their proportion in area or location. The conductivities of the materials are known exactly, and correspond to the values $C_1, C_2 \dots C_{\mathcal{C}}$.

Model B The types of possible materials (labeled $c = 1, 2 \dots \mathcal{C}$) are known. The conductivities of these materials vary from one instance to another, but are assumed to have a single uniform value in the sample at hand. The conductivity values C_c, $c = 1, 2 \dots \mathcal{C}$, are supposed to be normally distributed random variables.

Model C We may, in addition, have knowledge about the structure, shape, and arrangement of objects in the regions. For example the region may be known to contain man-made objects, with outlines that are composed of simple geometric structures, such as line-segments and circular arcs. In this case one may use an intermediate or higher level prior, as described, for example, in Grenander and Miller (1993).

In this paper we report results for Model A only.

5 Markov Chain Monte Carlo

We now define a Markov chain $\{X_t\}_{t=0}^{\infty}$ of random variables X_t taking values in Σ_Ω, and having a t-step distribution $\Pr(X_t = \sigma | X_0 = \sigma^{(0)})$ which tends to the posterior distribution, $\Pr(\sigma|v)$, as t tends to infinity. Such a process will generate the samples we require. We use the Metropolis-Hastings method for our Markov chain Monte Carlo as Gibbs sampling is not practicable, on account of the form of the likelihood. In order to get ergodic behavior on useful time scales we found it necessary to include several different types of update in the MCMC.

The state is composed of regions of uniform conductivity. Referring ahead, to Figure 1d, we see that much of the uncertainty in the reconstructed state occurs at change points in the conductivity. We therefore focus updates on pixels in regions where the conductivity is non-uniform. This idea can be found in Ripley (1987). There is a second reason why we want to do this. The likelihood is a very sensitive function of pixel conductivity. Changes at conductivity boundaries often make only small changes to the likelihood, and therefore stand a better chance of being accepted. Move 3 below is an example.

The simplest way to build multiple moves into a Metropolis-Hastings sampler is to define a separate reversible transition probability *for each move type*, as in Green (1995). The method is favored over using a single transition probability: in particular, when several moves are allowed to connect the same pair of states, the condition for reversibility is simplified. Let P be the number of moves, and let $\{\Pr^{(p)}(X_{t+1} = \sigma_{t+1} | X_t = \sigma_t)\}_{p=1}^{P}$ be a set of P transition probabilities, reversible with respect to the distribution $\Pr(\sigma|v)$. Let $\zeta_p, p = 1, 2 \ldots P, \sum \zeta_p = 1$, be the probability to choose move p. We select a transition probability at random and use it to generate a MCMC update. The overall transition probability is

$$\Pr(X_{t+1} = \sigma_{t+1} | X_t = \sigma_t) = \sum_{p=1}^{P} \zeta_p \; \Pr{}^{(p)}(X_{t+1} = \sigma_{t+1} | X_t = \sigma_t).$$

If at least one of the moves is irreducible on Σ_Ω, then the equilibrium distribution of the chain is $\Pr(\sigma|v)$, independent of the choice, $\sigma^{(0)}$, of starting state.

We now list the moves used. The current state is $X_t = \sigma$, and the move is used to determine a value for X_{t+1}. The candidate state proposed in the Metropolis-Hastings scheme is labeled σ'. A pixel n is a *near-neighbor* of pixel m if their lattice distance is less than $\sqrt{8}$. A pixel has twenty four near-neighbors, unless it is on or next to the boundary. An *update-edge* is a pair of near-neighboring pixels on the lattice. Let $\mathcal{N}^*(\sigma)$ be the set of update-edges connecting two pixels of unequal conductivity in σ. Let $\mathcal{N}^*_m(\sigma)$ be the set of edges in $\mathcal{N}^*(\sigma)$ involving the pixel m. We maintain a list of the update-edges in $\mathcal{N}^*(\sigma)$, for the current state σ of the Markov chain. We use this to select at random pairs of pixels with unequal conductivities. Let $|S|$ be the number of elements in the set S.

Move 1 *Flip a pixel*. Select a pixel m uniformly at random from $1, 2 \ldots M$ and assign σ_m a new conductivity chosen uniformly at random from the other $\mathcal{C} - 1$ possible conductivity values. Call this new value σ'_m and the new conductivity distribution σ'. σ' differs from σ at just one pixel. The probability to generate the move is $q_1(\sigma \to \sigma') = 1/M(\mathcal{C} - 1)$. The acceptance probability is

$$\alpha_1(\sigma \to \sigma') = \min\left\{1, e^{2\theta(H_m(\sigma') - H_m(\sigma))} \times e^{L(\sigma') - L(\sigma)}\right\}.$$

Move 2 *Flip a pixel near a conductivity boundary*. Pick an update-edge at random from $\mathcal{N}^*(\sigma)$. Pick one of the two pixels in that edge at random, pixel m say. Pixel m can be selected in $|\mathcal{N}^*_m(\sigma)|$ possible ways. Proceed as in Move 1, generating a candidate state σ'. The overall probability to generate the move is $q_2(\sigma \to \sigma') = |\mathcal{N}^*_m(\sigma)|/2|\mathcal{N}^*(\sigma)|(\mathcal{C} - 1)$. The acceptance probability is

$$\alpha_2(\sigma \to \sigma') = \min\left\{1, e^{2\theta(H_m(\sigma') - H_m(\sigma))} \times e^{L(\sigma') - L(\sigma)} \times \frac{|\mathcal{N}^*_m(\sigma')||\mathcal{N}^*(\sigma)|}{|\mathcal{N}^*_m(\sigma)||\mathcal{N}^*(\sigma')|}\right\}.$$

Move 3 *Swap conductivities at a pair of pixels*. Pick an update-edge at random from $\mathcal{N}^*(\sigma)$. Suppose it connects pixels m and n with conductivities $\sigma_m = a$ and $\sigma_n = b$. The candidate state σ' is equal to σ except that $\sigma'_m = b$ and $\sigma'_n = a$. The generation probability is $q_3(\sigma \to \sigma') = 1/|\mathcal{N}^*(\sigma)|$. The acceptance probability is

$$\alpha_3(\sigma \to \sigma') = \min\left\{1, e^{2\theta(H_m(\sigma') - H_m(\sigma) + H_n(\sigma') - H_n(\sigma))} \times e^{L(\sigma') - L(\sigma)} \times |\mathcal{N}^*(\sigma)| \,/\, |\mathcal{N}^*(\sigma')|\right\}.$$

Account has been taken, in the formula for α_3, of the possibility that m and n may be neighbors.

Notice that Move 2 duplicates Move 1, but allows us to focus updates on the interface between regions of differing conductivity as in Ripley (1987). Move 1 is still needed, for irreducibility. Note also that there is no useful simplification of the likelihood ratio $\exp(L(\sigma') - L(\sigma))$, as occurs whenever the posterior random field is Markov. This is because the log-likelihood is not an additive function of σ, but depends on the conductivity at each pixel through the non-linear field functions $\phi(\sigma)$.

6 Markov Chain Monte Carlo with an EIT Likelihood

In order to compute the acceptance probabilities, α, at an update, we need $L(\sigma')$. This, as was shown in Section 3, involves solving Equations (11) and (12) for the Green's functions $g(\mathbf{x}|\mathbf{x}_k, \sigma')$ in the candidate conductivity σ', and computing $\phi(\mathbf{x}|n, \sigma')$ and $L(\sigma')$ using Equations (15) and (6). We have described how we solve for the Green's functions. This operation takes too much time to be feasible at each step of the Markov chain. Notice, however, that the two conductivities σ and σ' differ only slightly, and that we have already calculated the Green's functions $g(\mathbf{x}|\mathbf{x}_k, \sigma)$ in the current conductivity σ. Let $\delta g(\mathbf{x}|\mathbf{x}_k, \sigma, \sigma') \equiv g(\mathbf{x}|\mathbf{x}_k, \sigma') - g(\mathbf{x}|\mathbf{x}_k, \sigma)$ and $\delta\sigma(\mathbf{x}) \equiv \sigma'(\mathbf{x}) - \sigma(\mathbf{x})$. In terms of these quantities we have, from Equation (11), approximately,

$$\nabla_x \cdot (\sigma \nabla_x \delta g(\mathbf{x}|\mathbf{x}_k, \sigma, \sigma')) \simeq \nabla_x \cdot (\delta\sigma \nabla_x g(\mathbf{x}|\mathbf{x}_k, \sigma)) \,. \tag{17}$$

The Green's functions for the operator on the left will satisfy Equations (11) and (12) as before. The term on the right hand side of Equation (17) may be regarded as a current source. If we take $Q(\mathbf{x}|\mathbf{x}_k\sigma, \sigma') \equiv \nabla_x \cdot (\delta\sigma \nabla_x g(\mathbf{x}|\mathbf{x}_k, \sigma))$, we have, as in Equation (14),

$$\delta g(\mathbf{x}|\mathbf{x}_k, \sigma, \sigma') \simeq \int_\Omega g(\mathbf{y}|\mathbf{x}, \sigma) Q(\mathbf{y}|\mathbf{x}_k, \sigma, \sigma') d\nu(\mathbf{y}) \tag{18}$$

We have exploited the symmetry $g(\mathbf{x}|\mathbf{y}, \sigma) = g(\mathbf{y}|\mathbf{x}, \sigma)$ of our Green's functions in the final step. The important point is that, in order to compute the likelihood, we need the field values at electrodes only. It follows that we need only compute the perturbing field, $\delta g(\mathbf{x}|\mathbf{x}_k, \sigma, \sigma')$ for $\mathbf{x} \in \mathcal{E}_\Omega$, ie, for $\mathbf{x}$ corresponding to an electrode position. But then the Green's functions on the right hand side of Equation (18) are precisely the Green's functions $g(\mathbf{x}|\mathbf{x}_k, \sigma)$ we have available already. Moreover, if the conductivity change $\delta\sigma(\mathbf{x})$ is local, as is usual in MCMC, $Q(\mathbf{y}|\mathbf{x}_k, \sigma, \sigma')$ is zero over much of Ω. When the integral is reduced to a sum on the lattice variables, this sum is local, involving just a few variables around the point where the change in conductivity is made.

The local linearization we have used is good if the term $\nabla_x \cdot (\delta\sigma \nabla_x \delta g(\mathbf{x}|\mathbf{x}_k, \sigma, \sigma'))$ dropped from Equation (17) is small. Similar approximation schemes are possible for Models of type B and C. However, if the conductivity levels C_c, C_{c+1} differ too greatly, for a given PDE mesh, the approximation will be poor. We have only tried this straightforward linearization, however we note that other authors have found that other parameterizations of σ give linear approximations with a greater range of accuracy (Martin and Idier 1997) and these may improve our method.

The MCMC algorithm is as follows. Given a state $X_t = \sigma$, given the Green's functions $g(\mathbf{x}|\mathbf{x}_k, \sigma)$, for $\mathbf{x} \in \Omega$ and $\mathbf{x}_k \in \mathcal{E}_\Omega$, $k = 1, 2 \ldots K$, and given $L(\sigma)$, the log-likelihood of the state, proceed as follows.

Step I Select a move p with probability ζ_p. Generate σ' by sampling $q_p(\sigma \to \sigma')$.

Step II Using the local linearization of the present section, estimate $L(\sigma')$.

Step III Compute the acceptance probability for the state, $\alpha_p(\sigma \to \sigma')$.

Step IV Accept or reject the candidate state σ':

- **IV.i** If σ' is rejected, set $X_{t+1} = \sigma$. Since there is no change to the state, the Green's functions and log-likelihood are unchanged, available for the next iteration.
- **IV.ii** If σ' is accepted, set $X_{t+1} = \sigma'$. We must now solve Equations (11) and (12) to compute $g(\mathbf{x}|\mathbf{x}_k, \sigma')$, for $k = 1, 2 \ldots K$, and update the log-likelihood to $L(\sigma')$.

To summarize, given the Green's functions for current sources in a conductivity σ, we may estimate the likelihood of a candidate conductivity σ' generated in the Metropolis-Hastings procedure, without solving the PDE again. Metropolis-Hastings updates which lead to rejection are local operations, involving the conductivity values at just a few pixels neighboring the point at which a change is made. However, if σ' is accepted, we must solve the PDE, as we now wish to perturb about the fields of the conductivity σ'.

7 Experiments

The algorithm described in Section 6 was implemented in C. The MCMC software was tested (to an accuracy of about 1 part in 1000) by taking an independent sequence of MCMC runs, and varying between runs those parameters of the Markov process which leave the equilibrium distribution undisturbed. In particular we can use different sets of ζ-values in different runs. In this way we vary the proportions of the moves used. We can then check, by comparing the values of statistics estimated from the output, that all the moves preserve the same stationary distribution. We can also vary the number of update-edges at each pixel. We experimented with four, eight and twenty four update-edges, corresponding to near-neighbors out to 1, $\sqrt{2}$ and $\sqrt{8}$ lattice units.

In the following example we start with the known conductivity distribution, $s(\mathbf{x})$, shown in Figure 1a, along with the locations of the sixteen electrodes. The conductivity model and the finite difference PDE solver use the same 25×25 conductivity lattice. We condition on the value of the conductivity at the boundary itself, so $X_t(\mathbf{x}) = s(\mathbf{x})$ for $\mathbf{x} \in \partial\Omega$ and all t. The two grey-levels correspond to two conductivities ($\mathcal{C} = 2$) of $C_1 = 1$ and

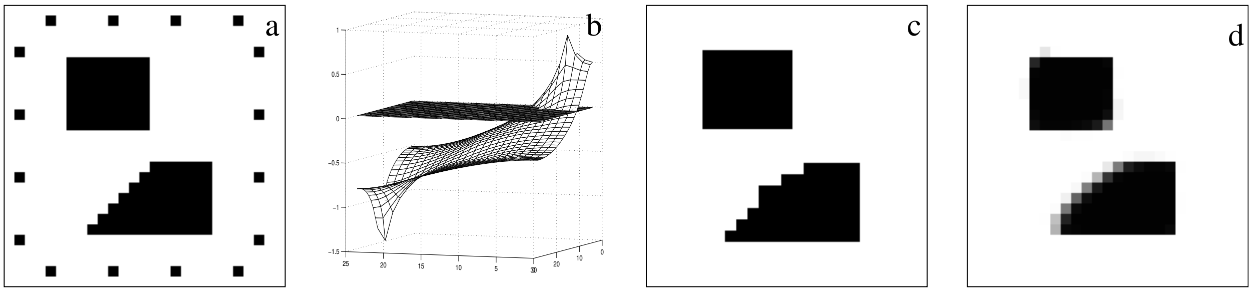

Figure 1: Recovering conductivity information for a synthetic data set, computed from (a). The dots around the boundary in (a) are the electrode locations. The graph in (b) is the potential field $\phi(\mathbf{x}|n, s)$ for the conductivity in (a), for one of the $n = 1, 2 \ldots 120$ fields. The state in (c) was obtained after about 10 minutes time on the reference machine, (see below). The mean conductivity from a long run is shown in (d).

$C_2 = 2$ conductivity units. A binary Markov random field is a reasonable prior for distributions of this kind, and we choose a prior parameter value $\theta = 1$. The potential fields $\phi(\mathbf{x}|n, s), n = 1, 2 \ldots N$ are computed for $\mathbf{x} \in \Omega$. The computed potential for one such electrode pair is shown in Figure 1b. Artificial data v is generated by adding Gaussian noise to the potential observed at each electrode, as in Equation (5). We choose to test observations with standard deviation $d = 0.005$, a signal to noise ratio of around $1500 : 1$, taking account of the $(K - 3)$-fold over-measurement. Hua *et al* (1991) find such SNR's physically achievable. We run MCMC using moves 1, 2 and 3 in the proportions 1:9:81, with the chain initialized to a random conductivity at each pixel. The state shown in Figure 1c was reached after 100 sweeps (1 sweep = M Markov chain steps), which took less than about 10 minutes on readily available hardware (a Pentium Pro 200 costing about three thousand pounds in 1997). The posterior mean conductivity at each pixel from a run of 50,000 sweeps is shown in Figure 1d. We might just as easily plot the variance at each pixel, or any other summarizing statistic of interest. Has the MCMC converged ? The log-likelihood is plotted against sweeps in Figure 2. Standard tests shown no sign of slow mode-hopping, and give an integrated auto-covariance time (for the number of pixels at a given conductivity level) of $27(\pm 5)$ sweeps. We are getting an effective independent sample each 90 seconds. Other statistics were considered and have similar behavior. The acceptance rate for updates is low, only 0.6%. The high computational cost of an acceptance is balanced against the more common but faster rejections. The approximation Equation (18) is therefore most important, as the PDE need only be solved in the event that an update is accepted. The approximation leads to an error rate of 14%: that is, 14% of the acceptances would not occur if the likelihood was calculated without approximation.

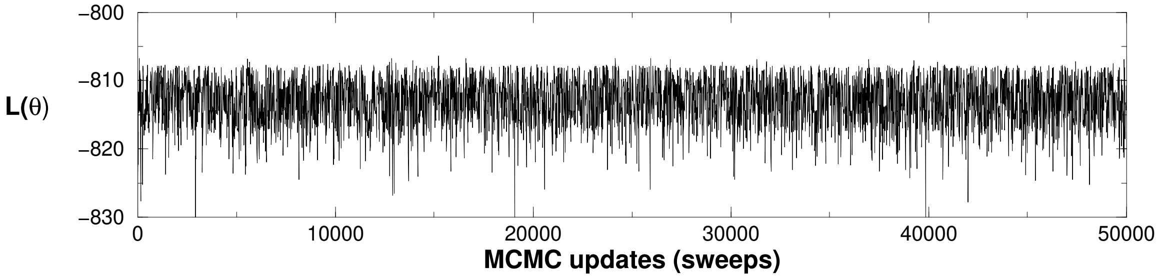

Figure 2: The sample log-likelihood (up to a constant independent of σ) plotted against the update number.

8 Conclusions

We found that a Bayesian approach to inverting EIT data, with analysis based on posterior samples obtained via MCMC simulation, was feasible, for simple two-level synthetic problems. We have experimented with conductivity distributions in which there are three or more conductivity levels, using a multi-level Markov random field or Potts model, and again, found it possible to treat simple structures with the methods described here. However, several problems remain.

At noise levels low enough to define more delicate structures, the MCMC sticks in unrepresentative states. It seems to be possible, in many cases, to recover ergodicity by adding new types of moves to the MCMC. Viewing the state of the MCMC, X_t, evolving as t increases can be helpful. Pixels near the boundary cause particular difficulty at the convergence stage. Indeed, a Bayesian/sampling approach may be inappropriate for reconstruction near the boundary, where there is little uncertainty in the conductivity. However, we expect that higher level models, of the kind described in Section 4, Model C, and explicitly in this workshop, will resolve the problem in a more general way, for at least some electrode arrangements and conductivity distributions. Higher level models remove or change the character of the ambiguity in the reconstruction. This ambiguity corresponds to multi-modality in the posterior, which is the immediate difficulty.

It is clearly necessary to test the approach on real data. With synthetic data, the effects of discretization are canceled, since the forward and inverse problems are solved under the same approximation. Real data introduces a number of new problems, such as modeling electrode structure, and allowing for uncertainty in electrode position, which will greatly increase the computational burden of computing a likelihood ratio. However there remain a number of ways in which our computational procedure may be accelerated, and so we consider it at least possible that the sampling problem will continue to be tractable.

References

Barber, D.C. (1989). A review of image reconstruction techniques for electrical impedance tomography *Med. Phys.*, 16:162–169.

Green, P. (1995). Reversible jump MCMC and Bayesian model determination. *Biometrika*, 82:711–732.

Grenander, U. and Miller, M. (1993). Representation of knowledge in complex systems (with discussion). *Journal of the Royal Statistical Society B*, 56:549–604.

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. Imag.*, 621–628.

Martin, T. and Idier, J. (1997). A Bayesian non-linear approach for electrical impedance tomography Technical Report, Lab. des Signaux at Systems, CNRS, Plateau de Moulon, 91192 Gif-sur-Yvette Cedex, France.

Ripley, B. D. (1987). *Stochastic Simulation*. Cambridge University Press, Cambridge.