Prior Modelling and Posterior sampling in Impedance Imaging

G. K. Nicholls and C. Fox

Mathematics Department, The University of Auckland, Auckland, New Zealand

ABSTRACT

We examine sample based Bayesian inference from impedance imaging data. We report experiments employing low level pixel based priors with mixed discrete and continuous conductivities. Sampling is carried out using Metropolis-Hastings Markov chain Monte Carlo, employing both large scale, Langevin updates, and state-adaptive local updates. Computing likelihood ratios of conductivity distributions involves solving a second order linear partial differential equation. However our simulation is rendered computationally tractable by an update procedure which employs a linearization of the forward map and thereby avoids solving the PDE for those updates which are rejected.

Keywords: Impedance tomography, Markov chain Monte Carlo, Bayesian, Langevin, Inverse problem

1. INTRODUCTION

Conductivity imaging, also known as electrical impedance tomography (EIT), is not yet established as a workable imaging technique. The best reconstructions obtained to date[1] have very low resolution and are hard to interpret. The question is, is this low resolving power an inherent weakness of the physical imaging process, or is it an artifact of the data analysis. We are unable to give a final answer, since we have not yet obtained real data for trial analysis. However our experiments with synthetic data, which we present below, indicate substantial improvements on existing inversion methods are to be made.

We begin by defining the physical imaging process. Referring to Figure 1(a), electrodes are placed in contact with an object, and a small current is applied through some of these electrodes. One of the non-current electrodes is selected as a reference electrode, and the potential difference between each of the remaining non-current electrodes and the reference is measured in volts. These measured potential differences depend on the unknown conductivity distribution of material inside the object. In conductivity imaging we attempt to reconstruct the conductivity at each point in the object using the known currents and the observed potentials.

The definitions which follow are from Fox and Nicholls.[2] Consider the two-dimensional region, Ω with boundary $\partial\Omega$ shown in Figure 1(a). The conductivity σ (units ohms^{-1}) is an unknown function of position, $\sigma(x), x \in \Omega$, in the region. It may in practice be bounded above and below so that $0 < \sigma_{\min} \leq \sigma(x) \leq \sigma_{\max} < \infty$. Electrodes producing fixed current source distributions $J_{\partial\Omega} = J_{\partial\Omega}(x)$ for $x \in \partial\Omega$ and $J_\Omega = J_\Omega(x)$ for $x \in \Omega$ 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. Within the region the potential satisfies the partial differential equation

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

If the region is otherwise insulated, the current crossing $\partial\Omega$ is just $J_{\partial\Omega}$. Hence if n_x is the unit outward normal on the boundary $\partial\Omega$, the boundary current density is

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

Suppose K electrodes are attached at the boundary of a region Ω at points $x_k \in \partial\Omega, k = 1, 2 \ldots K$. Let $\mathcal{E}_\Omega = \{x_1, x_2 \ldots x_K\}$ 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, and let $e_n = \{e_n^a, e_n^c, e_n^r\}$. Since there are $N = K(K-1)/2$ distinct

CF: E-mail: fox@math.auckland.ac.nz, GKN: E-mail: nicholls@math.auckland.ac.nz
Further info at `http://www.math.auckland.ac.nz/~nicholls/linkfiles/papers.html`

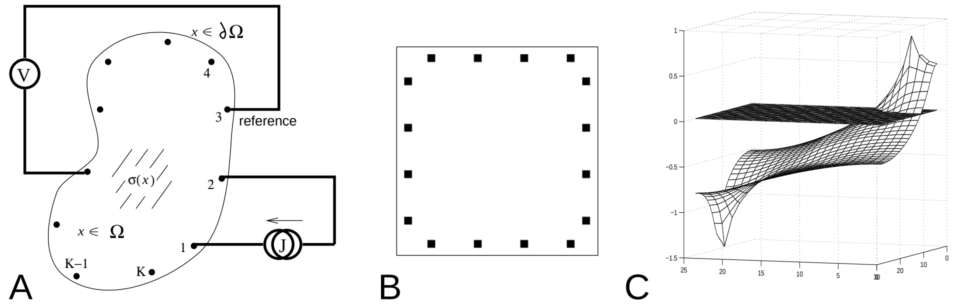

Figure 1. (a) Region Ω of spatially varying conductivity and with bounding curve $\partial\Omega$. The choice of anode, cathode and reference electrode shown corresponds to the choice $e_n^a = 1, e_n^c = 2$ and $e_n^r = 3$. The full measurement set is made up of all electrode potentials, measured in turn for all distinct choices of anode and cathode. The index n indicates a set of measurements associated with some particular choice of anode and cathode. (b) Electrode positions used in the experiments section. The graph in (c) is the potential field $\phi(x|n, s)$ for the conductivity of Figure 3(a), for one of the $n = 1, 2 \dots 120$ electrode arrangements.

anode-cathode choices, we measure N, $K-3$ component vectors, $v_n, n = 1, 2 \dots N$. Let the current due to the n'th electrode pair be $J_n = \{J_{\Omega,n}, J_{\partial\Omega,n}\}$. The electrodes are modelled as point sources of magnitude I_0, located in the interior, so $J_{\partial\Omega,n}(x) = 0$ and

$$J_{\Omega,n}(x) = I_0\delta(x - x_{e_n^a}) - I_0\delta(x - x_{e_n^c}). \tag{3}$$

We measure one vector v_n for each of the current distributions $J_1, \dots J_N$ generated by the N distinct current electrode pairs. Note that these measurements are not all independent. We have a total of $(K-3) \times K(K-1)/2$ 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.

Let $s(x)$ be the true conductivity at each point $x \in \Omega$, and let $\phi(x|n, s)$ be the potential at x generated by current distribution J_n in the conductivity s. Refer to Figure 1(c) for the potential function of an applied current distribution in the conductivity of Figure 3(a). Let $\phi(s)$ denote the N component vector of functions $(\phi(x|1, s), \dots \phi(x|N, s))$. We suppose the voltage measurements are made in noise that is independent, additive and Gaussian, mean zero and standard deviation d, whilst currents are measured exactly. Thus if $\epsilon_n(k)$ is a Gaussian random variable, $\epsilon_n(k) \sim N(0, d^2)$, we observe

$$v_n(k) = \phi(x_k|n, s) + \epsilon_n(k), \qquad \text{for each } n = 1..N, k = 1..K, k \notin e_n. \tag{4}$$

The potential at the reference electrode is $\phi(x_{e_n^r}|n, s) \equiv 0$. The solution to the PDE, Equation (1), is determined only up to an additive constant, and by setting $\phi(x_{e_n^r}|n, s) = 0$ we fully determine the solution. Now, given applied currents $J = (J_1, \dots J_N)$ and data $v = (v_1, \dots v_N)$, the likelihood of a conductivity distribution $\sigma(x)$ is proportional to $\exp(L(\sigma|v, J))$, where $L(\sigma) \equiv L(\sigma|v, J)$ is the log-likelihood, up to a term independent of σ,

$$L(\sigma) = -|v - \phi(\sigma)|/2d^2, \qquad \text{with} \qquad |v - \phi(\sigma)| \equiv \sum_{n=1}^{N} \sum_{\substack{k=1 \\ k \notin e_n}}^{K} (v_n(k) - \phi(x_k|n, \sigma))^2. \tag{5}$$

In order to compute the likelihood, we must solve Equations (1) and (2) for $\phi(x|n, \sigma)$,

$$\nabla_x \cdot (\sigma \nabla_x \phi(x|n, \sigma)) = I_0\delta(x - x_{e_n^a}) - I_0\delta(x - x_{e_n^c}) \tag{6}$$

with

$$\sigma \partial\phi(x|n, \sigma)/\partial n_x = 0, \tag{7}$$

for each $n = 1, 2 \dots N$. The posterior for a conductivity distribution σ, given observations v, generated by currents J, is then $\Pr\{s \in d\sigma|v, J\} \propto \Pr\{s \in d\sigma\} \exp(L(\sigma|v, J))$. We intend to sample the distribution $\Pr\{s \in d\sigma|v, J\}$ and use these samples to answer questions concerning the unknown true conductivity s.

2. PROPERTIES OF THE INVERSE PROBLEM

For a fixed conductivity $\sigma(x)$ and for interior sources J_Ω, the *Dirichlet* boundary value problem $\nabla \cdot (\sigma \nabla \phi) = J_\Omega$ in Ω subject to boundary conditions $\phi|_{\partial\Omega} = V$ has a unique solution, ϕ, which has current crossing the boundary equal to $J_{\partial\Omega} = \sigma \frac{\partial \phi}{\partial n}$. The map $\Gamma_\sigma : V \mapsto J_{\partial\Omega}$ is called the Dirichlet to Neumann map for given σ (and J_Ω). Note that the differential equation is linear in ϕ and, hence, so is Γ_σ. However, the relationship between σ and Γ_σ is not linear. The *forward problem* is then to determine Γ_σ given σ. One simply solves the PDE. The *inverse problem* is: determine σ from incomplete measurements $(J_{\Omega,n}, J_{\partial\Omega,n}, v_n)_{n=1}^{N}$ of Γ_σ.

A number of existence and uniqueness results[3] hold for the inverse image of the nonlinear forward map, $\sigma \mapsto \Gamma_\sigma$. However, the issue of uniqueness of the inverse, *per se*, is not of practical importance and greater insight may be gained by examining the singular value decomposition (SVD) of the the linearised forward problem. The right- and left-singular vectors give the bases of conductivity and data, respectively, in which the linearised forward map is diagonalised and hence reduced to component-wise multiplication by the singular values. With respect to those bases, the singular values give the magnitude of the directional derivative of the forward map in each coordinate direction and hence give the sensitivity of the forward map to small changes in conductivity.

The linearised forward map for imaging the conductivity of a disc, assuming *complete* measurement of Γ_σ, has been shown to have singular values that are arbitrarily close to zero[4] and hence the inverse is discontinuous. So the idealised inverse problem is ill-posed. In the practical case of reconstructing a pixelised version of σ from a finite number of data, the decreasing singular values cause the inverse problem to be very badly conditioned. For imaging the conductivity of a disc which is close to homogeneous, with measurement accuracy of 0.1%, about 140 singular values exceed the threshold mentioned above and thus there are about 140 independent structures of the conductivity about which measurements can be made.[4] Thus if direct inversion of the forward problem is attempted, any image with more than 140 pixels must have correlations between pixels that are solely functions of the particular noise realisation – and those spurious structures will tend to be large as a result of the small singular values. On the other hand, reconstructing an image of conductivity with few pixels, e.g. 140, will not give good reconstructions as the forward problem is poorly approximated unless the true conductivity happens to be piece-wise constant with the same structure as the pixilation.

3. PREVIOUS APPROACHES TO CONDUCTIVITY IMAGING

Despite the inverse problem being very badly conditioned, with a theoretically discontinuous inverse, the vast majority of effort into conductivity imaging has been to implement algorithms that seek to simply invert the forward operator. It is not surprising, therefore, that the literature contains a conspicuous paucity of reconstructed images.

A review paper[5] in 1989 Barber outlined a number of the direct-inversion schemes tried to that date. All of those schemes sought to directly invert using a single, or iterated, inversion of an approximation to the nonlinear forward map, $\sigma \mapsto \Gamma_\sigma$, with the approximation relying on a single trial value of conductivity. An obviously good choice for an approximation to the forward map is a local linearisation about a known conductivity as given by the Fréchet derivative. However, even though this linearisation was known early on,[6] many early algorithms used *ad hoc* approximate inverses often loosely based on 'back-projection' as used in x-ray tomography.

More promising have been the direct inversion schemes using a *local* linearisation of the nonlinear forward map to iterate towards a solution. Barber concludes that Yorkey's[7] method was the best. Yorkey's method does not employ a *global* regulariser on the reconstructed solution, but only a *local* regulariser on each linearised step. When examined in detail, almost all of the iterative algorithms can be found to regularise the inversion of the forward map implicitly in the numerical scheme used as opposed to the desirable situation of having an explicitly chosen means of dealing with the ill-posedness.

In the statistical framework, the direct solution of the inverse problem is a ML estimate and suffers from all the usual problems when the forward mapping has a large range of sensitivities such as the reconstructed image consisting mainly of amplified noise. But a deeper problem with the algorithms outlined can be seen through each of them being equivalent to some gradient-based algorithm for maximising the likelihood – or some implicitly modified likelihood. Gradient-based methods require the space of allowable solutions to be connected and convex because iteration is performed by moving along search directions derived from local information. Thus it is very difficult to include high-level prior information such as "the conductivity at each point is one of five values" or "the image is a finite number of circular regions", etc, as would be desirable.

Since 1990 a number of authors have recognised that the inverse problem requires global regularisation and have optimised an appropriately modified objective function. An example of classical regularisation within an iterative scheme is given by Hua *et al.*[8] who propose using Sobolev-type norms as regularisers and provide results using the first of these, namely the square-norm of the conductivity, using a Newton-Raphson technique to perform the resulting optimization. Recently, Martin and Idier[9] have compared Hua *et al.*'s method with one using a Markov random field as prior within a Bayesian statement of the inverse problem. They used the sum of the convex Huber penalty function of nearest-neighbour differences for their log-prior and calculated a MAP estimate (from an improper posterior) using a gradient based optimiser. Since their prior prefers local smoothness with discontinuities, the reconstructions provided of conductivities of this type are improved. As far as the present authors are aware, the few other Bayesian formulations of conductivity imaging have been restricted to calculating MAP estimates with low-level priors and are, therefore, not fundamentally different to regularised inversion.

4. 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 $\phi(x|n,\sigma)$ for $x \in \mathcal{E}_\Omega$. That is, we need the field at the electrodes only. Let $g(x|y,\sigma)$ be the Neumann Green's function for the potential at $x \in \Omega$ given a point-like current source located at $y \in \Omega$, with a uniform current sink along the boundary. That is, $g(x|y,\sigma)$ satisfies the PDE

$$\nabla_x \cdot (\sigma \nabla_x g(x|y,\sigma)) = \delta^{(2)}(x-y) \tag{8}$$

with

$$\sigma \partial g(x|y,\sigma)/\partial n_x = 1/|\partial\Omega|, \tag{9}$$

where $|\partial\Omega|$ is the length of the boundary. The Green's function is determined up to an additive constant. We fix this by referring our Green's functions to a mean zero voltage over the boundary, that is, we impose

$$\int_{\partial\Omega} g(x|y,\sigma)\, d\ell(x) = 0. \tag{10}$$

where $d\ell(x)$ is the element of length measure along $\partial\Omega$. It is straightforward to show that the Green's functions we have defined are symmetric, *ie*, $g(y|x,\sigma) = g(x|y,\sigma)$, although this depends delicately on the boundary conditions, Equation (9), and on the choice of reference potential made in Equation (10).

Since the potential fields, $\phi(x|n,\sigma)$, $n = 1, 2 \ldots N$, satisfy Equations (6) and (7), they are given in terms of the Green's functions by

$$\begin{aligned}\phi(x|n,\sigma) &= \int_\Omega g(x|y,\sigma) J_{\Omega,n}(y)\, d\nu(y) + \int_{\partial\Omega} g(x|y,\sigma) J_{\partial\Omega,n}(y)\, d\ell(y) - \phi_0(x_{e_n^r}|n,\sigma) \\ &= I_0 g(x|x_{e_n^a},\sigma) - I_0 g(x|x_{e_n^c},\sigma) - \phi_0(x_{e_n^r}|n,\sigma)\end{aligned} \tag{11}$$

using Equation (3), where $d\nu(x)$ is the element of area measure in Ω, and

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

We subtract $\phi_0(x_{e_n^r}|n,\sigma)$ so that the potential at the reference electrode located at $x_{e_n^r}$ is zero. It follows that, in order to compute the $N = K(K-1)/2$ fields $\phi(x|n,\sigma)$, we must compute K Green's functions, $g(x|x_k,\sigma), k = 1, 2 \ldots K$, one for each point, 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 (11). 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 discretise the region in an M-element regular square array (mesh) of pixels labelled $m = 1, 2 \ldots M$ assigning the pixel conductivity σ_m according to conductivity at the center of the pixel. Thus the conductivity $\sigma(x)$, the fields $\phi(x|n,\sigma)$, and the Green's functions $g(x|y,\sigma)$ have corresponding discrete representations σ_m, $\phi_m(n,\sigma)$ and $g_m^{m'}(\sigma)$ satisfying difference equations and summation conditions derived in the simplest possible way from Equations (8), (9) and Equation (10). Note that 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, $m \in 1, 2 \ldots M$ and the distribution function $\sigma(x)$, $x \in \Omega$. For the results presented a simple finite difference scheme was implemented in C and called to solve for the Green's functions.

5. MODELLING THE CONDUCTIVITY

We wish to specify a prior distribution for the probability that some trial conductivity σ coincides with the unknown true conductivity s in Ω. We must specify model variables, their state space, and probability distribution or density. These prior models may also be regarded as representing varying states of prior knowledge. In this paper we present results for Models A and B. See the conclusions for further remarks concerning Model C.

Model A There are $\mathcal{C}$ possible types of material in the region. The conductivities of the materials are known exactly, and correspond to the values $C = \{C_1, C_2 \dots C_{\mathcal{C}}\}$. In this case $\sigma_m \in C$. We model σ on the lattice as an Ising or Potts binary or higher order Markov random field, with distribution

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

The sum over $m' \sim m$ is a sum over the four immediate neighbours m' of m on the square lattice. The function $\delta_{a,b}$ is the indicator function for the event $a = b$. Let $\Sigma_\Omega^A = C^M$ denote the set of all possible pixel-based conductivity distributions σ in model A so that in model A, $\sigma \in \Sigma_\Omega^A$.

Model B The types of possible materials (labelled $1, 2 \dots \mathcal{C}$) making up the object are known. However, the conductivities of these materials vary from pixel to pixel within Ω. Let $\tau(x), x \in \Omega$ be some estimate of the material type field. Corresponding lattice variables τ_m, $m = 1, 2 \dots, M$ with state space $\tau_m \in \{1, 2 \dots \mathcal{C}\}$ give the estimated material type at pixel m. Let $t(x), x \in \Omega$ denote the unknown true material type field in Ω. Let $\Sigma_\Omega^B = \{\Re^+\}^M$ and $T_\Omega^B = \{1, 2 \dots \mathcal{C}\}^M$ denote the state space of σ and τ in model B, so that in model B, $\sigma \in \Sigma_\Omega^B$ and $\tau \in T_\Omega^B$. Let $d\sigma = d\sigma_1 d\sigma_2 \dots d\sigma_M$. Conditional on the material type τ, we suppose that the prior distribution of σ is multivariate normal, subject to the condition $\sigma_m > 0, m = 1, 2 \dots, M$, with mean $\mu(\tau)$ and variance $\rho(\tau)$ depending on the material type, and correlation parameter $\tilde{\theta}$. Let $N_0(r; \mu, \rho)$ denote the normal density, mean μ, variance ρ, but normalised over $r > 0$ only. Conditional on material type the prior distribution is then

$$\Pr\{s \in d\sigma | \tau\} \propto \exp\left(-\tilde{\theta} \sum_{m=1}^{M} G_m(\sigma)\right) \times \prod_{m'=1}^{M} N_0(\sigma_{m'}; \mu(\tau_{m'}), \rho(\tau_{m'}))\, d\sigma \quad \text{where} \quad G_m(\sigma) \equiv \sum_{m'' \sim m} (\sigma_m - \sigma_{m''})^2$$

In this model the $\mathcal{C}$ parameters $\mu(c), c \in \{1, 2 \dots \mathcal{C}\}$ replace the $\mathcal{C}$ parameters $\{C_1, C_2 \dots C_{\mathcal{C}}\}$ of model A. The field τ representing the spatial distribution of material type has a Potts or Ising prior $\Pr\{t = \tau\}$, identical to $\Pr\{s = \sigma\}$ in model A. The overall prior is $\Pr\{t = \tau, s \in d\sigma\} = \Pr\{t = \tau\}\Pr\{s \in d\sigma | \tau\}$.

Model C We may, in addition, have knowledge about the structure, shape, and arrangement of objects in the region. For example it may be known to contain man-made objects, with outlines that are composed of simple polygonal structures. An intermediate level prior, modelling the underlying continuum process[10,11] is appropriate.

6. MARKOV CHAIN MONTE CARLO

We wish to sample the posterior distributions obtained under models A and B. Model A may be regarded as a special case of B in which $\tilde{\theta} = 0$ and $\rho(\tau_i)$ is very small compared to $|\mu(\tau_i) - \mu(\tau_{i-1})|$ and $|\mu(\tau_{i+1}) - \mu(\tau_i)|$ for each $i = 1, 2 \dots \mathcal{C}$. We define a Markov chain $\{X_n\}_{n=0}^{\infty}$ of random variables X_n taking values in $\Sigma_\Omega^B \times T_\Omega^B$, and having an n-step distribution with the property $\lim_{n\to\infty} \Pr\{X_n \in (\tau, d\sigma) | X_0 = (\tau^{(0)}, \sigma^{(0)})\} = \Pr\{t = \tau, s \in d\sigma | v\}$. Such a process will generate the samples we require. We use Metropolis-Hastings updates. In order to get ergodic behaviour on useful time scales we found it necessary to include several different types of update or move-types in the MCMC. A move is a stochastic operation used to determine a value for X_{n+1} given a value for X_n.

The simplest way to build multiple moves into a Metropolis-Hastings sampler is to define a separate reversible transition probability *for each move type.*[12] Let P be the number of moves, and let $\{\Pr^{(p)}\{X_{n+1} \in (\tau_{n+1}, d\sigma_{n+1}) | X_n = (\tau_n, \sigma_n)\}\}_{p=1}^{P}$ be a set of P transition probabilities, reversible with respect to the posterior distribution $\Pr\{\tau, d\sigma | v, J\}$. Let $\zeta_p, p = 1, 2 \dots P$ be the probability to choose move p. Necessarily, $\sum \zeta_p = 1$. We select move p with probability ζ_p and use it to generate a MCMC update. The overall transition probability is

$$\Pr\{X_{n+1} \in (\tau_{n+1}, d\sigma_{n+1}) | X_n = (\tau_n, \sigma_n)\} = \sum_{p=1}^{P} \zeta_p \Pr{}^{(p)}\{X_{n+1} \in (\tau_{n+1}, d\sigma_{n+1}) | X_n = (\tau_n, \sigma_n)\}.$$

If at least one of the moves is $\Pr\{\tau, d\sigma|v, J\}$-irreducible on $\Sigma_\Omega^B \times T_\Omega^B$, then the equilibrium distribution of the chain is $\Pr\{\tau, d\sigma|v, J\}$,[13] independent of X_0, the MCMC initialisation. We list the moves used in an Appendix A.

Assertion For each $p = 1, 2 \ldots 7$ let

$$\Pr{}^{(p)}\{X_{n+1} \in (\tau_{n+1}, d\sigma_{n+1})|X_n = (\tau_n, \sigma_n)\} \equiv q_p((\tau, \sigma) \to (\tau', d\sigma')) \times \alpha_p((\tau, \sigma) \to (\tau', \sigma'))$$

for $q_p((\tau, \sigma) \to (\tau', d\sigma'))$ and $\alpha_p((\tau, \sigma) \to (\tau', \sigma'))$ the proposal and acceptance distributions defined in Appendix A. For each $p = 1, 2 \ldots 7$, the transition kernel $\Pr^{(p)}\{X_{n+1} \in (\tau_{n+1}, d\sigma_{n+1})|X_n = (\tau_n, \sigma_n)\}$ is reversible with respect to $\Pr\{\tau, d\sigma|v, J\}$. Also, the transition kernel for move 1 defines a process $\Pr\{\tau, d\sigma|v, J\}$-irreducible on $\Sigma_\Omega^B \times T_\Omega^B$.

7. MARKOV CHAIN MONTE CARLO WITH AN EIT LIKELIHOOD

In this section we refer to a "current" process state, $X_n = (\tau, \sigma)$, and a "candidate" state, labelled (τ', σ'), proposed in the course of a Metropolis-Hastings update. In order to compute the acceptance probabilities, α_p, at an update, we need $L(\sigma')/L(\sigma)$. As we saw in Section 4, this involves solving Equations (8) and (9) for the Green's functions $g(x|x_k, \sigma')$ in the candidate conductivity σ', and computing $\phi(x|n, \sigma')$ and $L(\sigma')$ using Equations (11) and (5). 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 for moves 1-4, at least, the two conductivities σ and σ' differ only slightly, and that we have already calculated the Green's functions $g(x|x_k, \sigma)$ in the current conductivity σ. In a previous paper[2] we have shown how the Green's functions for σ may be used to estimate the Green's functions in σ', and therefore the log-likelihood $L(\sigma')$. Our estimate is good when σ' and σ differ only slightly.

The MCMC algorithm is as follows. Given a state $X_n = (\tau, \sigma)$, given the Green's functions $g(x|x_k, \sigma)$, for $x \in \Omega$ and $x_k \in \mathcal{E}_\Omega$, $k = 1, 2 \ldots K$, and given $L(\sigma)$, the log-likelihood of the state, X_{n+1} is determined in the following way.

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

Step II Using the local linearization,[2] estimate $L(\sigma')$.

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

Step IV Accept or reject the candidate state (τ', σ'):

- **IV.i** If (τ', σ') is rejected, set $X_{t+1} = \sigma$. There is no change to the state so the Green's functions and log-likelihood are unchanged, available for the next iteration.
- **IV.ii** If (τ', σ') is accepted, set $X_{t+1} = \sigma'$. We must now solve Equations (8) and (9) K times in order to compute $g(x|x_k, \sigma')$, for $k = 1, 2 \ldots K$, and update the log-likelihood to $L(\sigma')$.

To summarise, 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 neighbouring 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'.

We conclude with a comment on computing the Metropolised Langevin update. Referring to Equation (13), the key computation is of $\partial\phi(x_k|n, \sigma)/\partial\sigma_m$. We must compute the change in the field at electrodes caused by a small change in the conductivity within the region. This is effected without solving the PDE, in a computation of order constant in time for each σ_m, using the same linearisation as before[2] to compute the response of the Green's functions to a small change in σ at each pixel $m = 1, 2 \ldots M$ in turn. In a process in which the computation of likelihood ratios is expensive, "large" updates, involving changes to a large number of pixels, become attractive. Usually such moves are not feasible, as randomly generated candidate states tend to have low probability, and are therefore rejected. However Langevin updates generate candidate states which step in the direction of locally decreasing posterior potential, leading to acceptable acceptance rates if the jump size h is chosen appropriately.[14] In the process we define, the Langevin update has a very low acceptance rate for reasonable h values, probably on account of the extreme sensitivity of the likelihood to variations in the conductivity near the boundary. It may be appropriate to normalise the Langevin update vector $\nabla \log(f(\tau, \sigma|v))$, or work with a spatially varying parameter h. Either revision is straightforward in a Metropolis adjusted diffusion. Other schemes exist for selecting large updates which have good acceptance rates, notably micro-canonical updates.[15]

8. EXPERIMENTS

The algorithm described in Section 7 was implemented in C. The MCMC software was tested in the usual way.[2] In each of the following examples we start with a known, or "true", conductivity distribution, $s(x)$. This is imaged with 16 electrodes located as shown in Figure 1(b). The conductivity model and the finite difference PDE solver use the same 25×25 conductivity lattice. We condition on the value of the conductivity around the electrodes, so $X_n(x) = s(x)$ for all $n = 0, 1, 2 \ldots$ and all x inside a strip of width 2 pixels neighbouring the border. Details of all the experiments we present are shown in Table 1. Artificial data v is generated by adding Gaussian noise to the potential observed at each electrode, as in Equation (4). 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. Measurements of this quality are physically achievable.

We sample the posterior with Markov Chain Monte Carlo, initialising the chain in all cases with a random pattern of conductivity (and material types, Model B and Figure 4). We compute the mean and variance of the sampled conductivity at each pixel. These statistics are displayed as images, along with a sample from the posterior, in the figures which follow. The sample we present is simply the state of the MCMC process at the end of a long sequence of updates of fixed length. A number of statistics are recorded along the run, including the log-likelihood $L(\sigma)$, the posterior potential, and the proportion of pixels of each conductivity value (Model A) or type (Model B). Convergence is rather slow, though it seems to be reliable, as can be judged from the output statistics plotted in Figure 2. The integrated autocorrelation time (IACT) of the log-likelihood (τ_{LL} say) was computed. This quantity is small in efficient MCMC. Roughly speaking, τ_{LL} is the number of correlated MCMC samples from the posterior with the variance-reducing power of one independent sample. We find that the IACT does indeed drop when moves 5 and 6 are included alongside moves 1-4.

In the experiment from the synthetic data of Figure 3(a), three conductivity values ($C_1 = 1$, $C_2 = 2$ and $C_2 = 3$) appear. Referring to the variance image, Figure 3(d), the reconstruction of the square region with conductivity value 2 units is somewhat uncertain. The data may now be fitted by replacing part of this region by a small region in which the conductivity value is 2 units, and shrinking it somewhat. In an earlier experiment (presented elsewhere[2]), treating synthetic *binary* conductivity patterns, observed under the same conditions as here, and reconstructed using the prior knowledge of Model A, we found we were able to recover $\sigma(x)$ essentially perfectly, and with negligible variation in the posterior conductivities at pixels. The uncertainty in the recovered conductivity distribution increases as the number of allowed conductivity levels increases, as Figure 3 shows. However, the mean conductivity remains a good estimator for the true conductivity $s(x)$.

When we move to the more realistic state of prior knowledge represented by Model B, the ambiguity in the reconstruction increases again. The simulation summarised in Figure 4 illustrates this observation. The prior

Table 1. Run parameters for figures. Blank entries are not applicable or not of interest. The second row corresponds to a run sampling the same posterior distribution as in the first row, but with varied move proportions.

Figure	Model	move fractions % $\zeta_1 : \zeta_2 : \zeta_3 : \zeta_4 : \zeta_5 : \zeta_6$	$\theta, \tilde{\theta}$	$\mu(1), \mu(2), \mu(3)$ or C_1, C_2, C_3	$\rho(1), \rho(2), \rho(3)$
3	A	1:9:50:40:0:0	0.75,-	1.0, 2.0, 3.0	- - -
(no image)	A	1:9:40:40:5:5	0.75,-	1.0, 2.0, 3.0	- - -
4	B	1:9:40:40:0:0	0.0,0.5	1.0, 2.0, 3.0	0.2 0.2 0.2
5	A	1:9:40:40:5:5	0.75,-	1.0, 2.0, 3.0	- - -

Table 2. Output statistics for Figures. Missing entries are not applicable or not of interest. Bracketed quantities indicate standard errors. The first two rows repeat the same simulation with varied move proportions. This is done to test the algorithm implementation. Estimated means, $\mathsf{E}\{L(\sigma)\}$, should agree. 1 sweep is 441 updates, the number of pixels in the lattice which are subject to updates. τ_{LL} is the IACT of the log-likelihood, see text. Timing is given for a simple serial C-implementation of the sampling algorithm on a machine with SPECfp95 equal 12.4

Figure	τ_{LL} sweeps	τ_{LL} minutes	$\mathsf{E}\{L(\sigma)\}$
3 (a-d)	100(10) sweeps	4	-810.9 (0.1)
3 (no image)	45(10) sweeps	10	-811.0 (0.2)
4 (g,h,i,j)	170(40) sweeps	5.5	-

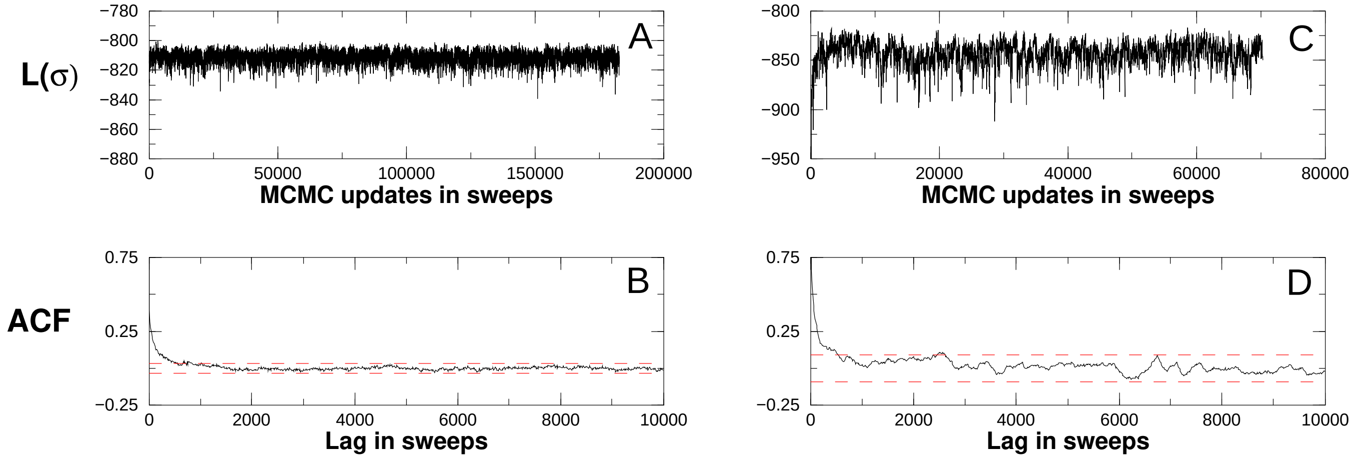

Figure 2. The sample log-likelihood (up to a constant independent of σ) plotted against the update number, along with the autocorrelation function (ACF) of the MCMC output series. (A) Model A, series used to estimate results in Figure 3. (B) ACF of series in (A). Dashed lines indicate variance of ACF asymptotic in the lag. (C) Model B, for Figure 4g-j. (D) ACF of series in (C). Otherwise as for (B). The quality of convergence in Figure 4g-j is poor as can be seen from (D).

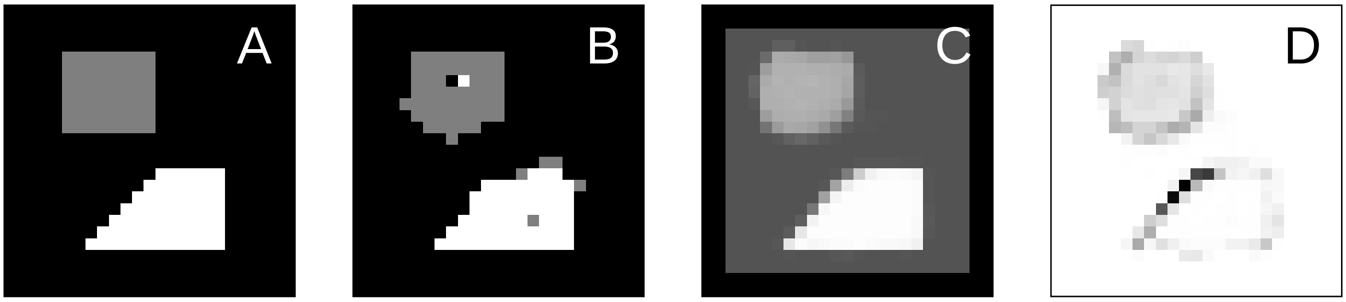

Figure 3. Recovering conductivity information for a synthetic data set with three conductivity levels, computed from (a), electrode positions as in Figure 1a. Conductivities C_1, C_2 and C_3 are black grey and white respectively. The state in (b) is a sample from the posterior, the mean conductivity is shown in (c), and pixel conductivity variance is shown in (d), where darker shades indicate greater variance.

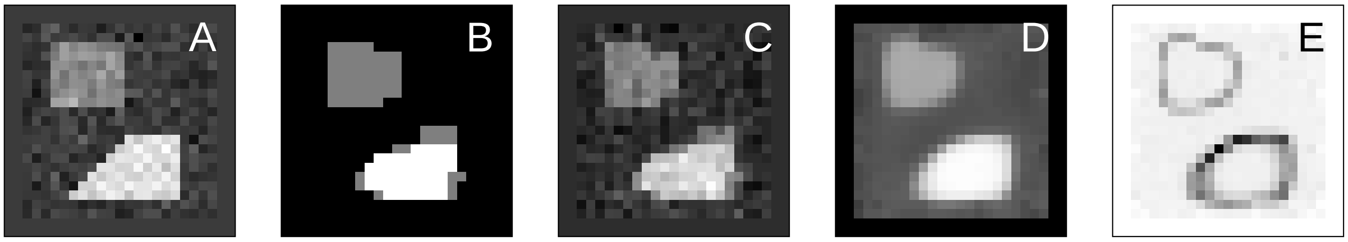

Figure 4. Recovering conductivity information for a synthetic data set, computed from (a) with continuous, varying, conductivity pixel variables. Lighter shades indicate higher conductivity. Material type as for Figure 3(a). The state in (b) gives sampled reconstructed material type classifications while (c) gives sampled reconstructed conductivity values. Mean conductivity values are shown in (d), and pixel conductivity variance is shown in (e).

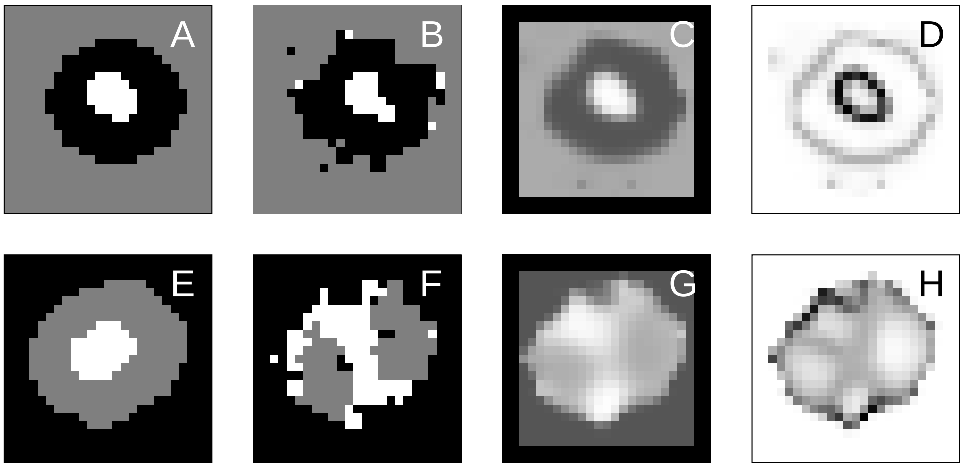

Figure 5. Indications of shielding leading to increased uncertainty in sampled reconstructed conductivity. Recovering conductivity information for a synthetic data set with three conductivity levels. (a) gives true conductivity for (b)-(d), (e) gives true conductivity for (f)-(h). (a,e) true conductivity, (b,f) sample reconstructed conductivity from posterior, (c,g) mean conductivity and (d,h) pixel conductivity variance.

knowledge of material type correlation between adjacent pixels becomes important in lowering the probability of reconstructions in which the border of blocks of uniform material type loose definition. We find we are able to sample the posterior adequately using moves 1-4. However, the state-adaptivity of moves 1-4 is important. Simpler move sets not focusing on the discontinuity set of the type field τ prove to be inadequate.[16]

Finally, certain patterns of conductivity can screen internal structure. The true conductivity distributions shown in Figures 5(a) and (e) are similar, except that in (a) the conductivity levels on an inward bound path go medium-low-high (*ie*, $2.0-1.0-3.0$) whilst in (e) the sequence is low-medium-high (*ie*, $1.0-2.0-3.0$). The lack of contrast in the interior makes the conductivity distribution of Figure 5(e) relatively difficult to reconstruct. The posterior variance is large in the interior, and the structure there is determined largely by the prior.

9. CONCLUSIONS

We find that a Bayesian approach to inverting EIT data, with analysis based on posterior samples obtained via MCMC simulation, is feasible, for simple synthetic problems. This is a promising result. It suggests that it will be possible to explore real EIT data using well developed statistical methods and visualisation devices. A relatively sophisticated adaptive sampling algorithm is required. This algorithm is modular, *ie* built of separate move types. Its main elements will be useful in MCMC outside Bayesian EIT. The high computational cost of an acceptance in our MCMC algorithm is balanced against the more common but faster rejections. The PDE need only be solved in the event that an update is accepted. Thus we found moves 1-4 adequate for sampling despite an overall acceptance rate around 2%. Moves 5 and 6, which use sub-chains of MCMC updates, are expensive to compute. They solve initial convergence problems in certain difficult (low-noise) examples we have explored.

We expect that higher level models, of the kind described in Section 5, Model C, and explicitly elsewhere in this conference, will resolve these problems 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. Ambiguities caused by screening also may be ameliorated if extra high level information is available. Our own initial experiments in this line with relatively complex intermediate level (continuum mosaic[11,17]) models have themselves shown poor convergence. We expect that simpler continuum models[18] will be more effective.

APPENDIX A. MARKOV CHAIN MONTE CARLO UPDATES

In this Appendix we detail the moves used to define the MCMC. Each move defines a transition kernel reversible with respect to the posterior distribution. The state is composed of regions of uniform conductivity. Referring to

Figure 3(d), we see that much of the uncertainty in the reconstructed state occurs at change points in the conductivity. We therefore construct moves which focus updates on pixels in regions where the conductivity is non-uniform.[19] The expressions for the acceptance probabilities α_p, $p = 1, 2 \ldots 7$ given for the following moves are calculated using the usual Metropolis-Hastings formula.[20]

A pixel n is a *near-neighbour* of pixel m if their lattice distance is less than $\sqrt{8}$. A pixel has twenty four near-neighbours, unless it is on or near the boundary. An *update-edge* is a pair of near-neighbouring 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 $|\mathcal{N}^*(\sigma)|$ be the number of elements in set $\mathcal{N}^*(\sigma)$. Let $S_4 \equiv \zeta_1 + \zeta_2 + \zeta_3 + \zeta_4$.

Move 1 *Flip a pixel.* Select a pixel m uniformly at random from $1, 2 \ldots M$ and assign τ'_m a conductivity type chosen uniformly at random from the $\mathcal{C} - 1$ conductivity types not equal to the original type τ_m. Next, a conductivity value $\sigma'_m \sim N_0(\sigma'_m; \mu(\tau'_m), \rho(\tau'_m))$ is selected. Let (τ', σ') be the candidate state thus generated. τ' differs from τ at just one pixel; σ' and σ likewise. The probability to generate the move is

$$q_1((\tau, \sigma) \to (\tau', d\sigma')) = 1/M(\mathcal{C} - 1) \times N_0(\sigma'_m; \mu(\tau'_m), \rho(\tau'_m))d\sigma'$$

The acceptance probability is

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

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 (τ', σ'). The overall probability to generate the move is

$$q_2((\tau, \sigma) \to (\tau', d\sigma')) = \frac{|\mathcal{N}^*_m(\sigma)|}{2|\mathcal{N}^*(\sigma)|(\mathcal{C} - 1)} \times N_0(\sigma'_m; \mu(\tau'_m), \rho(\tau'_m))d\sigma'.$$

The acceptance probability is

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

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 conductivity types $\tau_m = a$ and $\tau_n = b$. The candidate state τ' is equal to τ except that $\tau'_m = b$ and $\tau'_n = a$. The conductivities σ_m and σ_n are likewise swapped to generate candidate state σ'. The generation probability is

$$q_3((\tau, \sigma) \to (\tau', \sigma')) = 1/|\mathcal{N}^*(\sigma)|.$$

The acceptance probability is

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

The formula for α_3 remains correct as it stands when m and n are neighbours.

Move 4 *Assign new conductivities to each of a pair of pixels.* Pick an update-edge at random from $\mathcal{N}^*(\sigma)$. Suppose it connects pixels m and n with conductivity types $\tau_m = a$ and $\tau_n = b$. Assign new random conductivity types $\tau'_m = a'$ and $\tau'_n = b'$ chosen uniformly at random from $\{1, 2 \ldots \mathcal{C}\}^2$, subject to the conditions $a' \neq a$, $b' \neq b$ (avoid duplicating move 2) and $a' \neq b'$ (for reversibility), at each pixel. Assign new conductivities distributed like $\sigma'_m \sim N(\mu(\tau'_m), \rho(\tau'_m))$ and $\sigma'_n \sim N(\mu(\tau'_n), \rho(\tau'_n))$. The probability to generate the move is

$$q_4((\tau, \sigma) \to (\tau', d\sigma')) = \frac{1}{|\mathcal{N}^*(\sigma)|(\mathcal{C}^2 - 3\mathcal{C} + 2)} \times N_0(\sigma'_m; \mu(\tau'_m), \rho(\tau'_m)) \times N_0(\sigma'_n; \mu(\tau'_n), \rho(\tau'_n))d\sigma'_m d\sigma'_n.$$

If pixel m and n are not neighbours then $\alpha_4((\tau, \sigma) \to (\tau', \sigma')) = \alpha_3((\tau, \sigma) \to (\tau', \sigma'))$. If m and n are neighbours,

$$\alpha_4((\tau, \sigma) \to (\tau', \sigma')) = \min\left\{1, \frac{e^{2\theta(H_m(\tau') + H_n(\tau'))}}{e^{2\theta(H_m(\tau) + H_n(\tau))}} \times \frac{e^{-2\tilde{\theta}(G_m(\sigma') - (\sigma'_m - \sigma'_n)^2 + G_n(\sigma'))}}{e^{-2\tilde{\theta}(G_m(\sigma) + (\sigma_m - \sigma_n)^2 - G_n(\sigma))}} \times e^{L(\sigma') - L(\sigma)} \times \frac{|\mathcal{N}^*(\sigma)|}{|\mathcal{N}^*(\sigma')|}\right\}.$$

Move 5 *Multiple prior updates, likelihood acceptance*[16]. Let $\Pr^{(p,\mathrm{PRIOR})}\{X_{n+1} \in (\tau', d\sigma')|X_n = (\tau,\sigma)\}$ denote the transition probability obtained for each move $p = 1$ to 4 by omitting the likelihood ratio $e^{L(\sigma')-L(\sigma)}$ in α_p. Simulate a subchain $Y_0, Y_1 \dots Y_A$ of length $A+1$ starting with $Y_0 = (\tau,\sigma)$ and stepping with transition probability

$$Q^{\mathrm{PRIOR}}\{Y_{a+1} \in (\tau', d\sigma')|Y_a = (\tau,\sigma)\} \equiv \sum_{p=1}^{4} \frac{\zeta_p}{S_4} \Pr^{(p,\mathrm{PRIOR})}\{Y_{a+1} \in (\tau', d\sigma')|Y_a = (\tau,\sigma)\}.$$

Now Q^{PRIOR} is reversible with respect to the prior, $\Pr\{\tau, d\sigma\}$, and so the A-step transition probability in the chain $Y_0, Y_1 \dots Y_A$ satisfies

$$Q^{\mathrm{PRIOR}}\{Y_A \in (\tau', d\sigma')|Y_0 = (\tau,\sigma)\}\Pr\{\tau, d\sigma\} = Q^{\mathrm{PRIOR}}\{Y_A \in (\tau, d\sigma)|Y_0 = (\tau',\sigma')\}\Pr\{\tau', d\sigma'\}.$$

Let $Y_A = (\tau', \sigma')$, so that the subchain generates our candidate state (τ', σ'). The candidate state is accepted with probability

$$\alpha_5((\tau,\sigma) \to (\tau',\sigma')) = \min\left\{1, e^{L(\sigma')-L(\sigma)}\right\}.$$

Move 6 *Multiple likelihood updates, prior acceptance.* For each move $p = 1$ to 4 let

$$\Pr^{(p,L)}\{X_{n+1} \in (\tau', d\sigma')|X_n = (\tau,\sigma)\} = q_p((\tau,\sigma) \to (\tau', d\sigma')) \times \min\left\{1, e^{L(\sigma')-L(\sigma)} \times \frac{q_p((\tau',\sigma') \to (\tau, d\sigma))}{q_p((\tau,\sigma) \to (\tau', d\sigma'))}\right\}$$

denote the transition probability obtained by omitting in α_p factors associated with the prior. Simulate a subchain $Y_0, Y_1 \dots Y_A$ of length $A+1$ starting with $Y_0 = (\tau,\sigma)$ and stepping with transition probability

$$Q^{L}\{Y_{a+1} \in (\tau', d\sigma')|Y_a = (\tau,\sigma)\} \equiv \sum_{p=1}^{4} \frac{\zeta_p}{S_4} \Pr^{(p,L)}\{Y_{a+1} \in (\tau', d\sigma')|Y_a = (\tau,\sigma)\}.$$

Now the A-step transition probability in the chain $Y_0, Y_1 \dots Y_A$ satisfies

$$Q^{L}\{Y_A \in (\tau', d\sigma')|Y_0 = (\tau,\sigma)\} \times e^{L(\sigma)} = Q^{L}\{Y_A \in (\tau, d\sigma)|Y_0 = (\tau',\sigma')\} \times e^{L(\sigma')}.$$

Let $Y_A = (\tau', \sigma')$, so that the subchain generates our candidate state (τ', σ'). The acceptance probability is

$$\alpha_5((\tau,\sigma) \to (\tau',\sigma')) = \min\left\{1, \frac{\Pr\{\tau', d\sigma'\}}{\Pr\{\tau, d\sigma\}}\right\},$$

where $\Pr\{\tau, d\sigma\}$ is the full prior.

Move 7 *Metropolised Langevin updates.* This update modifies the conductivities only, leaving the conductivity types unchanged. Let ∇_σ be the gradient operator in σ, let h be a small constant (chosen small enough to give an acceptance rate around 0.6[14]) and let r be a realisation of an M-component random vector with iid components distributed like $N(0, h^2)$. Let $f(\tau,\sigma|v)$ denote the posterior density, *ie* $\Pr\{\tau, d\sigma|v, J\} = f(\tau,\sigma|v,J)d\sigma$ and set

$$\sigma' = \sigma + h * r + \frac{h^2}{2}\nabla_\sigma \log(f(\tau,\sigma|v,J)). \tag{12}$$

In model B, for $m = 1, 2 \dots M$ a pixel index,

$$\frac{\partial}{\partial\sigma_m}\log(f(\tau,\sigma|v)) = \sum_{n=1}^{N}\sum_{\substack{k=1\\k\notin e_n}}^{K} \frac{\partial\phi(x_k|n,\sigma)}{\partial\sigma_m}(v_n(k) - \phi(x_k|n,\sigma))/2d^2 - 4\tilde{\theta}\sum_{i\sim m}(\sigma_m - \sigma_i) - (\sigma_m - \mu(\tau_m))/\rho(\tau_m). \tag{13}$$

This choice emerges as the natural discretisation of the Langevin diffusion[21] reversible with respect to $\Pr\{\tau, d\sigma|v, J\}$. We now regard σ' defined by Equation (12) as the candidate state in a Metropolis-Hastings update.[14] The generation probability is simply

$$q_7((\tau,\sigma) \to (\tau, d\sigma')) = \frac{1}{\sqrt{2\pi h^2}}\, e^{-\sum_{m=1}^{M}(\sigma'_m - \sigma_m)^2/2h^2}\, d\sigma'.$$

Now, let

$$\sigma'' = \sigma' + \frac{h^2}{2}\nabla_\sigma \log(f(\tau, \sigma'|v, J))$$

The acceptance probability is

$$\alpha_3((\tau,\sigma) \to (\tau,\sigma')) = \min\left\{1, \frac{\Pr\{\tau, d\sigma'\}}{\Pr\{\tau, d\sigma\}} \times \exp(-\sum_{m=1}^{M}(\sigma''_m - \sigma_m)^2/2h + \sum_{m=1}^{M}(\sigma'_m - \sigma_m)^2/2h)\right\}.$$

Notice that Move 2 duplicates Move 1, but allows us to focus updates on the interface between regions of differing conductivity.[19] Move 1 is still needed, for irreducibility. Note also that there is no useful simplification of the likelihood ratio $\exp(L(\sigma') - L(\sigma))$. 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 $\phi(\sigma)$.

REFERENCES

1. P. Metherall, D. C. Barber, R. H. Smallwood, and B. H. Brown, "Three-dimensional electrical impedance tomography," *Nature* **380**, pp. 509–512, 1996.
2. C. Fox and G. K. Nicholls, "Sampling conductivity images via MCMC," in *The Art and Science of Bayesian Image Analysis*, K. Mardia, R. Ackroyd, and C. Gill, eds., Leeds Annual Statistics Research Workshop, pp. 91–100, University of Leeds, 1997.
3. J. Sylvester and G. Uhlman, "A global uniquness theorem for an inverse boundary value problem," *Ann. Math.* **125**, pp. 643–667, 1987.
4. C. Fox, *Conductance Imaging.* PhD thesis, University of Cambridge, 1989.
5. D. C. Barber, "A review of image reconstruction techniques for electrical impedance tomography," *Med. Phys* **16**(2), pp. 162–169, 1989.
6. A. P. Calderon, "On an inverse boundary value problem," in *Seminar on Numerical Analysis and its Applications*, W. H. Meyet and M. A. Ranpp, eds., Brazilian Math. Soc., (Rio de Janerio, Brazil), 1980.
7. T. J. Yorkey, *Comparing reconstruction methods for electrical impedance tomography.* PhD thesis, University of Wisconson, Madison, WI, 1986.
8. P. Hua, E. J. Woo, J. G. Webster, and W. J. Tompkins, "Iterative reconstruction methods using regularization and optimal current patterns in electrical impedance tomography," *IEEE Transactions on Medical Imaging* **10**, pp. 621–628, 1991.
9. T. Martin and J. Idier, "A Bayesian non-linear approach for electrical impedance tomography," tech. rep., CNRS, `martin@lss.supelec.fr`, Lab. des Signaux at Systems, CNRS, Plateau de Moulon, 91192 Gif-sur-Yvette Cedex, France, 1997.
10. U. Grenander and M. Miller, "Representation of knowledge in complex systems (with discussion)," *Journal of the Royal Statistical Society B* **56**, pp. 549–604, 1993.
11. G. K. Nicholls, "Bayesian image analysis with Markov chain Monte Carlo and colored continuum triangulation mosaics," *Journal of the Royal Statistical Society B* **60**, pp. 643–659, 1998.
12. P. Green, "Reversible jump MCMC and Bayesian model determination," *Biometrika* **82**, pp. 711–732, 1995.
13. L. Tierney, "Markov chains for exploring posterior distributions," *Annals of Statistics* **22**, pp. 1701–1762, 1994.
14. G. O. Roberts and J. S. Rosenthal, "Optimal scaling of discrete approximations to langevin diffusions," *Journal of the Royal Statistical Society B* , p. to appear, 1998.
15. M. Creutz, "Microcanonical monte carlo simulation," *Phys. Rev. Lett.* **50**, pp. 1411–1414, 1983.
16. E. Clark, "Bayesian impedance imaging." Seminar, Physics Department, Auckland University, November 1997.
17. G. K. Nicholls, "Modelling planar change-point problems using colored continuum triangulation models," in *The art and science of Bayesian Image analysis*, K. Mardia, R. Ackroyd, and C. Gill, eds., Leeds Annual Statistics Research Workshop, pp. 143–150, 1997.
18. R. M. West, "Image reconstruction for electrical resistance tomography by updating local parameters," tech. rep., Camborne School of Mines, University of Exeter, UK, 1998.
19. B. D. Ripley, *Stochastic Simulation*, Cambridge University Press, Cambridge, 1987.
20. W. K. Hastings, "Monte Carlo sampling methods using Markov chains and their applications," *Biometrika* **57**, pp. 97–109, 1970.
21. P. J. Rossky, J. D. Doll, and H. L. Friedman, "Brownian dynamics as smart monte-carlo simulation," *J. Chem. Phys* **69**, pp. 4628–4633, 1978.