Recovery of blocky images in electrical impedance tomography

David C. Dobson

Department of Mathematics, Texas A&M University, College Station, TX 77843-3368 USA dobson@math.tamu.edu

1 Introduction

The techniques of electrical impedance tomography (EIT) have been widely studied over the past several years, for applications in both medical imaging and nondestructive evaluation. The goal is to find the electrical conductivity of a spatially inhomogeneous medium inside a given domain, using electrostatic measurements collected at the boundary.

One of the primary difficulties with EIT is the instability of the reconstruction problem. This instability has been characterized in various ways by studying linearized problems [1, 11, 15]. In essence, the difficulty is that the measured data are very insensitive to certain features in the conductivity profile. For instance, two profiles that differ primarily in the high-frequency Fourier components may yield boundary data that are nearly indistinguishable in the presence of measurement noise. This is a manifestation of the information content of the data. Any reconstruction procedure will be unstable to data errors unless the problem is properly regularized. Unfortunately, most common regularization techniques yield reconstructed conductivity images that are often smeared, blurred, or otherwise distorted.

In an effort to reduce the blurring effect, an "image enhancement" technique was applied to the linearized EIT problem in [13]. The technique is essentially a nonlinear regularization by total variation, following ideas introduced by Rudin, Osher, and Fatemi [23] for image denoising. Measuring the total variation has long been recognized as an effective way to quantify the "simplicity" of a given signal or function. It measures the oscillations of a given function, but still admits discontinuities. Ideas incorporating total variation based methods in reconstruction problems have been used successfully in applications such as image reconstruction [5, 6, 7, 8, 17, 22, 23, 24, 27], inverse problems [25, 26], and optimal design [11].

The motivation for [13] was to try to incorporate *a priori* information about the unknown conductivity into the problem formulation. The *a priori* information is roughly the knowledge that the unknown conductivity is "blocky". By "blocky" we mean that the conductivity is a piecewise constant function for

which the set of discontinuities has relatively small measure in codimension 1. Such conductivity functions have small total variation, provided the discontinuities are not large. The problem is formulated as that of finding the conductivity of minimal total variation which fits the data. The conductivity reconstructions obtained in [13] were very good for blocky images. However as with other known techniques, reconstructions of fine details and highly oscillatory features were difficult or impossible to achieve.

One of the key questions arising from the work in [13] is: how does one characterize the set of images which can be recovered reasonably well by minimization of total variation? This question was explored in a general context of linear image processing problems in [14]. As will be outlined in the present paper, the results have some relevance to EIT and give some insight into the kinds of results one can expect.

In this paper we describe some aspects of the application of total variation minimization techniques to the linearized EIT problem. In Section 2, we formulate a reconstruction problem, describing one approach to total variation regularization by constrained minimization. In Section 3, a stabilization strategy for the constraints is described. Section 4 discusses some results on characterizing conductivity images which can be completely recovered, under the assumption that the instability in the problem is restricted to a limited range of frequency components in the image. Section 5 motivates other conditions which are favorable for recovering images when the instabilities are not bandlimited. In Section 6 we describe a very simple minimization scheme for the constrained regularized problem. Finally, some representative numerical results are presented in Section 7.

2 EIT and minimal total variation regularization

Given some domain $\Omega \subset I\!R^d$, $d = 2$ or 3, the idealized model problem is to find the conductivity distribution in the interior of Ω from electrostatic measurements on the boundary $\partial\Omega$. In all of the following, we assume $d = 2$. A spatially distributed current flux density pattern f which satisfies $\int_{\partial\Omega} f = 0$ is applied to $\partial\Omega$. The voltage potential u inside Ω then satisfies

$$\nabla \cdot (\sigma \nabla u) = 0 \quad \text{in } \Omega, \tag{1a}$$

where σ is the conductivity of the medium, with the Neumann boundary condition

$$\sigma \frac{\partial u}{\partial \nu} = f \quad \text{on } \partial\Omega. \tag{1b}$$

With the additional normalization constraint $\int_{\partial\Omega} u = 0$, and the assumption that $\sigma \in L^\infty(\Omega)$ is uniformly bounded away from zero, equations (1) uniquely determine the voltage potential u.

For every "current pattern" f, we measure the corresponding voltage potential u on the boundary $\partial\Omega$. Hence, the data in the ideal problem can be viewed as the so-called Neumann-to-Dirichlet map $\Lambda(\sigma; f)$ which takes current patterns

to voltage measurements, defined by $\Lambda(\sigma; f) = u|_{\partial\Omega}$, where u satisfies (1). The operator Λ is linear in f but nonlinear in σ.

It is infeasible in practice to measure all of $u|_{\partial\Omega}$. Instead, one generally obtains a finite number of "samples" of u on $\partial\Omega$. For a given current pattern f, let us assume that rather than measuring all of $u|_{\partial\Omega}$, we have access to the finite set of measurements

$$g_i = \int_{\partial\Omega} uq_i \quad i = 1, \ldots, m, \tag{2}$$

where u satisfies (1). Each distribution $q_i \in H^{-1/2}(\partial\Omega)$ can be chosen to represent a measurement process. For example, the q_i's could be concentrated at m distinct points on $\partial\Omega$ to approximate a set of point measurements, or concentrated on a set of characteristic functions to approximate spatially averaged measurements. Naturally, one would wish to choose the q_i's to be orthogonal, or at least linearly independent.

The forward map, which we denote $F(\sigma; f)$, can then be viewed as taking a current pattern f and a conductivity distribution σ to an $\mathbb{R}^m$ vector of measurements. In practice, one is limited to performing only a finite set of experiments, generally carried out with very specific f, such as those generated by a given number of electrodes in fixed positions. If we make n experiments, that is, we apply n different current patterns

$$f^{(1)}, f^{(2)}, \ldots, f^{(n)},$$

and make m measurements for each pattern

$$g^{(1)}, g^{(2)}, \ldots, g^{(n)},$$

(where $g^{(j)} = (g_1^{(j)}, \ldots, g_m^{(j)})$ according to (2)), then the inverse problem is to determine σ such that

$$F(\sigma; f^{(j)}) = g^{(j)}, \quad j = 1, \ldots, n.$$

A common technique to simplify the analysis of the problem is to linearize F with respect to σ about about a constant background. Thus, we let

$$\sigma = 1 + \delta\sigma$$

where $\|\delta\sigma\|_{L^\infty(\Omega)}$ is "small", and assume that the voltage potential is

$$u = U + \delta u.$$

By considering terms of the same order in (1a), we see that the background potential U satisfies

$$\triangle U = 0 \quad \text{in } \Omega, \tag{3a}$$

$$\frac{\partial U}{\partial \nu} = f \quad \text{on } \partial\Omega. \tag{3b}$$

The linear part δu of the perturbational voltage potential satisfies

$$\triangle \delta u = -\nabla \cdot \delta\sigma \nabla U \quad \text{in } \Omega, \tag{4a}$$

$$\frac{\partial \delta u}{\partial \nu} = 0 \quad \text{on } \partial\Omega, \tag{4b}$$

where we again enforce the normalization $\int_{\partial\Omega} \delta u = 0$. Since U depends linearly on the current pattern f, we see that δu is also linearly dependent on f.

For a given current pattern f, the linearized forward map DF takes a conductivity perturbation $\delta\sigma$ to perturbational measurements on the boundary. The i-th component of this map is

$$(DF(f)\delta\sigma)_i := \int_{\partial\Omega} \delta u \, q_i. \tag{5}$$

Assuming that $\delta\sigma$ is supported in $\Omega' \subset\subset \Omega$, we can view the map as

$$DF(f) : L^2(\Omega') \to I\!R^m. \tag{6}$$

In the linearized inverse problem, the goal is to find $\delta\sigma$ from knowledge of the *differences* between the measured voltages and the background voltages $\int_{\partial\Omega} U \, q_i$ for a set of current patterns $f^{(j)}$, $j = 1, \ldots, n$. Let us denote the background voltage due to current pattern $f^{(j)}$ by $U^{(j)}$. The linearized inverse problem is to determine $\delta\sigma$ in the equation

$$DF(f^{(j)})\delta\sigma = g^{(j)} - TU^{(j)} =: \delta g^{(j)}, \quad \text{for} \quad j = 1, \ldots, n, \tag{7}$$

where $TU^{(j)}$ is the $I\!R^m$-vector whose components are $\int_{\partial\Omega} U^{(j)} \, q_i$. To simplify notation, let us assume that the current patterns $f^{(1)}, \ldots, f^{(n)}$ have been chosen and are henceforth *fixed*. Forming the $I\!R^{n\cdot m}$ data vector

$$\delta g = (\delta g^{(1)T}, \ldots, \delta g^{(n)T})^T$$

and the corresponding operator

$$DF = (DF(f^{(1)})^T, \ldots, DF(f^{(n)})^T)^T,$$

the linearized inverse problem can be written

$$DF \; \delta\sigma = \delta g. \tag{8}$$

Of course, $\delta\sigma \in L^2(\Omega')$ is not determined due to the finite amount of data. Furthermore, even if $\delta\sigma$ is discretized in a straightforward way to a finite number (say $n \cdot m$) of unknowns, the conductivity perturbation $\delta\sigma$ is generally not well-determined by the data due to instability [1, 11, 15]. In particular, information corresponding to high frequency components of $\delta\sigma$ is easily corrupted by noise.

To motivate the use of total variation regularization, assume for a moment that the unknown $\delta\sigma$ can be described as a sum of characteristic functions:

$$\delta\sigma(x) = \sum_{k=1}^{N} a_k \chi(\Omega_k)(x). \tag{9}$$

Here N is unknown but finite, the coefficients a_k are unknown, the subdomains $\Omega_k \subset\subset \Omega$ are unknown, and $\chi(\Omega_k)$ denotes the characteristic function on Ω_k. For convenience, assume that each Ω_k has a C^2 boundary $\partial\Omega_k$. Equation (9) approximates a very wide variety of conductivity profiles which could be encountered in medical imaging and non-destructive testing.

Let $\nabla\delta\sigma$ denote the gradient of $\delta\sigma$ in the sense of distributions; $\nabla\delta\sigma$ is a vector valued Radon measure and

$$TV(\delta\sigma) := \int_\Omega |\nabla\delta\sigma| \tag{10}$$

is the total variation of $\delta\sigma$. We denote by $BV(\Omega)$ the space of functions of bounded variation in Ω, equipped with the norm

$$\|\delta\sigma\|_{BV(\Omega)} = \|\delta\sigma\|_{L^1(\Omega)} + \int_\Omega |\nabla\delta\sigma|$$

(see for example [18]). We observe that $\delta\sigma$ defined by (9) is in $BV(\Omega)$.

Now suppose that we are trying to reconstruct $\delta\sigma$ by solving $DF\ \delta\sigma = \delta g$. Recall that the data is finite-dimensional. Because of the ill-posedness of the problem, we are only able to obtain an approximation $\tilde{\delta\sigma} = \delta\sigma + \eta$, where η is some unknown "error" from the nullspace of DF. Notice that $\nabla\delta\sigma$ is zero except on $\cup_{k=1}^N \partial\Omega_k$. Thus if η happens to have bounded derivatives in the parts of Ω where $\nabla\delta\sigma$ is singular, then necessarily

$$\int_\Omega |\nabla(\delta\sigma + \eta)| \geq \int_\Omega |\nabla\delta\sigma|.$$

This suggests that one might find a better approximation to $\delta\sigma$ by choosing from all possible solutions $\tilde{\delta\sigma} = \delta\sigma + \eta$, one with smallest total variation. Intuitively, at the very least, this strategy rules out a large class of error components η which have oscillations in the parts of Ω where $\delta\sigma$ is constant. We will describe in Section 4 specific (and rather stringent) conditions which ensure that $\delta\sigma$ can be recovered exactly.

Thus our starting point is the optimization problem

$$\min_{\delta\sigma\in\mathcal{A}} \int_\Omega |\nabla\delta\sigma|, \tag{11a}$$

$$\text{subject to: } DF\ \delta\sigma = \delta g, \tag{11b}$$

where $\mathcal{A}$ is some admissible class of conductivity perturbations. The constraint ensures that a feasible conductivity perturbation is one that matches the observed data. A natural choice for the admissible set is $\mathcal{A} = \{\delta\sigma \in BV(\Omega) : \delta\sigma|_{\partial\Omega} = 0\}$. Since the cost functional is convex and the constraints are linear, any local minimizer of problem (11) must be a global minimizer. However, the cost functional is not *strictly* convex, so uniqueness of solutions is not guaranteed in general.

For existence of solutions, one can show the following.

Theorem 1. *Given any data vector $\delta g \in \mathbb{R}^{n \cdot m}$ in the range of DF, the problem (11) admits at least one solution $\delta\sigma \in \mathcal{A}$.*

The proof follows the direct method in the calculus of variations, using compactness properties of BV. Details can be found in [13].

3 Stabilizing the constraints

There are two main difficulties associated with the constraint $DF\ \delta\sigma = \delta g$. The first is that if a measured data vector δg is not in the range of DF, the constraints cannot be satisfied and so the problem has no solution. Consider the following "near-null" set associated with DF:

$$V_{\text{near-null}} = \{v\ :\ v \in L^2(\Omega), ||DF\ v|| \leq \epsilon\},$$

for some small ϵ. The second difficulty is that even if δg is in the range of DF, it is still possible for δg to have a component in the set $(DF\ v)$ where $v \in V_{\text{near-null}}$. In this case, small changes in the measurements δg could lead to large changes in the subspace defined by the constraints $\{\delta\sigma\ :\ DF\ \delta\sigma = \delta g\}$, and hence cause instability in the minimization problem.

One remedy for both difficulties can be obtained in the following way. Let ϵ be the "noise level" in the data. Consider the singular value decomposition (SVD):

$$DF = \mathbf{U}\Sigma\mathbf{V}^T,$$

where $\mathbf{U}$ is an $nm \times nm$ orthogonal matrix, $\mathbf{V}$ is an orthogonal operator mapping $L^2(\Omega) \to L^2(\Omega)$, and Σ is a "diagonal" operator with diagonal $\{s_1, \ldots, s_{nm}\}$, where $s_1 \geq s_2 \geq \ldots \geq s_{nm}$ are the singular values of DF. Given any discrete approximation to $L^2(\Omega)$, the corresponding SVD can be calculated numerically by standard methods, see [19]. A slightly different approach is needed if we choose not to discretize. Such an approach is outlined in [15]. The SVD has been used previously to analyze instability in the EIT problem [3].

Let s_p be the smallest singular value greater than ϵ. Form the new operators

$$\Sigma' = \text{diag}\{s_1, \ldots s_p\},$$

$$\mathbf{U}' = (\mathbf{u}_1, \ldots, \mathbf{u}_p),$$

and

$$\mathbf{V}' = (\mathbf{v}_1, \ldots, \mathbf{v}_p),$$

where $\mathbf{u}_j$ denotes the jth column of $\mathbf{U}$, and similarly for $\mathbf{v}_j$. Now we can form the "reduced rank" operator

$$M = \Sigma'\mathbf{V}'^T.$$

The range space of M is p-dimensional, and M has a well-defined pseudo-inverse $M^\dagger$ whose norm is less than or equal to $1/\epsilon$. The "stabilized" linear inverse problem can then be posed as

$$M\delta\sigma = \mathbf{U}'^T\delta g := \delta g'.$$

The effect of this manipulation is that δg has been projected onto the subspace spanned by range vectors of DF whose singular values are greater than ϵ, thus eliminating both difficulties described above. The constraint $M\delta\sigma = \delta g'$ defines a subspace of conductivity perturbations which are consistent with the observed data to within the noise level ϵ.

We emphasize that there are alternate ways of handling the instability in the constraints. The penalty method [27, 25, 26], augmented Lagrangian formulations [20], and the method in [24] are such alternatives. These methods are for the most part more computationally efficient for large problems because they do not require calculating the SVD. However, the approach presented here is useful from a conceptual standpoint.

4 The case of band-limited instability

In this section we describe some conditions under which blocky images $\delta\sigma_0$ can be fully recovered, under the assumption that there is no instability in recovering the low-frequency Fourier components of $\delta\sigma_0$. This assumption is certainly not true in general, but we outline below that it may be a useful approximation.

From (4a) it follows that

$$(DF(f)\delta\sigma)_i := \int_{\partial\Omega} \delta u \; q_i = \int_\Omega \delta\sigma \nabla U \cdot \nabla V_i,$$

where $\triangle V_i = 0$, and $\partial V_i/\partial\nu = q_i$. We may assume without loss of generality that Ω is contained within a unit ball centered at the origin. Based on the approach of Calderon [4], consider the harmonic functions

$$U_\xi(x) = e^{(i\xi+\zeta)\cdot x - |\xi|}, \quad V_\xi(x) = e^{(i\xi-\zeta)\cdot x - |\xi|}$$

where $\xi, \zeta \in I\!R^n$, with $\xi \cdot \zeta = 0$ and $|\xi| = |\zeta|$ in the expression above to obtain

$$\int_{\partial\Omega} \delta u \; q_\xi = e^{-2|\xi|}|\xi|^2 \int_\Omega \delta\sigma \; e^{i\xi\cdot x}, \tag{12}$$

where $q_\xi = \partial V_\xi/\partial\nu$. In other words, one can recover the Fourier coefficients of $\delta\sigma$ by choosing the "current patterns" $f_\xi = \partial U_\xi/\partial\nu$, and appropriate measurement weights. In practice, the weights q_ξ could be approximated by taking appropriate linear combinations of point measurements of δu. Notice that the magnitude of f_ξ and q_ξ is bounded with increasing ξ.

Now suppose as in the previous section that we make a finite number n of experiments, each corresponding to a different frequency ξ_j, $j = 1, \ldots, n$. For each current pattern f_{ξ_j}, we make one measurement with weight q_{ξ_j}. For these particular experiments and measurements,

$$(DF \; \delta\sigma)_j = e^{-2|\xi_j|}|\xi_j|^2 \int_\Omega \delta\sigma \; e^{i\xi_j\cdot x}. \tag{13}$$

The operator DF is thus diagonalized and the SVD is simply U = identity, $\Sigma = \mathrm{diag}(e^{-2|\xi_j|}|\xi_j|^2)$, and V is the Fourier transform.

Since the magnitude of the measurements is exponentially decreasing with increasing frequency $|\xi_j|$, one would expect that with a fixed amount of background noise, the relative noise level for low frequency measurements would be very low compared to high frequency measurements. Under this assumption, we use the data as equality constraints to define a set of feasible images. Recall the stabilized operator M as defined in Section 3. The feasible set of images is described by

$$V = \{\delta\sigma \ : \ M\delta\sigma = \delta g'\}.$$

Since M is simply a weighted and truncated Fourier transform, the set V contains all images whose spectrum for frequencies $|\xi_j| \leq R$, for some R defined by the noise level, agree with that of the true image.

The problem is then to solve

$$\min_{\delta\sigma} \ TV(\delta\sigma) \quad \text{subject to } \ \delta\sigma \in V. \tag{14}$$

Note that knowledge about the noise is used to determine R, and hence V.

If the low frequency measurements $\delta g'$ corresponding to a given image $\delta\sigma_0$ are not corrupted by noise, then clearly $\delta\sigma_0 \in V$, since then $M\ \delta\sigma_0 = \delta g'$. Also notice that every $\delta\sigma \in V$ can be written $\delta\sigma = \delta\sigma_0 + \alpha$, where α is from the null space of M. In order to establish that the minimization (14) will successfully recover $\delta\sigma_0$, we must show that any other image $\delta\sigma = \delta\sigma_0 + \alpha \in V$, necessarily has larger total variation than $\delta\sigma_0$. In other words we must show that

$$TV(\delta\sigma_0 + \alpha) > TV(\delta\sigma_0)$$

for all α in the null space of M. Due to the form of M this means simply that high frequency perturbations to $\delta\sigma_0$ should increase its total variation.

4.1 Recovery of discrete images

Recall that Ω is assumed to be contained in a unit ball centered at the origin. Assume that the true image $\delta\sigma_0$ is piecewise constant on a uniform square grid on $(-1,1)^2$ with $N \times N$ cells and denote the value of $\delta\sigma_0$ on the (n,m)-th cell by $(\delta\sigma_0)_{n,m}$. To calculate the total variation of $\delta\sigma_0$, we define the horizontal and vertical difference operators D_H and D_V by

$$\begin{aligned}(D_H\delta\sigma_0)_{n,m} &= (\delta\sigma_0)_{n+1,m} - (\delta\sigma_0)_{n,m},\\ (D_V\delta\sigma_0)_{n,m} &= (\delta\sigma_0)_{n,m+1} - (\delta\sigma_0)_{n,m},\end{aligned}$$

where, since $\delta\sigma_0$ is supported away from the boundary, $(\delta\sigma_0)_{N+1,m} \equiv (\delta\sigma_0)_{1,m}$ and $(\delta\sigma_0)_{n,N+1} \equiv (\delta\sigma_0)_{n,1}$. Defining $\|v\|_1 = \sum_{n,m=1}^{N} |v_{n,m}|$, we see that

$$\int_\Omega |\nabla\delta\sigma_0| = \frac{2}{N}\left(\|D_H(\delta\sigma_0)\|_1 + \|D_V(\delta\sigma_0)\|_1\right), \tag{15}$$

gives the total variation of $\delta\sigma_0$. Let us define the set of all edges upon which $\delta\sigma_0$ is not constant:

$$S = \{(n,m) : (D_H\delta\sigma_0)_{n,m} \neq 0 \ \text{ or } \ (D_V\delta\sigma_0)_{n,m} \neq 0\}$$

and assume that there are exactly ν elements in S, numbered (n_r, m_r) for $r = 1, \ldots, \nu$. Notice that ν/N^2 indicates the proportion of edges upon which $\delta\sigma_0$ is not constant.

Assume that α is a *band-limited* perturbation, i.e, we can express

$$\alpha_{n,m} = \frac{1}{N^2} \sum_{(k,l)\in\mathcal{K}} \hat{\alpha}_{k,l} \exp(\frac{-2\pi i(n-1)(k-1)}{N}) \exp(\frac{-2\pi i(m-1)(l-1)}{N}) \tag{16}$$

where $\mathcal{K}$ is some finite set with K elements. By a direct calculation and an application of the triangle inequality, one finds that

$$\sum_{(n,m)\in S} |D_H(\delta\sigma_0 + \alpha)_{n,m}| \geq \sum_{(n,m)\in S} |(D_H\delta\sigma_0)_{n,m}| - \frac{K\nu}{N^2}\|\hat{\alpha}\|_\infty,$$

where $\|\hat{\alpha}\|_\infty = \max_{n,m} |\hat{\alpha}_{n,m}|$. With a similar calculation, one can show that

$$\left(1 - \frac{K\nu}{N^2}\right) \|\hat{\alpha}\|_\infty \leq \sum_{(n,m)\notin S} |(D_H(\delta\sigma_0 + \alpha)_{n,m}|.$$

Applying the previous two inequalities, it follows that

$$\|D_H(\delta\sigma_0 + \alpha)\|_1 \geq \|D_H\delta\sigma_0\|_1 + \left(1 - \frac{2K\nu}{N^2}\right) \|\hat{\alpha}\|_\infty.$$

The same derivation yields an analogous inequality for the the vertical difference operator D_V. Adding the horizontal and vertical contributions, we find that

$$\int |\nabla(\delta\sigma_0 + \alpha)| \geq \int |\nabla\delta\sigma_0| + \frac{4}{N}\left(1 - \frac{2K\nu}{N^2}\right) \|\hat{\alpha}\|_\infty, \tag{17}$$

recalling that $\int |\nabla\delta\sigma_0|$ is given by (15).

The bound (17) indicates that as long as

$$\frac{2K\nu}{N^2} < 1 \tag{18}$$

and $\alpha \neq 0$, the perturbed signal $\delta\sigma_0 + \alpha$ has higher total variation than the original signal $\delta\sigma_0$. This implies that $\delta\sigma_0$ can be recovered by minimizing $\int |\nabla(\delta\sigma_0 + \alpha)|$ over all perturbations of the form (16).

Estimate (18) indicates a tradeoff between the spectral content of the noise (measured by K) and the proportion of edges upon which $\delta\sigma_0$ is not constant (measured by ν/N^2). This is a form of the uncertainty principle as discussed in Donoho [16]. It is interesting that neither the intensity of any of the frequency components of α, nor distribution of the spectrum of α in important in this condition. Only the *number* of frequencies K allowed in the noise matters.

In terms of the minimization problem (14), if the noise level in the problem is such that the subspace V spans K or fewer frequencies, and the number of non-constant edges ν in $\delta\sigma_0$ satisfies (18), then solving the minimization will recover $\delta\sigma_0$. Note that inherent in this statement is the assumption that the problem is discrete; otherwise V is infinite dimensional.

4.2 Sharpness of the spectral bounds

The estimate (18) gives a condition which guarantees that a given image can be recovered from a bandlimited perturbation. It is worthwhile to ask how "sharp" is the bound, e.g., what is the fewest number of frequencies K which can be perturbed in order to decrease the total variation of a given image? The inequality (18) states that if $K < N^2/2\nu$, then the total variation cannot be decreased.

In [14] several numerical experiments were carried out to test the bound. Generally speaking, for "blocky" images $\delta\sigma_0$, it was found to be very difficult to lower the total variation with simple filtering operations. Low-pass and high-pass filters were applied to blocky images. It was generally necessary for the filter to alter substantially more than $K = N^2/2\nu$ frequencies before the total variation of the image was lowered. This indicates that for blurring operators like DF applied to blocky images, the band limit suggested by (18) is too conservative. However, this is only supported by numerical evidence.

Images for which the bound is "least conservative" seem to be composed of a few pure frequencies. For example, images $\delta\sigma_0$ consisting of vertical or horizontal "stripes" of uniform width can be perturbed by $\alpha = -\delta\sigma_0$, thus reducing the total variation to zero with $K = N^2/\nu$, thereby violating the bound (18) by only a factor of two. Similarly, other images with concentrated frequency content are easily diminished by changing a few frequency components, thereby decreasing the total variation.

4.3 Recovery of nearly piecewise-constant images

Estimates analogous to (17) can be established in an infinite-dimensional setting. Such estimates have implications for denoising and image recovery problems when the spectral content of the noise can be controlled. Here we merely state the main result; the reader is referred to [14] for details.

Let Ω be the unit square and assume as before that $\delta\sigma_0$ is supported away from the boundary. Assume for convenience that $\delta\sigma_0 \in C^1(\Omega)$ and define the set upon which $\delta\sigma$ is not constant:

$$B = \{x \in \Omega : \nabla\delta\sigma_0(x) \neq 0\}.$$

Let the signal perturbation α have the Fourier series representation

$$\alpha(x) = \sum_{k \in Z^2} \hat{\alpha}_k e^{2\pi i x \cdot k},$$

where $Z = \{0, \pm 1, \pm 2, \ldots\}$ and

$$\hat{\alpha}_k = \int_\Omega \alpha(x) e^{-2\pi i x \cdot k}\, dx.$$

Theorem 2. *The total variation of $\delta\sigma_0 + \alpha$ is strictly greater than the total variation of $\delta\sigma_0$ provided that*

$$2|B|\ \|k\hat{\alpha}_k\|_{l^1} < \|k\hat{\alpha}_k\|_{l^\infty}. \tag{19}$$

Condition (19) says that if $\|k\hat{\alpha}_k\|_{l^1}$ is finite, then as long as the nonconstant set B is small enough, $\delta\sigma_0$ can be recovered. On the other hand, if $|B|$ is fixed, condition (19) indicates that $\delta\sigma_0$ can be recovered as long as the spectrum of the noise is not too spread out compared to the magnitude of any one of its Fourier components. So for example, a "nearly piecewise constant" image polluted by a few pure frequencies would be relatively easily recoverable. On the other hand, a smoothly varying function polluted by a spatially localized perturbation (implying that the spectrum of the perturbation is "spread out") may be impossible to recover.

Without additional information about the asymptotic behavior of the spectrum of the noise, the estimate (19) does not give particularly useful information about EIT. On the other hand, suppose we knew that the spectrum of α is bounded by some sequence $\hat{h}_k$. If $|\hat{h}_k|$ decreases quickly enough as $|k| \to \infty$, it follows that $\|k\hat{\alpha}_k\|_1 / \|k\hat{\alpha}_k\|_\infty$ must decrease as the frequency cutoff $R \to \infty$. Then (19) shows that $\delta\sigma_0$ can be recovered provided $|B|$ is small enough, and furthermore the constraint on $|B|$ becomes less restrictive as R increases.

In case α is bandlimited, i.e. all but K Fourier components $\hat{\alpha}_k$ are zero, condition (19) becomes

$$2K|B| < 1, \tag{20}$$

in analogy with the discrete case (18).

In the limiting case $|B| = 0$, under the assumption that $\delta\sigma_0$ can be described as a sum of characteristic functions (9), It can easily be shown that if $\alpha \in W^{1,1}(\Omega)$, then

$$\int_\Omega |\nabla(\delta\sigma_0 + \alpha)| = \int_\Omega |\nabla\delta\sigma_0| + \int_\Omega |\nabla\alpha|. \tag{21}$$

Thus, provided $\delta\sigma_0$ is piecewise constant, α has integrable derivatives, and the minimization problem has a unique solution, $\delta\sigma_0$ can be recovered by minimizing the total variation of $\delta\sigma_0 + \alpha$.

It would be interesting to see how results such as these might be combined with other regularization techniques designed to limit the spectral content of α, in order to obtain recoverability results under reasonable assumptions on noise and discretization. This is in effect what took place in Section 4.1, where regularization by discretization truncates the spectrum of α.

5 The case of non-bandlimited instability

The case of recovering $\delta\sigma$ in the presence of general measurement noise is now considered. We consider the unconstrained minimization

$$\min_u \; TV(\delta\sigma) + \lambda||DF\ \delta\sigma - \delta g||^2_{L^2(\Omega)}. \tag{22}$$

The penalty coefficient λ is chosen so that at minimum, the data mismatch $||DF\ \delta\sigma - \delta g||$ is within the estimated signal to noise ratio. If the noise is Gaussian with variance η is γ, then we choose λ to be such that

$$||DF\ \delta\sigma - \delta g|| \leq \gamma. \tag{23}$$

As one might expect, without the assumption of band-limited noise in the problem, results on recoverability are mostly negative. Let $\delta\sigma_0$ be the image we wish to recover; for now assume that $\delta\sigma_0$ is any BV function. First notice that there are *always* perturbations α such that $(\delta\sigma_0 + \alpha)$ has lower total variation than $\delta\sigma_0$ while at the same time $(\delta\sigma_0 + \alpha)$ satisfies the data fit requirement

$$||DF(\delta\sigma_0 + \alpha) - g||_2 \leq \gamma,$$

for some "noise level" $\gamma > 0$. For convenience, assume that $DF\delta\sigma_0 = \delta g$. Then the set of feasible perturbations is

$$P_\gamma = \{\alpha : ||DF\alpha||_2 \leq \gamma\}. \tag{24}$$

One then wishes to determine if

$$\min_{\alpha \in P_\gamma} \int_\Omega |\nabla(\delta\sigma_0 + \alpha)| \tag{25}$$

can be less than $\int_\Omega |\nabla\delta\sigma_0|$. Notice that P_γ contains rough (non-bandlimited) perturbations. We can set $a = \min(\gamma/\|DF\delta\sigma_0\|_2, 1)$ so that

$$\|DF(a\delta\sigma_0)\|_2 \leq \gamma,$$

i.e., $a\delta\sigma_0 \in P_\gamma$. Then by taking the perturbation $\alpha = -a\delta\sigma_0$, we see that

$$\min_{\alpha \in P_\gamma} \int_\Omega |\nabla\delta\sigma_0 + \nabla\alpha| \leq (1-a) \int_\Omega |\nabla\delta\sigma_0| < \int_\Omega |\nabla\delta\sigma_0|. \tag{26}$$

Thus one cannot recover $\delta\sigma_0$ by solving problem (25).

On the other hand, if $(1-a)\delta\sigma_0$ *were* the solution to (25), we would probably be quite happy since it differs from the true image only by a multiplicative constant. In this sense, we could say that the image $(1-a)\delta\sigma_0$ is a "qualitatively correct" reconstruction of $\delta\sigma_0$ (provided $(1-a) \neq 0$).

However, one is generally not guaranteed a qualitatively correct solution either. Using a blurring operator for which explicit calculations could be made, an example was constructed in [14] of a perturbation which lies inside the constraint set and which lowers the total variation, while making the image qualitatively different from the original image. The same construction could be easily repeated for EIT using the currents and measurements measurements f_{ξ_j}, q_{ξ_j} as used in Section 4. A prominent feature of the example from [14] is that the total variation of the true image $\delta\sigma_0$ is very large relative to its "mass" $\int |\delta\sigma_0|$. This condition seems to characterize most of the images we found hard to reconstruct accurately in the numerical experiments.

6 Minimization of the total variation functional

The main computational difficulty associated with problems (14) and (22) is mainly due to the nondifferentiability of the total variation functional. Rudin and Osher overcome this difficulty by using a nonlinear diffusion equation whose steady state is the Euler-Lagrange equation of the minimization problem [23].

Vogel [27, 28] has introduced a fixed point method to solve this type of problem. In this approach, the total variation seminorm is mollified slightly to get around the nondifferentiability. This method has exhibited rapid convergence in numerical experiments.

Ito and Kunisch [20] have introduced an active set strategy based on the augmented Lagrangian formulation for similar problems in image restoration.

A new class of methods to solve optimization problems with nondifferentiable functional, such as (14), has been proposed by Coleman and Li [9]. An image enhancement method based on the ideas of Coleman and Li has been developed by Li and Santosa. [21]. The method, which works on the original nondifferentiable functional, is based on an affine scaling strategy and appears to be quite efficient for large problems. The code minimizes the total variation while searching in the feasible set (23). To solve the optimization (22), in [14] a quadratic programming method of Coleman and Liu [10] was employed.

Rather than discussing the different algorithmic approaches to this problem, we will describe here a simple gradient-descent type method used in [13], and based on the original scheme of Rudin, Osher and Fatemi [23].

For arbitrary current patterns and measurements, the minimal total variation reconstruction problem with stabilized linear constraints can be formulated:

$$\min_{\delta\sigma\in X} J(\delta\sigma) = \int_\Omega |\nabla\delta\sigma|, \tag{27a}$$

$$\text{subject to } M\delta\sigma = \delta g', \tag{27b}$$

where $M : L^2(\Omega) \to I\!R^p$ is the linear operator constructed in Section 3, and $X = \{\delta\sigma \in BV(\Omega) : \delta\sigma|_{\partial\Omega} = 0\}$. In the two-dimensional case we consider here, $BV(\Omega)$ imbeds in $L^2(\Omega)$, so $M\delta\sigma$ is well-defined.

The fact that the cost functional is not smooth creates certain complications from both practical and theoretical standpoints. A common approach is to "mollify" the cost functional J with a small smoothing parameter and solve the resulting problem by straightforward methods. In principle the smoothing parameter can be taken arbitrarily small, so that in the limit one should obtain a solution to problem (27). The problem we try to solve can be written

$$\min_{\delta\sigma\in X} J_\epsilon(\delta\sigma) = \int_\Omega h_\epsilon(|\nabla\delta\sigma|), \tag{28a}$$

$$\text{subject to } M\delta\sigma = \delta g', \tag{28b}$$

where

$$h_\epsilon(s) = \begin{cases} s & \text{if } s > \epsilon, \\ \frac{s^2}{2\epsilon} + \frac{\epsilon}{2} & \text{if } s \le \epsilon. \end{cases}$$

Thus h_ϵ is C^1 for $\epsilon > 0$. The effect of h_ϵ is to "round off" the corner in the absolute value function.

The Euler-Lagrange equations for problem (28) are formally given by

$$\nabla \cdot (q_\epsilon(|\nabla \delta\sigma|)\nabla \delta\sigma) - M^T\lambda = 0, \quad \text{in } \Omega, \tag{29a}$$

$$\delta\sigma = 0, \quad \text{on } \partial\Omega, \tag{29b}$$

$$M\delta\sigma = \delta g', \tag{29c}$$

where $\lambda \in I\!R^p$, M^T denotes the L^2 adjoint of M, and the function q_ϵ is defined by

$$q_\epsilon(s) = \frac{h'_\epsilon(s)}{s} = \begin{cases} 1/s \text{ if } s > \epsilon, \\ 1/\epsilon \text{ if } s \leq \epsilon. \end{cases}$$

In [23] the approach to solving the Euler-Lagrange equations is to use time as an evolution (iteration) parameter. Applying this idea to our problem, we would solve

$$\delta\sigma_t = \nabla \cdot (q_\epsilon(|\nabla \delta\sigma|)\nabla \delta\sigma) - M^T\lambda, \quad \text{in } \Omega \times (0,\infty), \tag{30a}$$

$$\delta\sigma = 0, \quad \text{on } \partial\Omega \times (0,\infty), \tag{30b}$$

$$\delta\sigma(x,0) = \delta\sigma_0(x), \quad x \in \Omega, \tag{30c}$$

where $\delta\sigma_0$ is a solution of $M\delta\sigma_0 = \delta g'$, given, say, by the pseudo-inverse solution $\delta\sigma_0 = M^\dagger \delta g'$.

To solve the initial boundary value problem in (30), one could apply an explicit time-stepping scheme, that is, starting with the initial step $\delta\sigma_0$, apply the iteration

$$\delta\sigma_{i+1} = \delta\sigma_i + \tau \left[\nabla \cdot (q_\epsilon(|\nabla \delta\sigma_i|)\nabla \delta\sigma_i) - M^T\lambda^{(i)}\right], \tag{31}$$

where τ is the "step length". To solve for the Lagrange multiplier approximation $\lambda^{(i)}$, let $b_j \in L^2(\Omega)$ be a solution of

$$Mb_j = e_j, \quad j = 1, \ldots, p,$$

where e_j denotes the standard unit basis vectors for $I\!R^p$. Assuming that the SVD has already been calculated as in the last section, we can set $b_j = \mathbf{v}_j/s_j$, where $\mathbf{v}_j$ is the j-th column of the orthogonal operator $\mathbf{V}$. We see that if (29a) is satisfied then

$$\int_\Omega \nabla \cdot (q_\epsilon(|\nabla \delta\sigma|)\nabla \delta\sigma) b_j - \int_\Omega (M^T\lambda) b_j = 0,$$

and hence

$$\lambda_j = \langle \lambda, Mb_j \rangle = \int_\Omega (M^T\lambda) b_j = \int_\Omega \nabla \cdot (q_\epsilon(|\nabla \delta\sigma|)\nabla \delta\sigma) b_j. \tag{32}$$

Using equation (32), $\lambda^{(i)}$ can be calculated from $\delta\sigma_i$ and the vectors b_j. This procedure can be viewed as projecting the L^2 gradient

$$D_{\delta\sigma} J_\epsilon = -\nabla \cdot (q_\epsilon(|\nabla \delta\sigma|)\nabla \delta\sigma) \tag{33}$$

of the cost functional J_ϵ onto the linear subspace $\{\delta\sigma : M\delta\sigma = 0\}$. Thus the entire iteration scheme (31) can be viewed as a "projected gradient method". Taking this viewpoint, define the negative "projected gradient"

$$G_i = -D_{\delta\sigma} J_\epsilon(\delta\sigma_i) - M^T \lambda^{(i)}.$$

This leads to the following simple minimization procedure.

1. Choose an initial steplength τ_0, a parameter $\epsilon > 0$, and an initial iterate $\delta\sigma_0$ satisfying the constraint (27b).
2. For $i = 1, \ldots,$ convergence do
3. If $J_\epsilon(\delta\sigma_i + \tau_i G_i) < J_\epsilon(\delta\sigma_j)$ then
   $\delta\sigma_{i+1} = \delta\sigma_i + \tau_i G_i$
   $\tau_{i+1} = \tau_i$
   else
   $\tau_i = \tau_i/2$
   If $\tau_i \|G_i\|$ is too small then stop
   otherwise go to step 3
   end do

A more complicated backtracking strategy that allows an increase as well as a decrease in the step size can also be implemented, however experiments in [13] indicated very little improvement in speed of convergence with such a scheme. Faster convergence is generally observed for larger mollification ϵ, but the scheme seems to converge even for very small positive values of ϵ.

In the next section, we discretize the problem by making the simplifying assumption that $\delta\sigma$ is a piecewise constant function over a square array of square pixels. We note that one does *not* obtain convergence in BV with this discretization as the number of pixels is increased. The intuitive reason for the lack of convergence is that the resulting discrete total variation measure is anisotropic. For example, a discretized function with a discontinuity along a straight line oriented parallel to the gridlines would have smaller total variation than one oriented at a 45 degree angle to the gridlines (independent of the level of discretization). The primary advantage of this discretization is its simplicity. The cost functional $J_\epsilon(\delta\sigma)$ can be calculated in closed form. This leads to a very simple exact formula for the gradient of $J_\epsilon(\delta\sigma)$ and avoids the difficulty of choosing an appropriate approximation of the partial derivatives in G_i.

7 Implementation and numerical examples

In this section we describe some numerical computations carried out for a particular linearized EIT problem. We take Ω to be unit disk in $\mathbb{R}^2$. In polar coordinates, $\Omega = \{(r,\theta) : r < 1\}$, and $\partial\Omega = \{(r,\theta) : r = 1\}$. Furthermore, we assume that the current patterns f are generated by a finite number n of fixed electrodes. Thus, in (1b) we can write f in the form

$$f(\theta) = \sum_{i=1}^{n} f_i \chi(\theta, \theta_i), \tag{34}$$

where χ is the characteristic function

$$\chi(\theta, \theta_i) = \begin{cases} 1/h \text{ for} & |\theta - \theta_i| \leq h/2, \\ 0 \quad \text{otherwise,} \end{cases}$$

h is the electrode width, and $\{\theta_i\}_{i=1}^n$ are the electrode centers. For simplicity set $\theta_i = 2\pi i/n$, $i = 1, \cdots, n$. We identify the current pattern function f with the n-vector $(f_i)_{i=1}^n$. Finally, the measurements are taken as voltage drops between adjacent electrodes. That is, let u_i be the voltage potential at electrode i, then the data, corresponding to a current pattern f, is the $\mathbb{R}^n$ vector

$$g = (g_1, g_2, \cdots, g_n)^T,$$

where $g_i = u_{i+1} - u_i$ (electrode $n+1$ is identified with electrode 1), and it is understood that u solves (1) with (34). We make n experiments, that is, we apply n different current patterns

$$f^{(1)}, f^{(2)}, \ldots, f^{(n)},$$

and measure the corresponding voltage drops

$$g^{(1)}, g^{(2)}, \ldots, g^{(n)}.$$

The problem is then linearized to obtain $DF\ \delta\sigma = \delta g$, as described in Section 2. We note that this arrangement of current patterns and measurements is similar to one which has been implemented in medical imaging [2].

Extending the domain to the unit square

$$Q = (-1, 1) \times (-1, 1),$$

we discretize the conductivity perturbation $\delta\sigma$ to lie the space of piecewise constant functions on a uniform square grid with $N \times N$ cells. The linearized map DF acting on a $\delta\sigma$ defined over Q is simply defined to be DF acting on the restriction of $\delta\sigma$ to Ω. In this case $h_\epsilon(|\nabla\delta\sigma|)$ is a measure supported on the lines of the grid. Thus if we denote the value of $\delta\sigma$ on the (i,j)-th cell by $\delta\sigma_{i,j}$ then

$$\int_Q h_\epsilon(|\nabla\delta\sigma|) = \frac{1}{N} \sum_{i,j=1}^{N-1} h_\epsilon(|\delta\sigma_{i+1,j} - \delta\sigma_{i,j}|) + h_\epsilon(|\delta\sigma_{i,j+1} - \delta\sigma_{i,j}|).$$

Also discretizing the domain of DF, the stabilized constraint operator M becomes a $p \times N^2$ matrix. From now on we denote the $\mathbb{R}^{N^2}$ vector $(\delta\sigma_{i,j})_{i,j=1}^N$ by $\delta\sigma$. Problem (28) can then be written

$$\min_{\delta\sigma \in X} J_\epsilon^N(\delta\sigma) = \frac{1}{N} \sum_{i,j=1}^{N-1} h_\epsilon(|\delta\sigma_{i+1,j} - \delta\sigma_{i,j}|) + h_\epsilon(|\delta\sigma_{i,j+1} - \delta\sigma_{i,j}|), \quad (35a)$$

$$\text{subject to } M\delta\sigma = \delta g', \quad (35b)$$

where X is the subset of $\mathbb{R}^{N^2}$ with zero boundary values, that is,

$$X = \{\delta\sigma \in \mathbb{R}^{N^2} : \delta\sigma_{1,k} = \delta\sigma_{N,k} = \delta\sigma_{k,1} = \delta\sigma_{k,N} = 0, \quad k = 1, \ldots, N\}.$$

The gradient of the functional $J(\delta\sigma)$ is easy to calculate. For $1 < k, l < N$,

$$\begin{aligned}\frac{\partial J_\epsilon^N}{\partial \delta\sigma_{k,l}} &= q_\epsilon(|\delta\sigma_{k,l} - \delta\sigma_{k-1,l}|)(\delta\sigma_{k,l} - \delta\sigma_{k-1,l}) \\ &- q_\epsilon(|\delta\sigma_{k+1,l} - \delta\sigma_{k,l}|)(\delta\sigma_{k+1,l} - \delta\sigma_{k,l}) \\ &+ q_\epsilon(|\delta\sigma_{k,l} - \delta\sigma_{k,l-1}|)(\delta\sigma_{k,l} - \delta\sigma_{k,l-1}) \\ &- q_\epsilon(|\delta\sigma_{k,l+1} - \delta\sigma_{k,l}|)(\delta\sigma_{k,l+1} - \delta\sigma_{k,l})\ .\end{aligned} \tag{36}$$

With the convention that $\partial J_\epsilon^N / \partial \delta\sigma_{k,l} = 0$ for k or l equal to 1 or N, the gradient of $J_\epsilon^N(\delta\sigma)$ over X may be written

$$\operatorname{grad}_{\delta\sigma} J_\epsilon^N = \left(\frac{\partial J_\epsilon^N}{\partial \delta\sigma_{k,l}}\right)_{k,l=1}^N .$$

With this formula, one can easily compute the projected gradient G_i as described in the previous section. The gradient descent method is then applied to solve (35).

In the following experiments we use a 32×32 grid. The number of electrodes is fixed at 20 and we make 20 measurements, so the data vector δg consists of 400 data points. We remark that the method is feasible for finer discretizations. Most of the computational work in the problem is expended computing DF and its singular value decomposition. The SVD can be avoided by considering other equivalent constraints or by incorporating the constraints as penalty as in (22). The data error tolerance is set at 0.0001, which resulted in $p = 150$ "useable" singular values in the SVD of DF. Thus the constraint operator M is a 150×1024 matrix and the reduced data vector $\delta g'$ has 150 elements. The smoothing parameter $\epsilon = 0.001$ for all experiments.

In the first experiment we attempt to recover the separated block profile pictured in Figure 1. The pseudo-inverse solution $\delta\sigma_0 = M^\dagger\ \delta g'$ was taken as the starting point in the minimization. Figure 2 compares the initial iterate $\delta\sigma_0$ with the final approximate minimizer. As the figure indicates, a near-exact reconstruction is obtained. The reconstruction not only sharpens the edges of the image, but also recovers nearly exactly the true values of the image on the "blocks". This example confirms our expectations from Sections 4 and 5 that good reconstructions can be obtained for images with a relatively small number of non-constant edges ν, and hence relatively small total variation (compared to "mass").

Unfortunately it is not always possible to recover images as accurately as in the previous example. As described in Sections 4 and 5, one should not expect good reconstructions for images $\delta\sigma_0$ with high total variation (or number of non-constant edges) and low "mass". This is illustrated in the following example. Figure 3 shows the true image consisting of a thin curved line. Figure 4 shows the comparison between the pseudo-inverse solution and the minimal total variation reconstruction. As it turns out, in this example the reconstruction has much

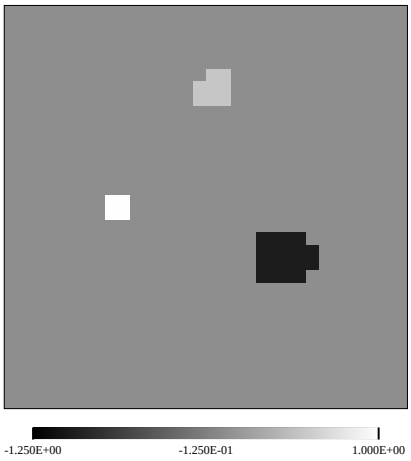

Fig. 1. True profile for the first example.

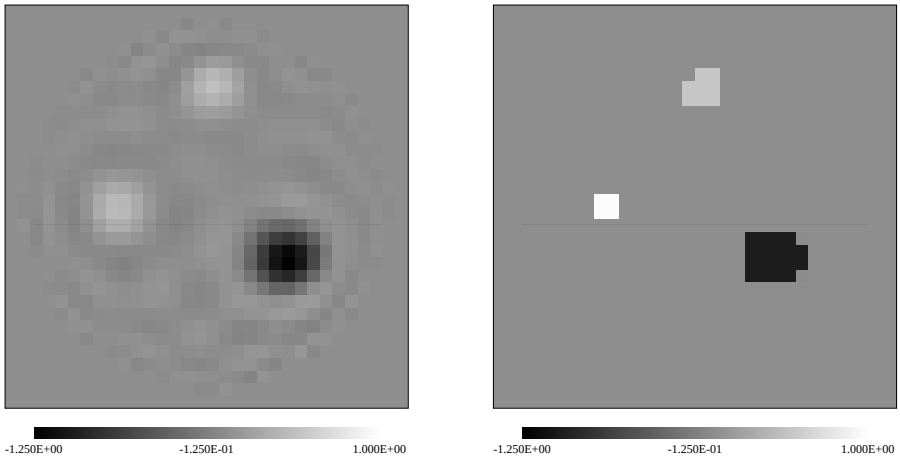

Fig. 2. Comparison of pseudo-inverse solution and minimal total variation reconstruction for the first example. a.) Pseudo-inverse. b.) Reconstruction.

smaller total variation than the original image: approximately half as much. In fact, even the pseudo-inverse has less total variation than the original image.

The preceding two examples were in some sense extreme cases. For the last example, we try to indicate the behavior of the method in a slightly more realistic situation. A caricature of a cross-section of the human torso is pictured in Figure 5. The large light-colored areas on the right and left indicate the lungs, the slightly darker area between the lungs represents muscle and other tissue, and the two dark areas in the lower center represent heart chambers filled with

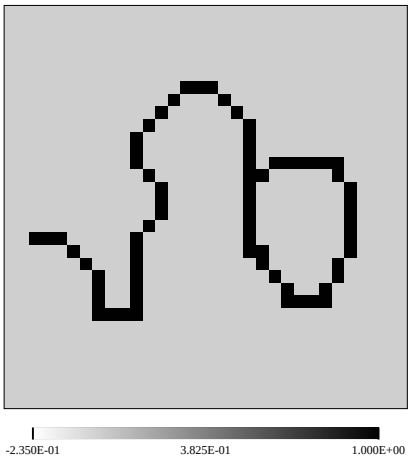

Fig. 3. True profile for the second example.

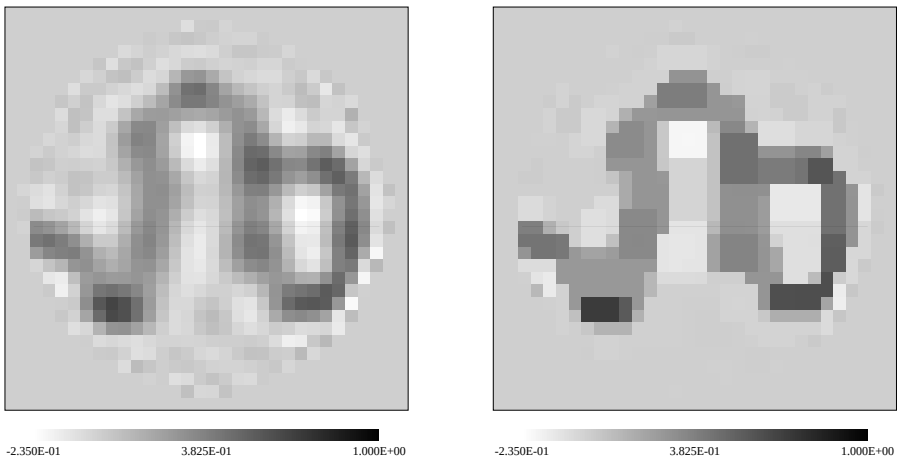

Fig. 4. Comparison of pseudo-inverse solution and minimal total variation reconstruction for the second example. a.) Pseudo-inverse. b.) Reconstruction.

blood. Many details have been omitted from the image, of course. Since the underlying problem is linearized, one might think of Figure 5 as the difference between the true image and a background conductivity.

As can be seen from Figure 6, the reconstruction recovers some features not visible in the pseudo-inverse, most noticeably the lungs and the tissue between the lungs. In this example as in the previous example, the reconstruction had slightly lower total variation than the original image, however in this case the pseudo-inverse had higher total variation.

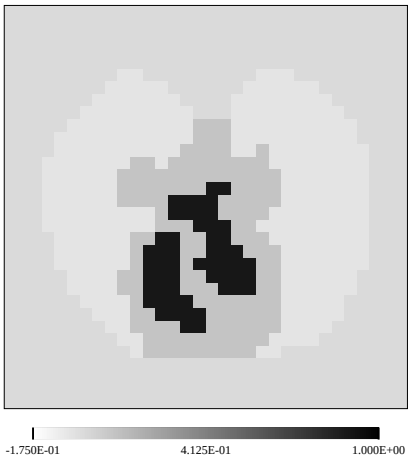

Fig. 5. True profile for the third example.

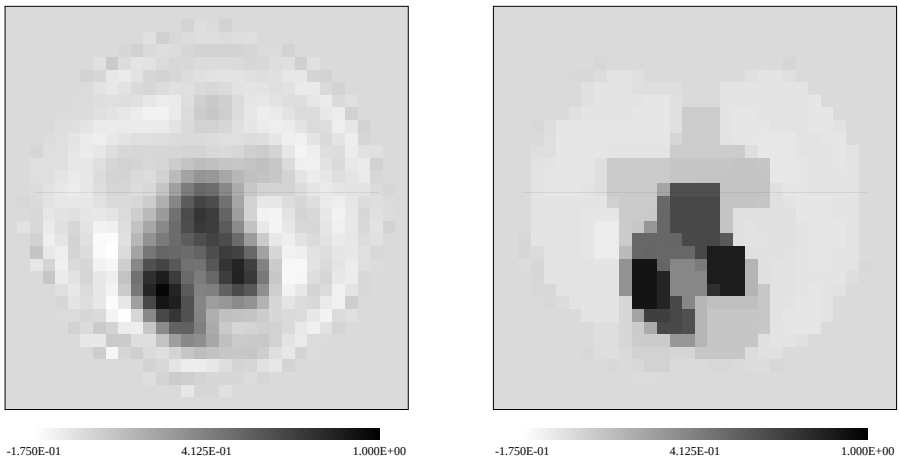

Fig. 6. Comparison of pseudo-inverse solution and minimal total variation reconstruction for the third example. a.) Pseudo-inverse. b.) Reconstruction.

In conclusion, we believe that total variation based methods and other edge-preserving regularizations have great potential in electrical impedance tomography, and in numerous other applications involving inverse problems, optimal design, and image processing. Further research is necessary to develop efficient computational techniques, applications to fully nonlinear problems, and to better understand questions of recoverability, such as those touched upon here.

Acknowledgements

Most of the work reported here was done jointly with Fadil Santosa. The author wishes to express sincere thanks the organizers of the Oberwolfach conference. This work was partially supported by AFOSR Grant number F49620-95-1-0497.

References

1. A. Allers and F. Santosa, *Stability and resolution analysis of a linearized problem in electrical impedance tomography*, Inverse Problems, 7 (1991), pp. 515–533.
2. D. Barber, B. Brown, and J. Jossinet, *Electrical impedance tomography*, Clinical Physics and Physiological Measurements, 9 (1988). Supplement A.
3. W.R. Breckon and M.K. Pidcock, *Some mathematical aspects of electrical impedance tomography*, in Mathematics and Computer Science in Medical Imaging, M.A. Viergever and A. Todd-Pokropek, eds., NATO ASI Series, Springer Verlag (1987), pp. 351–362.
4. A.P. Calderón, *On an inverse boundary value problem*, in Seminar on Numerical Analysis and its Applications to Continuum Physics (Soc. Brasileira de Matematica, Rio de Janerio (1980).
5. F. Catté, P.L. Lions, J. Morel, and T. Coll, *Image selective smoothing and edge detection by nonlinear diffusion*, SIAM J. Numer. Anal., 29 (1992), pp. 182–193.
6. T. F. Chan, H. M. Zhou, and R. H. Chan, *A Continuation Method for Total Variation Denoising Problems*, UCLA CAM Report 95-18.
7. A. Chambolle and P. L. Lions, *Image recovery via total variation minimization and related problems*, Research Report No. 9509, CEREMADE, Universite de Paris-Dauphine, 1995.
8. P. Charbonnier, L. Blanc-Feraud, G. Aubert, and M. Barlaud, *Deterministic edge-preserving regularization in computed imaging*, Research Report no. 94-01, Univ. of Nice–Sophia Antipolis, 1994.
9. T. Coleman and Y. Li, *A globally and quadratically convergent affine scaling method for linear l_1 problems*, Mathematical Programming, 56 (1992), pp. 189–222.
10. T. Coleman and J. Liu, *An interior Newton method for quadratic programming*, Cornell University Department of Computer Science Preprint TR 93-1388, 1993.
11. D. Dobson, *Estimates on resolution and stabilization for the linearized inverse conductivity problem*, Inverse Problems, 8 (1992), pp. 71–81.
12. D. Dobson, *Exploiting ill-posedness in the design of diffractive optical structures,* in "Mathematics in Smart Structures", H. T. Banks, ed., SPIE Proc. **1919** (1993), pp. 248–257.
13. D. Dobson and F. Santosa, *An image-enhancement technique for electrical impedance tomography*, Inverse Problems 10 (1994) pp. 317–334.
14. D. Dobson and F. Santosa, *Recovery of blocky images from noisy and blurred data*, SIAM J. Appl. Math., to appear.
15. D. Dobson and F. Santosa, *Resolution and stability analysis of an inverse problem in electrical impedance tomography—dependence on the input current patterns.* SIAM J. Appl. Math, 54 (1994) pp. 1542-1560.
16. D. Donoho, *Superresolution via sparsity constraints*, SIAM J. Math. Anal., 23 (1992), pp. 1309–1331.

17. D. Geman and C. Yang, *Nonlinear image recovery with half-quadratic regularization*, IEEE Trans. Image Proc., vol. 4 (1995), pp. 932-945.
18. E. Giusti, *Minimal Surfaces and Functions of Bounded Variation*, Birkhauser, Boston, 1984. Monographs in Mathematics, Vol. 80.
19. G. Golub and C. V. Loan, *Matrix Computations*, Johns Hopkins, 1983.
20. K. Ito and K. Kunisch, *An active set strategy based on the augmented Lagrangian formulation for image restoration*, preprint (1995).
21. Y. Li and F. Santosa, *An affine scaling algorithm for minimizing total variation in image enhancement*, Cornell Theory Center Technical Report 12/94, submitted to IEEE Trans. Image Proc.
22. S. Osher and L.I. Rudin, *Feature-oriented image enhancement using shock filters*, SIAM J. Numer. Anal., 27 (1990), pp. 919–940.
23. L.I. Rudin, S. Osher, and E. Fatemi, *Nonlinear total variation based noise removal algorithms*, Physica D., 60 (1992), pp. 259–268.
24. L.I. Rudin, S. Osher, and C. Fu, *Total variation based restoration of noisy blurred images*, SIAM J. Num. Anal., to appear.
25. F. Santosa and W. Symes, *Linear inversion of band-limited reflection seismograms*, SIAM J. Sci. Stat. Comput., 7 (1986), pp. 1307–1330.
26. F. Santosa and W. Symes, *Reconstruction of blocky impedance profiles from normal-incidence reflection seismograms which are band-limited and miscalibrated*, Wave Motion, 10 (1988), pp. 209–230.
27. C.R. Vogel and M. E. Oman, *Iterative methods for total variation denoising*, SIAM J. Sci. Comput., to appear.
28. C.R. Vogel and M. E. Oman, *Fast numerical methods for total variation minimization in image reconstruction*, in SPIE Proc. Vol. 2563, Advanced Signal Processing Algorithms, July 1995.

This article was processed using the LaTeX macro package with LMAMULT style