Geophys. J. Int. (1998) **134,** 213–227

Non-linear inversion using general measures of data misfit and model structure

Colin G. Farquharson and Douglas W. Oldenburg

UBC-Geophysical Inversion Facility, Department of Earth & Ocean Sciences, University of British Columbia, Vancouver, BC, Canada.
E-mail: farq@geop.ubc.ca

Accepted 1998 January 28. Received 1998 January 16; in original form 1997 April 23

SUMMARY

We investigate the use of general, non-l_2 measures of data misfit and model structure in the solution of the non-linear inverse problem. Of particular interest are robust measures of data misfit, and measures of model structure which enable piecewise-constant models to be constructed. General measures can be incorporated into traditional linearized, iterative solutions to the non-linear problem through the use of an iteratively reweighted least-squares (IRLS) algorithm. We show how such an algorithm can be used to solve the linear inverse problem when general measures of misfit and structure are considered. The magnetic stripe example of Parker (1994) is used as an illustration. This example also emphasizes the benefits of using a robust measure of misfit when outliers are present in the data. We then show how the IRLS algorithm can be used within a linearized, iterative solution to the non-linear problem. The relevant procedure contains two iterative loops which can be combined in a number of ways. We present two possibilities. The first involves a line search to determine the most appropriate value of the trade-off parameter and the complete solution, via the IRLS algorithm, of the linearized inverse problem for each value of the trade-off parameter. In the second approach, a schedule of prescribed values for the trade-off parameter is used and the iterations required by the IRLS algorithm are combined with those for the linearized, iterative inversion procedure. These two variations are then applied to the 1-D inversion of both synthetic and field time-domain electromagnetic data.

Key words: inversion, non-linear problems.

1 INTRODUCTION

In this paper we investigate the solution of the non-linear inverse problem using measures other than the typical sum-of-squares, or l_2, measure. Our motivation is two-fold. First, we would like our solution to be less influenced by outliers and other non-Gaussian noise, that is to be more robust, than when an l_2 measure of data misfit is used. Second, we would like to be able to construct models that are piecewise constant, or 'blocky', and therefore more in accord with the geology of certain regions than the smooth, 'smeared-out' images obtained using an l_2 measure of model structure. These two aims can be achieved with the use of judiciously chosen measures of data misfit and model structure, for example an l_1 measure, or the measure related to the 'most robust' of the M-estimators of Huber (1964). Traditional linearized, iterative inversion procedures can be modified to handle such non-l_2 measures by the inclusion of an iteratively reweighted least-squares algorithm.

There have been a number of reports of the use of robust measures of misfit in geophysical inverse problems. Claerbout & Muir (1973) espouse the merits of an l_1 measure when a set of observations contains outliers or 'infinite blunders'. In the context of the 1-D seismic inverse problem, Gersztenkorn, Bednar & Lines (1986) illustrate the superiority of a solution obtained using an l_1 measure over one obtained using an l_2 measure when spikes (i.e. outliers) are present in the seismic trace. Scales, Gersztenkorn & Treitel (1988) provide a similar illustration of the robustness of an l_1 measure in the context of seismic tomography. Both sets of authors use an iteratively reweighted least-squares algorithm. Crase *et al.* (1990) provide another illustration of the benefits of robust measures of misfit when they compare the results of seismic waveform inversion using a number of robust measures and the traditional l_2 measure. Amundsen (1991) likewise provides a comparison for seismic frequency–wavenumber inversion. In an application to gravity data, Beltrão, Silva & Costa (1991) use a robust measure when fitting polynomials to observations to determine

the regional field. Their method uses an iteratively reweighted least-squares procedure. Finally, to make their procedure more robust in the presence of realistic noise, LaBrecque *et al.* (1997) use an iteratively reweighted least-squares scheme to down-weight the importance of outliers in their 3-D inversion of electrical resistance tomography data.

The use of robust measures of misfit has also led to significant improvements in the processing of magnetotelluric and geomagnetic depth sounding data (see Egbert & Booker 1986; Chave, Thomson & Ander 1987; Chave & Thomson 1989; Sutarno & Vozoff 1991; Larsen *et al.* 1996). The time-series of the observed electric and magnetic fields from which frequency-domain response functions are computed contain noise that can rarely be considered as purely Gaussian. Robust measures of misfit therefore give much more reliable estimates of the response functions than does an l_2 measure.

Mention has also been made, although to a much lesser extent and predominantly for linear inverse problems, of using non-l_2 measures to construct models possessing some character other than being 'smeared-out' and smooth. Claerbout & Muir (1973) put forward the possibility of obtaining sparse or piecewise-constant models by minimizing an l_1 measure of model structure. Levy & Fullagar (1981) obtain sparse models in seismic deconvolution by minimizing an l_1 norm of the model. Oldenburg (1984a) constructs minimum-structure models when inverting lead isotope data by minimizing an l_1 norm of the model's gradient. Oldenburg (1984b) discusses the use of l_1 measures of model structure and data misfit for the general linear inverse problem. Alliney & Ruzinsky (1994) investigate the solution of the general linear problem when an objective function comprising an l_1 measure of model structure and an l_2 measure of data misfit is minimized. In the context of seismic data processing, Sacchi & Ulrych (1995, 1996) obtain sparse solutions to Radon and Fourier transforms by considering the respective inverse transforms as linear inverse problems and then solving by minimizing an objective function made up of a suitable measure of model structure and an l_2 measure of misfit.

As yet very little use has been made of non-l_2 measures of model structure in the solution of the non-linear geophysical inverse problem. There are some notable exceptions, however. Dosso (1990) minimizes the l_1 norm of the model's vertical derivative in order to construct piecewise-constant models for the 1-D magnetotelluric problem. Dosso also uses an l_1 measure of data misfit and solves the linearized inverse problem using linear programming. Oldenburg & Ellis (1993) and Ellis, Farquharson & Oldenburg (1993) use an l_1 measure to construct blocky models for the 2-D magnetotelluric problem. These authors also use linear programming.

Here we present a means of using general measures of both data misfit and model structure in the solution of the non-linear geophysical inverse problem. The technique is based on the standard linearized, iterative approach (see e.g. Constable, Parker & Constable 1987; Smith & Booker 1988; deGroot-Hedlin & Constable 1990), but with the solution of the system of equations at each iteration, which is non-linear for general non-l_2 measures, being obtained using an iteratively reweighted least-squares algorithm. We first introduce the measures of misfit and model structure that particularly interest us. We then show how the solution to a linear inverse problem can be obtained using an iteratively reweighted least-squares procedure, and illustrate this with an example. Finally, we present and discuss two procedures for solving the non-linear inverse problem, and illustrate their behaviour and the results they produce for the 1-D inversion of time-domain electromagnetic data.

2 GENERAL MEASURES, THE LINEAR INVERSE PROBLEM AND ITERATIVELY REWEIGHTED LEAST SQUARES

There are numerous functions that can be used as a measure of the 'size' of a vector, whether the elements of that vector are the misfits between observed and predicted data, or the parameters describing a model. Consider a vector $\mathbf{x}$ with elements x_j, $j = 1, \ldots, N$. Then a general measure of its size is given by

$$\phi(\mathbf{x}) = \sum_{j=1}^{N} \rho(x_j). \tag{1}$$

It can easily be seen that when $\rho(x) = x^2$ this general expression gives (the square of) the l_2 norm of the vector:

$$\|\mathbf{x}\|_2^2 = \sum_{j=1}^{N} x_j^2. \tag{2}$$

The measure in eq. (1) also includes the more general l_p norm:

$$\|\mathbf{x}\|_p^p = \sum_{j=1}^{N} |x_j|^p, \tag{3}$$

where $1 \leq p < \infty$. Two other measures that we consider in this paper are one related to the particular M-estimator that was considered the 'most robust' by Huber (1964), and one proposed by Ekblom (1973, 1987). The first of these corresponds to

$$\rho(x) = \begin{cases} x^2 & |x| \leq c, \\ 2c|x| - c^2 & |x| > c, \end{cases} \tag{4}$$

where c is a positive constant that separates the elements of the vector $\mathbf{x}$ into those that are considered large and those that are considered small. The large values then contribute less to this measure than they do to the l_2 norm. This measure is also linked to a probability density function for x_j that is Gaussian in the centre and exponential in the tails (Huber 1964). For the purposes of this paper we shall call the measure given by eqs (1) and (4) the 'M-measure' because of its origin amongst Huber's M-estimators. The measure proposed by Ekblom (1973, 1987) that we have found useful, especially from a numerical point of view, is a perturbed version of the l_p norm:

$$\rho(x) = (x^2 + \varepsilon^2)^{p/2}, \tag{5}$$

where ε is some positive number. As ε becomes small, this measure tends towards the l_p norm given in eq. (3). For large ε, the above measure behaves like a scaled sum-of-squares measure. This can be seen from the following series expansion of eq. (5) valid for large ε (Ekblom 1973):

$$\rho(x) = \varepsilon^p \left[1 + \frac{p}{2} \frac{x^2}{\varepsilon^2} + O\left(\frac{x^4}{\varepsilon^4} \right) \right]. \tag{6}$$

Consider now a linear inverse problem that we want to solve using the general expression in eq. (1) as a measure of both data misfit and model structure. We assume that the model is the spatial distribution of some physical property of

the Earth, $m(\mathbf{r})$, and that it is discretized as follows:

$$m(\mathbf{r}) = \sum_{j=1}^{N} m_j \psi_j(\mathbf{r}), \tag{7}$$

where $\mathbf{r}$ is the position vector and ψ_j are basis functions. Usually these basis functions are pulse basis functions that are equal to unity within a cuboidal cell and zero everywhere else. In such a case m_j is the value of the physical property in the jth cell. Let the forward problem of determining the set of observations for a particular model be expressed as

$$\mathbf{d}^{\mathrm{prd}} = \mathbf{G}\mathbf{m}, \tag{8}$$

where $\mathbf{m} = (m_1, \ldots, m_N)^{\mathrm{T}}$, $\mathbf{d}^{\mathrm{prd}}$ is the vector of M predicted data for this model, and $\mathbf{G}$ is an $M \times N$ matrix that contains essentially all the physics of this linear problem. The inverse problem consists of finding the one model from the collection that reproduce the observations to an appropriate level of misfit that is the most plausible model given any prior knowledge of the region under investigation.

To solve the inverse problem we construct an objective function that we are going to minimize:

$$\begin{aligned}\mathbf{\Phi} = {} & \alpha_s \phi_s(\mathbf{W}_s(\mathbf{m} - \mathbf{m}^{\mathrm{ref}})) + \alpha_x \phi_x(\mathbf{W}_x \mathbf{m}) \\ & + \beta[\phi_d(\mathbf{W}_d(\mathbf{d}^{\mathrm{prd}} - \mathbf{d}^{\mathrm{obs}})) - \phi_d^{\mathrm{tar}}]. \end{aligned} \tag{9}$$

In this expression, the three instances of the general measure function, namely ϕ_s, ϕ_x and ϕ_d, are intended to provide measures, respectively, of (1) how close the model, $\mathbf{m}$, is to a reference model $\mathbf{m}^{\mathrm{ref}}$, (2) the amount of structure in the model, and (3) the misfit between the predicted data, $\mathbf{d}^{\mathrm{prd}}$, and the observations, $\mathbf{d}^{\mathrm{obs}}$. What the three terms in eq. (9) are actually measures of depends not only on the particular form used for each measure, but also on the weighting matrices $\mathbf{W}_s$, $\mathbf{W}_x$ and $\mathbf{W}_d$. The particular model character that interests us in this paper is simplicity of structure. We therefore use the first-order finite-difference operator for $\mathbf{W}_x$. For a 1-D inverse problem in which $m(\mathbf{r}) = m(x)$, and for pulse basis functions of equal length, this matrix has the form

$$\mathbf{W}_x = \begin{pmatrix} -1 & 1 & 0 & \ldots & 0 \\ 0 & -1 & 1 & \ldots & 0 \\ \vdots & & \ddots & & \vdots \\ 0 & \ldots & 0 & -1 & 1 \\ 0 & \ldots & 0 & 0 & 0 \end{pmatrix}. \tag{10}$$

If the pulse basis functions are of varying lengths, $\mathbf{W}_x$ can be modified accordingly. We also assume for the purposes of this paper that when a reference model is included it is of equal appropriateness over the complete spatial extent of the model. We therefore take $\mathbf{W}_s$ to be the identity matrix. We also assume that the noise in the observations is not correlated between observations. The data-weighting matrix, $\mathbf{W}_d$, is therefore diagonal. We also expect, for the examples in this paper, that the ith element on the diagonal of this matrix is the reciprocal of the estimate of the uncertainty in the ith datum. If the noise is Gaussian with zero mean this estimate of uncertainty is the usual standard deviation. Returning to eq. (9), the constant, ϕ_d^{tar}, is the desired target misfit appropriate for the amount of noise in the observations. As a final comment, the three constants α_s, α_x and β, which determine the balance of the three terms in $\mathbf{\Phi}$, are not independent. Three are used only to provide a certain aesthetic symmetry to the objective function.

We now solve the linear inverse problem by finding the model $\mathbf{m}$ that minimizes the objective function $\mathbf{\Phi}$ given in eq. (9). To do this we follow the standard procedure of differentiating $\mathbf{\Phi}$ with respect to the model parameters and equating the resulting equations to zero. Consider first the differentiation of the general measure given in eq. (1) with respect to the kth model parameter, assuming that $\mathbf{x}$ is some function of the model $\mathbf{m}$:

$$\frac{\partial \phi(\mathbf{x})}{\partial m_k} = \sum_{j=1}^{N} \rho'(x_j) \frac{\partial x_j}{\partial m_k}. \tag{11}$$

If all values of k are considered, this relationship can be written as

$$\frac{\partial \phi(\mathbf{x})}{\partial \mathbf{m}} = \mathbf{B}^{\mathrm{T}} \mathbf{q}, \tag{12}$$

where $\partial\phi/\partial\mathbf{m} = (\partial\phi/\partial m_1, \ldots, \partial\phi/\partial m_N)^{\mathrm{T}}$, $B_{ij} = \partial x_i/\partial m_j$ and $\mathbf{q} = [\rho'(x_1), \ldots, \rho'(x_N)]^{\mathrm{T}}$. The explicit form of $\mathbf{q}$ will depend on the particular form of ρ. However, in order to construct a system of equations upon which an iteratively reweighted least-squares routine can be used, which is our intention, consider the substitution $\mathbf{q} = \mathbf{R}\mathbf{x}$ where $\mathbf{R}$ is a diagonal matrix given by

$$\mathbf{R} = \mathrm{diag}\{\rho'(x_1)/x_1, \ldots, \rho'(x_N)/x_N\}. \tag{13}$$

Eq. (12) therefore becomes

$$\frac{\partial \phi(\mathbf{x})}{\partial \mathbf{m}} = \mathbf{B}^{\mathrm{T}} \mathbf{R} \mathbf{x}. \tag{14}$$

The elements of $\mathbf{R}$ for the three particular measures that we consider in this paper are as follows. First, for the l_p norm given by eq. (3),

$$R_{ii} = p|x_i|^{p-2}. \tag{15}$$

For $p < 2$ this expression is singular for $x_i = 0$. A small cut-off value, γ, is therefore used in any numerical implementation (see e.g. Gersztenkorn *et al.* 1986):

$$R_{ii} = \begin{cases} p\gamma^{p-2} & |x_i| \le \gamma, \\ p|x_i|^{p-2}, & |x_i| > \gamma. \end{cases} \tag{16}$$

Second, for the M-measure given by eqs (4) and (1),

$$R_{ii} = \begin{cases} 2, & |x_i| \le c, \\ 2c/|x_i|, & |x_i| > c. \end{cases} \tag{17}$$

Note that, unlike the l_p norm, there is now no difficulty with small values of x_i. Finally, for Ekblom's measure given by eqs (5) and (1),

$$R_{ii} = p(x_i^2 + \varepsilon^2)^{p/2-1}. \tag{18}$$

Let us now return to the minimization of the whole objective function $\mathbf{\Phi}$ in eq. (9). Consider the differentiation of the term involving ϕ_s. For this term, $\mathbf{x} = \mathbf{W}_s(\mathbf{m} - \mathbf{m}^{\mathrm{ref}})$. The jth element of this vector, written in summation notation, is

$$x_j = \sum_{l=1}^{N} [W_s]_{jl}(m_l - m_l^{\mathrm{ref}}). \tag{19}$$

From the definition given after eq. (12), the matrix $\mathbf{B}$ for this vector $\mathbf{x}$ is given by $B_{ij} = [W_s]_{ij}$. Hence, from eq. (14), the

derivative of ϕ_s with respect to the model $\mathbf{m}$ is

$$\frac{\partial \phi_s}{\partial \mathbf{m}} = \mathbf{W}_s^{\mathrm{T}} \mathbf{R}_s \mathbf{W}_s (\mathbf{m} - \mathbf{m}^{\mathrm{ref}}). \tag{20}$$

$\mathbf{R}_s$ is the matrix given in eq. (13) for the appropriate form of ρ for ϕ_s. Identical reasoning leads to

$$\frac{\partial \phi_x}{\partial \mathbf{m}} = \mathbf{W}_x^{\mathrm{T}} \mathbf{R}_x \mathbf{W}_x \mathbf{m} \tag{21}$$

for the derivative of the second term in $\mathbf{\Phi}$. For the derivative of the third term on the right-hand side of eq. (9), let $\mathbf{x} = \mathbf{W}_d(\mathbf{d}^{\mathrm{prd}} - \mathbf{d}^{\mathrm{obs}})$. Using the definition of the forward modelling given in eq. (8), this becomes $\mathbf{x} = \mathbf{W}_d(\mathbf{G}\mathbf{m} - \mathbf{d}^{\mathrm{obs}})$. The matrix $\mathbf{B}$ for this $\mathbf{x}$ is $\mathbf{B} = \mathbf{W}_d\mathbf{G}$. Hence, from eq. (14),

$$\frac{\partial \phi_d}{\partial \mathbf{m}} = \mathbf{G}^{\mathrm{T}} \mathbf{W}_d^{\mathrm{T}} \mathbf{R}_d \mathbf{W}_d (\mathbf{G}\mathbf{m} - \mathbf{d}^{\mathrm{obs}}). \tag{22}$$

From eqs (20) to (22), the system of equations obtained by differentiating the objective function $\mathbf{\Phi}$ with respect to the model parameters and equating to zero is therefore

$$[\alpha_s \mathbf{W}_s^{\mathrm{T}} \mathbf{R}_s \mathbf{W}_s + \alpha_x \mathbf{W}_x^{\mathrm{T}} \mathbf{R}_x \mathbf{W}_x + \beta \mathbf{G}^{\mathrm{T}} \mathbf{W}_d^{\mathrm{T}} \mathbf{R}_d \mathbf{W}_d \mathbf{G}] \mathbf{m}$$
$$= \beta \mathbf{G}^{\mathrm{T}} \mathbf{W}_d^{\mathrm{T}} \mathbf{R}_d \mathbf{W}_d \mathbf{d}^{\mathrm{obs}} + \alpha_s \mathbf{W}_s^{\mathrm{T}} \mathbf{R}_s \mathbf{W}_s \mathbf{m}^{\mathrm{ref}}. \tag{23}$$

Our solution to the linear inverse problem is the model $\mathbf{m}$ that satisfies this system of equations.

Eq. (23) is very similar to the system of equations that would be obtained if all the measures in the objective function $\mathbf{\Phi}$ were sum-of-squares measures. The difference is the presence of the diagonal matrices $\mathbf{R}_s$, $\mathbf{R}_x$ and $\mathbf{R}_d$. These matrices actually depend on the model, $\mathbf{m}$, so eq. (23) is a non-linear system of equations. This system can be solved with an iterative procedure:

$$\mathbf{m}^{k+1} = \mathbf{M}_k^{-1} \mathbf{y}^k, \tag{24}$$

where, from eq. (23),

$$\mathbf{M}_k = [\alpha_s \mathbf{W}_s^{\mathrm{T}} \mathbf{R}_s^k \mathbf{W}_s + \alpha_x \mathbf{W}_x^{\mathrm{T}} \mathbf{R}_x^k \mathbf{W}_x + \beta \mathbf{G}^{\mathrm{T}} \mathbf{W}_d^{\mathrm{T}} \mathbf{R}_d^k \mathbf{W}_d \mathbf{G}] \tag{25}$$

and

$$\mathbf{y}^k = \beta \mathbf{G}^{\mathrm{T}} \mathbf{W}_d^{\mathrm{T}} \mathbf{R}_d^k \mathbf{W}_d \mathbf{d}^{\mathrm{obs}} + \alpha_s \mathbf{W}_s^{\mathrm{T}} \mathbf{R}_s^k \mathbf{W}_s \mathbf{m}^{\mathrm{ref}}, \tag{26}$$

where $\mathbf{R}_s^k = \mathbf{R}_s(\mathbf{m}^k)$, etc. The starting model in this procedure, $\mathbf{m}^1$, is given by the solution of eq. (23) with $\mathbf{R}_s = \mathbf{R}_x = \mathbf{R}_d = \mathbf{I}$, where $\mathbf{I}$ is the identity matrix. This starting model is the traditional least-squares solution to the linear inverse problem. The matrices $\mathbf{R}_s$, $\mathbf{R}_x$ and $\mathbf{R}_d$ are computed for this model using eq. (13) or whichever of eqs (16) to (18) may be appropriate. The matrix equation (23) is then formed and solved to give the new model. The matrices $\mathbf{R}_s$, $\mathbf{R}_x$ and $\mathbf{R}_d$ are recomputed for this new model and the process repeated. The iterations are terminated once the model is no longer changing by a significant amount. We use LU decomposition, in particular the routine 'sgesv' from LAPACK (Anderson *et al.* 1995), to solve the matrix equation given by eq. (24).

The above iterative procedure for solving the non-linear system of equations is commonly referred to as the iteratively reweighted least-squares (IRLS) procedure. It is very similar to Newton's method for minimizing $\mathbf{\Phi}$. For a comparison of these two methods when a single l_p measure is being minimized, see e.g. Watson (1980) or Ekblom (1987). There have been various theoretical and empirical discussions of the convergence of the IRLS procedure when a single, particular measure is considered (see e.g. Holland & Welsch 1977; Wolke & Schwetlick 1988; O'Leary 1990). Alliney & Ruzinsky (1994) prove the convergence of a descent algorithm when a combination of an l_2 measure of misfit and l_1 measure of model structure is minimized. We have found, as we shall proceed to demonstrate, that the IRLS procedure is successful in the minimization of a general, composite objective function such as the one given in eq. (9), and hence enables us to solve the linear inverse problem in the way we would like.

3 AN EXAMPLE LINEAR INVERSE PROBLEM

Before proceeding to the non-linear inverse problem, which is our ultimate goal, we illustrate the use of the general objective function and the IRLS procedure described in the preceeding section in the solution of a linear inverse problem. The example that we use is the linear magnetic stripe problem of Parker (1994). This is a simplification of the actual geophysical inverse problem of determining the magnetization of the oceanic crust from the magnetic anomalies aligned parallel to mid-ocean ridges.

To provide an instructive, linear problem it is assumed that (1) the magnetization of the oceanic crust occurs within an infinitesimally thin layer at a depth h below the ocean surface, (2) that this magnetization is in the vertical direction, and (3) that its intensity varies only with distance, x, from the ridge axis. The model, m, to be found is the distribution with x of the intensity of the magnetization. The observations are of the vertical component of the magnetic field at the ocean surface at various distances from the ridge axis. If we discretize the model using pulse basis functions (see eq. 7),

$$m(x) = \sum_{j=1}^{N} m_j \psi_j(x), \tag{27}$$

where

$$\psi_j(x) = \begin{cases} 1, & X_{j-1} \le x < X_j, \\ 0, & \text{elsewhere}, \end{cases} \tag{28}$$

then the forward modelling (see eq. 8) is given by

$$\mathbf{d}^{\mathrm{prd}} = \mathbf{G}\mathbf{m}, \tag{29}$$

with

$$G_{ij} = \frac{\mu_0}{2\pi} \left[\frac{x - x_i}{(x - x_i)^2 + h^2} \right]_{x = X_{j-1}}^{X_j}, \tag{30}$$

where x_i is the location of the ith observation, and μ_0 is the magnetic permeability of free space.

We generated synthetic data from the model shown in Fig. 1(b). As done by Parker, the plotted values of magnetization have been scaled by 0.001 and given the units A m^{-1}. This mimics the effect of a magnetized layer that is 1 km thick. Data were computed using the formulae in eqs (29) and (30). Initially only Gaussian noise, of 8 nT standard deviation, was added to these data. The resulting synthetic data are shown in Fig. 1(a). We present the inversion, using l_2 measures, of these outlier-free data in Section 3.1. Outliers were then introduced by perturbing two of the data. The inversion, using various measures, of these perturbed data is described in Sections 3.2 to 3.4.

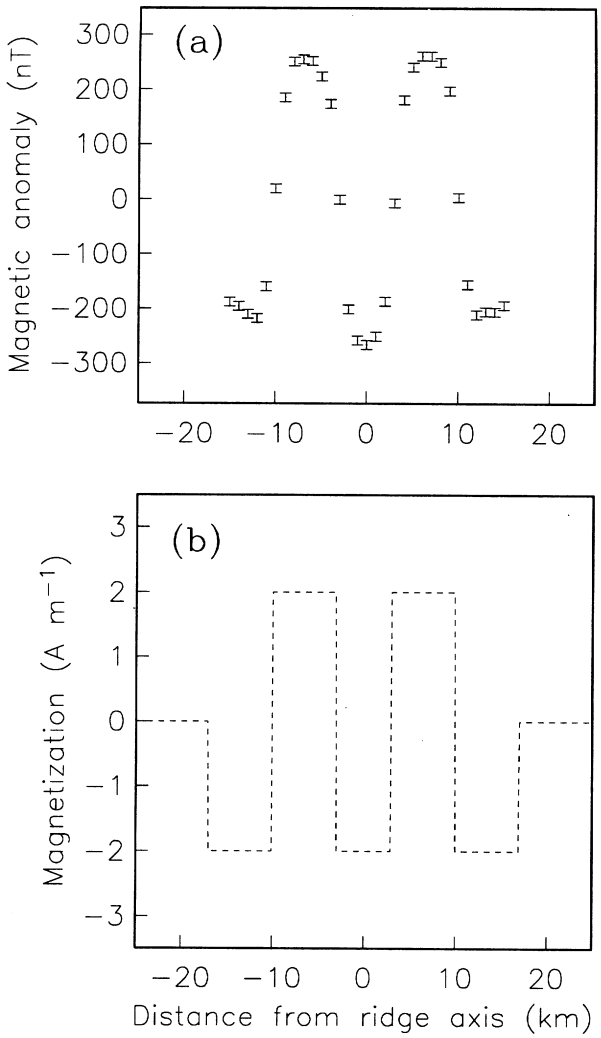

Figure 1. (a) The synthetic data for the linear magnetic stripe example discussed in Section 3.1. The error bars indicate the standard deviation of the noise added to the data. (b) The model from which these data were computed.

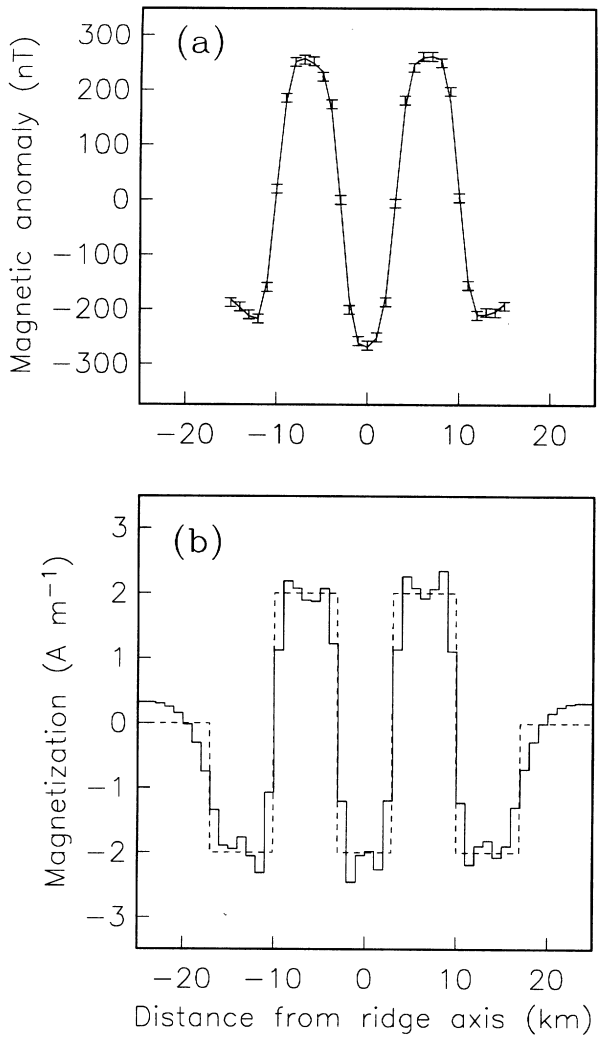

Figure 2. (a) The predicted data (solid line) for the model produced by the inversion of the synthetic magnetic stripe data in Fig. 1(a) using l_2 measures for ϕ_s, ϕ_x and ϕ_d. The error bars indicate the synthetic data. (b) The model (solid line) produced by this l_2 inversion. The dashed line represents the true model.

3.1 Gaussian noise and l_2 inversion

Consider first an inversion of the data contaminated with only Gaussian noise. We used the square of the l_2 norm for each of ϕ_s, ϕ_x and ϕ_d in the objective function $\mathbf{\Phi}$. A minimum-structure model was desired, so we used the first-order finite-difference operator given in eq. (10) for $\mathbf{W}_x$. The reference model was a constant magnetization of zero, and $\alpha_s = 0.1$ and $\alpha_x = 1$. The model was discretized into 60 equally sized cells extending from $x = -30$ km to $x = 30$ km. Outwith these limits the model was fixed and equal to zero. The elements of $\mathbf{W}_d$ were the reciprocals of the standard deviations of the noise in the data. Since the noise in the data is solely Gaussian, the expected value of the sum-of-squares measure of misfit is equal to 31 (the number of observations for this example). This value was therefore used as the target misfit. The first iteration of the procedure described in Section 2 was used to give the solution for this particular example. The target misfit of 31 was obtained with $\beta = 2.4 \times 10^5$. The corresponding model and predicted data are shown in Fig. 2.

The inclusion of a reference model was required in this and all other inversions of the magnetic stripe data to fix the background value of the models. This is because the kernels [that is the term $(x - x_i)/[(x - x_i)^2 + h^2]$ in eq. 30] for this inverse problem have zero area, meaning the data are insensitive to the background value of magnetization.

3.2 Outliers and l_2 inversion

To illustrate the benefits of using robust measures of misfit, we perturbed two of the synthetic data discussed in the previous section. The datum at $x = -8$ km was changed from 251 to 300 nT, and the datum at $x = 3$ km was changed from -6.7 to 181 nT. (This new value is equal to the neighbouring datum at $x = 4$ km.) The estimated uncertainties for these two data were not changed from 8 nT. The modified data set is shown in Fig. 3(a).

The inversion described in Section 3.1 involving sum-of-squares measures for each of ϕ_s, ϕ_x and ϕ_d was repeated for

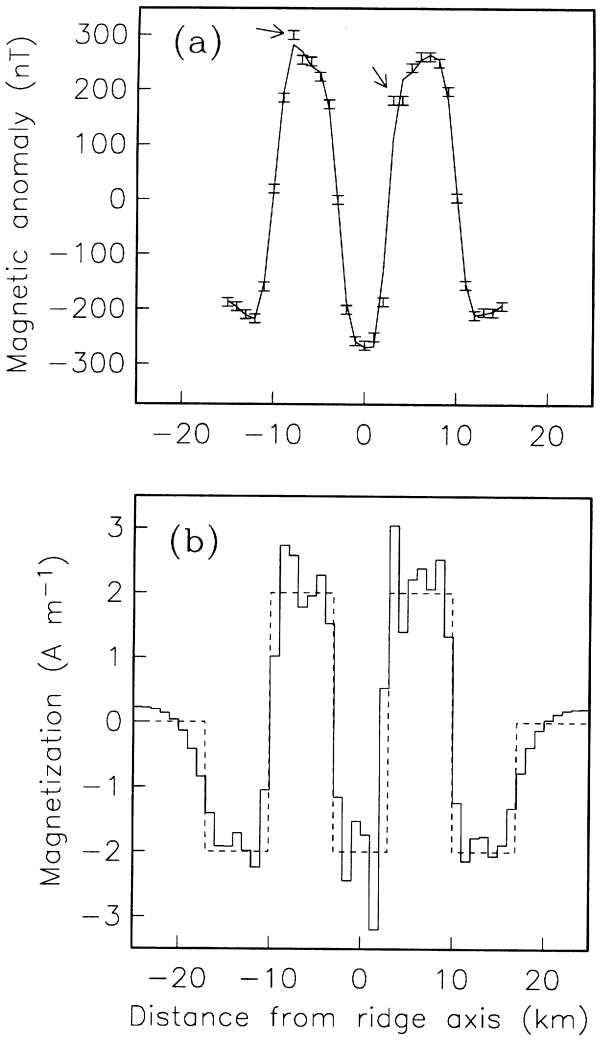

Figure 3. (a) The perturbed synthetic magnetic stripe data. The two outliers are indicated by the arrows. The solid line represents the predicted data for the model obtained using l_2 measures for both data misfit and model structure. This model is shown as the solid line in panel (b). The dashed line in panel (b) represents the model from which the synthetic data were generated.

the perturbed data set. The model deemed to be the most appropriate is shown in Fig. 3(b). The corresponding predicted data are shown in Fig. 3(a). The value of β giving this model was 4.27×10^5. The corresponding misfit was $\phi_d = 186 = 6M$. This is six times greater than what one would attempt to achieve if the data were not thought to contain outliers. However, models corresponding to lower values of misfit contained more structure than seemed justified by the data. Also, the corresponding predicted data showed a large amount of variation around the perturbed datum at $x = 3$ km. Such models were therefore deemed unacceptable. Note that even though the chosen model gives a misfit six times greater than that achieved for the inversion of the outlier-free data, it still contains significantly more structure than the model for that inversion (see Fig. 2b). This is because such a large contribution to the sum-of-squares measure of misfit comes from the two outliers that the other data have to be fit too closely in order to get the misfit even as small as $6M$.

3.3 Outliers and l_1 inversion

The preceeding example hints at the unsuitability of sum-of-squares measures of misfit when the noise in the data is non-Gaussian in general, and when outliers are present in particular. We repeated the inversion of the previous section using a robust measure of misfit. The chosen measure was that of Ekblom given by eqs (5) and (1) with $p = 1$. We also used this measure for ϕ_x in order to produce a piecewise-constant model, and for ϕ_s. The value of ε for ϕ_x was 1, and for ϕ_d, 0.01. These values were chosen as follows. The measures ϕ_x, ϕ_s and ϕ_d were computed for a range of ε using the final model and predicted data for the l_2 inversion in the previous section. The resulting curves are shown in Fig. 4. We recognize that the curves for the model and predicted data produced by the inversion in this section will not be identical to those in Fig. 4. However, the curves in Fig. 4 provide a means of estimating values of ε, prior to carrying out the inversion, that are small enough that the final values of ϕ_x, ϕ_s and ϕ_d will not differ significantly from their values with $\varepsilon = 0$. Values of ε lying about two orders of magnitude to the left of each point of maximum curvature were small enough that ϕ_x, ϕ_s and ϕ_d differed only in the third significant figure from their values with ε equal to zero, and yet large enough that the IRLS procedure converged. The inversion in this section was therefore minimizing, in essence, l_1 norms of misfit and model structure.

From Parker & McNutt (1980), an appropriate target value for an l_1 norm misfit is $\sqrt{2/\pi}\,M$ if there is only Gaussian noise present in the data. The value of this target misfit for this example is 24.7. However, the model that gave this misfit possessed an unacceptable amount of structure. The model chosen as the solution corresponded to an l_1 norm misfit of 39.5 (the value of ϕ_d was 39.6). This model is shown in Fig. 5(b), and the predicted data in Fig. 5(a). A comparison of Figs 3 and 5 highlights two important points. First, the fit to the perturbed data obtained using essentially an l_1 norm

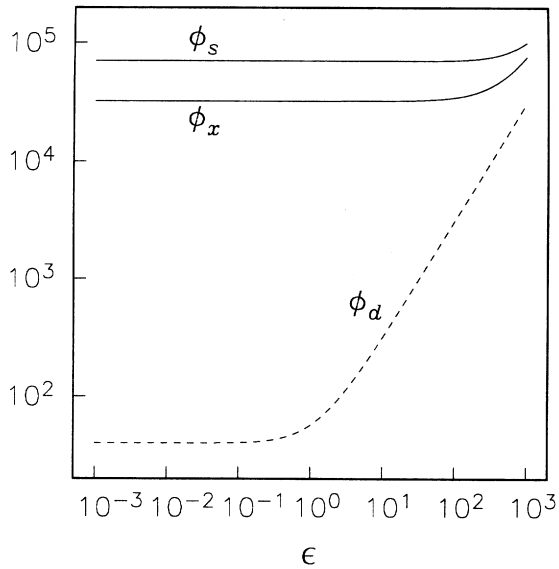

Figure 4. The variation with ε of Ekblom's measure (with $p = 1$) for ϕ_x, ϕ_s and ϕ_d computed using the final model and corresponding predicted data from the l_2 inversion in Section 3.2.

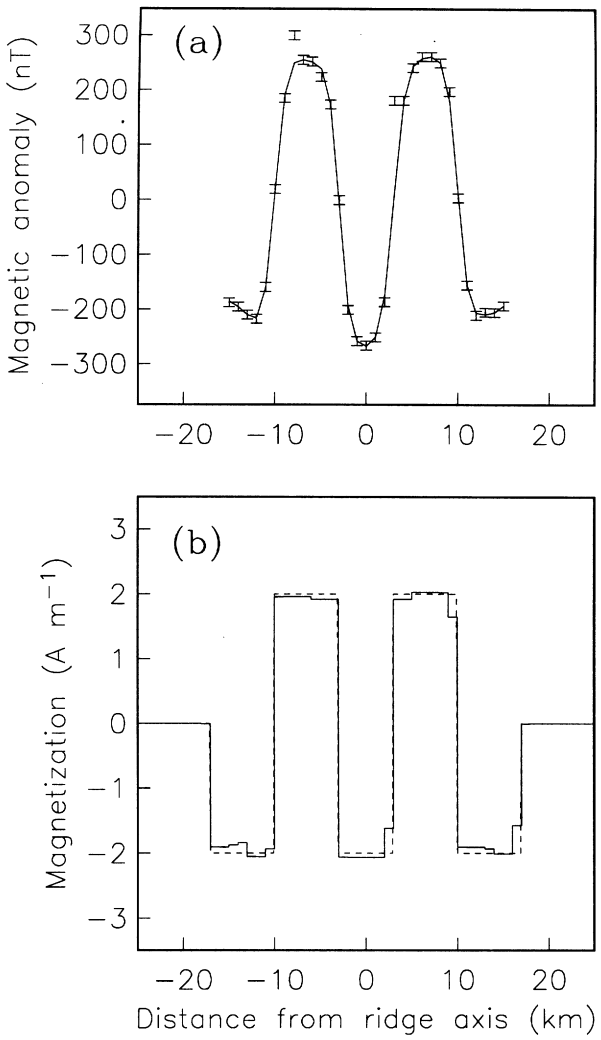

Figure 5. The predicted data (solid line, a) and model (solid line, b) resulting from the inversion of the perturbed synthetic magnetic stripe data (shown by the error bars in a) using Ekblom's measure with $p = 1$ for ϕ_s, ϕ_x and ϕ_d.

measure of misfit is more appropriate than that obtained using a sum-of-squares measure. The two outliers in Fig. 5(a) are in effect ignored by the predicted data. This is not the case in Fig. 3(a). Second, a model produced by minimizing the l_1 norm of its gradient has a piecewise-constant, or 'blocky', appearance which can make it preferable in many situations to the smooth, smeared-out models produced using l_2 measures.

The IRLS loop required 93 iterations for the results shown in Fig. 5. When the values of ε were increased by an order of magnitude, 59 iterations were required for convergence. (The convergence criterion was that no model parameter had changed by more than 0.001 A m^{-1} between successive iterations.) The resulting model did differ slightly from the one shown in Fig. 5, although differences in the predicted data were not discernable. It is only when p is equal to or very close to 1 that a large number of IRLS iterations is required. For other values of p, the number of IRLS iterations is typically 20 or 30.

Finally, the inversion in this section was also carried out for true l_1 norms using (1) linear programming and (2) the IRLS procedure with weights given by eq. (16) and $\gamma = 0.01$. The results produced by both of these approaches were essentially identical to those in Fig. 5. The IRLS procedure required 79 iterations for this example.

3.4 Outliers and M-measure/l_2 inversion

As a final example based on the linear magnetic stripe data, we inverted the perturbed data using an M-measure (see eqs 1 and 4) as the measure of misfit and Ekblom's measure with $p = 1$ for ϕ_s and ϕ_x. As before, the value of ε in the Ekblom measure was 1. The weighting matrices $\mathbf{W}_s$, $\mathbf{W}_x$ and $\mathbf{W}_d$, and the values of α_s and α_x were also the same as before. If a value of $c = 1.5$ is taken as the constant in eq. (4), and only Gaussian noise contaminates the data, then the expected value of the M-measure for 31 data is 30.3. Also, the average number of normalized residuals, that is elements of the vector $\mathbf{W}_d(\mathbf{d}^{\text{prd}} - \mathbf{d}^{\text{obs}})$, whose absolute values are less than and greater than the constant c are 24.2 and 6.8, respectively. These two values, and the value of the target misfit, were obtained by generating 10^5 realizations of a set of 31 normally distributed random variables and calculating the M-measure of misfit for each realization considering the random variables to be the normalized residuals. However, just as for the previous inversions of the perturbed data, the model that gave the target misfit contained an unacceptable amount of structure. The target misfit was therefore increased to 90.9 (three times that for the outlier-free data). The resulting model is shown in Fig. 6(b), and the corresponding predicted data in Fig. 6(a). The number of data with normalized residuals lower than and greater than $c = 1.5$ was 28 and 3 respectively. The IRLS loop required 49 iterations to converge.

3.5 Summary for the example linear inverse problem

It can be seen from the inversions of the synthetic magnetic stripe data that robust measures of misfit enable results to be obtained that are significantly better when outliers are present than those obtained using the sum-of-squares measure. The examples also illustrate the ability to construct piecewise-constant, or 'blocky', models if an l_1 measure of model structure is minimized.

4 GENERAL MEASURES AND THE NON-LINEAR INVERSE PROBLEM

The standard approach for solving the non-linear inverse problem is to use an iterative procedure in which the problem is linearized at each iteration. This linearized problem can then be solved using any of the plethora of methods available for linear inverse problems. It is then mostly assumed that such an iterative procedure will converge to the solution of the non-linear problem. Building on the discussion of the linear problem in Section 3, we now describe how general measures of misfit and model structure can be incorporated into an iterative, linearized solution to the non-linear inverse problem.

As for the linear problem, we solve the non-linear inverse problem by finding the model that minimizes an objective function (*cf.* eq. 9):

$$\mathbf{\Phi} = \alpha_s \phi_s(\mathbf{W}_s(\mathbf{m}^n - \mathbf{m}^{\text{ref}})) + \alpha_x \phi_x(\mathbf{W}_x \mathbf{m}^n) + \beta[\phi_d(\mathbf{W}_d(\mathbf{d}[\mathbf{m}^n] - \mathbf{d}^{\text{obs}})) - {}^n\phi_d^{\text{tar}}]. \quad (31)$$

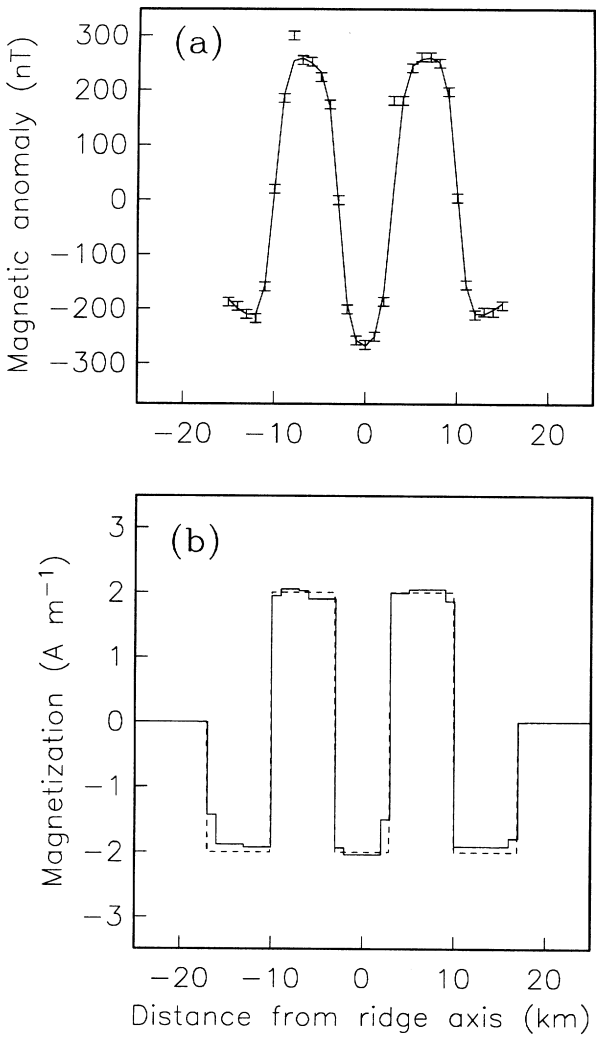

Figure 6. The results of inverting the perturbed magnetic stripe data, shown by the error bars in (a), by minimizing Ekblom's measure with $p = 1$ for ϕ_s and ϕ_x, and an M-measure for the data misfit. The solid line in (b) represents the model produced by the inversion, and the solid line in (a) represents the corresponding predicted data. The dashed line in (b) indicates the model from which the synthetic data were generated.

We recognize that an iterative process is now needed. We therefore require that the model resulting from the nth iteration, that is $\mathbf{m}^n$, will minimize this objective function. Note that the target misfit, ${}^n\phi_d^{\text{tar}}$, depends on the iteration.

The data-misfit term in eq. (31) involves the predicted data for the model that will result from the iteration $\mathbf{d}[\mathbf{m}^n]$. There is no longer a linear relationship, such as the one in eq. (8), between these data and the model. However, these data can be approximated with the first two terms in a Taylor series expansion about the model obtained from the previous iteration:

$$\mathbf{d}[\mathbf{m}^n] \approx \mathbf{d}[\mathbf{m}^{n-1}] + \mathbf{J}[\mathbf{m}^{n-1}](\mathbf{m}^n - \mathbf{m}^{n-1}), \tag{32}$$

where $\mathbf{J}$ is the Jacobian matrix of sensitivities:

$$J_{ij}[\mathbf{m}^n] = \frac{\partial d_i[\mathbf{m}^n]}{\partial m_j}. \tag{33}$$

The data misfit in eq. (31) can therefore be written as

$$\phi_d(\mathbf{W}_d(\mathbf{d}^{n-1} + \mathbf{J}^{n-1}\mathbf{m}^n - \mathbf{J}^{n-1}\mathbf{m}^{n-1} - \mathbf{d}^{\text{obs}})), \tag{34}$$

where $\mathbf{d}^n = \mathbf{d}[\mathbf{m}^n]$ and $\mathbf{J}^n = \mathbf{J}[\mathbf{m}^n]$.

To minimize the objective function in eq. (31) we differentiate with respect to the elements of $\mathbf{m}^n$ and equate the resulting equations to zero, exactly as discussed in Section 2 for the linear problem. The only significant difference is in the differentiation of the data misfit. Instead of eq. (22) for the linear case, differentiation of eq. (34) gives

$$\frac{\partial \phi_d}{\partial \mathbf{m}} = \mathbf{J}^{\mathrm{T}}\mathbf{W}_d^{\mathrm{T}}\mathbf{R}_d\mathbf{W}_d(\mathbf{J}\mathbf{m}^n - \mathbf{J}\mathbf{m}^{n-1} + \mathbf{d}^{n-1} - \mathbf{d}^{\text{obs}}), \tag{35}$$

where $\mathbf{J} = \mathbf{J}^{n-1}$. By analogy with eq. (23), the system of equations to be solved is therefore

$$[\alpha_s\mathbf{W}_s^{\mathrm{T}}\mathbf{R}_s\mathbf{W}_s + \alpha_x\mathbf{W}_x^{\mathrm{T}}\mathbf{R}_x\mathbf{W}_x + \beta\mathbf{J}^{\mathrm{T}}\mathbf{W}_d^{\mathrm{T}}\mathbf{R}_d\mathbf{W}_d\mathbf{J}]\mathbf{m}^n = \beta\mathbf{J}^{\mathrm{T}}\mathbf{W}_d^{\mathrm{T}}\mathbf{R}_d\mathbf{W}_d(\mathbf{d}^{\text{obs}} - \mathbf{d}^{n-1} + \mathbf{J}\mathbf{m}^{n-1}) + \alpha_s\mathbf{W}_s^{\mathrm{T}}\mathbf{R}_s\mathbf{W}_s\mathbf{m}^{\text{ref}}. \tag{36}$$

We solve this system of equations using the IRLS procedure described in Section 2. There are therefore two iterative procedures involved in the solution of the non-linear problem: the one associated with the linearization of the non-linear problem, and the one associated with the iteratively reweighted least-squares procedure. There are a number of ways in which the two can be combined. We discuss two possibilities in the next section when we apply the preceeding ideas to the 1-D inversion of synthetic time-domain electromagnetic data.

5 AN EXAMPLE NON-LINEAR INVERSE PROBLEM: SYNTHETIC DATA

We now illustrate the solution of the non-linear inverse problem using general measures of misfit and model structure, and the IRLS procedure described in Section 2. Our chosen example is the 1-D inversion of time-domain electromagnetic (TEM) data. The TEM method is used extensively for mineral exploration, and is seeing ever-increasing use in environmental and engineering surveys. It has also been used for hydrocarbon exploration. For an overview of the TEM method and its uses, see Nabighian & Macnae (1991). We have chosen this as our example because of our previous experience with this inverse problem: see Farquharson & Oldenburg (1993), in which we present an inversion procedure for obtaining, from TEM data, horizontally layered, minimum-structure models of the Earth's electrical conductivity. We used the traditional sum-of-squares measure for both the data misfit and the measure of model structure. For details of the forward modelling of TEM data over a horizontally layered model, and the computation of the Jacobian matrix of sensitivities for such a model, we refer the reader to this paper.

In this section we consider the inversion of synthetic TEM data. Synthetic data were computed for the three-layered model shown in Fig. 8(b). The data were computed for a 5×5 m horizontal transmitter loop on the surface of the model, and a horizontal receiver loop, also on the surface of the model, of unit area at a distance of 20 m from the centre of the transmitter loop (see Fig. 7). The transmitter current waveform was a pure step-off with a pre-shut-off current of 12 A. The range of delay times over which the voltage in the receiver loop was computed can be seen from Fig. 8(a). Gaussian noise was added to the

Figure 7. The transmitter–receiver geometry for the synthetic TEM data used as the first non-linear example. The origin is at the centre of the transmitter loop, and the coordinates are in metres. **T** and **R** indicate the transmitter and receiver, respectively, both of which lie on the surface of the model, that is in the plane $z = 0$.

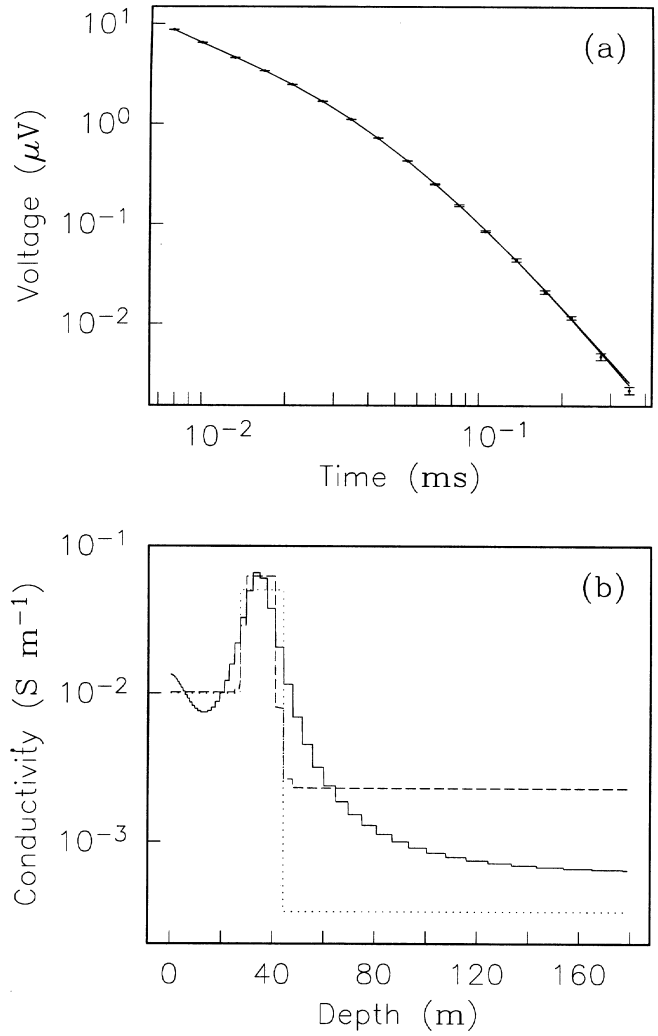

Figure 8. (a) The synthetic TEM data (shown by the error bars) contaminated only with Gaussian noise that are used for the first examples of non-linear inversion using general measures of misfit and model structure. The two almost coincident solid lines indicate the predicted data for the model produced using the traditional l_2 measures of misfit and model structure, and for the model produced using l_1 measures. (b) The model (solid line) produced by the inversion of the synthetic data shown in panel (a) using l_2 measures, and the model (dashed line) produced using l_1 measures. The dotted line indicates the model from which the synthetic data were computed.

voltages. The standard deviation of this noise increased in steps from 0.5 to 8 per cent over the range of delay times. The resulting synthetic data are shown in Fig. 8(a). We now present the results of inverting these data using the linearized, iterative procedure summarized in the previous section, and using non-l_2 measures of misfit and model structure. The same model parametrization was used for all inversions: the model was made up of 50 layers whose thicknesses increased exponentially with depth. The parameter to be found was the natural logarithm of the electrical conductivity in each layer.

We restrict our consideration here to only one measure: the l_1 measure, or more precisely, Ekblom's measure given by eqs (1) and (5) with $p = 1$ and ε small. We chose this measure because it enables us to produce piecewise-constant models, and because it is a robust measure of misfit. It also allows comparison with the alternative approach for l_1 measures of using linear programming. In the following inversions of the synthetic TEM data we also present two possible ways in which the two iterative procedures, the one associated with the linearization of the non-linear problem, and the one associated with the IRLS algorithm, can be combined.

5.1 Gaussian noise, l_2 inversion and a line search over β

Consider first the inversion of the synthetic TEM data using the traditional sum-of-squares measure for ϕ_s, ϕ_z ('z' since our model is now a function of depth) and ϕ_d. A purely minimum-structure model was desired, so $\alpha_s = 0$ and $\alpha_z = 1$, and $\mathbf{W}_z$ was the first-order finite-difference operator. The data-weighting matrix, $\mathbf{W}_d$, was a diagonal matrix, the elements of which were the reciprocals of the standard deviations of the noise in the data. The first iteration in the IRLS procedure was used to solve eq. (36), which is a linear system for this example.

A line search was performed at each iteration associated with the linearization of the inverse problem to determine the appropriate value of β. Since the emphasis of these examples is illustration, we used the comprehensive but inefficient process of solving eq. (36) for 100 values of β spread over a suitably wide range, and computing the misfit associated with each β. A bisection search was then performed once the two values of β bracketing the desired value had been determined. The desired value of β was the smallest value giving a misfit equal to the target misfit for that iteration, or if the misfit could not be reduced that far, then the chosen value was the one that gave the largest decrease in misfit.

The target misfit at each iteration was taken as ${}^n\phi_d^{\text{tar}} = \max(0.66 \times {}^{n-1}\phi_d^{\text{tar}}, \phi_d^{\text{tar}})$, where ϕ_d^{tar} was the ultimate target misfit for the non-linear inversion. Since the data misfit for this example is a sum-of-squares measure, ϕ_d^{tar} was chosen to be equal to the number of observations (= 17). The final model is shown in Fig. 8(b), and the corresponding predicted data in Fig. 8(a). The values of β chosen at each iteration are shown in Fig. 9(a), along with the corresponding values of ϕ_z and ϕ_d. Note the characteristic decrease in ϕ_d at each iteration and increase in ϕ_z until the target misfit of $\phi_d = 17$ is reached after 21 iterations. The starting model was a homogeneous half-space of conductivity 0.01 S m^{-1}. Also note the characteristic smooth appearance of the final model for this minimum-structure inversion using an l_2 measure of model structure.

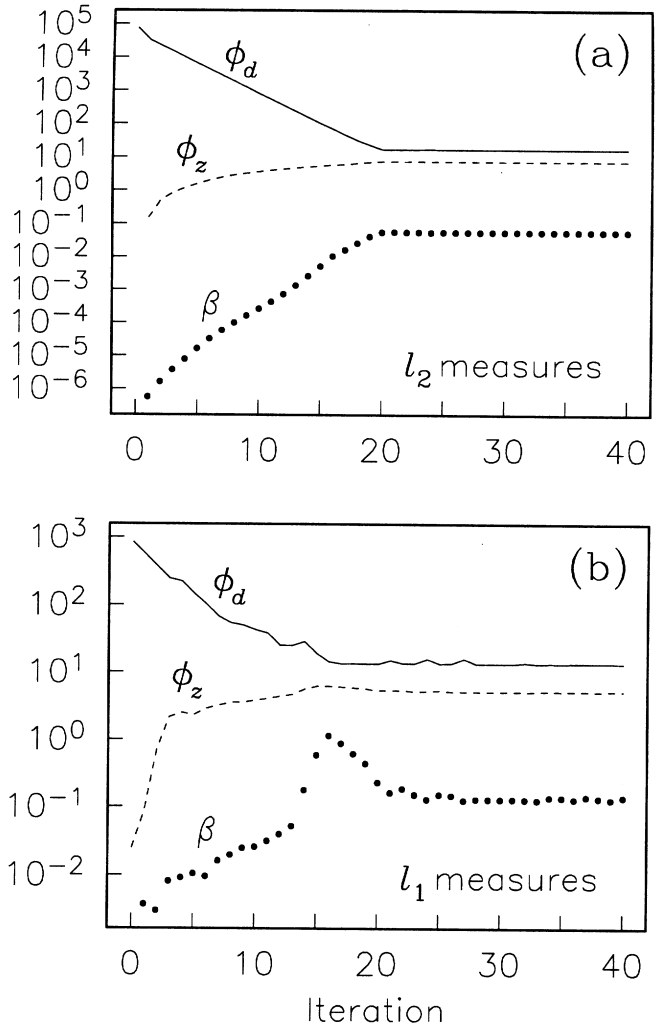

Figure 9. (a) The variation with iteration of the trade-off parameter β, the measure of model structure ϕ_z, and the data misfit ϕ_d for the inversion, using l_2 measures for both ϕ_z and ϕ_d, of the synthetic TEM data shown in Fig. 8. (b) The corresponding variation during the inversion using l_1 measures for both ϕ_z and ϕ_d, and the IRLS procedure. Note the different scales on the vertical axes.

5.2 Gaussian noise, l_1 inversion and a line search over β

The previous inversion was repeated using the same parameters and weighting matrices, but using Ekblom's measure, given by eqs (1) and (5), for both ϕ_z and ϕ_d. The value of p was 1, and ε was equal to 5×10^{-4} in ϕ_z and 10^{-1} in ϕ_d. The IRLS solution to eq. (36) was needed, typically requiring 20 to 30 iterations. As in the previous section, a line search was performed to find the most appropriate value of β for each iteration associated with the linearization of the non-linear problem. The target misfit at each iteration was taken to be ${}^n\phi_d^{\text{tar}} = \max(0.66 \times {}^{n-1}\phi_d^{\text{tar}}, \phi_d^{\text{tar}})$. Following Parker & McNutt (1980) and the linear l_1 example described in Section 3.3, the ultimate target misfit was taken to be $\phi_d^{\text{tar}} = \sqrt{2/\pi} \times 17 = 13.6$.

The final model for this inversion is indicated by the dashed line in Fig. 8(b). The predicted data are shown by the solid line in Fig. 8(a). The chosen values of β at each iteration, along with the corresponding values of ϕ_z and ϕ_d, are shown in Fig. 9(b). The variation of these three quantities is not as smooth and regular as during the inversion using l_2 measures (see Fig. 9a). In particular, the desired decrease in misfit at each iteration was not always achieved, and there is considerable variation in β for the iterations immediately before and after the ultimate target misfit of 13.6 is achieved.

The inversion was repeated using true l_1 norms for ϕ_z and ϕ_d, and using linear programming to solve the linearized problem. Instead of constructing the system given in eq. (36), the following overdetermined system of equations was formed:

$$\begin{pmatrix} \beta \mathbf{W}_d \mathbf{J}^{n-1} \\ \alpha_s \mathbf{W}_s \\ \alpha_z \mathbf{W}_z \end{pmatrix} \mathbf{m}^n = \begin{pmatrix} \beta \mathbf{W}_d (\mathbf{d}^{\text{obs}} - \mathbf{d}^{n-1} + \mathbf{J}^{n-1} \mathbf{m}^{n-1}) \\ \alpha_s \mathbf{W}_s \mathbf{m}^{\text{ref}} \\ 0 \end{pmatrix}. \quad (37)$$

The model that minimized the l_1 norm of the difference between the two sides of this equation was found using the routine of

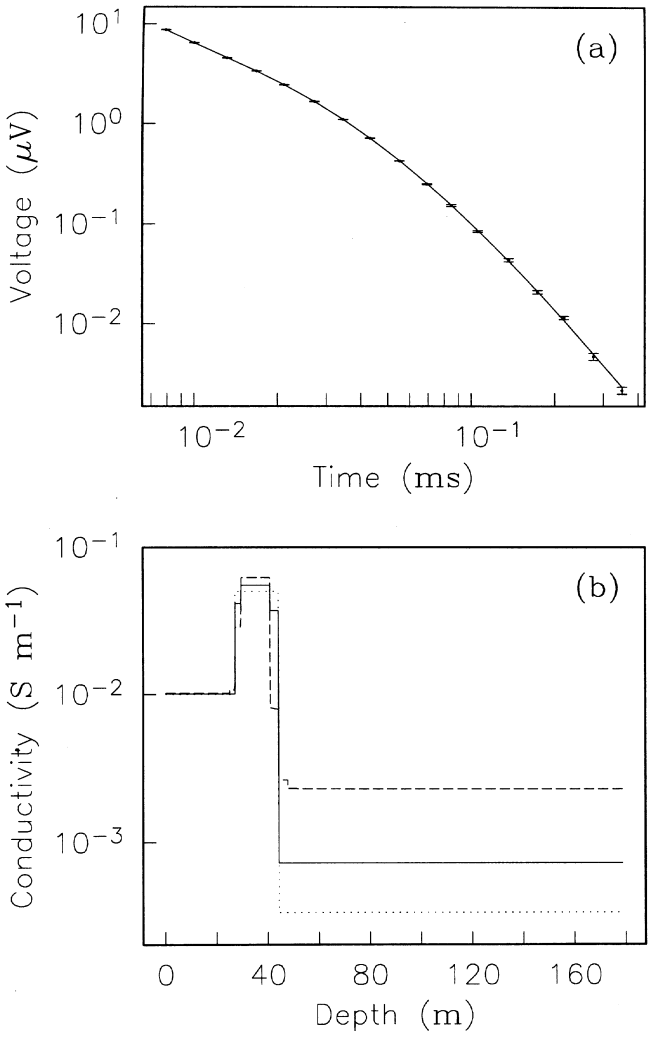

Figure 10. The final model (solid line in b) produced by the inversion of the synthetic TEM data (shown by the error bars in a) using l_1 norms for ϕ_z and ϕ_d, and using linear programming to solve the system of equations at each iteration in the linearized inversion procedure. The predicted data for this model are indicated by the solid line in (a). The dashed line in (b) indicates the model obtained from the l_1 inversion using Ekblom's measure and the IRLS procedure, already shown in Fig. 8(b). The dotted line in (b) indicates the model from which the synthetic data were generated.

Abdelmalek (1980a,b), available from the TOMS collection of software at NETLIB (http://www.netlib.org/). A line search was performed, just as for the other inversions of the synthetic TEM data so far discussed. The final model is shown in Fig. 10(b), and the predicted data in Fig. 10(a). The chosen value of β at each iteration and the corresponding variation in ϕ_z and ϕ_d are shown in Fig. 11. Note that the variation in these three quantities is not as smooth and uncomplicated as the behaviour for the l_2 inversion (see Fig. 9a). Also note that the value of β, and to a lesser extent ϕ_z and ϕ_d, never converge to a stable value but continue to oscillate as the number of iterations increases beyond the low teens. It is also intriguing that the values of β that are reached are over an order of magnitude greater than in the supposedly equivalent inversion using Ekblom's measure with $p = 1$. There are also differences between the final models for the two inversions, most notably the conductivity of the basement (see Figs 10b and 8b).

The differences between the inversion using Ekblom's measure and the IRLS procedure, and the inversion using the true l_1 norm and linear programming, develop through the iterative solution to the non-linear problem, and the choice of β at each iteration. Fig. 12 shows the value of the data misfit as a function of β for the final iteration in the three inversions of the synthetic TEM data so far discussed. The curve for the l_2 inversion, shown in panel (a), is smoothly varying and simple, as one would hope. By contrast, the 'curve' for the l_1 inversion using linear programming, shown in panel (c), is predominantly piecewise-constant: ranges of β for which the misfit is not changing separated by discontinuities in the value of misfit. In fact, it is generally the case that the target misfit for a particular iteration cannot be achieved because it falls between the values on either side of a discontinuity. The search procedure described in Section 5.1 was modified for this situation to choose simply the value of β that gave the misfit closest to the target value. The curve shown in panel (b) is that for Ekblom's measure with $p = 1$ and a small but non-zero value for ε, and for solution via the IRLS procedure. This curve is significantly smoother than the one for the l_1 norm inversion using linear programming. It is the difficulty in carrying out a line search along a curve such as the one shown in Fig. 12(c) compared to the one in Fig. 12(b), and the modification of the line search procedure required to address this issue, that led to different values of β being chosen during the two corresponding inversions, and different final results being attained.

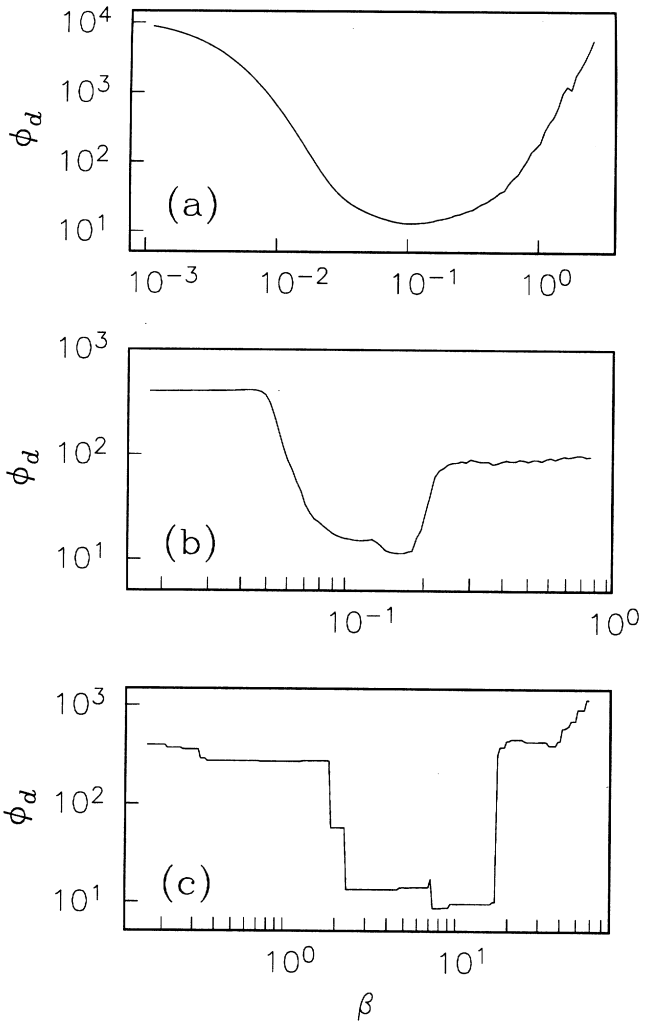

Figure 12. The data misfit, ϕ_d, as a function of the trade-off parameter β for the final iteration in (a) the l_2 inversion, (b) the l_1 inversion using Ekblom's measure and the IRLS procedure, and (c) the l_1 inversion using the true l_1 norm and linear programming. Note the different scales in all three plots.

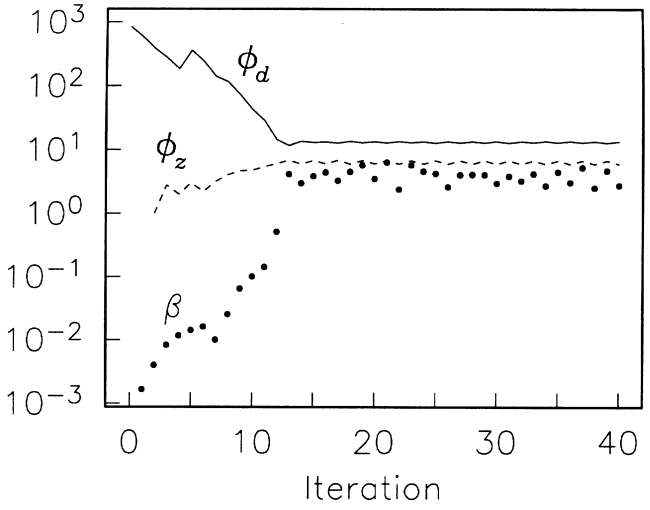

Figure 11. The variation of β, ϕ_z and ϕ_d during the l_1 inversion of the synthetic TEM data for which the solution at each iteration was found using linear programming.

5.3 Gaussian noise, l_1 inversion and a schedule for β

The inversion in the previous section using Ekblom's measure involved two distinct iterative loops: the familiar loop associated with the linearization of a non-linear problem, and the loop within the IRLS procedure used to obtain the solution for each linearized problem. A line search over β was also incorporated into the solution of each linearized problem in order to find the β that gave a misfit equal to some target value. In this line search, the IRLS solution to the system of equations has to be found for each value of β, and a forward

modelling performed to determine the misfit. These multiple applications of the forward modelling and IRLS procedures in the solution of each linearized problem make this inversion procedure relatively time-consuming.

We now describe a variation of this inversion procedure which in itself is significantly quicker. In this variation a schedule of values of β is prescribed (Oldenburg & Ellis 1993; Ellis *et al.* 1993). This means that only a single system of equations has to be solved and a single forward modelling carried out at each iteration. However, a trial-and-error process involving re-running the inversion is required to determine an appropriate schedule.

This inversion procedure can be made even quicker by absorbing the iterations required by the IRLS procedure into the iterations associated with the linearization of the non-linear problem. At each iteration associated with the linearization of the problem, and for the single prescribed value of β, the system of equations given in eq. (36) is constructed and solved once using the weighting matrices $\mathbf{R}_s$, $\mathbf{R}_z$ and $\mathbf{R}_d$ computed for the model obtained from the previous iteration. The two iterative procedures, the one required to overcome the non-linearity of the inverse problem, and the one required to overcome the non-linearity introduced by non-l_2 measures, therefore proceed in unison. This is similar to the approach used by LaBrecque *et al.* (1997).

The inversion described in Section 5.2 using Ekblom's measure was repeated using exactly the same parameters and weighting matrices, but using the procedure described in the previous two paragraphs. The resulting model and predicted data are shown in Fig. 13. The prescribed value of β for each iteration, and the resulting values of ϕ_z and ϕ_d, are shown in Fig. 14. The final value of ϕ_d was 13.6, just as for the related inversion in Section 5.2. Comparison of Figs 13 and 8 shows that essentially the same model has been produced by the two different inversions. The 50 iterations shown in Fig. 14 required significantly less time that the 40 shown in Fig. 9. However, several runs were required to determine the schedule of values of β that gave the desired value for the final misfit.

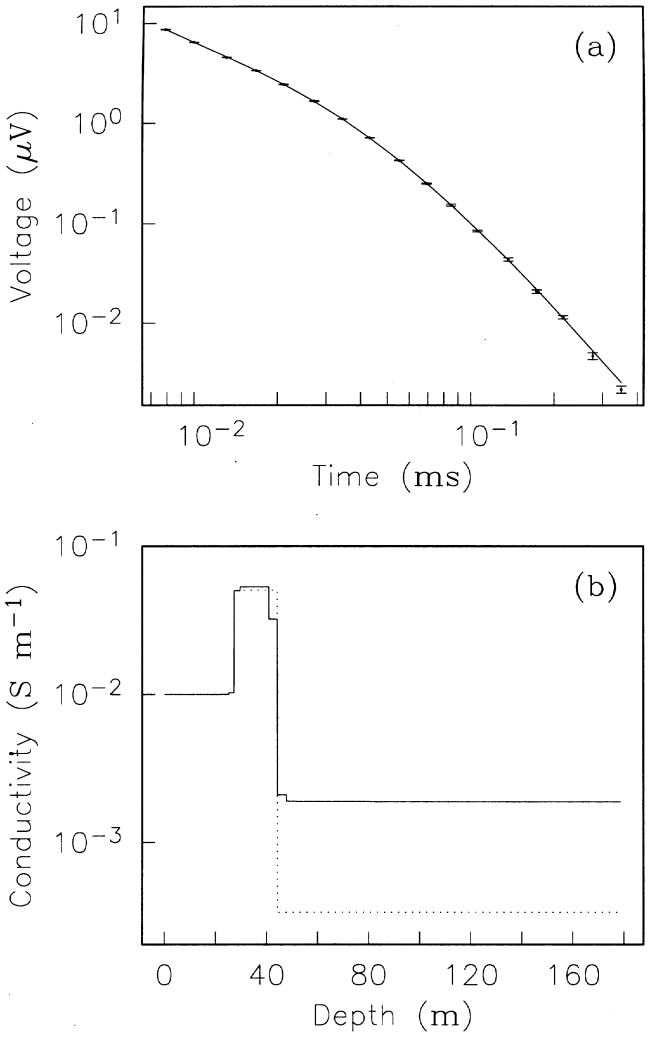

Figure 13. The final model produced by the l_1 inversion of the synthetic TEM data using Ekblom's measure, and for which a prescribed schedule of values of β was used, is shown as the solid line in (b). The predicted data for this model are shown by the solid line in (a). The error bars in (a) represent the synthetic data, and the dotted line in (b) represents the model from which the synthetic data were computed.

5.4 Outliers, and l_2 and l_1 inversions

To conclude the example inversions of the synthetic TEM data set, we perturbed two of the data to create outliers. The value of the voltage at the third measurement time was decreased by ten per cent, and the values of the voltage and measurement uncertainty at the fourth-last measurement time were set equal to the voltage and uncertainty at the preceeding time. These two perturbed data are indicated by the arrows in Fig. 15(a).

The perturbed data were inverted using the l_2 inversion described in Section 5.1, and the procedure using Ekblom's measure with $p = 1$ described in Section 5.2. Both of these procedures involved a line search over β. Not surprisingly, neither inversion could fit the outlier-contaminated data to the same values of ϕ_d (17 and 13.6 for the l_2 and l_1 inversions) as when only Gaussian noise contaminated the data. Final models were therefore chosen on the basis of giving the smallest

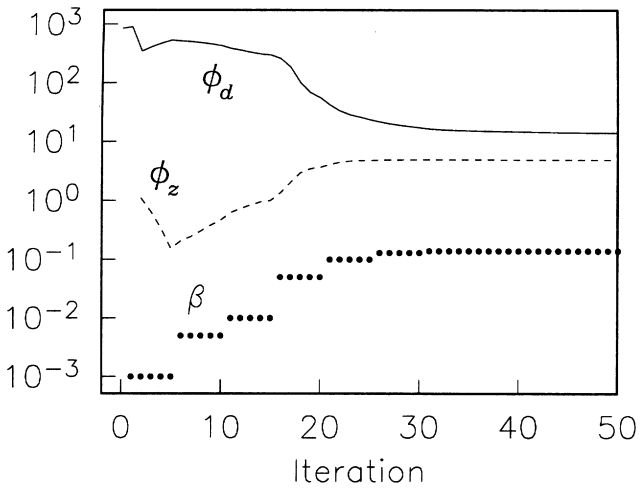

Figure 14. The prescribed values of β for the l_1 inversion of the synthetic TEM data using Ekblom's measure, the results of which are shown in Fig. 13. The corresponding behaviour of the data misfit, ϕ_d, and the measure of model structure, ϕ_z, are also shown.

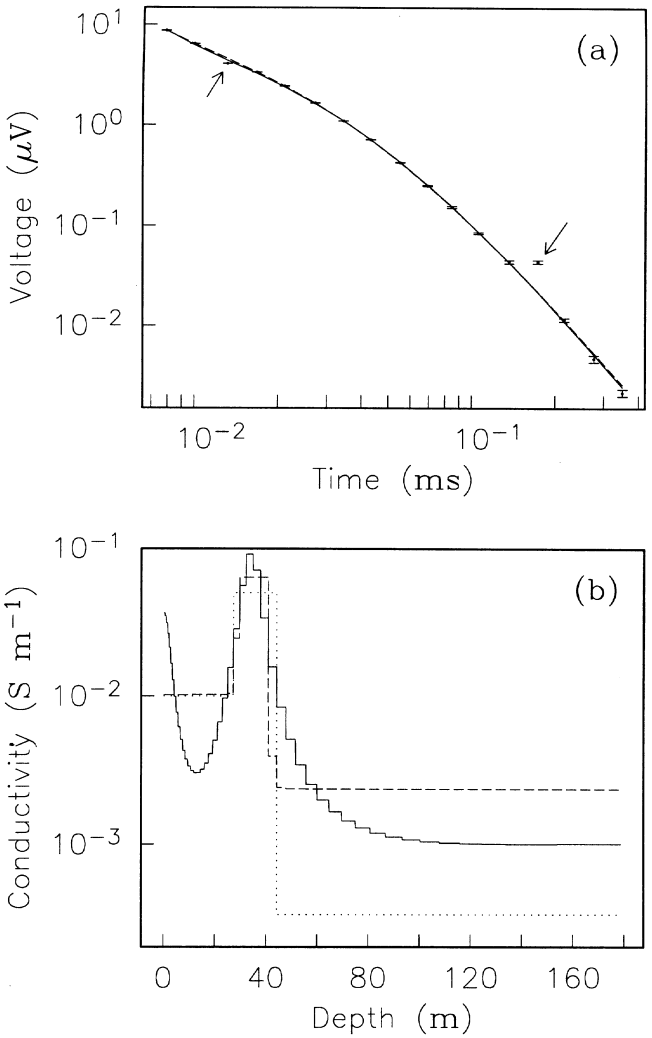

Figure 15. The perturbed synthetic TEM data are shown by the error bars in (a). The two data that were perturbed are indicated by the arrows. The solid and dashed lines in (a) are the predicted data for the final models produced by the l_2 and l_1 inversions respectively. The corresponding models are shown in (b): the solid line is the model for the l_2 inversion, and the dashed line is the model for the l_1 inversion. The dotted line in (b) is the model from which the synthetic data were generated.

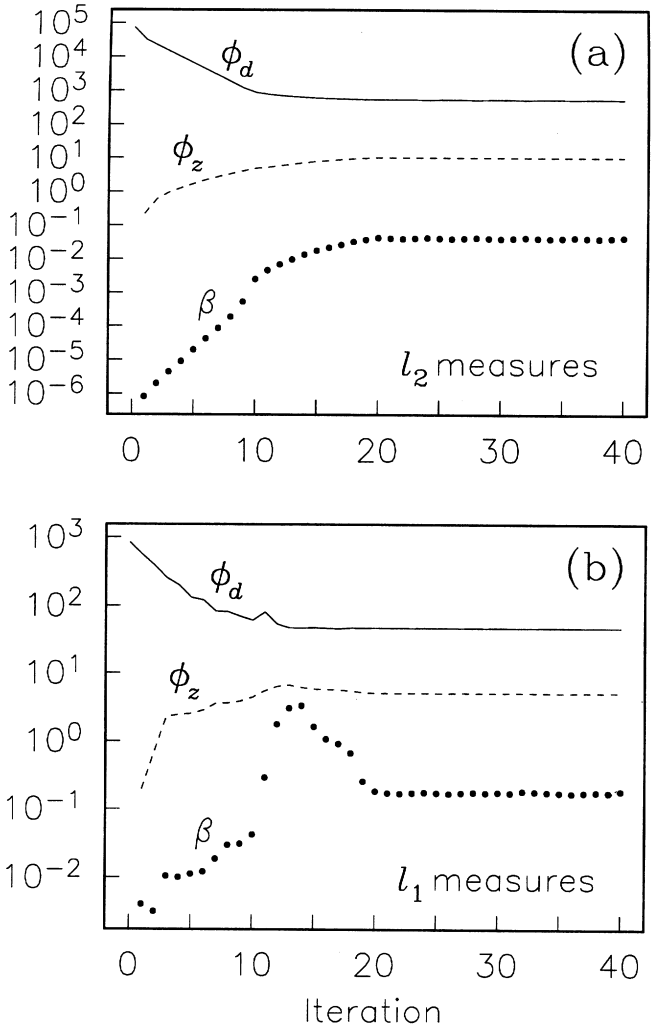

Figure 16. (a) The variation in the chosen value of β, and in the values of ϕ_z and ϕ_d during the l_2 inversion of the synthetic TEM data containing outliers. The results of this inversion are shown in Fig. 15. (b) The variation of β, ϕ_z and ϕ_d during the l_1 inversion of the perturbed data set. The results of this inversion are also shown in Fig. 15.

misfit without possessing an unreasonable amount of structure. The misfits for the chosen models were 544 for the l_2 inversion, and 47 for the l_1 inversion. The models are shown in Fig. 15, along with the corresponding predicted data. The behaviour of β, ϕ_z and ϕ_d during the two inversions is shown in Fig. 16.

There are two important points that arise from the comparison of the models in Figs 15 and 8. First, there is very little difference between the models for the l_1 inversions of the outlier-free data (Fig. 8b) and the outlier-contaminated data (Fig. 15b). This is because the l_1 measure of misfit used in these inversions is a robust measure, and, as such, is not greatly affected by the outliers. Second, and in obvious contrast to the first point, there are significant differences between the models for the l_2 inversions of the two data sets. This is because the l_2 measure is not a robust measure.

5.5 Summary for the synthetic TEM examples

We have given illustrations, using synthetic TEM data, of the solution of the non-linear inverse problem using non-l_2 measures of misfit and model structure. We have presented two possible methods for combining the iterative loops associated with the linearization of the non-linear problem and the IRLS procedure needed to accommodate non-l_2 measures. The first method involves a line search over β at each iteration in the former of the two loops, and a complete solution of the system of equations for each β using the IRLS procedure. The second method uses a prescribed schedule of the values of β, and absorbs the IRLS iterations into the iterations associated with the linearization of the non-linear problem. The second of these methods is by far the quicker if an appropriate schedule is known, or can be readily established. However, the first method is the more robust, and can be expected to attain the solution to a particular problem at the first attempt.

6 AN EXAMPLE NON-LINEAR INVERSE PROBLEM: FIELD DATA

As a final illustration of the solution of the non-linear inverse problem using general measures of misfit and model structure we present the results of an inversion of field TEM data supplied by CRA Exploration Party. The data were collected using a 500×500 m transmitter loop with the receiver at its centre. Two overlapping sweeps were used to record the data, one from 0.68 to 27.8 ms, and one from 4.37 to 446 ms. Linear ramp turn-offs were used, of duration 0.3 ms for the earlier sweep, and 0.5 ms for the later sweep. The data are shown by the dots in Fig. 17(a). The uncertainties in the observations were estimated to be of 2 per cent for all measurement times. These estimates were derived from the repeatability of the measurements. The error bars are smaller than the dots in Fig. 17(a), so are not shown.

The data were inverted using the procedure involving the line search over β described in Sections 5.1 and 5.2. The model was discretized into 100 layers of exponentially increasing thicknesses. A blocky, minimum-structure model was desired, so Ekblom's measure with $p = 1$ and ε small ($= 5 \times 10^{-4}$) was used as a measure of model structure, and the first-order finite-difference operator used for $\mathbf{W}_z$. No proximity to a reference model was required, so $\alpha_s = 0$ and $\alpha_z = 1$. The data shown in Fig. 17(a) are of high quality with very small uncertainties. The noise that is present was considered to be Gaussian. The sum-of-squares measure was therefore used as the measure of misfit.

The model produced by the inversion is shown in Fig. 17(b), and the corresponding predicted data are shown in Fig. 17(a). Given the number of data, the initial value of the ultimate target misfit was 69. However, the minimum misfit that could be attained was 149. The model shown in Fig. 17(b) was produced by the inversion with a target misfit of 155. The behaviour of β, ϕ_d and ϕ_z during the inversion is shown in Fig. 18.

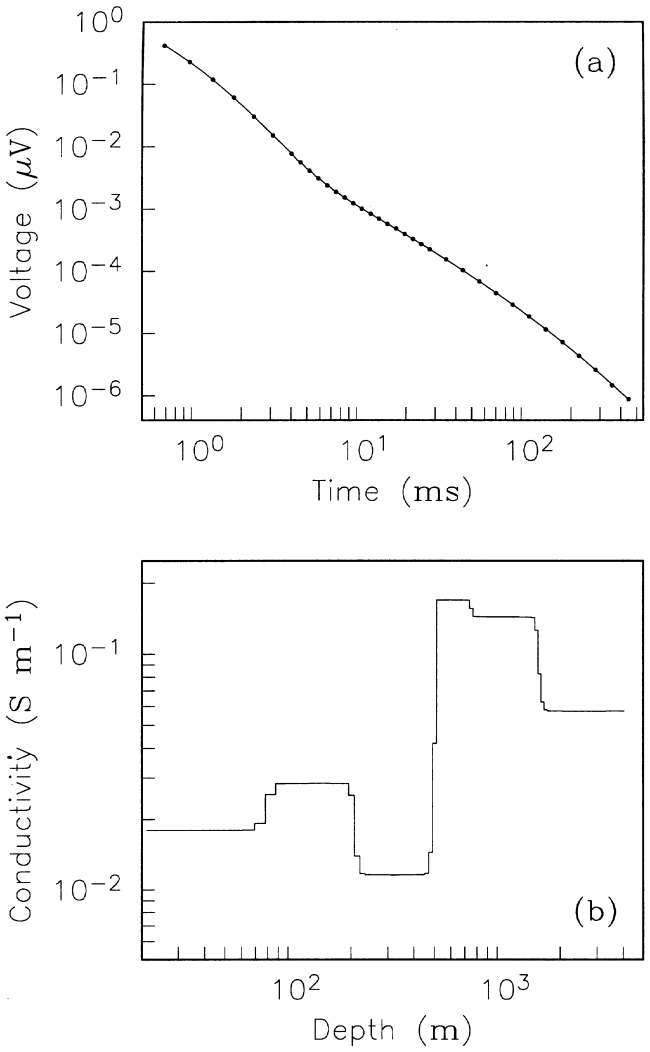

Figure 17. The field TEM data, shown by the dots in (a), the inversion of which is discussed in Section 6. The final model, shown by the solid line in (b), produced by this inversion. Note that both axes of (b) are logarithmic. The predicted data for the final model are shown by the solid line in (a).

7 CONCLUSIONS

The traditional linearized, iterative approach to solving the non-linear inverse problem can be readily generalized to incorporate non-l_2 measures of data misfit and model structure. This can be done with the inclusion of an iteratively reweighted least-squares procedure. The successful implementation of this approach has been illustrated for a number of possible measures of misfit and model structure. The illustrations also serve to demonstrate the benefits of the particular measures that we chose to use, namely the robustness of the l_1 measure and the M-measure when used as measures of data misfit, and the ability to construct piecewise-constant, or 'blocky', models that an l_1 measure of model structure provides.

ACKNOWLEDGMENTS

The work for this paper was supported by an NSERC IOR grant and the Consortium for the Joint and Cooperative Inversion of Geophysical and Geological Data (the 'JACI'

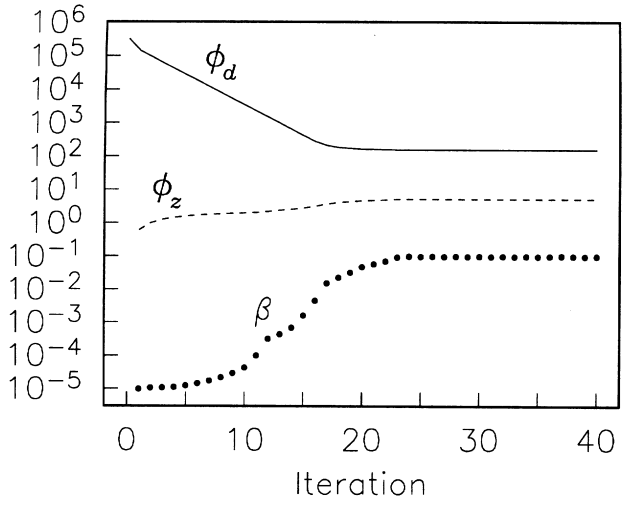

Figure 18. The variation in β, ϕ_d and ϕ_z during the inversion of the field TEM data shown in Fig. 17(a).

Consortium). The following are participants in this consortium: Placer Dome, BHP Minerals, Noranda Exploration, Cominco Exploration, Falconbridge, INCO Exploration and Technical Services, Hudson Bay Exploration and Development, Kennecott Exploration Company, Newmont Gold Company, WMC Exploration, and CRA Exploration Party. The authors are grateful for their participation. We are further indebted to CRA Exploration Party for supplying the field TEM data. We would also like to thank Drs Yaoguo Li, Mauricio Sacchi and Eldad Haber for many ever-useful discussions.

REFERENCES

Abdelmalek, N.N., 1980a. L_1 solution of overdetermined systems of linear equations, *ACM Trans. math. Software,* **6,** 220–227.

Abdelmalek, N.N., 1980b. ALGORITHM 551: a Fortran subroutine for the L_1 solution of overdetermined systems of linear equations, *ACM Trans. math. Software,* **6,** 228–230.

Alliney, S. & Ruzinsky, S.A., 1994. An algorithm for the minimization of mixed l_1 and l_2 norms with application to Bayesian estimation, *IEEE Trans. Signal Processing,* **42,** 618–627.

Amundsen, L., 1991. Comparison of the least-squares criterion and the Cauchy criterion in frequency-wavenumber inversion, *Geophysics,* **56,** 2027–2035.

Anderson, E. *et al.,* 1995. LAPACK Users' Guide, 2nd edn, SIAM, Philadelphia, PA.

Beltrão, J.F., Silva, J.B.C. & Costa, J.C., 1991. Robust polynomial fitting method for regional gravity estimation, *Geophysics,* **56,** 80–89.

Chave, A.D. & Thomson, D.J., 1989. Some comments on magnetotelluric response function estimation, *J. geophys. Res.,* **94,** 14 215–14 225.

Chave, A.D., Thomson, D.J. & Ander, M.E., 1987. On the robust estimation of power spectra, coherences, and transfer functions, *J. geophys. Res.,* **92,** 633–648.

Claerbout, J.F. & Muir, F., 1973. Robust modeling with erratic data, *Geophysics,* **38,** 826–844.

Constable, S.C., Parker, R.L. & Constable, C.G., 1987. Occam's inversion: a practical algorithm for generating smooth models from electromagnetic sounding data, *Geophysics,* **52,** 289–300.

Crase, E., Pica, A., Noble, M., McDonald, J. & Tarantola, A., 1990. Robust elastic non-linear waveform inversion: application to real data, *Geophysics,* **55,** 527–538.

deGroot-Hedlin, C. & Constable, S., 1990. Occam's inversion to generate smooth, two-dimensional models for magnetotelluric data, *Geophysics,* **55,** 1613–1624.

Dosso, S.E., 1990. Inversion and appraisal for the one-dimensional magnetotellurics problem, *PhD thesis,* University of British Columbia, Vancouver, Canada.

Egbert, G.D. & Booker, J.R., 1986. Robust estimation of geomagnetic transfer functions, *Geophys. J. R. astr. Soc.,* **87,** 173–194.

Ekblom, H., 1973. Calculation of linear best L_p-approximations, *BIT,* **13,** 292–300.

Ekblom, H., 1987. The L_1-estimate as limiting case of an L_p- or Huber-estimate, in *Statistical Data Analysis Based on the* L_1*-Norm and Related Methods,* pp. 109–116, ed. Dodge, Y., Elsevier, Amsterdam.

Ellis, R.G., Farquharson, C.G. & Oldenburg, D.W., 1993. Approximate inverse mapping inversion of the COPROD2 Data, *J. Geomag. Geoelectr.,* **45,** 1001–1012.

Farquharson, C.G. & Oldenburg, D.W., 1993. Inversion of time-domain electromagnetic data for a horizontally layered Earth, *Geophys. J. Int.,* **114,** 433–442.

Gersztenkorn, A., Bednar, J.B. & Lines, L.R., 1986. Robust iterative inversion for the one-dimensional acoustic wave equation, *Geophysics,* **51,** 357–368.

Holland, P.W. & Welsch, R.E., 1977. Robust regression using iteratively reweighted least-squares, *Communications in Statistics—Theory and Methods,* **A6,** 813–827.

Huber, P.J., 1964. Robust estimation of a location parameter, *Ann. math. Stat.,* **35,** 73–101.

LaBrecque, D., Morelli, G., Daily, B., Ramirez, A. & Lundegard, P., 1997. Occam's inversion of 3D ERT data, in *Three-Dimensional Electromagnetics,* eds Oristaglio, M.L. & Spies, B.R., Soc. Expl. Geophysicists, Tulsa, OK, in press.

Larsen, J.C., Mackie, R.L., Manzella, A., Fiordelisi, A. & Rieven, S., 1996. Robust smooth magnetotelluric transfer functions, *Geophys. J. Int.,* **124,** 801–819.

Levy, S. & Fullagar, P.K., 1981. Reconstruction of a sparse spike train from a portion of its spectrum and application to high-resolution deconvolution, *Geophysics,* **46,** 1235–1243.

Nabighian, M.N. & Macnae, J.C., 1991. Time-domain electromagnetic prospecting methods, in *Electromagnetic Methods in Applied Geophysics,* Vol. 2, pp. 427–479, ed. Nabighian, M.N., Soc. Expl. Geophysicists, Tulsa, OK.

Oldenburg, D.W., 1984a. The inversion of lead isotope data, *Geophys. J. R. astr. Soc.,* **78,** 139–158.

Oldenburg, D.W., 1984b. An introduction to linear inverse theory, *IEEE Trans. Geosci. Remote Sensing,* **GE-22,** 665–674.

Oldenburg, D.W. & Ellis, R.G., 1993. Efficient inversion of magnetotelluric data in two dimensions, *Phys. Earth planet. Inter.,* **81,** 177–200.

O'Leary, D.P., 1990. Robust regression computation using iteratively reweighted least squares, *SIAM J. Matrix Anal. Appl.,* **11,** 466–480.

Parker, R.L., 1994. *Geophysical Inverse Theory,* Princeton University Press, Princeton, NJ.

Parker, R.L. & McNutt, M.K., 1980. Statistics for the one-norm misfit measure, *J. geophys. Res.,* **85,** 4429–4430.

Sacchi, M.D. & Ulrych, T.J., 1995. Improving resolution of Radon operators using a model re-weighted least-squares procedure, *J. seis. Expl.,* **4,** 315–328.

Sacchi, M.D. & Ulrych, T.J., 1996. Estimation of the discrete Fourier transform, a linear inversion approach, *Geophysics,* **61,** 1128–1136.

Scales, J.A., Gersztenkorn, A. & Treitel, S., 1988. Fast l_p solution of large, sparse, linear systems: application to seismic travel time tomography, *J. comput. Phys.,* **75,** 314–333.

Smith, J.T. & Booker, J.R., 1988. Magnetotelluric inversion for minimum structure, *Geophysics,* **53,** 1565–1576.

Sutarno, D. & Vozoff, K., 1991. Phase-smoothed robust *M*-estimation of magnetotelluric impedance functions, *Geophysics,* **56,** 1999–2007.

Watson, G.A., 1980. *Approximation Theory and Numerical Methods,* John Wiley, Chichester.

Wolke, R. & Schwetlick, H., 1988. Iteratively reweighted least squares: algorithms, convergence analysis, and numerical comparisons, *SIAM J. Sci. Stat. Comput.,* **9,** 907–921.