Physics of the Earth and Planetary Interiors, 81 (1993) 177–200
Elsevier Science Publishers B.V., Amsterdam

Efficient inversion of magnetotelluric data in two dimensions

D.W. Oldenburg *, R.G. Ellis

UBC — Geophysical Inversion Facility, Department of Geophysics and Astronomy, University of British Columbia, Vancouver, B.C., V6T 1Z4 Canada

(Received 14 November 1992; revision accepted 15 March 1993)

ABSTRACT

Two algorithms are presented to invert magnetotelluric data over 2D conductivity structures. Both algorithms use approximate sensitivities that arise from the 1D conductivity profile beneath each station; this avoids the large computations normally required to calculate the true 2D sensitivities. The first algorithm produces 'blocky' models by minimizing the l_1 norm of the conductivity, whereas the second algorithm produces smoother models by minimizing the l_2 norm of the model. The inversion of a large matrix in the l_2 norm minimization is obviated by using a sub-space solution. Our presentation is motivated by three goals. The first is to provide details of these algorithms and to show their efficiency compared with more commonly used techniques. The second goal is to illustrate the non-uniqueness inherent in these inversions and to illuminate the importance of choice of norm that is minimized. To accomplish this we perform l_1 and l_2 norm inversions on a synthetic model. The third goal is to show the utility, the practical importance, and the limitations of inverting determinant average data. The utility is demonstrated by the similarity that is often observed when comparing the inversion result obtained by inverting transverse electric (TE) and transverse magnetic (TM) mode data jointly with that obtained from inverting determinant average data. The importance of inverting determinant average data in field experiments arises because determinant average data are insensitive to rotations of the impedance matrix and perhaps other artefacts of processing the data. The limitations are shown by the loss of resolution in the model obtained by inverting the determinant average data compared with joint inversion of TE and TM modes. We illustrate our algorithms by inverting synthetic data and the COPROD2 data set.

1. Introduction

The approach for solving the magnetotelluric (MT) inverse problem for 1D conductivity structures is well in hand. The Earth is first discretized into a number of layers (generally more than there are data) of constant conductivity. The equations are linearized, sensitivities calculated, an an iterative algorithm is used to generate a conductivity for each layer so that the data are adequately reproduced. Because the inverse problem is non-unique the inversion is solved by minimizing a functional of the conductivity subject to the data constraints. The conductivity functional is usually designed to penalize roughness. This procedure, used by Constable et al. (1987), Smith and Booker (1988), Dosso and Oldenburg (1989) and others has been shown to work well. Moreover, because calculation of the forward modelled responses and the sensitivities (or Frechet derivatives) are easily carried out, and because the $M \times M$ matrix (where M is the number of layers) is easily invertible, the inversion result can be rapidly achieved on a small workstation.

The above inversion method can, in principle, be used to solve 2D and 3D inverse problems. An example of this is the inversion of MT data in two dimensions by deGroot-Hedlin and Constable (1990). The practical issue, however, is the amount of computation required. In typical 2D problems the number of cells and the number of data may both be of the order of a few thousand. The computational roadblocks are: (1) computations

* Corresponding author.

required to carry out a forward modelling; (2) calculation of the sensitivities; (3) solution of a large matrix. As the size of the problem increases, the computational burdens in these three areas are exacerbated and eventually the researcher tires of waiting for the output, even when he or she has access to a larger computer.

This paper proposes an alternative solution to the problem. Our goal is to obtain a conductivity distribution that acceptably reproduces the observations but at the same time keeps the computations to a minimum. With reference to the three computational roadblocks listed above, our general philosophy is to use accurate forward modelling but to keep the number of forward modellings to a minimum, and to use approximate sensitivities in place of accurate values. The difficulty with inverting a large matrix is addressed in two ways. An efficient linear programming solution is advocated for mid-sized problems and a generalized subspace technique is recommended for large-scale problems. In implementing the above philosophy, we have generated two algorithms, both of which make use of sensitivities derived from a 1D linearization of the electromagnetic (EM) equations with respect to the z coordinate. The first inversion algorithm is a modification of the AIM-DS algorithm of Oldenburg and Ellis (1991) (hereafter referred to as P1). It produces a block model having minimum variation subject to an l_1 norm misfit criterion. The primary modification to the algorithm in P1 is the introduction of prespecified values for the Lagrange multiplier controlling the misfit at successive iterations. This keeps the number of forward modellings to one per iteration. The second algorithm minimizes an l_2 norm objective function of the model subject to an l_2 norm misfit in the data. The inversion of a large matrix is obviated by appealing to a generalized sub-space method (Skilling and Bryan, 1984; Kennett and Williamson, 1988; Oldenburg et al., 1993). In the sub-space methods the perturbation to the conductivity at each iteration is restricted to be a linear combination of q search vectors. The efficiency of the method lies in the fact that only a $q \times q$ matrix needs to be inverted and the efficacy of the method lies in the appropriate choice of vectors. We follow closely the work presented by Oldenburg et al. (1993) and henceforth refer to this paper as P2.

The l_1 and l_2 inversion algorithms can invert individually or jointly transverse electric (TE) and transverse magnetic (TM) mode responses over a 2D conductivity model. We illustrate the algorithms by inverting data obtained from the test model of Smith (1988) and the COPROD2 data of Jones and Savage (1986) and Jones (1988). For the COPROD2 data set we illustrate the extra resolution obtained by inverting the TE and TM modes jointly compared with the inversion of the determinant average data.

2. Inversion algorithms

Great flexibility exists in setting up any inverse problem. We begin by presenting details of our choices for: (1) forward modelling; (2) data; (3) inversion model; (4) sensitivities. These items are the same for both algorithms. Explicit details on the model norm to be minimized and the method of solution will be given after these four items have been considered.

To effect our forward mapping, the 2D conductivity model is first cellularized with rectangular elements. The model is partitioned into n_y horizontal cells and n_z vertical cells and the 2D conductivity $\sigma(y, z)$ is partitioned to an $n_z \times n_y$ array σ_{ij}, $i = 1, \ldots, n_z$; $j = 1, \ldots, n_y$. The thickness of the cells increases (usually logarithmically) with depth. Lateral partitioning in the survey region is dictated by the observation locations which are specified to be at the centre of each surface cell. This grid is terminated laterally by uniform layers and below by prisms elongated with depth. The conductivity is assumed to be constant in each cell and the 2D MT responses are computed using a transmission surface modelling code (Madden, 1972).

TE or TM impedances at n_y observation sites and at n_f frequencies can be inverted individually, jointly, or as determinant averages. Determinant average impedances are generated by

$$Z_{\text{det}} = \left(Z_{xx}Z_{yy} - Z_{xy}Z_{yx}\right)^{1/2} \tag{1}$$

(Berdichevsky and Dmitriev, 1976, p. 208). When data are provided in the form of apparent resistivity and phase, we transform these to the response

$$R_{jl} = \mu_0 \frac{H(y_j, 0, \omega_l)}{E(y_j, 0, \omega_l)}, \; j = 1, \ldots, n_y; \; l = 1, \ldots, n_f \tag{2}$$

and choose as data the amplitude and phase of this complex quantity. Errors are converted by numerical simulation, assuming that the errors on the initial data are Gaussian, unbiased and independent.

For the 2D MT inverse problem we choose $\ln(\sigma)$ as the 'model' for the inverse problem. We let m_{ij}, $i = 1, \ldots, n_z$; $j = 1, \ldots, n_y$ denote the cellularized array of $\ln(\sigma_{ij})$ values where the cellularization is the same as used for the forward modelling.

A major goal of this paper is to illustrate the practicality of using approximate sensitivities. Numerous approximations exist but the simplest is to generate the 1D sensitivities, as done in P1. We begin by linearizing, in the z direction, the exact forward mapping, $\mathcal{F}$, about a model $m^{(n)}$ which in this case represents $\sigma(y, z)$ at the nth iteration. This yields

$$\mathcal{F}[m; \omega]\big|_{y=y_0} = \mathcal{F}[m^{(n)}; \omega]\big|_{y=y_0} + \int_0^\infty g_{1D}[m^{(n)}; y_0, z, \omega] \left[m(y_0, z) - m^{(n)}(y_0, z)\right] \mathrm{d}z + \ldots \tag{3}$$

where g_{1D} is the 1D kernel function associated with the conductivity $\sigma(y, z)$ at model offset y_0. The expression of g_{1D} for the amplitude and phase responses used here has been given by Oldenburg (1979, Eq. (7)):

$$g_{1D}[m^{(n)}; y_0, z, \omega] = \mu_0 \sigma^{(n)}(y_0, z) \left[\frac{E(y_0, z, \omega)}{E(y_0, 0, \omega)}\right]^2 \tag{4}$$

where $E(y_0, z, \omega)$ is the electric field strength at the point (y_0, z). In the simplest approximation, $E(y_0, z, \omega)$ is taken to be the $E(z)$ field in a 1D Earth with conductivity $\sigma(y_0, z)$. Most of the work in this paper adopts this approximation. A better approximation to the 2D sensitivities (Smith and Booker, 1991) can be obtained by taking $E(y_0, z, \omega)$ to be the electric field at (y_0, z) in a 2D Earth $\sigma(y, z)$. Of course, the E fields are different for the TE and TM modes and so the approximate sensitivities will also be mode dependent in this approximation.

3. l_1 norm inversion algorithm

The l_1 norm inversion used here is essentially the AIM-DS algorithm in P1. The basic AIM-DS equation is

$$\tilde{\mathcal{F}}[m^{(n+1)}] = d^{\mathrm{obs}} + \tilde{\mathcal{F}}[m^{(n)}] - \mathcal{F}[m^{(n)}] \tag{5}$$

where $\mathcal{F}$ denotes a true forward mapping and $\tilde{\mathcal{F}}$ denotes an approximate forward mapping. Defining the approximate forward mapping by keeping only the first two terms in (3) yields

$$\tilde{\mathcal{F}}[m; \omega]\big|_{y=y_0} = \mathcal{F}[m^{(n)}; \omega]\big|_{y=y_0} + \int g_{1D}[m^{(n)}; y_0, z, \omega] \left[m(y_0, z) - m^{(n)}(y_0, z)\right] \mathrm{d}z \tag{6}$$

and substituting into (5) yields

$$\int g_{1D}[m^{(n)}; y_0, z, \omega] m^{(n+1)}(z) \, \mathrm{d}z = d^{\mathrm{obs}}(y_0, \omega) - d^{(n)}(y_0, \omega) + \int g_{1D}[m^{(n)}; y_0, z, \omega] m^{(n)}(z) \, \mathrm{d}z \tag{7}$$

Discretizing this equation produces

$$\sum_i^{n_z} A_{ijk} m_{ij}^{(n+1)} = d_{jk}^{\mathrm{obs}} - d_{jk}^{(n)} + \sum_i^{n_z} A_{ijk} m_{ij}^{(n)}, \quad j = 1, \ldots, n_y, \; k = 1, \ldots, 2n_f \tag{8}$$

where A_{ijk} is the integral of the kth 1D kernel function g_{1D} at model offset y_j over the ith depth partition. The model norm minimized is a discretization of the functional

$$\phi_m = \beta \int\int \left|\frac{\partial m}{\partial y}\right| \mathrm{d}y \, \mathrm{d}z + \gamma \int\int \left|\frac{\partial m}{\partial z}\right| \mathrm{d}y \, \mathrm{d}z \tag{9}$$

The inverse problem is solved by minimizing the objective function

$$\phi = \beta \sum_{i}^{n_z} \sum_{j}^{n_y - 1} \Delta z_i \left| m_{i,j+1}^{(n+1)} - m_{i,j}^{(n+1)} \right| + \gamma \sum_{i}^{n_z - 1} \sum_{j}^{n_y} \Delta y_j \left| m_{i+1,j}^{(n+1)} - m_{i,j}^{(n+1)} \right| + \mu \sum_{k}^{2n_f} \sum_{j}^{n_y} \left| \frac{R_{jk}^{\text{obs}} - R_{jk}^{(n)}}{\epsilon_{jk}} \right| \quad (10)$$

where ϵ_{jk} is the standard deviation of the datum R_{jk}^{obs}. The parameters β and γ control the relative weighting of the y and z variations and are fixed for each inversion. The Lagrange multiplier μ is generally sought so that the final misfit achieves a target value consistent with the errors associated with the data; however, rather than start with this value of μ, a schedule of μ values is found to be desirable. This schedule is a monotonically increasing sequence which asymptotically approaches a constant. The increasing nature of μ prevents unnecessary roughness being accumulated in the model at intermediate steps. Although this approach has the possible disadvantage that a non-optimum sequence of μ values may be selected, the major advantage is that only one forward modelling needs be carried out per iteration. This may be compared with the approach where a schedule of target misfits is chosen and then a line search performed to find the corresponding Lagrange multiplier μ at each iteration. Of course, here there is the possible disadvantage that a non-optimum sequence of target misfits may be selected, with the associated non-optimum number of line searches. Although the best approach is somewhat model dependent we find, in general, that the computational savings associated with a simple schedule of μ values can be substantial because line searches typically involved 3–8 forward modellings.

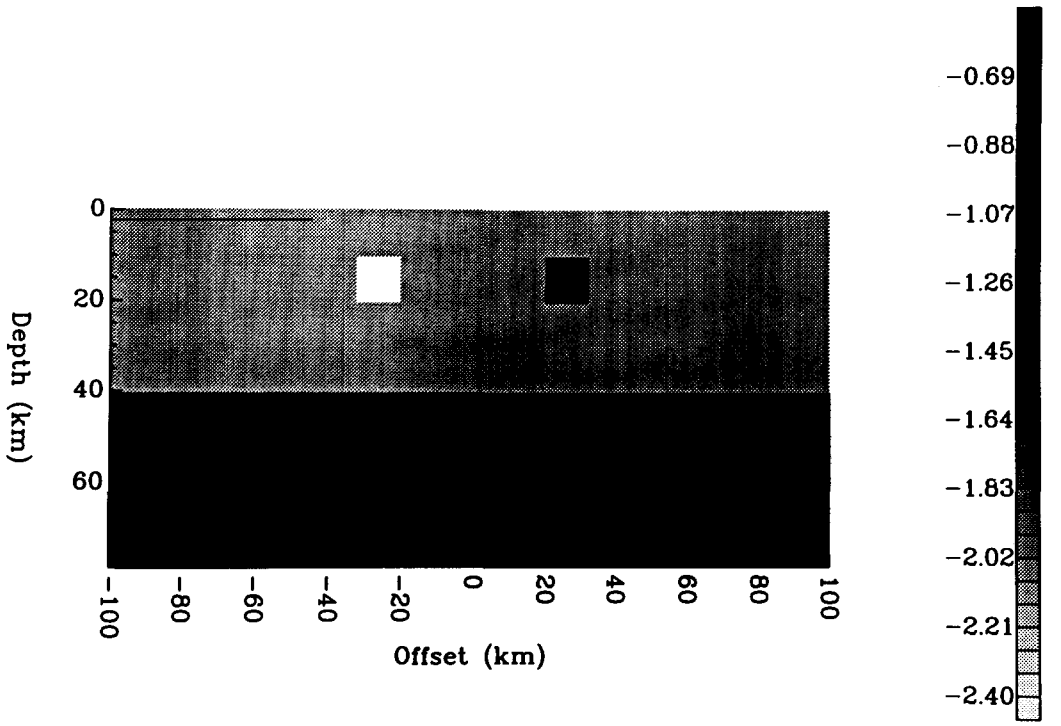

Fig. 1. Synthetic conductivity model consisting of a 1000 Ω m resistive prism and 10 Ω m conductive prism in a 100 Ω m background, all overlying a 10 Ω m conductive basement. The scale on the right is $\log_{10}\sigma$.

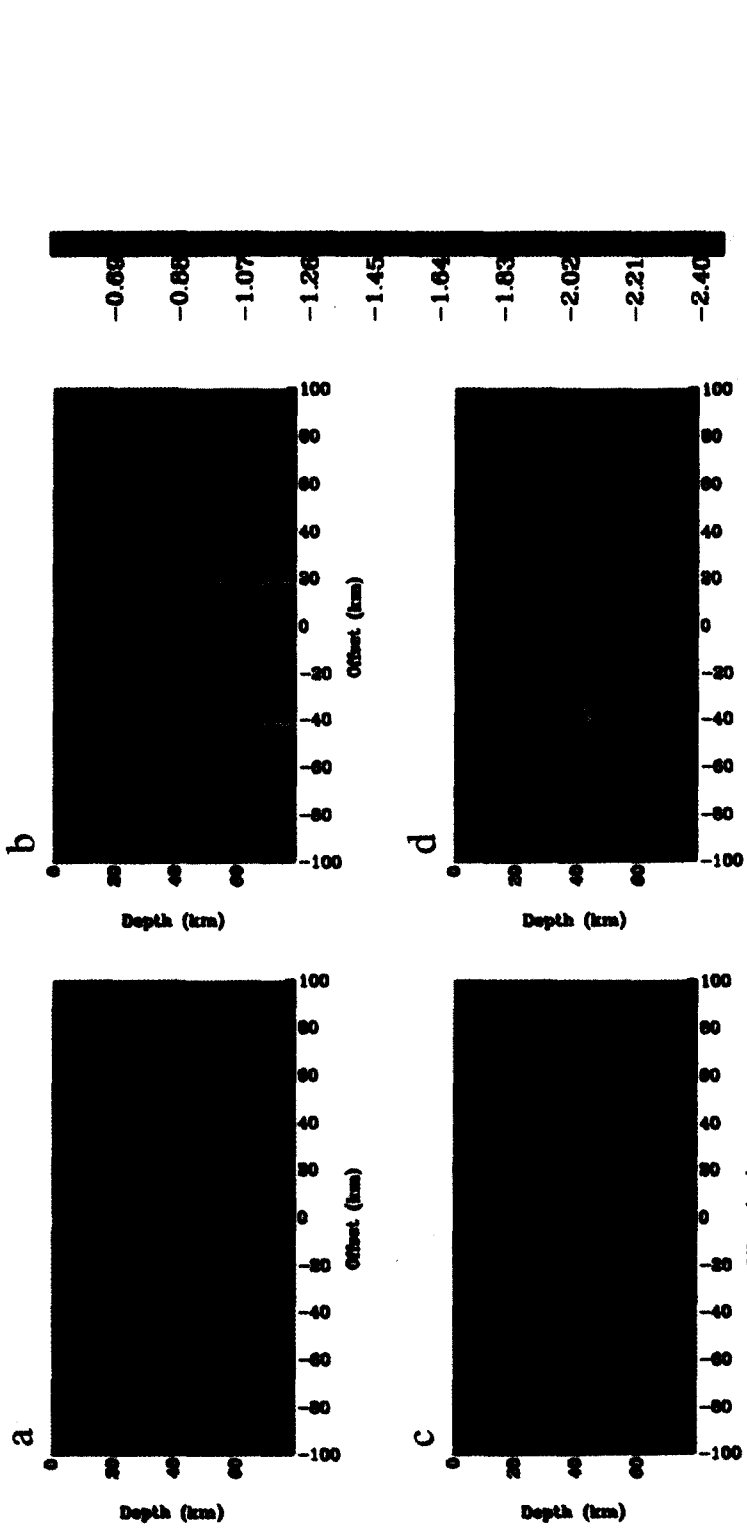

Fig. 2. (a)–(c), l_1 inversions, with $\beta:\gamma = 10:1$, 1:10 and 1:1, respectively, of the same noisy determinant data associated with Fig. 1. (d) Result of inverting the TE and TM mode data jointly, with $\beta:\gamma = 1:1$.

The use of l_1 norm requires the definition of an appropriate misfit measure. Here we follow Parker and McNutt (1980) and define

$$\chi_N^1 = \frac{1}{2n_f n_y}\left(\frac{\pi}{2}\right)^{1/2} \sum_{k}^{2n_f} \sum_{j}^{n_y} \left| \frac{R_{jk}^{\text{obs}} - R_{jk}^{(n)}}{-\epsilon_{jk}} \right| \qquad (11)$$

and note that $E[\chi_N^1] = 1$.

The problem of minimizing (10) subject to the constraints in (8) is solved using linear programming (LP) techniques and was outlined in P1. Here, we simply note that minimization of an l_1 model norm gives rise to models with large regions of constant conductivity, i.e. block type models, and the use of an l_1 data norm provides robustness in the presence of non-Gaussian noise on the data. These two characteristics of the l_1 norm are well suited to inverting MT data with static shift effects.

4. l_1 inversion of synthetic data

We now invert data from a model suggested by Smith and Booker (1991) and used in P1. The model has $n_y = 50$ horizontal cells and $n_z = 30$ vertical cells. It consists of a 1000 Ω m resistive prism and a 10 Ω m conductive prism in a 100 Ω m background, all overlying a 10 Ω m conductive basement. This model, shown in Fig. 1, was chosen mainly because of its similarity to models already used by other workers (Smith, 1988; deGroot-Hedlin and Constable, 1990). It is hoped that a comparison of the same model inverted by different techniques will provide information about the strengths and weaknesses of different methods. The cellularization of the model we use has logarithmically increasing cell thickness with depth and constant cell width in the region of interest. When referring to data associated with this model we consider nine frequencies in the range 1.0–0.00033 Hz, with ±5% unbiased Gaussian noise added to the apparent resistivities and ±2° noise to the phases.

It is well known that the inversion of a finite number of noisy data can yield a wide range of models depending on the choice of model norm regularization. To illustrate the significance of this fact we present the results of the inversion of the determinant data associated with the model shown in Fig. 1 for three choices of the ratios $\beta:\gamma$, 10:1, 1:10 and 1:1, in Figs. 2(a), 2(b) and 2(c), respectively. These three models all have the same misfit $\chi_N^1 = 1$ and have significantly different character. For comparison, we show in Fig. 2(d) the results of a joint inversion of the TE and TM mode data with $\beta:\gamma = 1:1$ and $\chi_N^1 = 1$. All four inversions were performed with 1D sensitivities (Eq. (4)), using the E field associated with the corresponding 1D Earth. As a test case (not shown), the joint inversion was repeated with 1D sensitivities (Eq. (4)), using the E field associated with the 2D Earth conductivity: only a minor improvement was observed in the final model and the rate of convergence. It should be noted that the joint inversion has twice the number of data compared with the determinant inversions and correspondingly produces a better representation of the true model than the determinant inversion. However, we remind the reader that the determinant inversion does not require knowledge of the TE and TM modes, is a smaller inverse problem, and yet yields a very good approximation of the true model. In practical terms, this makes the determinant inversion an appealing tool.

5. l_2 norm inversion algorithm

The blockiness in models produced by the l_1 norm inversion may be a desirable characteristic for some Earth geologies but other geologic environments might be better emulated by smoother models. Correspondingly, we present an algorithm which minimizes the l_2 norm of the conductivity model. In this algorithm we also allow for a base model in the objective function so that structure is minimized with respect to that background model.

In the l_1 norm inversion the recovered model is achieved by explicitly solving for the conductivity in each cell. Effectively, each cell is a basis element in model space and finding values of the conductivities requires solving a large matrix system. This is efficient for an LP solver (Marsten, 1981, D'Azevedo et al., 1991) for the size of

problem just considered. The l_2 norm inversion, however, requires the solution of an $M \times M$ matrix, where M is the number of cells, and this computation can become prohibitive as the number of cells increases. To obviate this difficulty we appeal to a generalized sub-space method. We adopt the formalism outlined in P2 but present the essence of the calculation here.

We first define a model objective function

$$\phi_m(\boldsymbol{m}) = \left\| \boldsymbol{W}_m(\boldsymbol{m} - \boldsymbol{m}_0) \right\|^2 \tag{12}$$

where $\boldsymbol{m}_0$ is a base model and $\boldsymbol{W}_m$ is a weighting matrix. The minimization of ϕ_m yields a model that is close to $\boldsymbol{m}_0$ with the metric defined by $\boldsymbol{W}_m$ and so the characteristics of the recovered model are directly controlled by these two quantities. The data misfit may be characterized by

$$\phi_d(\boldsymbol{d}) = \left\| \boldsymbol{W}_d(\boldsymbol{d} - \boldsymbol{d}^{\text{obs}}) \right\|^2 \tag{13}$$

where the relationship between the jth datum and the model is $d_j = \mathcal{F}_j[\boldsymbol{m}]$. The functionals $\mathcal{F}_j$ are assumed known. In (13), $\boldsymbol{W}_d$ is an $N \times N$ matrix. If the noise contaminating the jth observation is an uncorrelated Gaussian random variable having zero mean and standard deviation ϵ_j then an appropriate form for $\boldsymbol{W}_d$ is $\boldsymbol{W}_d = \text{diag}\{1/\epsilon_1, \ldots, 1/\epsilon_N\}$. With such assumptions, ϕ_d is a random variable distributed as χ^2 with N degrees of freedom. The expected value of ϕ_d is therefore approximately equal to N and, accordingly, the model sought from the inversion algorithm should reproduce the observations to about this value. The inverse problem of minimizing $\phi_m(\boldsymbol{m})$ such that $\phi_d(\boldsymbol{d}) = \phi_d^*$, a desired misfit, is solved by minimizing

$$\phi(\boldsymbol{m}) = \phi_m(\boldsymbol{m}) + \mu\left[\phi_d(\boldsymbol{d}) - \phi_d^*\right] \tag{14}$$

where μ is the Lagrange multiplier. The minimization problem is nonlinear and is addressed by linearizing (14) about the current model $\boldsymbol{m}^{(n)}$ and iterating. If $\delta\boldsymbol{m}$ is a model perturbation, then a truncated Taylor expansion which has terms only up to second order is

$$\begin{aligned}\phi(\boldsymbol{m}^{(n)} + \delta\boldsymbol{m}) &= \phi_m + \boldsymbol{\gamma}_m^T\delta\boldsymbol{m} + \tfrac{1}{2}\delta\boldsymbol{m}^T\boldsymbol{H}_m\delta\boldsymbol{m} \\ &\quad + \mu\left\{\phi_d + \boldsymbol{\gamma}_d^T\delta\boldsymbol{m} + \tfrac{1}{2}\delta\boldsymbol{m}^T\boldsymbol{H}_d\delta\boldsymbol{m} - \phi_d^*\right\}\end{aligned} \tag{15}$$

where $\boldsymbol{\gamma}_m = \nabla_m\phi_m$ and $\boldsymbol{\gamma}_d = \nabla_m\phi_d$ are gradient vectors, $\boldsymbol{H}_m = \nabla_m\nabla_m\phi_m$ and $\boldsymbol{H}_d = \nabla_m\nabla_m\phi_d$ are Hessian matrices and ∇_m is the operator $(\partial/\partial m_1, \partial/\partial m_2, \ldots, \partial/\partial m_M)^T$. In (15), ϕ_m is understood to be $\phi_m(\boldsymbol{m}^{(n)})$ and ϕ_d is $\phi_d(\boldsymbol{d}^{(n)})$.

In a sub-space approach, the 'model' perturbation $\delta m \in R^M$ is restricted to lie in a q-dimensional sub-space of R^M which is spanned by the vectors $\{\boldsymbol{v}_i\}$ $i = 1, q$. The model perturbation can be written as

$$\delta\boldsymbol{m} = \sum_{i=1}^{q} \alpha_i\boldsymbol{v}_i \equiv \boldsymbol{V}\boldsymbol{\alpha} \tag{16}$$

and is therefore specified once the parameters α_i are determined.

The sub-space equations are generated by substituting (16) into (15) to yield

$$\begin{aligned}\phi(\boldsymbol{m}^{(n)} + \boldsymbol{V}\boldsymbol{\alpha}) &= \phi_m + \boldsymbol{\gamma}_m^T\boldsymbol{V}\boldsymbol{\alpha} + \tfrac{1}{2}\boldsymbol{\alpha}^T\boldsymbol{V}^T\boldsymbol{H}_m\boldsymbol{V}\boldsymbol{\alpha} \\ &\quad + \mu\left\{\phi_d + \boldsymbol{\gamma}_d^T\boldsymbol{V}\boldsymbol{\alpha} + \tfrac{1}{2}\boldsymbol{\alpha}^T\boldsymbol{V}^T\boldsymbol{H}_d\boldsymbol{V}\boldsymbol{\alpha} - \phi_d^*\right\}\end{aligned} \tag{17}$$

This is a quadratic objective function to be solved for the parameter vector $\boldsymbol{\alpha}$. Differentiating (17) with respect to $\boldsymbol{\alpha}$ and setting the resultant equations equal to zero yields

$$\boldsymbol{V}^T(\boldsymbol{H}_m + \mu\boldsymbol{H}_d)\boldsymbol{V}\boldsymbol{\alpha} = -\mu\boldsymbol{V}^T\boldsymbol{\gamma}_d - \boldsymbol{V}^T\boldsymbol{\gamma}_m \tag{18}$$

The solution of this system requires that a line search be carried out to find the value of the Langrange multiplier μ so that a specific target value ϕ_d^* is achieved. This involves an initial guess for μ, solving (18) by singular value decomposition (SVD) for the vector $\boldsymbol{\alpha}$, computing the perturbation δ_m, carrying out a forward modelling to evaluate the true responses and misfit, and then adjusting μ. The estimation of an acceptable value of μ typically requires three or four forward modellings.

At each iteration in the inversion we desire a model perturbation which minimizes ϕ_m and alters ϕ_d so that it achieves a specific target value ϕ_d^*. To prevent the build-up of unnecessary roughness, the target misfit begins at an initial value (usually a fraction of the misfit generate by the starting model) and decreases with successive

iterations towards a final value selected by the interpreter. Convergence is reached when the data misfit reaches this final target and no further reduction in the model norm is obtained with successive iterations.

The success of the sub-space method depends strongly upon the choice of basis vectors. In P2, considerable success was achieved by subdividing the misfit objective function and using steepest descent vectors associated with each data gradient. Here we segment the data misfit objective function according to frequency, amplitude and phase. The ith search vector becomes

$$v_i = \left(W_m^T W_m\right)^{-1} \nabla_m \sum_{j=(i-1)*n_y+1}^{i*n_y} \left(\frac{d_j - d_j^{\text{obs}}}{\epsilon_j} \right)^2,$$

$$i = 1, \ldots, 2n_f \tag{19}$$

so that the datum in the sum is an amplitude (or phase) at a single frequency. Every inversion carried out contains these vectors, the constant vector and also $(W_m^T W_m)^{-1} \nabla_m \phi_m = m - m_0$, which is the steepest descent vector associated with the model objective function. In addition, we also segment the model norm objective function and form steepest descent vectors from each of the minor objective functions. The segmentation is invoked by choosing a rectangular region of the model and subdividing this region into smaller groups. The projection of $(m - m_0)$ onto any group of cells provides a new search direction. The rectangular region can be the entire model or a smaller portion. In the initial iterations we used each row of cells to make a basis vector and sometimes also subdivided the entire model into groupings of $n \times n$ cells, where n is typically 5×5. In later iterations, when features of interest appeared, we have centred the rectangle over the area to see if structure is enhanced or attenuated in a subsequent iteration. This process can be both interactive and dynamic. The attractive aspect of the dynamic use of additional search vectors is that in the sub-space inversion only a model perturbation is sought. At worst, a poor choice of vectors produces little benefit.

6. l_2 norm inversion of synthetic data

The l_1 norm inversion results illustrated effectively the dependence of the inversion result on the model norm. The same dependence exists for the l_2 inversions, but this aspect does not need further investigation here. As such, for the l_2 inversions, we select a single objective function. Our choice for ϕ_m is guided by a desire to find a model which has minimum structure in the vertical and horizontal directions and at the same time is close to a base model m_0. To accomplish this we minimize a discretized approximation to

$$\begin{aligned}
&\phi_m(m, m_0) \\
&\quad = \alpha_s \int\int w_s(y, z)(m - m_0)^2 \, \mathrm{d}y \, \mathrm{d}z \\
&\qquad + \int\int \left\{ \alpha_x w_y(y, z) \left[\frac{\partial (m - m_0)}{\partial y} \right]^2 \right. \\
&\qquad \left. + \alpha_z w_z(y, z) \left[\frac{\partial (m - m_0)}{\partial z} \right]^2 \right\} \mathrm{d}y \, \mathrm{d}z
\end{aligned} \tag{20}$$

where α_s, α_y and α_z are adjustable constants having suitable dimensions to make (20) dimensionless. The weighting functions $w_s(y, z)$, $w_y(y, z)$ and $w_z(y, z)$ provide additional flexibility to control the characteristic of the final model. The discrete form of (20) is

$$\begin{aligned}
\phi_m &\equiv \phi_s + \phi_y + \phi_z \\
&= \alpha_s \| W_s(m - m_0) \|^2 + \alpha_y \| W_y(m - m_0) \|^2 \\
&\quad + \alpha_z \| W_z(m - m_0) \|^2 \\
&= (m - m_0)^T \left\{ \alpha_s W_s^T W_s + \alpha_y W_y^T W_y + \alpha_z W_z^T W_z \right\} \\
&\quad \times (m - m_0) \\
&\equiv (m - m_0)^T W_m^T W_m (m - m_0)
\end{aligned} \tag{21}$$

In (21), W_s is a diagonal matrix with elements $(\Delta y \Delta z)^{1/2}$, where Δy is the length of the cell and Δz is its thickness, W_y has elements $\pm \Delta z / (\Delta z_1 \mathrm{d}y)^{1/2}$, where $\mathrm{d}y$ is the distance between the centres of horizontally adjacent cells, and W_z has elements $\pm (\Delta y \Delta z / \Delta z_1 \, \mathrm{d}z)^{1/2}$, where $\mathrm{d}z$ is the distance between the centres of vertically

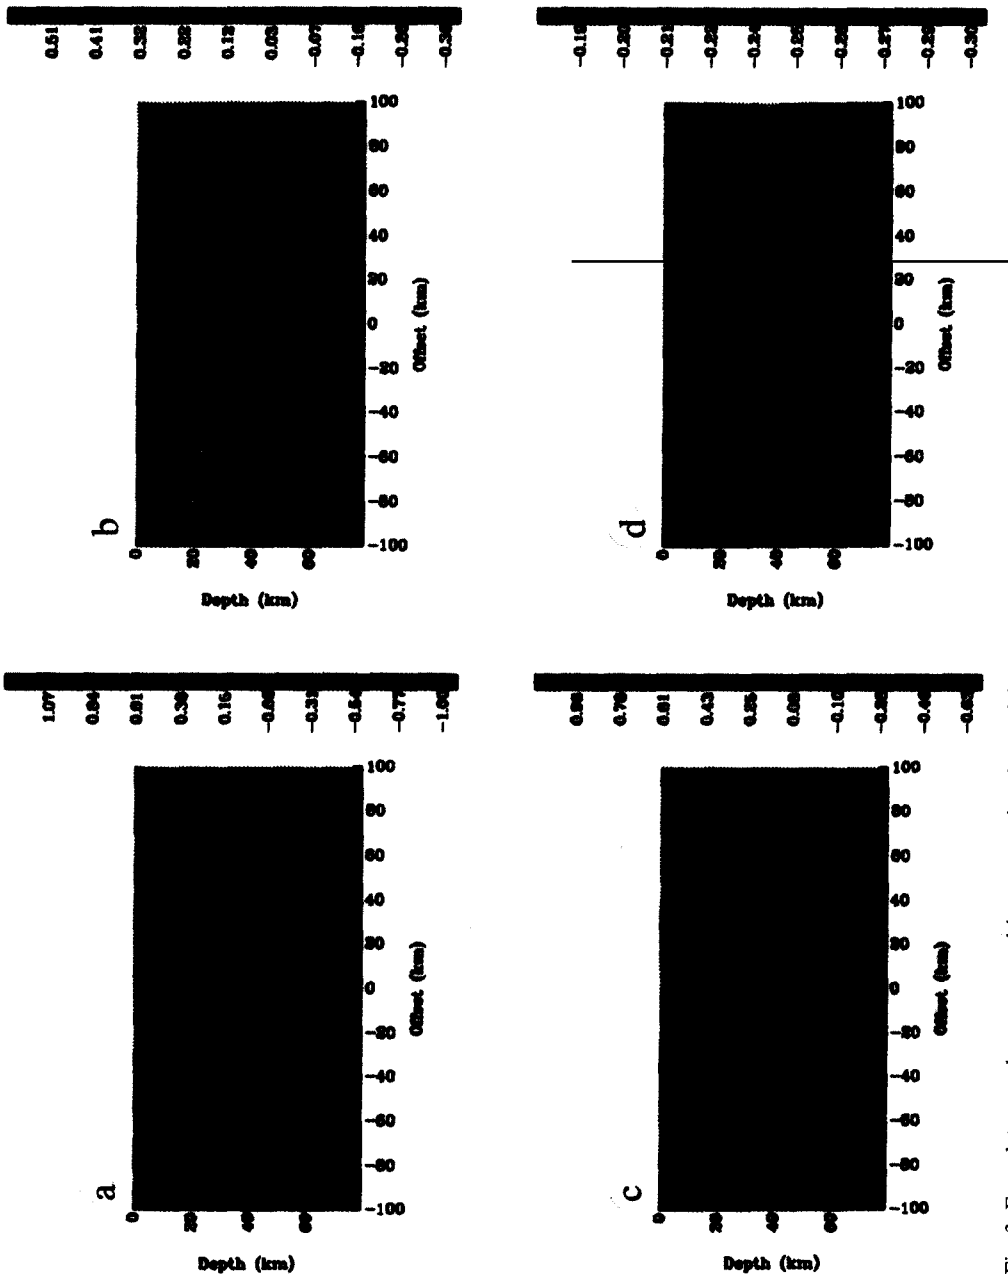

Fig. 3. Four data search vectors used in constructing the model perturbation at the first iteration. These are steepest descent vectors for (a) $f = 1.0$ Hz (phase); (b) $f = 0.05$ Hz (phase); (c) $f = 0.05$ Hz (amplitude) and (d) $f = 0.0033$ Hz (amplitude).

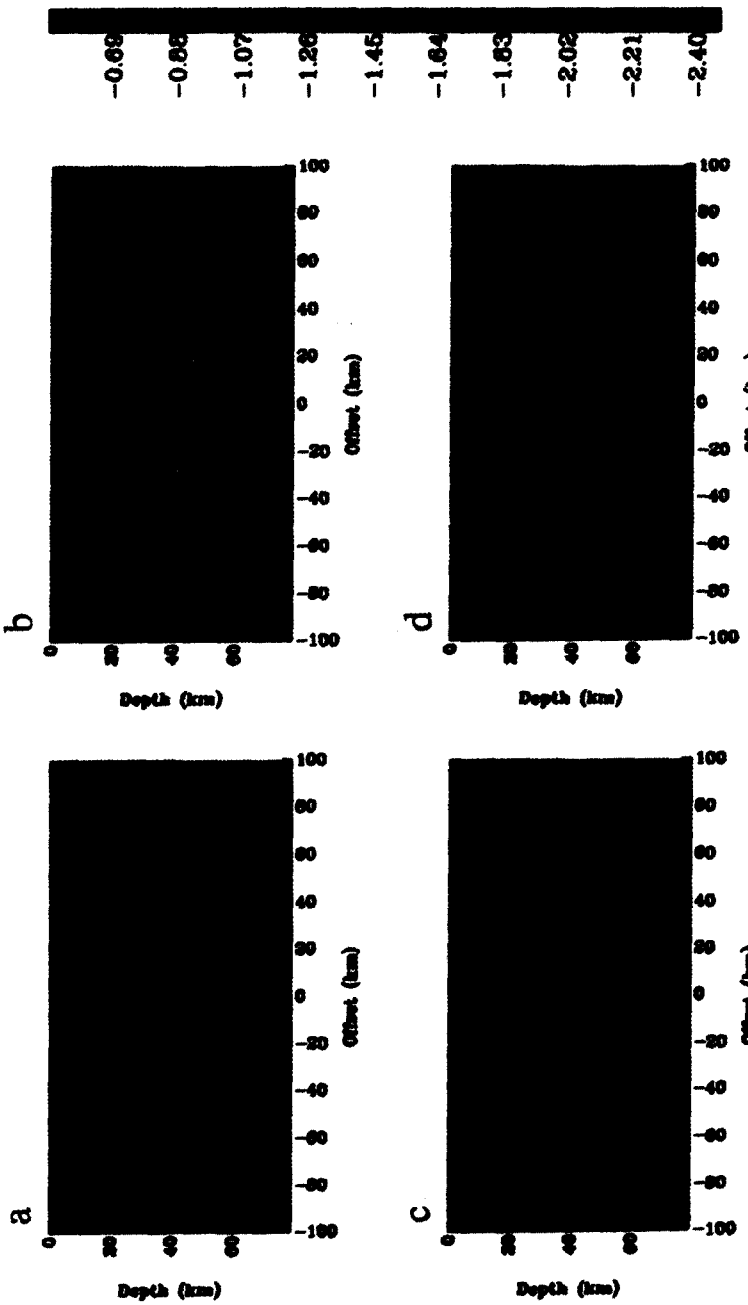

Fig. 4. (a) The synthetic model. (b) The l_2 norm model obtained at iteration 15 by inverting determinant average data. The reference model was a half-space of 0.03 S m^{-1}. (c) Result obtained by inverting the same data but using a 1D reference model obtained from (b). (d) Model produced by jointly inverting TE and TM mode data. This result was achieved at iteration 10.

adjacent cells. The factor $(\Delta z/\Delta z_1)^{1/2}$ in the specification of the $\boldsymbol{W}_y$ and $\boldsymbol{W}_z$ matrices accounts for non-constant $w_y(y, z)$ and $w_z(y, z)$ and causes the constructed model to discriminate against conductivity variation at depth. This is in accordance with the reality of strong surface conductivity variations that are often found in geophysical surveys and the reduced resolution with depth for EM investigations. We have chosen $\alpha_s = 10^{-8}$, $\alpha_y = 1.0$, and $\alpha_z = 1.0$.

For the first inversion of the Smith and Booker model we invert Berdichevsky determinant average data and use $m_0 = 0.03$ S m^{-1} as a reference model. The starting model is a half-space of 0.01 S m^{-1}. The inversion is carried out with 49 search vectors; there are 18 data gradient search vectors (two for each period), the constant vector and 30 steepest descent vectors for the model objective function (one for each row of cells in the model). Before showing the result we justify the grouping of the data in forming our search vectors. In Fig. 3 we show four selected vectors at the first iteration. These vectors are clearly useful search directions for they have manifestations of surface roughness, manifestations of the conductive blocks and have the potential for altering the half-space model toward the two-layer model.

The model shown in Fig. 4(b) was achieved after 15 iterations. The final fit to the data is $\phi_d = 988$, which is acceptably close the number of data, 900. The improvement in the data misfit is rapid in the initial iterations and decreases beyond the seventh iteration. The general features of the true model are visible; the 1D background is reasonably well defined and the sharp conductivity contrast at 40 km is seen as a smoothed transition zone; this is a typical result in an l_2 inversion. The resistive and conductive anomalies are fairly well determined in the horizontal direction but there are dropouts beneath both prisms. There is also a conductive overshoot visible beneath the conductive prism. The top of the conductive prism is close to that in the true model but the resistive prism has been pushed to greater depth. We notice also that the near-surface layer exhibits more lateral variability, which is in accordance with our weighting function.

The 1D background is often of limited interest geophysically and usually we are more interested in deviations from that background. Accordingly, we carry out another inversion where the base model has been altered from a half-space of 0.03 S m^{-1} to a 1D conductivity structure obtained by taking the lateral average of the conductivity shown in Fig. 4(b). The result after seven iterations is displayed in Fig. 4(c). The data misfit for this model is $\phi_d = 976$. The two prisms are visually enhanced and the conductivity overshoots (especially beneath the conductivity prism) are reduced from that shown in Fig. 4(b). We deem this to be a beneficial result.

The inversion above was performed using determinant average data. We now compare those results with that obtained by jointly inverting the TE and TM mode data. That inversion result, obtained after 10 iterations and using a reference model of 0.03 S m^{-1}, is shown in Fig. 4(d). The final misfit to the data was $\phi_d = 2487$. This is somewhat greater than the desired fit of 1800, but this seemed to be approximately the limit that could be achieved using these basis vectors. The joint inversion used 38 data gradient vectors, the constant vectors and 30 vectors associated with each row of cells, making a total of 69 vectors. The results are not substantially different from those in Fig. 4(c), and in fact, the result in Fig. 4(c) might be moderately superior.

7. COPROD2 data set l_1 norm inversion

We now invert a subset of the COPROD2 data set provided by A.G. Jones. These data were collected along an east–west traverse in southern Saskatchewan and Manitoba in Canada. There are 35 stations with spacings of approximately 10 km. The data have been analysed by Jones and Savage (1986) and Jones (1988), and have been used as test data for inversion algorithms. We now present the results of the application of the AIM l_1 inversion to a subset of the COPROD2 data. The subset chosen consists of the data at all 35 stations and at the eight frequencies 0.187, 9.37E − 02, 4.69E − 02, 2.33E − 02, 5.86E − 03, 2.93E − 03, and 1.46E − 03 Hz. From this subset we have extracted and inverted four data sets; the TE mode data, the TM mode data, the determi-

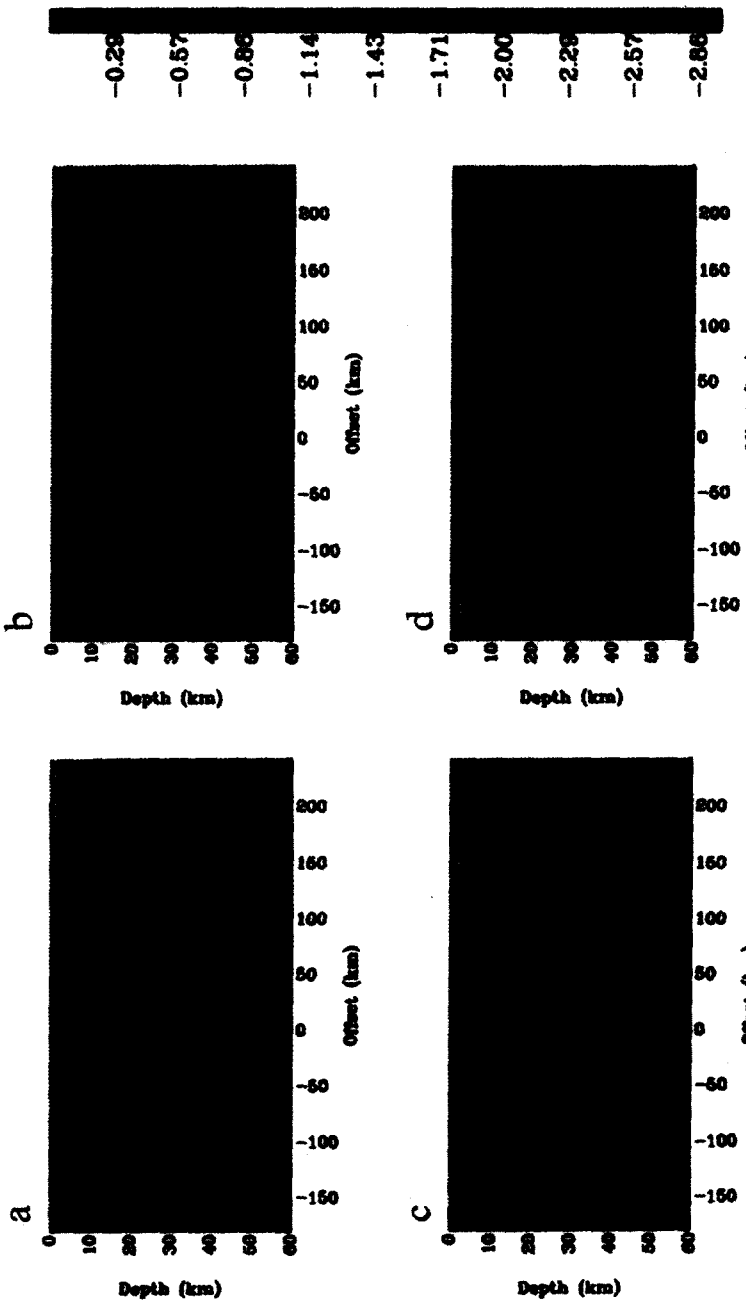

Fig. 5. The l_1 norm models obtained by inverting (a) TE, (b) TM, (c) determinant average, and (d) joint TE and TM modes for COPROD2 data.

nant average data and the joint TE and TM mode data. Following deGroot-Hedlin and Constable (1990), we attribute minimum errors of $\pm 10\%$ to the apparent resistivities and $\pm 5°$ to the phases.

The model produced by the AIM l_1 inversion using 1D sensitivities operating on the TE mode data with $\beta : \gamma = 5 : 1$ is shown in Fig. 5(a), and the observed and predicted data are shown in Fig. 6(a). The predicted data have $\chi_1^N = 0.97$. The TE mode apparent resistivity and phase data have a strong conducting signature in the vicinity of the NACP and the TOBE anomalies and this leads to compact conductors in the recovered model. We note that both conductivity anomalies are split vertically into two conductors. However, we emphasize that this is only one of many models which fit these data at this misfit level.

The result of inverting the TM mode data with $\beta : \gamma = 5 : 1$ and with 1D sensitivities gave the model shown in Fig. 5(b). The observed and predicted data are shown in Fig. 6(b). The data misfit is $\chi_1^N = 0.86$. The TM mode apparent resistivity and phase data show little manifestation of the conducting anomalies and consequently there is little structure in the recovered model.

Combining the TE and TM mode data to form the determinant average data and inverting ($\beta : \gamma$

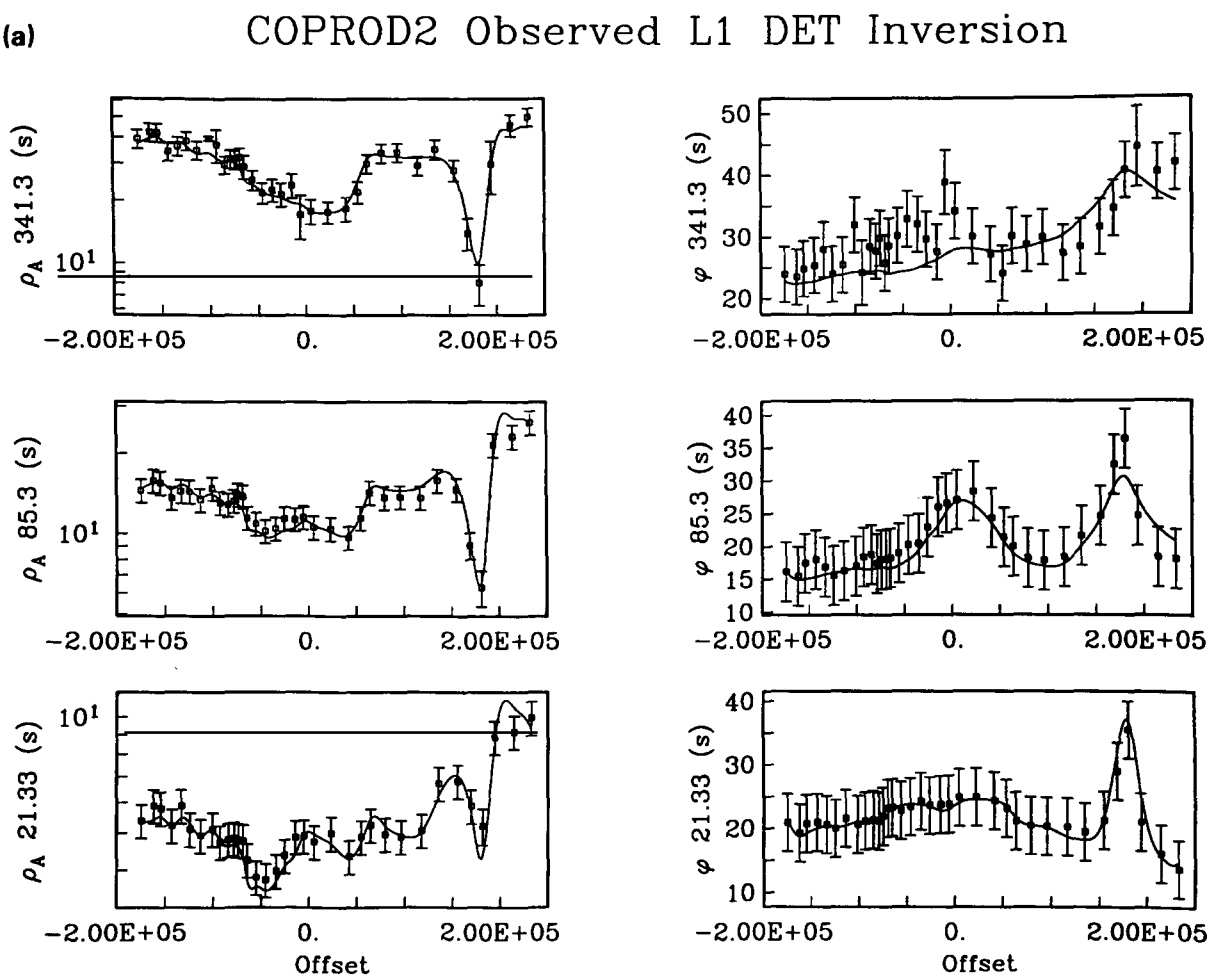

Fig. 6. The observed data (points with error bars) and the predicted data (curves) associated with the l_1 norm models produced by inverting (a) TE, (b) TM, (c) determinant average and (d) joint TE and TM modes for COPROD2 data.

= 5 : 1) produced the model shown in Fig. 5(c). This model gives rise to the predicted data shown in Fig. 6(c), which has $\chi_1^N = 0.67$. This inversion result clearly shows a blocky conductor associated with the NACP and a smaller, more compact conductor associated with the TOBE anomaly. The predicted data associated with this model overfit the observed data in a statistical sense. However, as the true data errors are somewhat uncertain, we do not consider this overfitting to be significant. More importantly, it has been our experience that optimal results are produced by inverting the determinant data when using 1D sensitivities, so that Fig. 5(c) represents our best representation, with 1D sensitivities using 1D E fields, of the conductivity structure along this profile.

After successful inversion of the determinant data, an attempt was made jointly to invert the TE and TM mode COPROD2 data. A satisfactory inversion result could only be achieved when the 2D E fields were used in the 1D sensitivity function (Eq. (4)); using the 1D E fields resulted in extremely poor convergence. The result of the successful inversion (Ellis et al., 1993) of the joint TE and TM mode data ($\beta : \gamma = 1 : 1$) is shown in Fig. 5(d), and the observed and predicted data are shown in Fig. 6(d). The predicted data have a normalized misfit $\chi_1^N = 1.15$. We were unable to reduce the misfit to $\chi_1^N = 1$ without increasing

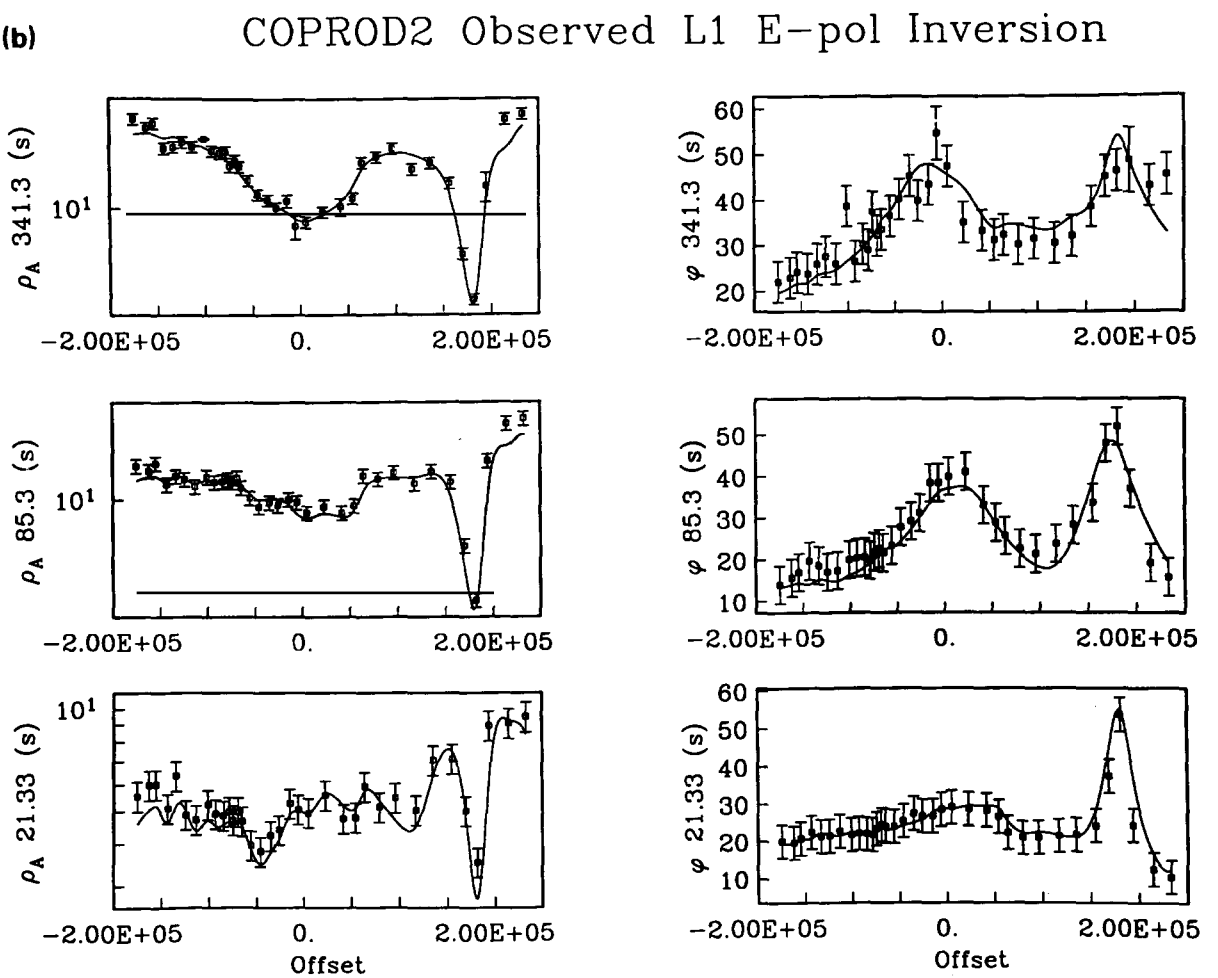

Fig. 6 (continued).

the number of cells in the model. Even at the $\chi_1^N = 1.15$ misfit level, the NACP anomaly is clearly resolved into three conductors at a depth of 20 km and dipping to the west. The TOBE anomaly is clearly resolved as a strong conductor at 10 km depth. These results confirm, in a general sense, the finding of other workers (e.g. deGroot-Hedlin and Constable 1990) who have inverted this data set.

8. COPROD2 data set l_2 norm inversion

The COPROD2 data set will now be inverted using the l_2 norm minimization and the sub-space method. We first invert the determinant average data. The model obtained after 18 iterations is shown in Fig. 7(a). The result was achieved in the following manner. The first six iterations were performed with a reference half-space model of 0.03 S m^{-1}. The resulting model was then averaged laterally to yield a 1D reference model for the remaining iterations. In these latter iterations, the model component of the objective function was divided in various ways (layers, 5×5 cell groupings, a large rectangular block centred on the NACP anomaly and a rectangular block centred on the TOBE anomaly). Steepest descent vectors associated with these cell groupings were used in conjunction with the data gradient vec-

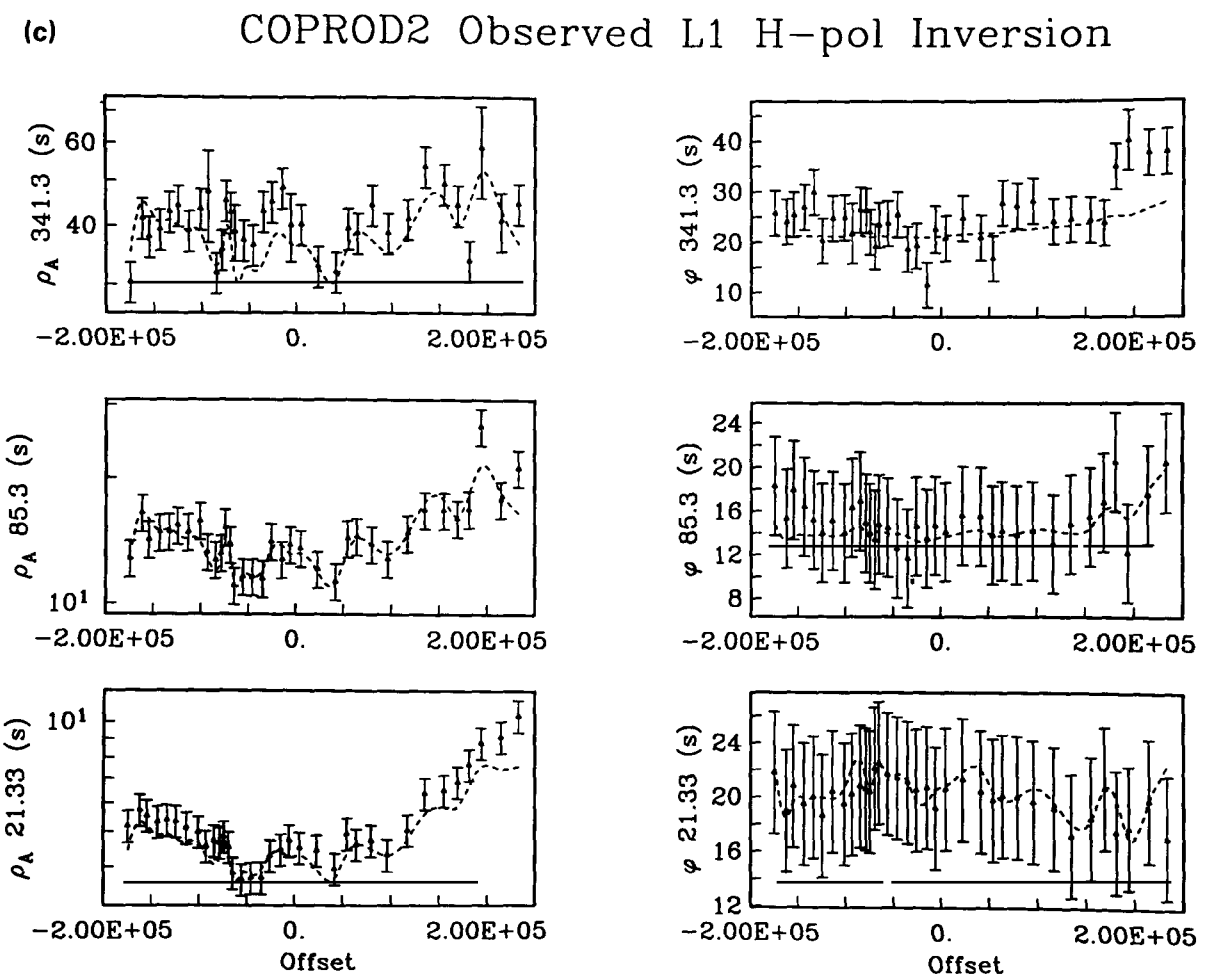

Fig. 6 (continued).

tors. For all iterations the number of vectors was less than 100. The constructed model in Fig. 7(a) shows that both the NACP and the TOBE anomalies are well defined. The data misfit for the constructed model is 666. This is somewhat more than the target misfit of 560. However, even at this level almost all of the data are well fitted. The major discrepancies and principal contributors to the global misfit are the phases at 5.33, 341 and 682 s. At these periods, the predicted phase curves are biased downward from the observations. The apparent resistivities at these frequencies, however, are well reproduced. The misfit information is summarized by the line plots in Fig. 8(a) and the plots in Fig. 9(a). These show the misfit contribution for amplitudes and phases for each frequency. The dashed line in Fig. 9(a) is the misfit expected if each data group contributed equally to the desired misfit of 560. Only the misfits from the phases at the three periods given above contribute more than this average value and, in fact, they contribute 514 to the total misfit of 666.

To look more closely at the information in the COPROD data set, we next inverted separately the TE and TM mode data. The reference model was the same 1D model as used when inverting the determinant average data. The starting model was the conductivity model shown in Fig. 7(a). The initial misfit for the TE mode data was

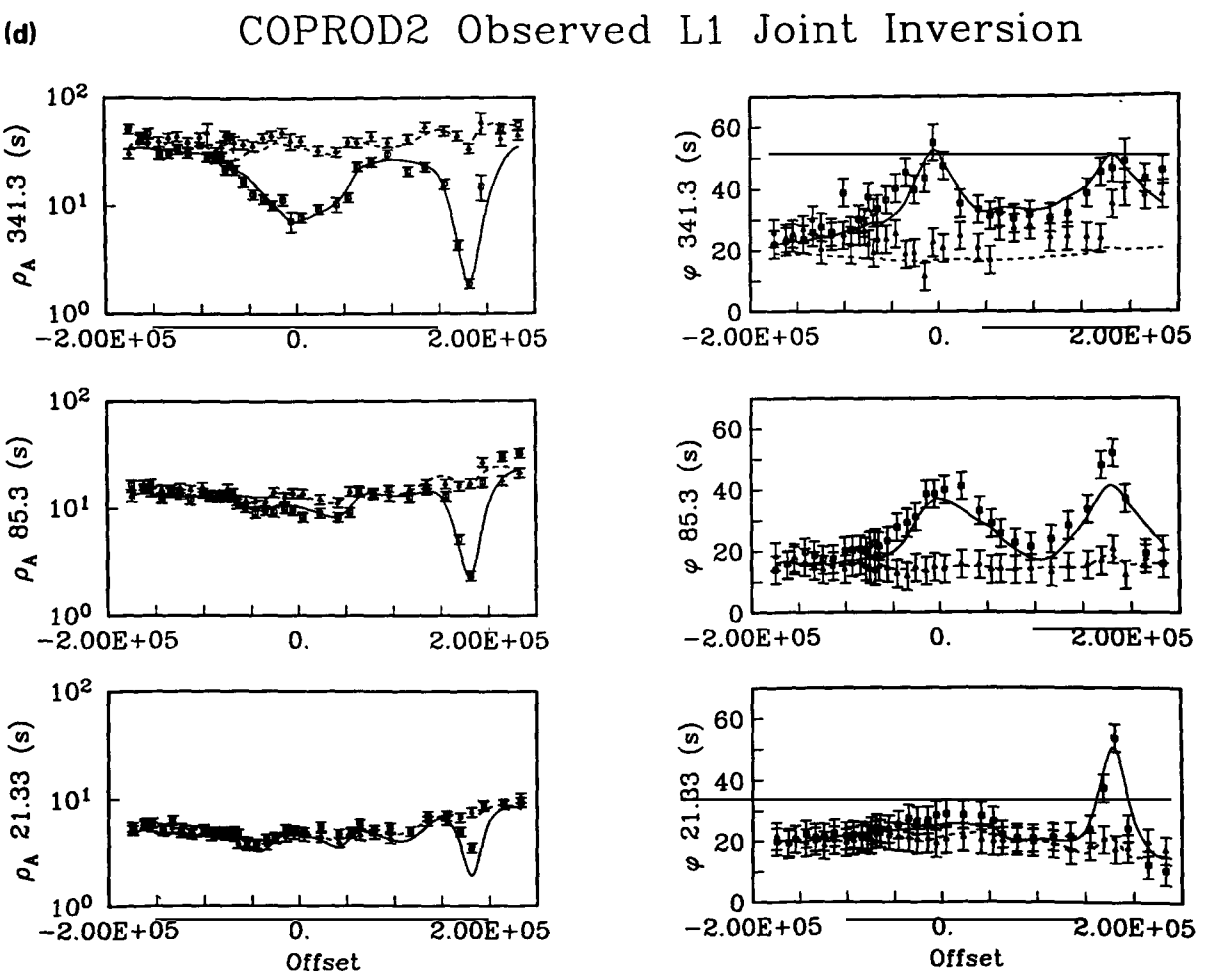

Fig. 6 (continued).

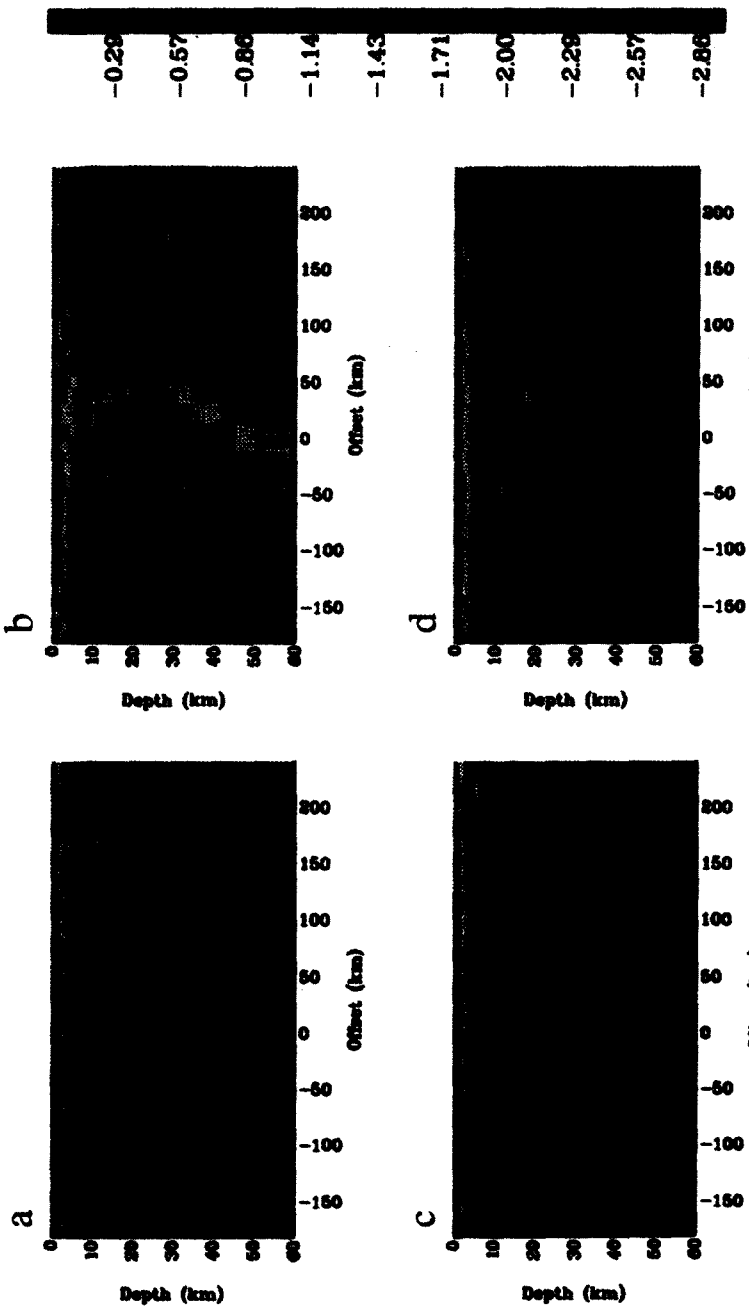

Fig. 7. The l_1 norm models obtained by inverting (a) determinant average, (b) TE, (c) TM and (d) joint TE and TE modes for the COPROD2 data.

$\phi_d = 2524$. The model achieved after seven iterations is shown in Fig. 7(b). It has a misfit of $\phi_d = 804$. Again, the greatest misfits occur at 5.33 and 682 s but the fit is distributed much more uniformly than was the fit for the determinant average data. The summary misfit plots are given in Figs. 8(b) and 9(b). The principal difference between the models in Fig. 7(b) and 7(a) is an enhancement of the magnitude of the NACP conductivity by a factor of five and the extension of the anomaly to greater depths. In addition, the inversion of the TE data has increased the physical size of the TOBE anomaly and generated a model with somewhat greater roughness.

The starting misfit for the TM mode inversion was $\phi_d = 2323$. The model achieved after 21 iterations, which had a misfit of $\phi_d = 877$, is shown in Fig. 7(c). There is no manifestation of a conductive NACP anomaly, even though that feature was well defined in the starting model. The other major difference between Figs. 7(c) and 7(a) is the character of the TOBE anomaly. Fig. 7(a) models the TOBE anomaly as a strong conductor adjacent to a large resistor. Fig. 7(c), however, models the anomaly as two separate conductors of different strength and there is no large resistor. The misfits are shown in Figs. 8(c) and 9(c). Again, the largest misfits are associated with the

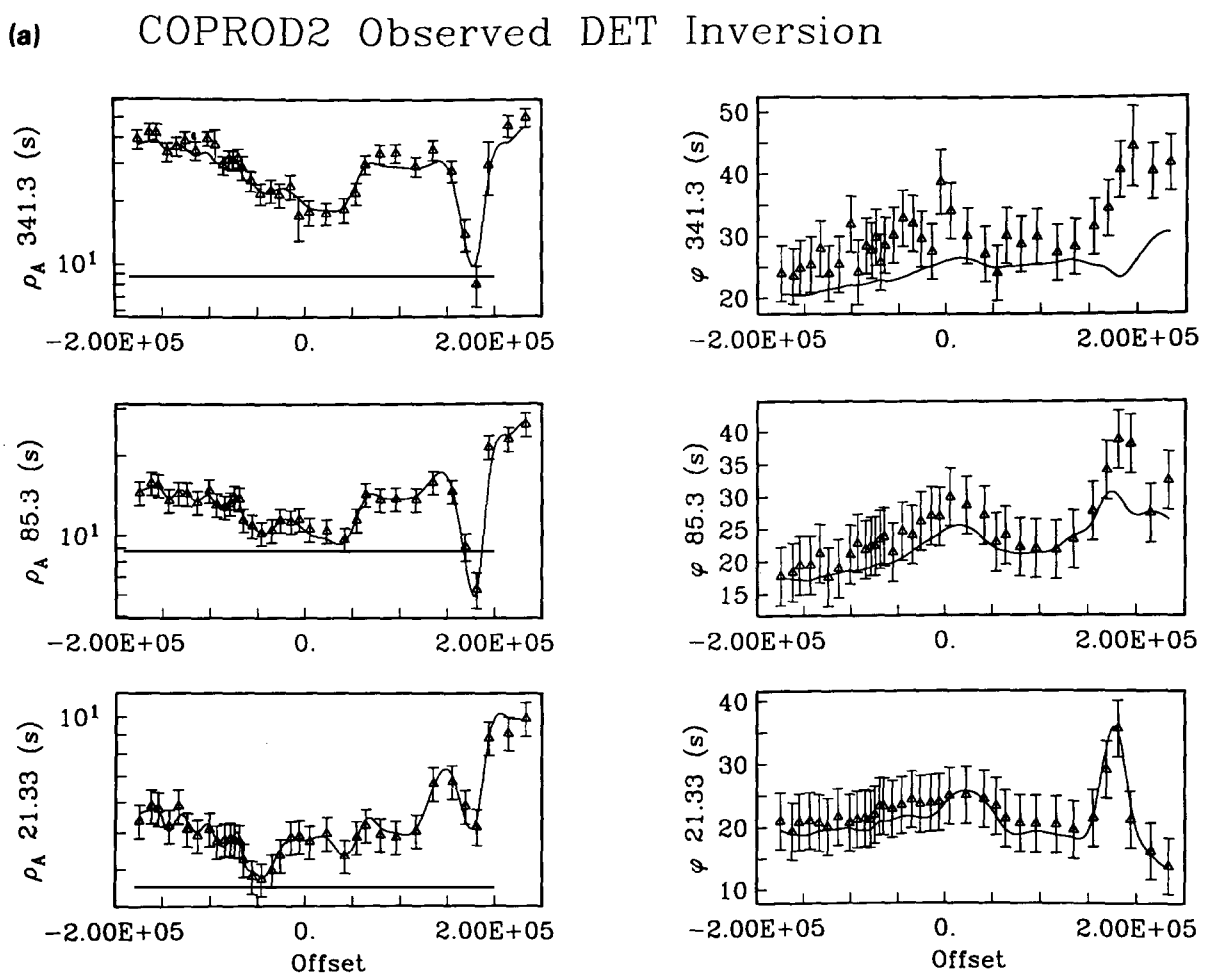

Fig. 8. The observed data (points with error bars) and the predicted data (curves) associated with the l_2 norm models produced by inverting (a) determinant average, (b) TE, (c) TM and (d) joint TE and TM modes for the COPROD2 data.

data at the shortest and longest periods and the phases are particularly erroneous. The phase misfits at 5.33 and 682 s account for 480 of the total misfit of 877.

The large differences observed between the determinant data inversions and separate inversion of the two modes is a strong motivation for carrying out a joint inversion. Several attempts were made to do this but the algorithm was not successful in finding an adequately fitting model. Using the same model norm and reference model as in the determinant data inversion and beginning with the model in Fig. 7(a), the smallest misfit achieved by the joint inversion was about 4500. It appears that the combination of only 1D sensitivities and the restrictions implicit in a subspace method have conspired to prevent the algorithm from working satisfactorily in the joint inversion. This might be expected, as the incorporation of 1D E fields in Eq. (4) means that the sensitivities for both TE and TM data are the same. This makes it extremely difficulty to model the mode splitting which is observed in the data. As in the l_1 inversion, we therefore use an improved sensitivity obtained by substituting the 2D E fields into Eq. (4). The resulting model, obtained with the algorithm presented here, but extracted from the study by Ellis et al. (1993), is shown in Fig. 7(d). Three well-defined conductors describe the NACP anomaly and there is a strong

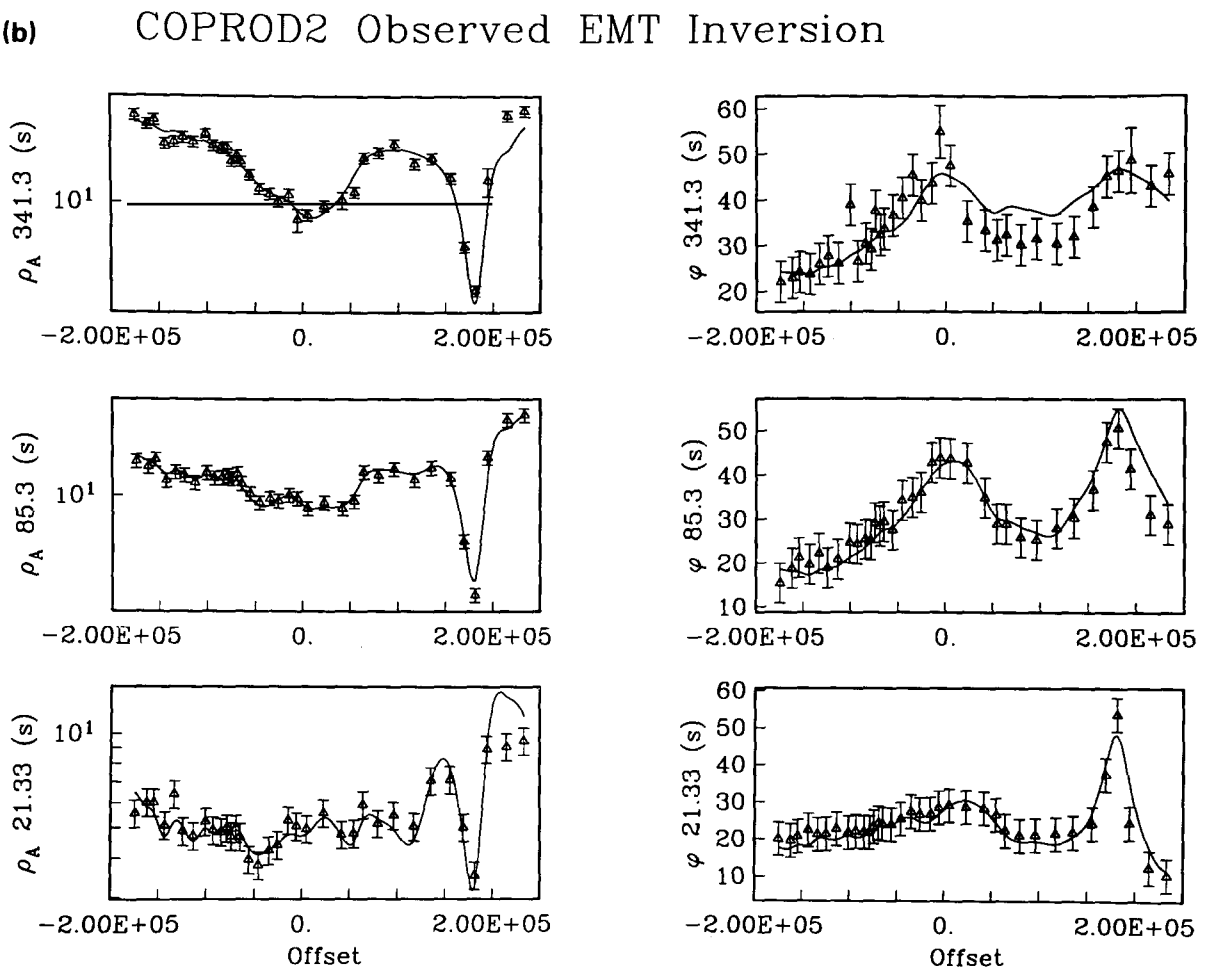

Fig. 8 (continued).

conductor corresponding to the TOBE anomaly. The inversion began with the model in Fig. 7(a) and used the same reference model as in that inversion. The final misfit is $\phi_d = 3200$, which is higher than desired, but again most of the misfit occurs at the longest period. The misfits are summarized in Figs. 8(d) and 9(d).

9. Discussion

The inversion of MT data to produce meaningful information about the Earth's conductivity structure is a difficult task. Any inversion method involves a sequence of decisions by the researcher so that a single model can be constructed. The main decision points concern the data to be inverted and the model objective function to be minimized. Processing and acquisition difficulties and 3D effects make it difficult in general to specify data and their error statistics. Yet when an inversion is carried out, the data, estimated errors and desired fit criterion are specified at the outset and assumed fixed. The next major hurdle faced by the researcher is to specify the model objective function to be minimized. Is a 'blocky' or 'smooth' model desired? Is there a base model, and if so, what is it? If a final model is to be close to a base model and yet smooth, what is the relative ratio for these quantities?

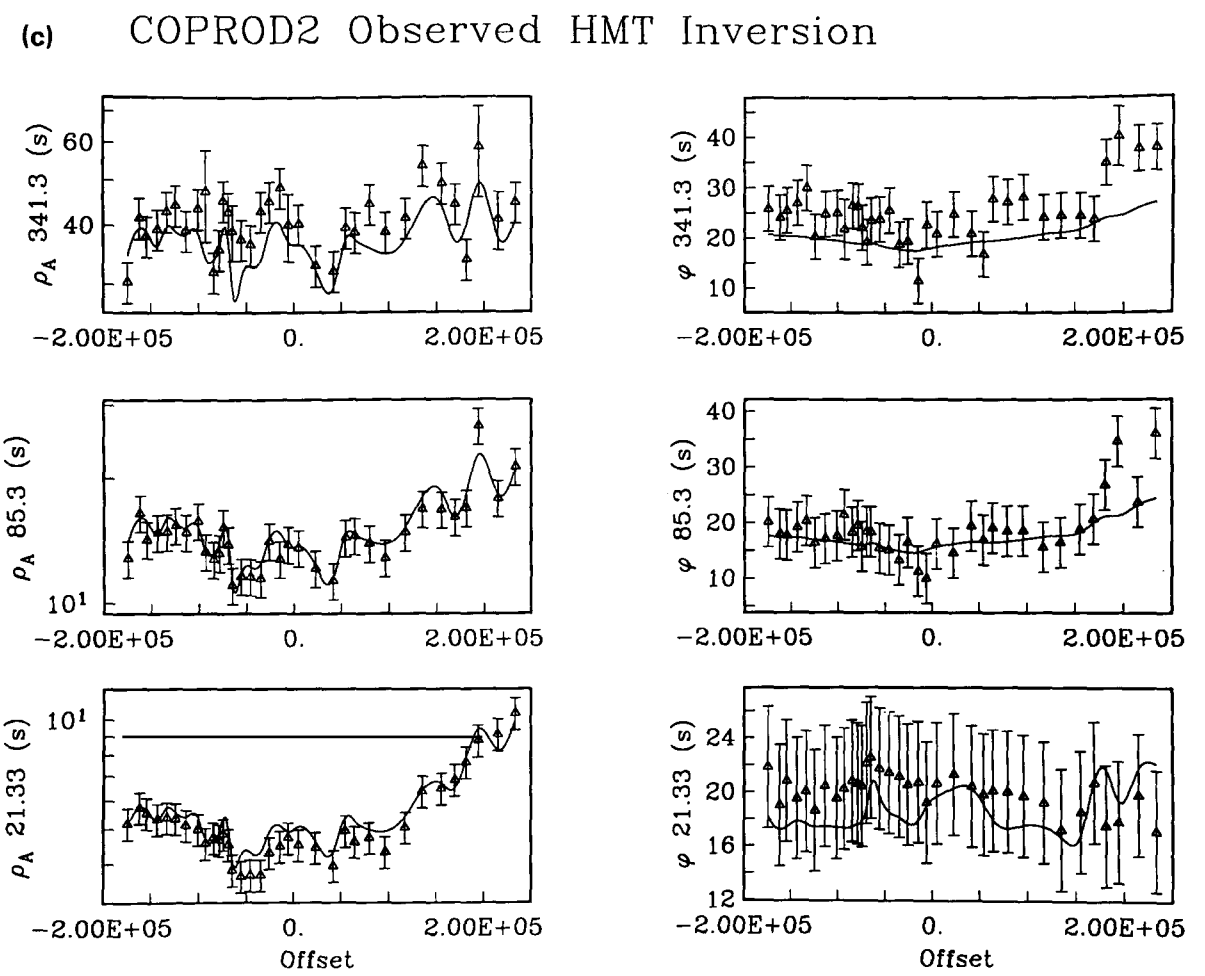

Fig. 8 (continued).

What are the additional multiplicative functions in the norm components needed to obtain the desired results (e.g. progressively more lateral smoothing with increased depth?). Alteration of any of the above factors dictates that the data be reinverted. Consequently, a particular data set must ultimately be inverted not once, but many times, and practicality demands that the inversion computation be efficient. This has been the goal of this paper. First, with respect to the data, we advocate working with determinant average data, at least in the formative stages of the inversion. Not only does this keep the problem smaller but it avoids the difficulties associated with inaccurate rotations of the impedance tensor and perhaps other processing difficulties as well. Also, in some of our inversions, the conductivity models from inverting determinant average data provided acceptable fits to TE and TM mode responses. In such cases, there is little to be gained from carrying out a joint inversion. Additionally, the inversion result from determinant average data can be used as the starting model and/or the reference model for a joint inversion if one is to be carried out. Exploration of the effects of details of model objective function on the inversion result can be explored by inverting determinant data. We have demonstrated that very different models can be produced by altering the objective function to be minimized. This explo-

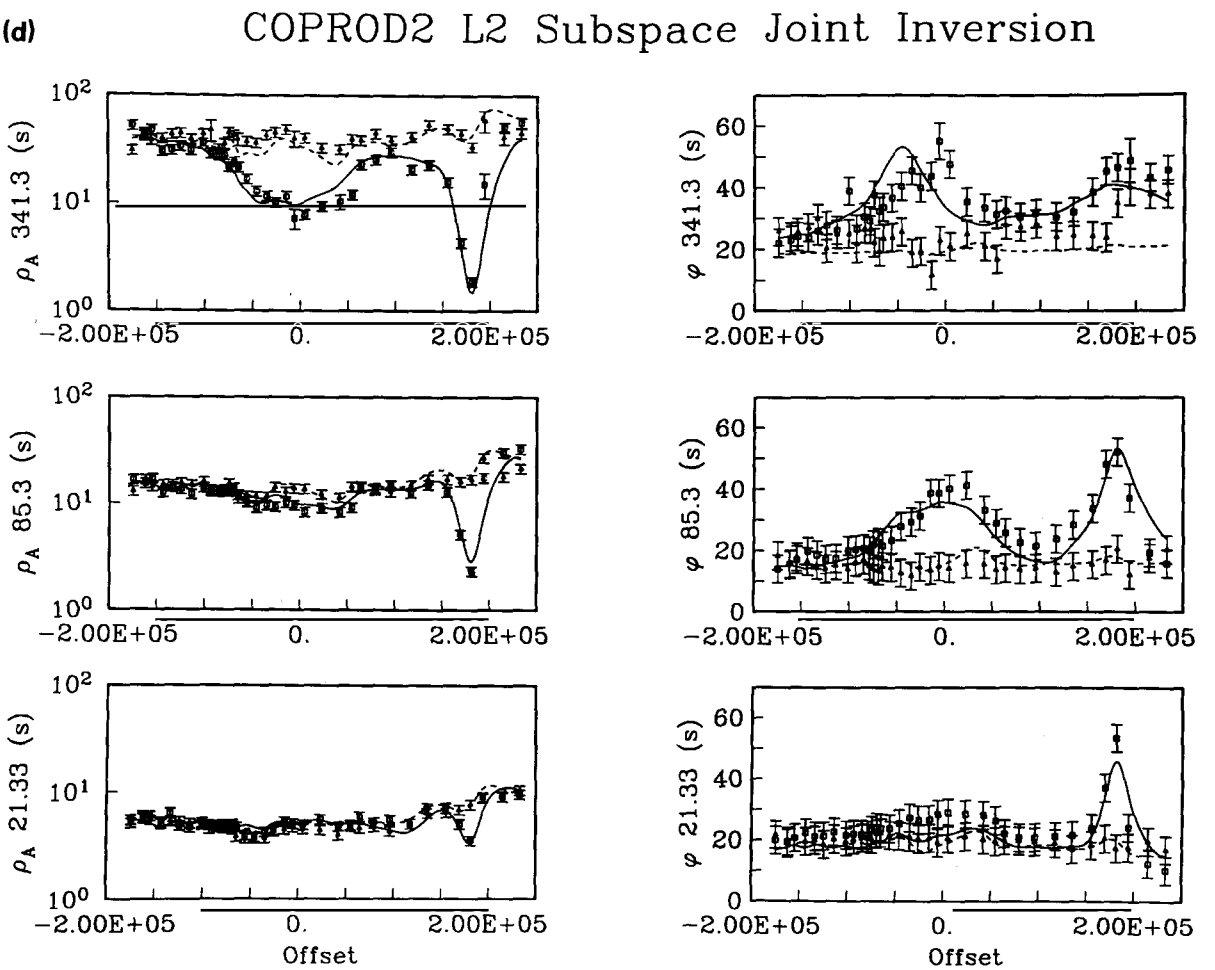

Fig. 8 (continued).

ration will provide the researcher with a greater understanding of the resolving power of the data and will also be useful in designing a final objective function to produce a conductivity model which best satisfies his or her personal prejudices and geologic intuition. There is no substitute for this exploration.

Numerical efficiency in both the l_1 and l_2 algorithms has been achieved through the use of 1D sensitivities. These sensitivities are surprisingly beneficial in 2D problems, at least if misfits of about 5% on the apparent resistivities and 2° on the phase are adequate. The fact that the 1D sensitivities work as well as they do suggests that rather crude approximations to 2D sensitivities may work extremely well. We note that for the successful joint inversion of the COPROD2 data, which have very different TE and TM mode responses, it is necessary to modify the purely 1D approximation to include the 2D field variation. The success of this simple improved approximation motivates further research into developing even better approximations which are still computationally less demanding than carrying out an accurate 2D linearization.

The other important aspect of numerical efficiency for the l_2 solution is the implementation of a sub-space approach to avoid the inversion of a large matrix. The crucial aspect here is the choice of basis vectors. There is still research to be carried out in this area. In this paper we have introduced a dynamic mode where the types of search vectors vary with iteration, changing from horizontal strips to subdivisions of localized areas of interest. The results were encouraging but we were not always able to achieve the desired target

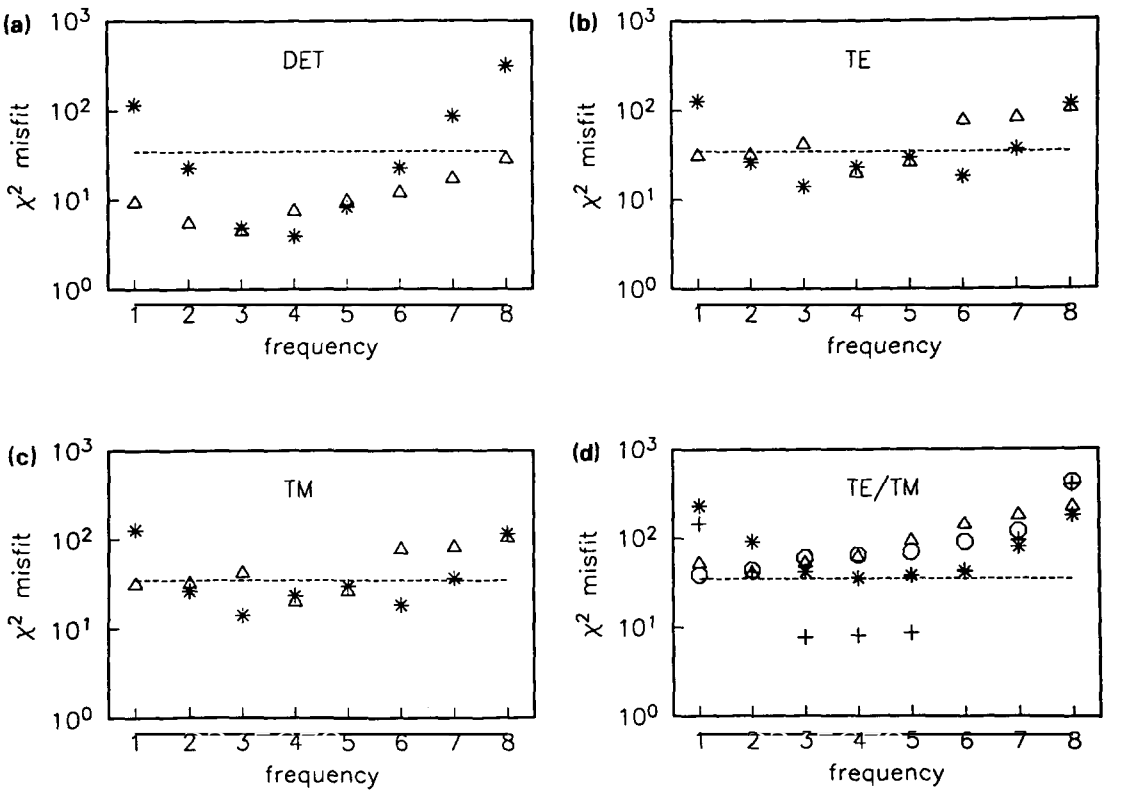

Fig. 9. Summary plots of the misfits. (a)–(d) Misfits for the models shown in Figs. 7(a)–7(d), respectively. In (a)–(c): △, cumulative amplitude misfits for all stations at each frequency; *, similar information for the phase misfits. In (d): △, TE amplitude; *, TE phase; ○,TM amplitude; +, TM phase.

misfit, and it is likely that the constructed model still contains some unnecessary roughness. The fact that the l_1 inversion, in which each cell is a basis vector, was able to generate somewhat lower misfits than that achieved with the sub-space approach suggests that the choice of search vectors may not have been optimum. On the other hand, the large improvement in the inverted results obtained by using the same sub-space method but substituting the 2D electric fields suggests that the focus for improving the inversion should be on obtaining better, but still approximate and quickly computable, sensitivities. These caveats aside, the algorithm presented here is still extremely valuable in performing many inversions of the data and exploring the effects of the details of the objective function on the recovered model.

The main computations for the algorithms presented here are connected with forward modelling. An average of about four forward modellings per iteration for the line search were used in the l_2 inversion and 8–20 iterations were required to achieve a final result. With the matrix to be inverted having a dimension of 100 or less, and with the sensitivities easily computed, the time required to carry out the inversion is only fractionally longer than the time required for the forward modellings. The l_1 solution is even more efficient because only one forward modelling is carried our per iteration. This requires, however, that a good strategy is implemented for specifying the value of the Lagrange multiplier at each iteration. Typically, 15–30 forward modellings are required for each l_1 inversion. The computational time for the l_1 solution is divided between the forward modelling and the need to solve a large sparse linear programming problem.

In conclusion, we emphasize that the fundamental difficulty with geophysical inverse problems is that they are ill-posed. Of particular concern is their inherent non-uniqueness, which can only be addressed by considering the class of models which fit the data to the desired level. One method for exploring the class of models is by performing a significant number of inversions with different model norms, and this can only be done if inversion algorithms are flexible and efficient. The algorithms presented in this paper were designed with flexibility and efficiency as a highest priority, and we have shown that they provide the means for a preliminary investigation of model space.

Acknowledgements

We thank Yaoguo Li for this assistance in programming the sub-space algorithm and A.G. Jones for supplying the COPROD2 data. We thank the University of Waterloo for making available the MAT1 Sparse Matrix Solver (D'Azevedo et al., 1991). Finally, we thank Torquil Smid and Greg Neumann for their comments on the original version of the paper.

References

Berdichevsky, M.N. and Dimitriev, V.I., 1976. Basic principles of interpretation of magnetotelluric sounding curves. In: A. Ádám (Editor), Geoelectric and Geothermal Studies. KAPG Geophysical Monograph. Akademiai Kiado, Budapest, pp. 165–221.

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

D'Azevedo, E.F., Kightley, J.R. and Forsyth, P.A., 1991. MAT1 Iterative Sparse matrix Solver: User's Guide. Department of Computer Science, University of Waterloo.

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

Dosso, S.E. and Oldenburg, D.W., 1989. Linear and nonlinear appraisal using extremal models of bounded variation. Geophys. J. Int., 99: 483–495.

Ellis, R.G., Farquharson, C. and Oldenburg, D.W., 1993. AIM Inversion of the COPROD2 data. J. Geomagn. Geoelectr., 45.

Jones, A.G., 1988. Static shift of magnetotelluric data and its removal in a sedimentary basin. Geophysics, 53: 967–978.

Jones, A.G. and Savage, P.J., 1986. North American Central Plains conductivity anomaly goes east. Geophys. Res. Lett., 13: 685–688.

Kennett, B.L.N. and Williamson, P.R., 1988. Subspace methods for large-scale nonlinear inversion. In: N.J. Vlaar, G. Nolet, M.J.R. Wortel and S.A.P.L. Cloetingth (Editors), Mathematical Geophysics: A Survey of Recent Developments in Seismology and Geodynamics. D. Reidel, Dordrecht, pp. 139–154.

Madden, T.R., 1972. Transmission systems and network analogies to geophysical forward and inverse problems. Tech. Rep. 72–3, Department of Earth and Planetary Sciences, Massachusetts Institute of Technology, Cambridge, MA.

Marsten, R.E., 1981. The design of the XMP linear programming library, ACM Trans. Math. Software, 7: 881–899.

Oldenburg, D.W., 1979. One dimension inversion of natural source magnetotelluric observations. Geophysics, 44: 1218–1244.

Oldenburg, D.W. and Ellis, R.G., 1991. Inversion of geophysical data using an approximate inverse mapping. Geophys. J. Int., 105; 325–353.

Oldenburg, D.W., McGillivray, P.R. and Ellis, R.G., 1993. Generalized subspace methods for large scale inverse problems. Geophys. J. Int., 114: 12–20.

Parker, R.L. and McNutt, M.K., 1980. Statistics for the one-norm misfit measure. J. Geophys. Res., 85: 4429–4430.

Skilling, J. and Bryan, R.K., 1984. Maximum entropy image reconstruction. general algorithm. Mon. Not. R. Astron. Soc., 211; 111–124.

Smith, J.T., 1988, Magnetotelluric inversion of multi-dimensional magnetotelluric data. Ph.D. Dissertation, University of Washington, Seattle.

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

Smith, J.T. and Booker, J.R., 1991. Rapid inversion of two- and three-dimensional magnetotelluric data. J. Geophys. Res., 96: 3905–3922.