ELSEVIER

Journal of Applied Geophysics 42 (1999) 71–80

JOURNAL OF APPLIED GEOPHYSICS

www.elsevier.nl/locate/jappgeo

An approximate inversion algorithm for time-domain electromagnetic surveys

Colin G. Farquharson *, Douglas W. Oldenburg, Yaoguo Li [1]

UBC-Geophysical Inversion Facility, Department of Earth and Ocean Sciences, University of British Columbia, 2219 Main Hall, Vancouver, BC, V6T 1Z4 Canada

Received 9 September 1998; accepted 2 June 1999

Abstract

Time-domain electromagnetic surveys typically comprise numerous soundings, and any useful interpretation procedure must be rapid enough to cope with the subsequent large amounts of data. Even though geological targets invariably display a degree of three-dimensionality, it is often possible to obtain information about their structure from an Earth model constructed from the results of one-dimensional inversions of each sounding. We derive from this process an approximate inversion procedure. The observations from all soundings are averaged to generate a representative data-set which is then inverted using a rigorous one-dimensional algorithm to produce a layered background model. As by-products of the inversion, the sensitivities for the background model are available, as well as the value of the trade-off parameter in the objective function being minimized. Model updates, which do not involve re-calculating the full sensitivities, are then carried out for each sounding. Each update requires only the solution of a small matrix equation, and a few forward modellings. Two or three updates generally result in a marked decrease in the objective function, and hence an improvement in the model for that sounding. The technique is illustrated with synthetic data generated from a three-dimensional model, and with field data collected in Venezuela. The inversion procedure is tailored to produce piecewise-constant models, and to use a robust measure of data misfit. For the field example, rigorous one-dimensional inversions provide the model for comparison. The approximate inversion is shown to provide much of the same information, but in a substantially reduced amount of time.

Keywords: Electromagnetic surveys; Inverse problem; Environmental geophysics; Mineral exploration

1. Introduction

A geophysical time-domain electromagnetic (TEM) survey aims to provide information about the electrical conductivity of the Earth. The conductivity in turn provides information about the composition of the subsurface, for example, the presence of conductive geological formations, of groundwater and its salinity, or of metallic debris.

To distinguish horizontal variations in the conductivity, a TEM survey will consist of numerous soundings distributed over the area of interest. A single TEM sounding involves measuring the magnetic field (or more usually its

* Corresponding author. Tel.: +1-604-822-4318; fax: +1-604-822-6047; E-mail: farq@geop.ubc.ca

[1] Now at: Department of Geophysics, Colorado School of Mines, Golden, CO 80401, USA.

0926-9851/99/$ - see front matter
PII: S0926-9851(99)00023-3

rate of change with time) associated with the diffusion and decay of electric currents induced in the subsurface by the termination of the current flowing in a transmitter loop (for a review of the TEM method in geophysics, see Nabighian and Macnae, 1987).

Ideally, all the soundings in a survey would be inverted together to produce a three-dimensional model of the Earth. This is not currently practical because of the large amounts of data commonly collected in a survey. One must therefore use short-cuts and approximations to speed up a multi-dimensional inversion algorithm (e.g., Christensen, 1997), or fuse together one-dimensional models obtained from each sounding. It is the latter option we pursue here.

Performing rigorous one-dimensional inversions of all soundings is quicker than a three-dimensional inversion, and results in a faithful representation of the Earth provided large lateral variations are not present (Newman et al., 1987; Auken, 1995). However, even this approach can require too much time, meaning that approximations have to be introduced to the one-dimensional inversions.

A number of rapid interpretation schemes have been published that generate a one-dimensional model for each sounding. Imaging techniques have been proposed (e.g., Macnae and Lamontagne, 1987; Nekut, 1987; Eaton and Hohmann, 1989; Fullagar, 1989) that map the voltage or magnetic field measurements into a conductivity-depth pseudo-section. Smith et al. (1994) combine a similar mapping technique with an approximate inversion procedure, and Christensen (1995) has developed an approximate inversion procedure that uses the Born approximation and approximations of the Fréchet derivatives. The goal of all these procedures is to produce a useful image of the subsurface in as short a time as possible. Here we propose a procedure that is also based on approximate one-dimensional inversions of each sounding, but which tries to retain many of the features of a rigorous inversion, and, as such, balances the need for speed with the desire to generate an accurate and reliable image of the subsurface.

We first provide a synopsis of our approximate inversion procedure, and then describe its various parts in more detail. We illustrate its performance by applying it to a synthetic data-set generated from a simple three-dimensional model, and to a field data-set collected in Venezuela.

2. The approximate inversion procedure

The most efficient way of performing a one-dimensional inversion of every sounding in a survey would be to start each inversion from a background model that is a good representation of the average layered structure of the region. This should minimize the number of iterations required to obtain the solution for each sounding, which in turn minimizes the number of times the sensitivities have to be computed and forward modellings carried out. It is these two operations that consume computation time. Going one stage further, a partial solution involving only a few judiciously chosen updates to the background model could be expected to provide much of the information contained in the complete solution, but at a reduced cost. This argument forms the basis for the approximate procedure we present here.

The approximate inversion procedure is as follows. Each stage is described in detail in the subsequent sections.

(1) The specific forms of the measures of data misfit and model character that comprise the objective function must be chosen.

(2) The observations are averaged over all soundings to create a representative voltage-vs.-time data-set.

(3) The representative data-set is inverted to produce a layered background model. The Jacobian matrix of sensitivities (or Fréchet derivatives) for the background model, and the value of the trade-off parameter in the objective function are saved and reused in stage 4.

(4) For each sounding:
(i) compute an update to the background model, for example, the linearized update or the AIM update described in this paper;
(ii) determine the step length that reduces the objective function the most, or, if no such step length can be found, exit this loop and proceed to the next sounding;
(iii) compute the new model, and, if linearized updates are being used, partially recompute the sensitivities;
(iv) repeat (i)–(iii) a prescribed number of times, or until the objective function can no longer be decreased for the sounding under consideration.

(5) The one-dimensional models obtained for each sounding are fused together to produce the requisite two- or three-dimensional image of the subsurface, and the misfits obtained for each sounding during the inversion procedure are plotted to indicate the confidence one should have in features present in the image.

2.1. The measures of data misfit and model character

In the inversion of the representative data-set, and during the model updates at each sounding, we strive to minimize the objective function

$$\Phi = \phi_{\mathrm{m}}(\mathbf{m}) + \beta\phi_{\mathrm{d}}(\mathbf{m}), \tag{1}$$

where $\mathbf{m}$ is the vector of model parameters (which we take to be the logarithms of the layer conductivities), ϕ_{m} is some measure of model character, ϕ_{d} is a measure of data misfit, and β is the trade-off parameter that balances their respective effects. Before the inversion procedure can begin, the specific forms of these two measures have to be chosen. There are many possibilities. Constructing ϕ_{m} to be a function of the spatial derivatives of the model parameters will result in a model which contains a minimum amount of structure. Furthermore, if ϕ_{m} is a sum-of-squares measure, or l_2-norm, a smooth, smeared-out model will be produced, whereas if ϕ_{m} is an l_1-norm, a blocky, piecewise-constant model will result. For ϕ_{d}, a sum-of-squares measure, or l_2-norm, is appropriate if the noise in the observations is Gaussian: if outliers are present, a robust measure such as an l_1-norm is more suitable.

2.2. The representative data-set

The next stage of the approximate inversion procedure is to obtain a voltage-vs.-time data-set that is representative of the majority of soundings in the survey, and which will give rise to an appropriate background model when inverted. This data-set must comprise representative values of the measurement uncertainties (which are assumed to be available for each sounding) as well as of the voltages themselves so that the inversion of this data-set behaves in a manner similar to an inversion of a typical sounding. The median values of the voltages at each time for all soundings are used as the representative voltages. The representative uncertainties are given by the median values of the measurement uncertainties. This gives a truer reflection of the size of the uncertainty in a typical sounding than the width of the distribution of the observed voltages about their median value.

2.3. One-dimensional inversion of the representative data-set

We invert the representative data-set using an algorithm developed previously (Farquharson and Oldenburg, 1993, 1998) and which we only summarize here. The Earth is discretized into several tens of layers, with the logarithms of the conductivities of the layers being sought in the inversion while the layer boundaries remain fixed. The techniques presented in this paper are applicable to any survey configuration. Here we consider a rectangular transmitter loop, along with measurements of the voltage induced in a horizontal receiver loop located anywhere on the surface. The inversion proceeds by minimizing the objective function Φ given in Eq. (1). Here we use the l_1-norm of the first-order finite-difference operator applied to $\mathbf{m}$ for ϕ_{m}, and a Huber M-measure for ϕ_{d}. These choices result in a piecewise-constant model and a robust fit to the observations.

The inverse problem is nonlinear. It is solved by constructing an iterative procedure in which a linearized approximation to the full nonlinear problem is solved at each iteration. At the j-th iteration the following equation has to be solved for the new model $\mathbf{m}^{(j)}$:

$$\left[\underline{\mathbf{W}}_{\mathrm{z}}^{\mathrm{T}}\underline{\mathbf{R}}_{\mathrm{z}}\underline{\mathbf{W}}_{\mathrm{z}} + \beta\underline{\mathbf{J}}^{\mathrm{T}}\underline{\mathbf{W}}_{\mathrm{d}}^{\mathrm{T}}\underline{\mathbf{R}}_{\mathrm{d}}\underline{\mathbf{W}}_{\mathrm{d}}\underline{\mathbf{J}}\right]\mathbf{m}^{(j)} = \beta\underline{\mathbf{J}}^{\mathrm{T}}\underline{\mathbf{W}}_{\mathrm{d}}^{\mathrm{T}}\underline{\mathbf{R}}_{\mathrm{d}}\underline{\mathbf{W}}_{\mathrm{d}}\left(\mathbf{d}^{\mathrm{obs}} - \mathbf{d}^{(j-1)} + \underline{\mathbf{J}}\mathbf{m}^{(j-1)}\right), \qquad (2)$$

where $\underline{\mathbf{W}}_{\mathrm{z}}$ is the first-order finite-difference operator, $\underline{\mathbf{W}}_{\mathrm{d}}$ is the diagonal matrix of the reciprocals of the measurement uncertainties, $\mathbf{d}^{\mathrm{obs}}$ is the vector of observations, and $\mathbf{d}^{(j-1)}$ is the forward-modelled data for $\mathbf{m}^{(j-1)}$, the model from the previous iteration. The matrix $\underline{\mathbf{J}}$ is the Jacobian matrix of sensitivities with respect to the logarithm of the layer conductivities, and given by $J_{mn} = \sigma_n^{(j-1)}\partial d_m/\partial\sigma_n(\mathbf{m}^{(j-1)})$, where $\sigma_n^{(j-1)}$ is the conductivity of the n-th layer in the model from the previous iteration, and $\partial d_m/\partial\sigma_n(\mathbf{m}^{(j-1)})$ is the sensitivity of the m-th datum with respect to the conductivity of the n-th layer in that model. When ϕ_{m} and ϕ_{d} are not l_2-norm measures, the diagonal matrices $\underline{\mathbf{R}}_{\mathrm{z}}$ and $\underline{\mathbf{R}}_{\mathrm{d}}$ are required (Farquharson and Oldenburg, 1998). They depend on $\mathbf{m}^{(j)}$, and hence Eq. (2) is a nonlinear system of equations. However, $\mathbf{m}^{(j)}$ can be obtained using an iteratively re-weighted least-squares procedure (e.g., Holland and Welsch, 1977; Gersztenkorn et al., 1986; Ekblom, 1987). The value of β in Eq. (2) is chosen using a line search so that the misfit is decreased to either half of what it was after the previous iteration, or to its target value. Assuming the uncertainties in the representative data-set are representative of the uncertainties in a typical sounding, the expected value of the data-misfit measure could be used as the target value. However, since the model is to be the starting point for the model updates for all soundings, it can be beneficial if a larger target misfit is used, resulting in a simpler background model. Some user input is therefore required to decide the best value for the target misfit.

2.4. Approximate inversion for each sounding

Once the background model has been obtained, we return to each sounding and carry out a few updates to the model. One possible update is the standard linearized step found within iterative solutions to nonlinear problems, and another is an AIM (approximate inverse mapping) update based on the work of Oldenburg and Ellis (1991) and Li and Oldenburg (1994). The objective function Φ given in Eq. (1) is monitored to ensure an improvement over the background model is obtained. Since sensitivities are expensive to compute, those for the background model are reused for all model updates, although a partial correction is incorporated for the linearized update. For consistency, all the soundings are inverted with the same objective function, using as the value of β its final value in the inversion of the representative data-set. This is appropriate if the noise characteristics of the observations, as well as the subsurface conductivity, do not vary greatly across the survey.

2.4.1. Linearized updates

The first model update that we consider has the same form as the step taken at each iteration in the iterative, linearized inversion of the representative data-set. The j-th update in a sequence of these model updates gives:

$$\mathbf{m}^{(j)} = \mathbf{m}^{(j-1)} + \alpha\left(\mathbf{m}_{\mathrm{trial}}^{(j)} - \mathbf{m}^{(j-1)}\right), \qquad (3)$$

where $\mathbf{m}_{\mathrm{trial}}^{(j)}$ is the solution to Eq. (2). The Jacobian matrix is partially re-computed for each update: the sensitivities used in Eq. (2) for the computation of $\mathbf{m}_{\text{trial}}^{(j-1)}$ are given by $J_{mn} = \sigma_n^{(j-1)} \partial d_m / \partial \sigma_n (\mathbf{m}^{\text{b}})$, where $\sigma_n^{(j-1)}$ is the conductivity of the n-th layer in the model from the $(j-1)$-th update, and $\partial d_m / \partial \sigma_n (\mathbf{m}^{\text{b}})$ is the sensitivity of the m-th datum with respect to the conductivity of the n-th layer in the background model. For each update a number of values of the step length $\alpha \in [0,1]$ are tried, with the chosen step being the one that decreases the objective function Φ by the greatest amount. If the objective function does not decrease for any of the step lengths tried, the model obtained from the previous update remains as the model for that sounding. The first update is applied to the background model, that is, $\mathbf{m}^{(0)} = \mathbf{m}^{\text{b}}$.

2.4.2. AIM updates

The second model update that we consider is an AIM update as developed by Oldenburg and Ellis (1991). We provide only a brief qualitative description here: for a derivation and extensive discussion of the AIM technique we refer the reader to the aforementioned paper.

An AIM update essentially tries to counter the failings of an approximate inverse operator. The approximate inverse operator we use corresponds to the first linearized model update (with $\alpha = 1$) described in Section 2.4.1. This update can be expressed as $\mathbf{m}^{(1)} = \tilde{F}^{-1}[\mathbf{d}^{\text{obs}}]$ where $\tilde{F}^{-1}$ is the approximate inverse operator:

$$\tilde{F}^{-1}\left[\mathbf{d}^{\text{obs}}\right] = \beta\left[\underline{\mathbf{W}}_{\text{z}}^{\text{T}}\underline{\mathbf{R}}_{\text{z}}\underline{\mathbf{W}}_{\text{z}} + \beta\underline{\mathbf{J}}^{\text{T}}\underline{\mathbf{W}}_{\text{d}}^{\text{T}}\underline{\mathbf{R}}_{\text{d}}\underline{\mathbf{W}}_{\text{d}}\underline{\mathbf{J}}\right]^{-1} \times \underline{\mathbf{J}}^{\text{T}}\underline{\mathbf{W}}_{\text{d}}^{\text{T}}\underline{\mathbf{R}}_{\text{d}}\underline{\mathbf{W}}_{\text{d}}\left(\mathbf{d}^{\text{obs}} - \mathbf{d}^{\text{b}} + \underline{\mathbf{J}}\mathbf{m}^{\text{b}}\right). \tag{4}$$

An AIM update to $\mathbf{m}^{(1)}$ tries to quantify and remove the discrepancy between this model and the one that would have resulted had the true nonlinear inverse operator been used. Suppose one linearized update has been applied to the background model to give $\mathbf{m}^{(1)}$, and that the corresponding forward-modelled data $\mathbf{d}^{(1)}$ have been computed. If the true inverse operator had been used to obtain $\mathbf{m}^{(1)}$ instead of the approximate operator, $\mathbf{d}^{(1)}$ would be equal to the observations $\mathbf{d}^{\text{obs}}$. The difference between $\mathbf{d}^{(1)}$ and $\mathbf{d}^{\text{obs}}$ is therefore due to the failings of the approximate operator. Assuming this difference would be the same for any $\mathbf{m}^{(1)}$, subtracting it from the observations will produce a vector $\tilde{\mathbf{d}}$ which, when acted upon by the approximate inverse operator, will give the desired model. The model resulting from the AIM update is therefore $\mathbf{m}^{\star} = \tilde{F}^{-1}[\tilde{\mathbf{d}}]$ where $\tilde{\mathbf{d}} = \mathbf{d}^{\text{obs}} + (\mathbf{d}^{\text{obs}} - \mathbf{d}^{(1)})$. See Fig. 1 for a geometrical description of this process.

Just as for the linearized model update, a number of different step lengths $\alpha \in [0,1]$ are actually tried, giving the final model as $\mathbf{m} = \mathbf{m}^{(1)} + \alpha(\mathbf{m}^{\star} - \mathbf{m}^{(1)})$. The chosen step is the one that results in the largest decrease in the objective function Φ. If none of the step lengths decrease Φ, $\mathbf{m}^{(1)}$ is retained as the model for that sounding.

2.5. Plots of model and misfits

Once the model updates have been carried out for every sounding, the resulting one-dimensional models can be fused together to produce the desired two- or three-dimensional image of the subsurface. Plotting the misfits obtained for each sounding is also important: this makes it clear which parts of the image give a good fit to the observations, which parts do not give such a

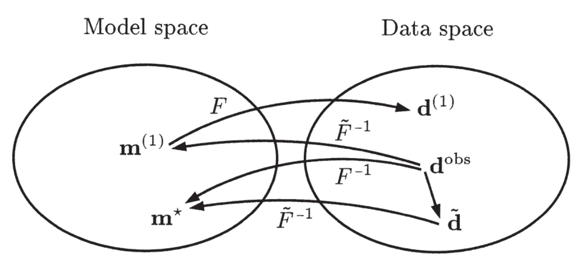

Fig. 1. A geometrical representation of the AIM update in which the new model is given by the approximate inverse operator applied to the corrected data: $\mathbf{m}^{\star} = \tilde{F}^{-1}[\tilde{\mathbf{d}}]$. F and F^{-1} are the true forward and inverse operators.

good fit but are nonetheless an improvement over the background model, and for which parts the background model could not be bettered.

3. Synthetic example

The first example we present uses synthetic data generated by Esben Auken (personal communication) for the same geometry as model b of Auken (1995), but for a conductive block within a resistive halfspace. In particular, a block of 0.1 S m^{-1}, whose centre is at a depth of 40 m and whose extents in the *X*-, *Y*- and *Z*-directions are 100, 500 and 40 m, respectively, in a halfspace of 0.01 S m^{-1} beneath a 20-m thick layer of 0.0333 S m^{-1}. The $Y = 0$ plane (for positive X) through this model is shown in Fig. 2(a). The voltages at 27 times from 0.005 to 2 ms at the centre of a 40×40 m transmitter loop were computed at 13 locations along the positive *X*-axis using the integral equation program of Newman et al. (1986). We then added Gaussian noise that was a combination of two parts: one whose standard deviation increased exponentially from 0.5% at the earliest time to 10% at the latest time, and one part which followed the noise model of Munkholm and Auken (1996) and whose standard deviation varied as $0.4 V_M \sqrt{t_M / t}$ where V_M is the noise-free voltage at the last observation time t_M. The standard deviations of the added noise were used as the measurement uncertainties for the data.

In the inversion of the representative data-set, the expected value of the misfit (equal to 26.4 given the 27 data) was attainable and produced a reasonable background model. The model produced using one AIM update (following an initial linearized update from the background model as described in Section 2.4.2) is shown in Fig. 2(b). (The final model for the sounding at $X = 60$ m gives a good indication of what the background model was since for this sounding the model changed little from the background model.) The misfit at each sounding after the three stages of the approximate inversion procedure (that is, for the background model, after the initial linearized update, and after the AIM update) are shown in Fig. 2(c). It is clear that the misfit has been significantly reduced for all but one sounding. At the apparently aberrant

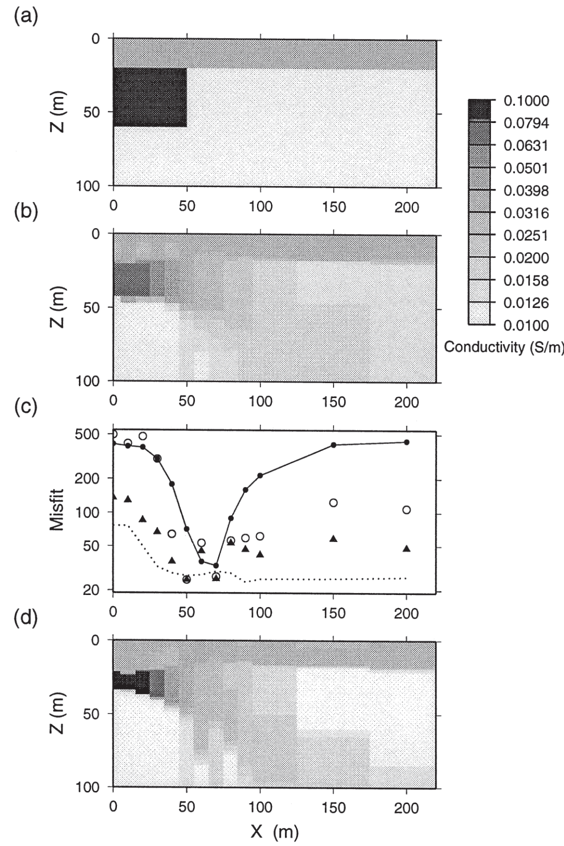

Fig. 2. (a) The model from which the synthetic data-set was created. The conductive block has a total extent of 500 m perpendicular to the page, and is centred beneath the observation locations. The model is symmetric about $X = 0$ m. (b) The model produced by the approximate inversion procedure using an AIM update. (c) The misfit after the three stages of the AIM update procedure: the solid line indicating the misfit for the background model, the open circles the misfit after the linearized update, and the solid triangles the misfit after the AIM update. The dotted line indicates the misfit for the model produced by rigorous one-dimensional inversions of each sounding. (d) The model produced by the rigorous one-dimensional inversions.

sounding, the whole objective function Φ was in fact decreased. Comparison of the model in Fig. 2(b) with that used to generate the data shows that the approximate inversion procedure has managed to obtain the true two-layered nature of the model away from the three-dimensional body, and to obtain, rather well, the shape and extent of this body, and its conductivity. The excessively high resistivity directly beneath the block is an artefact of the assumption of one-dimensionality, not of the approximate form of the model updates. No significant improvements were made to the model or in the misfits with further AIM updates. The inversion to produce the background model required 10 iterations and 25 min on a Sun Ultra 1 workstation. The sequence of updates applied at each sounding required 25 s per sounding.

A comparable model was produced using two linearized updates (with a search over the step length α for both), although the misfits for the soundings after these two updates were not quite as low as those shown in Fig. 2(c) for the AIM update. The misfits were only marginally reduced with more linearized updates.

Included in Fig. 2 (panel d) is the model obtained by rigorous one-dimensional inversions of each sounding. The model shown was produced using an l_1-norm as a measure of both model character and data misfit. The expected value of misfit could not be obtained for the first five soundings, but was for the remainder. The *M*-measure value of misfit for the final model at each sounding is shown by the dotted line Fig. 2(c). The model in Fig. 2(d) clearly displays the artefacts resulting from the assumption of one-dimensionality: the increased resistivity beneath the conductive block, the decreased depth-extent of the block, and a conductive zone extending outwards and downwards from the side of the block. Comparison with the model in Fig. 2(b), and between the corresponding values of misfit, shows that the approximate procedure has produced results very similar to those of the rigorous one-dimensional inversions. Equally as good results were obtained for model b of Auken (1995) in which the block is resistive (0.01 S m^{-1}) and the halfspace conductive (0.1 S m^{-1}).

4. Field example

The second example involves data from a survey carried out in Venezuela by Placer Dome Exploration of Vancouver. One of the main goals of the survey was to determine the thickness of a saprolite layer above Archean basement. One line of 61 soundings was used for this example. The data were collected using a 5×5 m transmitter loop, and a horizontal receiver loop offset 20 m along the survey line from the centre of the transmitter loop. The observed voltages are shown in Fig. 3(a). Estimates of the measurement uncertainties were provided with the voltages.

It was found that a target misfit of 98 in the inversion of the representative data-set resulted in a suitable background model: a smaller value gave a model with too much structure. This value is also commensurate with the values attained in rigorous one-dimensional inversions of each sounding. The model produced using an AIM update after a linearized update with $\alpha = 1$ is shown in Fig. 3(b). The corresponding forward-modelled data are shown in Fig. 3(c). The values of the misfit at the various stages of the approximate inversion procedure are shown in Fig. 3(d). It can be seen that the approximate procedure has managed to decrease the misfit for virtually every sounding. As with the synthetic example, two linearized updates gave misfits almost as low as those shown in Fig. 3(d) for the AIM update, and resulted in a very similar model.

The model obtained from the one-dimensional inversion of each sounding is shown in Fig. 3(e). The forward-modelled data for this model are shown in Fig. 3(f), and the misfits attained at each sounding are shown by the dotted line in Fig. 3(d). There is good agreement between the model produced by the ap-

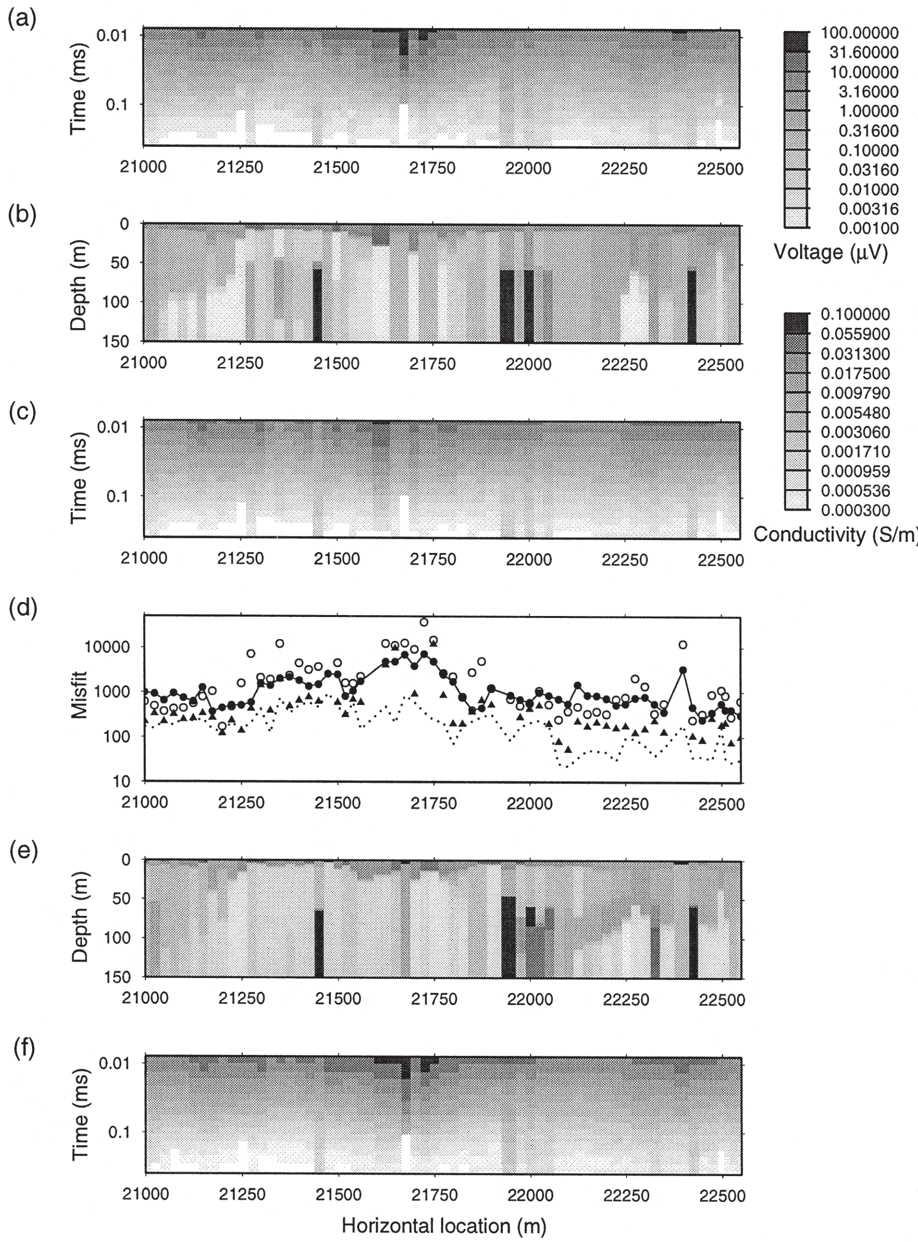

Fig. 3. (a) The observed voltages for the field example. (b) The model obtained using the approximate inversion procedure. (Note the vertical exaggeration of 2:1.) (c) The forward-modelled voltages for this model. (d) The misfit at each sounding during the approximate inversion process: the solid line and solid circles indicate the misfit for the background model, the open circles the misfit after the linearized update, and the solid triangles the final misfit after the AIM update. The dotted line is the misfit for the model shown in panel (e). (e) The model produced by fusing together the results of one-dimensional inversions of each sounding. (f) The forward-modelled voltages for the model in panel (e).

proximate inversion procedure and this model: the continuous surficial layer; the localized, very conductive features (the vertical extents of which have been greatly exaggerated by the one-dimensional nature of the inversion algorithms); and the thickness of the saprolite layer in the left halves of the models (thinning from about 100 m thickness at 21 000 m to 0 m at 21 250 m). However, some of the features present in the right half of Fig. 3(e), notably the structure within the saprolite layer between 21 900 and 22 300 m, are absent from Fig. 3(b). The difference in computation time is significant: 75 min to produce the model in Fig. 3(b) compared to the days required for the model in Fig. 3(e).

5. Discussion

We have presented an approximate inversion procedure for TEM data that is designed to cope with the large numbers of soundings typically collected during a survey. We have tried to construct an efficient procedure that retains many of the features of a rigorous inversion algorithm.

The procedure is considerably quicker than performing rigorous one-dimensional inversions of each sounding. The time-savings occur because only a small number of model updates are performed for each sounding, the sensitivities are not re-computed but are approximated by those for the background model, and a fixed regularization parameter β is used.

In the linearized model update presented in Section 2.4.1, the sensitivities required for each update, which are the sensitivities with respect to the logarithm of the conductivities, are partially re-computed. This gives a significant improvement in the sensitivities because they depend heavily on the scaling by the layer conductivities. This enhances the performance of the linearized updates.

From experience in comparing various approximate updates, one AIM update, as described in Section 2.4.2, often gives the best results, with the results produced by using two or three linearized updates almost as good. Performing more updates does not usually give a significant improvement upon these results. Also from experience, a step length of $\alpha = 1$ or $1/2$ most often gives the largest decrease in the objective function. No more than two or three step lengths need be tried.

Given the nature of the approximate inversion procedure presented here, it will work best when the subsurface is predominantly layered: the representative data-set will then be a good representation of the observations at each sounding, and the model updates will be capable of making the necessary changes to the background model to fit the observations at each sounding. However, the procedure also gives good results when the subsurface is considerably more complicated, as was shown by the examples in this paper.

The variation in the final values of the data misfit between soundings makes it clear where the approximate inversion procedure has done well, and where it has not done so well. This indicates which parts of the final image one can interpret, and which parts are missing features needed to fit the observations.

Finally, the one part of the procedure that we feel we have not been able to satisfactorily address is the choice of the target misfit in the inversion of the representative data-set. If good estimates of the uncertainties are known for the observations at all the soundings, and if the averaging process gives uncertainties in the representative data-set that are good representations of the uncertainties at the majority of soundings, then setting the target misfit equal to its expected value is appropriate. However, if these provisos are not met, then one has to resort to trial-and-error to determine the target misfit that gives a background model that in turn enables the model updates to perform as well as possible. This is the one part of the procedure that is not totally automated.

Acknowledgements

We would like to thank Dr. Esben Auken for providing us with the synthetic data-sets, and Peter Kowalczyk and Michael Ehling of Placer Dome Exploration for supplying us with the field data-set and for related discussions. This work was supported by an NSERC IOR grant and the Consortium for the Joint and Cooperative Inversion of Geophysical and Geological Data (the ''JACI'' Consortium). The following were 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, Newmont Gold, WMC Exploration, and CRA Exploration Party. The authors are grateful for their participation. We would also like to thank Roman Shekhtman for his assistance with all matters computer-related, and Dr. Gregory Newman and an anonymous reviewer for their constructive comments.

References

Auken, E., 1995. 1D time domain electromagnetic interpretations over 2D/3D structures. Proceedings of SAGEEP, Orlando, April, 1995, pp. 329–338.

Christensen, N.B., 1995. Imaging of central loop transient electromagnetic soundings. J. Environ. Eng. Geophys. 0, 53–66.

Christensen, N.B., 1997. Two-dimensional imaging of transient electromagnetic soundings. Proceedings of SAGEEP, Reno, March, 1997, pp. 397–406.

Eaton, P.A., Hohmann, G.W., 1989. A rapid inversion technique for transient electromagnetic soundings. Phys. Earth. Planet. Inter. 53, 384–404.

Ekblom, H., 1987. The L_1-estimate as limiting case of an L_p- or Huber-estimate. In: Dodge, Y. (Ed.), Statistical Data Analysis on the L_1-Norm and Related Methods. Elsevier, Amsterdam.

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

Farquharson, C.G., Oldenburg, D.W., 1998. Nonlinear inversion using general measures of data misfit and model structure. Geophys. J. Int. 134, 213–227.

Fullagar, P.K., 1989. Generation of conductivity-depth pseudosections from coincident loop and in-loop TEM data. Explor. Geophys. 20, 43–45.

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. Commun. Stat.: Theory Methods A6, 813–827.

Li, Y., Oldenburg, D.W., 1994. Inversion of 3-D DC resistivity data using an approximate inverse mapping. Geophys. J. Int. 116, 527–537.

Macnae, J., Lamontagne, Y., 1987. Imaging quasi-layered conductive structures by simple processing of transient electromagnetic data. Geophysics 52, 545–554.

Munkholm, M.S., Auken, E., 1996. Electromagnetic noise contamination on transient electromagnetic soundings in culturally disturbed environments. J. Environ. Eng. Geophys. 1, 119–127.

Nabighian, M.N., Macnae, J.C., 1987. Time domain electromagnetic prospecting methods. In: Nabighian, M.N. (Ed.), Electromagnetic Methods in Applied Geophysics, Vol. 2. Society of Exploration Geophysics, Tulsa.

Nekut, A.G., 1987. Direct inversion of time-domain electromagnetic data. Geophysics 52, 1431–1435.

Newman, G.A., Anderson, W.L., Hohmann, G.W., 1987. Interpretation of transient electromagnetic soundings over three-dimensional structures for the central-loop configuration. Geophys. J. R. Astron. Soc. 89, 889–914.

Newman, G.A., Hohmann, G.W., Anderson, W.L., 1986. Transient electromagnetic response of a three-dimensional body in a layered earth. Geophysics 51, 1608–1627.

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

Smith, R.S., Edwards, R.N., Buselli, G., 1994. An automatic technique for presentation of coincident-loop, impulse-response, transient electromagnetic data. Geophysics 59, 1542–1550.