

Constructing piecewise-constant models in multidimensional minimum-structure inversions

Colin G. Farquharson¹

ABSTRACT

A modification of the typical minimum-structure inversion algorithm is presented that generates blocky, piecewise-constant earth models. Such models are often more consistent with our real or perceived knowledge of the subsurface than the fuzzy, smeared-out models produced by current minimum-structure inversions. The modified algorithm uses l_1 -type measures in the measure of model structure instead of the traditional sum-of-squares, or l_2 , measure. An iteratively reweighted least-squares procedure is used to deal with the nonlinearity introduced by the non- l_2 measure. Also, and of note here, diagonal finite differences are included in the measure of model structure. This enables dipping interfaces to be formed. The modified algorithm retains the benefits of the minimum-structure style of inversion — namely, reliability, robustness, and minimal artifacts in the constructed model. Two examples are given: the 2D inversion of synthetic magnetotelluric data and the 3D inversion of gravity data from the Ovoid deposit, Voisey's Bay, Labrador.

INTRODUCTION

Minimum-structure inversion procedures are ones in which the model parameters being sought are the values of one or more physical properties in the cells in a fixed mesh and in which a measure of the structure in the model is minimized in combination with a measure of data misfit. The mesh should provide as fine a discretization of the subsurface as computational resources allow in order not to restrict possible models. This style of inversion procedure has proved successful, especially for geophysical methods of imaging the earth's subsurface that involve potential or diffusive fields, namely gravity, magnetic, electric, and electromagnetic methods (see, for example, Constable et al., 1987; Smith and Booker, 1988; de Groot-Hedlin and Constable, 1990; Oldenburg and Li, 1994; Li and Olden-

burg, 1996, 1998). The models produced have only enough features to reproduce the observations and have few, if any, artifacts arising from noise in the observations. Of arguably equal importance, minimum-structure inversion procedures generally are robust and reliable.

Traditional implementations of minimum-structure inversion procedures use a sum-of-squares, or l_2 , measure of model structure. This is because minimizing such a measure results in a linear system to be solved. However, the models produced typically have a smeared-out, fuzzy character. This can be at odds with the expected structure of the subsurface, which often comprises (rightly or wrongly) uniform geologic units separated by sharp, distinct interfaces. The purpose of the work presented here is to modify the traditional minimum-structure inversion procedure to give a technique that retains the benefits of the traditional method, but that can generate models with uniform regions separated by sharp interfaces.

The traditional minimum-structure procedure has been modified previously for general, non- l_2 measures using the iteratively reweighted least-squares (IRLS) algorithm. Last and Kubik (1983) minimize a measure of the total cross-sectional area of anomalous regions in 2D inversion of gravity data. Portniaguine and Zhdanov (1999), for 3D inversion of gravity and magnetic data, minimize a measure of the total volume within the model in which the physical property gradient is nonzero. These two approaches favor compact features and compact regions of property variation.

The size of the physical property jump across an interface is not inhibited. Farquharson and Oldenburg (1998) minimize an l_1 measure of the vertical derivative of the model in 1D inversion of electromagnetic data, and Loke et al. (2003) and Farquharson and Oldenburg (2003) minimize l_1 measures of the horizontal and vertical derivatives in 2D inversion of resistivity data. This is analogous to the usual implementation of minimum-structure inversions in which an l_2 measure of the spatial derivatives of the model is minimized. Minimizing an l_1 measure of a spatial derivative favors sharp interfaces over gradual, smeared-out gradations. Also, the size of the physical property jump across an interface does impact this measure, in contrast to those used by Last and Kubik (1983) and Portniaguine

Manuscript received by the Editor 27 June 2007; revised manuscript received 21 August 2007; published online 26 December 2007.

¹Memorial University of Newfoundland, Inco Innovation Centre, and Department of Earth Sciences, St. John's, Canada. E-mail: cgfarquh@mun.ca.

© 2008 Society of Exploration Geophysicists. All rights reserved.

and Zhdanov (1999). Ramos et al. (1999) minimize an entropy measure of the vector comprising the absolute values of the spatial finite differences of the model. They apply their technique to the 2D inversion of resistivity data as well as other geophysical data sets. However, despite the success of the preceding approaches in producing piecewise-constant, blocky models, the interfaces in the models are always either vertical or horizontal: The methods do not want to produce dipping interfaces (Auken and Christiansen, 2004). The work presented here attempts to remedy this problem.

A minimum-structure inversion procedure that uses l_1 measures for both the measure of model structure and the measure of data misfit can be solved via linear programming (see, for example, Dosso and Oldenburg, 1989, for 1D inversion of magnetotelluric data; Oldenburg and Ellis, 1991, for 2D inversion of magnetotelluric data; and van Zon and Roy-Chowdhury, 2006 for 2D and 3D gravity inversion). However, the flexibility afforded by the IRLS procedure to use non- l_2 measures other than the l_1 measure, and to use different measures for the model structure and data misfit, is preferred here.

Smith et al. (1999) and de Groot-Hedlin and Constable (2004) for magnetotelluric (MT) data and Auker and Christiansen (2004) for resistivity data have developed 2D inversion procedures that generate blocky, pseudolayered models. However, the approach of these authors has been to parameterize the subsurface in terms of a small number of cells beneath each observation location, to allow both the conductivities of the cells and the depths of their horizontal boundaries to vary in the inversion and to apply constraints on the lateral variability of the conductivities and depths of the cell boundaries. In some respects, these inversion procedures are hybrids of the minimum-structure approach and a parameter estimation approach; a full minimum-structure approach is of interest here.

The remainder of this paper is arranged as follows. First, the particulars of the modified minimum-structure inversion procedure are described, beginning with the general minimum-structure inversion strategy and continuing with examples of non- l_2 measures and their properties, the iteratively reweighted least-squares procedure for dealing with the nonlinearity introduced by non- l_2 measures, and the inclusion of diagonal finite differences in the measure of model structure for enabling dipping interfaces to be generated. The capabilities of the modified procedure are then illustrated by its application to the inversion of a synthetic 2D MT data set and to the 3D inversion of gravity data from the Ovoid deposit at Voisey's Bay, Labrador.

THEORY

General minimum-structure inversion strategy

The typical framework of a minimum-structure inversion procedure is used here. The subsurface is discretized using a fine mesh comprising uniform cells. The mesh is kept fixed during the inversion, with the values of the physical property (or properties) in the cells being the model parameters to be determined in the inversion. The solution strategy is to find the model parameters that minimize an objective function, which is a combination of a measure of how well the observations are reproduced and a measure of how complicated the model is:

$$\Phi = \phi_d + \beta \phi_m. \quad (1)$$

Here, ϕ_d is the measure of data misfit, having the general form

$$\phi_d = \phi_d(\mathbf{u}), \quad (2)$$

$$\mathbf{u} = \mathbf{W}_d(\mathbf{d}^{\text{obs}} - \mathbf{d}^{\text{prd}}), \quad (3)$$

where \mathbf{d}^{obs} is the vector of observations, \mathbf{d}^{prd} is the vector of data computed for the vector \mathbf{m} , of model parameters, and \mathbf{W}_d is a diagonal matrix whose elements are the reciprocals of the estimates of the standard deviations of the noise in the observations. Also, ϕ_m is a measure of the amount of structure in the model, having the form

$$\phi_m = \sum_k \alpha_k \phi_k(\mathbf{v}_k), \quad (4)$$

$$\mathbf{v}_k = \mathbf{W}_k(\mathbf{m} - \mathbf{m}_k^{\text{ref}}). \quad (5)$$

The summation in ϕ_m is over five terms for the 2D problem considered here and over 14 terms for the 3D problem. This will be treated below. The factor β in the objective function is the regularization, or trade-off, parameter, which controls the relative contributions of the data misfit term and the model complexity term.

General measures

A general form for ϕ_d and ϕ_k is

$$\phi(\mathbf{x}) = \sum_j \rho(x_j), \quad (6)$$

where x_j are the elements of the vector \mathbf{x} , which will be \mathbf{u} or \mathbf{v}_k from above, and the summation is over all elements in the vector. There are numerous possibilities for the specific form of the measure. One example is the l_p -norm:

$$\|\mathbf{x}\|_p^p = \sum_j |x_j|^p, \quad (7)$$

of which the traditional sum-of-squares measure, or l_2 -norm, is a special case. Other examples are the M -measure of Huber (1964), for which

$$\rho(x) = \begin{cases} x^2 & |x| \leq c, \\ 2c|x| - c^2 & |x| > c, \end{cases} \quad (8)$$

where c is the value of x at which the behavior of ρ changes from quadratic to linear and the perturbed l_p -norm measure of Eklblom (1987), for which

$$\rho(x) = (x^2 + \varepsilon^2)^{p/2}, \quad (9)$$

where ε is a small number. This measure is numerically more attractive than the l_p -norm because its derivative exists at $x = 0$ when $p = 1$. Furthermore, Last and Kubik (1983) and Portniaguine and Zhdanov (1999) use the measure with

$$\rho(x) = \frac{x^2}{x^2 + \varepsilon^2}. \quad (10)$$

For small ε , this measure essentially is proportional to the number of nonzero elements in \mathbf{x} and is analogous to the measure of the area or volume of support of a continuous function. There are innumerable other possibilities, allowing a measure to be chosen that produces a model (or data fit) with a specific desirable character. It is also conceivable to construct a measure function ρ that is optimal for a particular situation (Haber and Tenorio, 2003). Figure 1 illustrates how ρ varies for the measures described above.

The l_2 -norm traditionally has been used in inverse problems because its minimization results in a linear system of equations to be solved and because it is the most appropriate measure of misfit if the noise in the observations obeys a Gaussian distribution. However, squaring the elements of a vector in its l_2 -norm means that the contributions of large-valued elements to the norm are disproportionately large. Minimizing an l_2 -norm of a vector therefore results in a vector with no large, distinct elements. When minimizing a misfit, no single discrepancy between the forward-modeled and observed data will be dramatically larger than the others. Outliers will therefore drag the predicted data toward them and away from the remainder of the data set. When minimizing a measure of the spatial derivatives in a model, as is done in most minimum-structure inversions, using an l_2 -norm spreads a change in physical property from one region of the model to another over a number of cells. In other words, an abrupt, significant change from one cell to its neighbor is ruled out. In contrast, large-valued elements of a vector contribute proportionally to an l_1 measure of its size. Consequently, when an l_1 -type measure of misfit is minimized, outliers essentially are ignored. Also, when minimizing an l_1 measure of the spatial derivatives in a model, abrupt changes between regions of uniform physical property are not discriminated against and are, in fact, the kind of structure that is produced naturally.

Iterative solution procedure

Most geophysical inverse problems are nonlinear. The standard approach for this is used here: establishing a procedure, at each iteration of which a linearized approximation of the inverse problem is treated. This involves both the linearization of the relationship between the model parameters and the observed quantities and of the nonlinearity introduced by the use of non- l_2 measures.

The goal at the n th iteration is to find the model \mathbf{m} which minimizes the intermediate objective function:

$$\Phi^n = \phi_d^n + \beta^n \phi_m^n. \quad (11)$$

In this objective function,

$$\phi_d^n = \phi_d(\mathbf{u}), \quad (12)$$

$$\mathbf{u} = \mathbf{W}_d(\mathbf{d}^{\text{obs}} - \mathbf{d}^{n-1} - \mathbf{J}\delta\mathbf{m}), \quad (13)$$

where \mathbf{d}^{n-1} is the vector of data for the model \mathbf{m}^{n-1} obtained from the previous iteration, $\delta\mathbf{m} = \mathbf{m}^n - \mathbf{m}^{n-1}$, and \mathbf{J} is the Jacobian matrix of sensitivities for the linear approximation:

$$\mathbf{d}^n \approx \mathbf{d}^{n-1} + \mathbf{J}\delta\mathbf{m}, \quad (14)$$

$$J_{ij} = \frac{\partial d_i}{\partial m_j}. \quad (15)$$

Also in the intermediate objective function (equation 11),

$$\phi_m^n = \sum_k \alpha_k \phi_k(\mathbf{v}_k), \quad (16)$$

$$\mathbf{v}_k = \mathbf{W}_k(\mathbf{m}^{n-1} + \delta\mathbf{m} - \mathbf{m}_k^{\text{ref}}). \quad (17)$$

There are many ways to choose the trade-off parameter β^n ; these are not discussed here.

To minimize Φ^n , equation 11 is differentiated with respect to the perturbations of the model parameters, and the resulting derivatives are equated to zero. (For the specifics of the following summary that pertain to general measures, see Farquharson and Oldenburg, 1998, and references therein.) Differentiating the general form of the measures (equation 6) gives

$$\frac{\partial \phi(\mathbf{x})}{\partial \delta m_k} = \sum_j \rho'(x_j) \frac{\partial x_j}{\partial \delta m_k}, \quad (18)$$

that is,

$$\frac{\partial \phi(\mathbf{x})}{\partial \delta \mathbf{m}} = \mathbf{B}^T \mathbf{q}, \quad (19)$$

where $\partial \phi / \partial \delta \mathbf{m} = (\partial \phi / \partial \delta m_1, \dots, \partial \phi / \partial \delta m_N)^T$, $B_{ij} = \partial x_i / \partial \delta m_j$, and $\mathbf{q} = (\rho'(x_1), \dots, \rho'(x_N))^T$. Equation 19 can be reformulated by introducing a diagonal matrix:

$$\mathbf{R} = \text{diag}\{\rho'(x_1)/x_1, \dots, \rho'(x_N)/x_N\}, \quad (20)$$

which leads to

$$\frac{\partial \phi(\mathbf{x})}{\partial \delta \mathbf{m}} = \mathbf{B}^T \mathbf{R} \mathbf{x}. \quad (21)$$

For the measure of misfit in the intermediate objective function, \mathbf{B} is $-\mathbf{W}_d \mathbf{J}$. For the components of the measure of model structure, \mathbf{B} is \mathbf{W}_k . For the four examples of measures given in equations 7–10, the elements of the matrix \mathbf{R} are

$$R_{ii} = \begin{cases} p \gamma^{p-2} & |x_i| \leq \gamma, \\ p |x_i|^{p-2} & |x_i| > \gamma, \end{cases} \quad (22)$$

where γ is a small number so that \mathbf{R} does not become singular as $x_i \rightarrow 0$;




Figure 1. The behavior of the function ρ for various measures: (a) the l_2 -norm; (b) the l_1 -norm; (c) the Huber M -measure with $c = 0.5$; (d) the Eklom measure with $p = 1$ and $\varepsilon = 0.1$; (e) the support measure of equation 10 with $\varepsilon = 0.1$.

$$R_{ii} = \begin{cases} 2, & |x_i| \leq c, \\ 2c/|x_i|, & |x_i| > c; \end{cases} \quad (23)$$

$$R_{ii} = p(x_i^2 + \varepsilon^2)^{p/2-1}; \quad (24)$$

and

$$R_{ii} = \frac{2\varepsilon^2}{(x_i^2 + \varepsilon^2)^2}. \quad (25)$$

The linear system of equations to be solved at each iteration is therefore

$$\begin{aligned} & \left[\mathbf{J}^T \mathbf{W}_d^T \mathbf{R}_d \mathbf{W}_d \mathbf{J} + \beta^n \sum_k \alpha_k \mathbf{W}_k^T \mathbf{R}_k \mathbf{W}_k \right] \delta \mathbf{m} \\ & = \mathbf{J}^T \mathbf{W}_d^T \mathbf{R}_d \mathbf{W}_d (\mathbf{d}^{\text{obs}} - \mathbf{d}^{n-1}) \\ & + \beta^n \sum_k \mathbf{W}_k^T \mathbf{R}_k \mathbf{W}_k (\mathbf{m}_k^{\text{ref}} - \mathbf{m}^{n-1}). \end{aligned} \quad (26)$$

The matrices \mathbf{R} , as well as the Jacobian matrix, depend on the model. They are updated at each iteration. This technique is known as iteratively reweighted least squares (IRLS).

The examples presented below are for the 2D MT and 3D gravity inverse problems. For the 2D MT case, the forward modeling was done with the finite-difference procedure used to obtain the boundary conditions for the 3D program of Farquharson et al. (2002). The operation of the Jacobian matrix, or its transpose, on a vector was done by solving the forward-modeling equations with the appropriate right-hand side (see, for example, Mackie and Madden, 1993; Rodi and Mackie, 2001; Farquharson et al., 2002). Solution of equation 26 directly by LU decomposition and solution iteratively by incomplete LU decomposition preconditioned conjugate gradients were both implemented. For the 3D gravity case, forward modeling was done via a finite-difference solution of Poisson's equation. Although the relationship between observations and model parameters is linear in this case, the operation of what is essentially the Jacobian matrix, or its transpose, on a vector was computed using the same procedure as for the 2D MT example. Equation 26 was solved by the incomplete LU decomposition preconditioned conjugate gradient method.




Figure 2. A subsection of the 2D model mesh showing the possible directions of the finite differences, relative to cell i, j , in the measure of model structure.

Measure of model structure

In traditional minimum-structure inversion algorithms, the measure of model structure given by equations 4 and 5 would comprise three terms for a 2D model (four for a 3D model), one involving a diagonal weighting matrix and usually a reference model, and two involving finite differences in the x - and z -directions. The second two terms are sometimes considered together as the spatial gradient of the model. When l_1 measures are used in conjunction with the two finite-difference terms, large values of these x - and z -directed differences become acceptable in the constructed model. This produces sharp interfaces between essentially uniform regions. However, because the measure of model structure only explicitly involves finite differences in the x - and z -directions, irrespective of whether these differences are considered together as the gradient of the model, the sharp interfaces are normal only to one or the other of the finite differences in the measure — that is, they are either vertical or horizontal (Farquharson and Oldenburg, 2003; Auken and Christiansen, 2004). This is not sufficient to represent a general, piecewise-constant distribution of a physical property in the subsurface. The measure of model structure is therefore amended here in an ad hoc manner to include finite differences in diagonal directions. This encourages the formation of interfaces that are normal to these new diagonal directions. The two diagonal directions for the 2D case are shown in Figure 2. The corresponding measure of model structure in equation 4 now comprises five terms, with \mathbf{W}_1 a diagonal weighting matrix, \mathbf{W}_2 and \mathbf{W}_3 the usual x and z first-order finite-difference matrices, and \mathbf{W}_4 and \mathbf{W}_5 first-order finite-difference operators for diagonally up to the right and diagonally down to the right (Figure 2).

The measure of model structure for the 3D case is expanded in a similar manner. Here, it is expanded to include 14 terms: the diagonal weighting term; the differences in the x -, y -, and z -directions; six terms for diagonal differences in the xy -, yz -, and xz -planes; and four terms for diagonal differences directed from the center of a cell through its vertices. The areas of the cells are taken into account for the 2D case and the volumes of the cells for the 3D case.

Li and Oldenburg (2000) present a modification to their 2D minimum-structure inversion procedure for resistivity data that incorporates user-specified preferred dip directions into a typical l_2 measure of model structure. It is unclear whether their approach would generate sharp dipping interfaces if used with l_1 measures. The straightforward inclusion of diagonal finite differences as described above is favored here.

EXAMPLES

Synthetic 2D MT example

A synthetic MT data set was generated for the conductivity model shown in Figure 3. The data comprised the real and imaginary parts of the E-polarization impedances at the 17 locations shown by the triangles in Figure 3 and at the three frequencies 3, 10, and 30 Hz. (Only E-polarization data are included here because this gives a more dramatic demonstration of the differences between using l_2 and l_1 measures of model structure.) Gaussian random noise of standard deviation equal in magnitude to 1% of a datum was added to make the data set that is to be inverted. The data are shown in Figure 4. The mesh comprised 80×65 cells and extended from $x = -12$ km to $x = 12$ km and from $z = -18$ km to $z = 11$ km (with the earth-air interface at $z = 0$). The central region of the mesh, extending from $x = -5000$ m to $x = 5000$ m and from

$z = 0$ to $z = 4000$ m, comprised 50×40 , 200×100 m cells. The same mesh was used for the forward modeling and the inversions.

In the inversions described below, the l_2 -norm was used as the measure of misfit. The trade-off parameter was started at 100, a relatively large value in the sense that if an inversion were performed with the trade-off parameter fixed at this value, a model would be produced that had only a small amount of structure and gave data that substantially underfit the observations. At each iteration, the trade-off parameter was halved to give a slow but steady progression of models with increasing structure and decreasing data misfits. Once a value of the trade-off parameter was reached that resulted in a model with a data misfit close to the target misfit, the trade-off parameter was kept constant and further iterations were performed until the model and objective function no longer changed. The target misfit for this example is 102. This is the number of observations and hence the expectation of the l_2 measure of misfit being used. The data computed for the final models in the respective inversions are shown in Figure 4. All inversions started from a homogeneous half-space of 10^{-2} S/m. This was used also as the reference model $\mathbf{m}_1^{\text{ref}}$; reference models for $k = 2, \dots, 5$ were not included.

Figure 5 shows the model produced using the usual l_2 -norm for the measure of model structure and the x and z first-order finite-difference terms (specifically, in equation 4, $\alpha_1 = 0.001$, $\alpha_2 = \alpha_3 = 1$, and $\alpha_4 = \alpha_5 = 0$). The conductive region in this model corresponds nicely to the conductor in the true model in both location and general shape. However, the conductive zone has the smeared-out appearance typical of a model produced by minimizing an l_2 -norm measure of model structure.

Figure 6 shows three models produced using the Eklblom measure of equation 9 with $p = 1$ and $\varepsilon = 10^{-4}$ and with different combinations of the spatial finite differences. The model in Figure 6a was generated using only the horizontal and vertical finite differences (specifically, $\alpha_1 = 0.001$, $\alpha_2 = \alpha_3 = 1$, and $\alpha_4 = \alpha_5 = 0$). This model is piecewise constant, or blocky. However, the interfaces essentially are either vertical or horizontal: The cross section of the conductor is shaped more like an “L” than a triangle. Figure 6b shows the model produced when there was equal emphasis on all four of the finite-difference terms (i.e., $\alpha_1 = 0.001$, $\alpha_2 = \alpha_3 = 1$, and $\alpha_4 = \alpha_5 = 1$). The model contains sharp dipping, horizontal, and vertical interfaces. The center of the triangular conductor, where most of the induced current is flowing, is the dominant feature of the constructed model, with the points of the triangles being less well resolved. Figure 6c shows the resultant model when the diagonal finite differences were given more importance than the horizontal and vertical differences (i.e., $\alpha_1 = 0.001$, $\alpha_2 = \alpha_3 = 0.1$, and $\alpha_4 = \alpha_5 = 1$). (Setting $\alpha_2 = \alpha_3 = 0$ and $\alpha_4 = \alpha_5 = 1$ in the current implementation leads to a checked effect.) This model has sharp, dipping interfaces, exactly the kind of features that were not possible without the inclusion of the diagonal differences.

Figure 7 shows the values of data misfit, measure of model structure, and trade-off parameter as functions of iteration for the four inversions described above. The inversion with the l_2 model measure

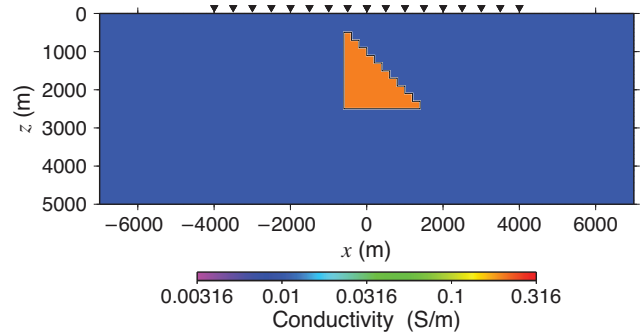


Figure 3. The 2D model from which the synthetic MT data set was generated. The triangles indicate the observation locations.




Figure 4. The synthetic observed data set and the predicted data sets for the example 2D MT inversions. The error bars indicate the synthetic data, and the lines indicate the data computed for the models shown in Figures 5 and 6.




Figure 5. The model produced using the usual l_2 measure of model structure. The black-and-white outline indicates the location of the conductive wedge in the true model (Figure 3).

converged in seven iterations. All of the inversions using the l_1 measure had effectively converged in 25 to 30 iterations.

Field 3D gravity example: Ovoid deposit, Voisey’s Bay, Labrador

The inversion procedure presented above was applied to the ground gravity data from over the Ovoid orebody at Voisey’s Bay (Figure 8). This is an internationally significant nickel-copper-cobalt deposit. For a description of its geology, see Naldrett (2000) and references therein. The gravity data set from the Ovoid has been used by a number of authors to test inversion methodologies (for example, Oldenburg et al., 1998; Ash et al., 2006).

The data set relevant to the Ovoid is shown in Figure 9. It consists of 89 data points distributed along four lines. The data are in the form of Bouguer anomaly values (relative to 2.67 g/cm^3). A regional trend has been removed via upward continuation (Ash, 2007).

The Ovoid deposit has been drilled extensively. Figure 10 shows the density model derived from downhole samples. The interpolation of densities between boreholes in this model was done by kriging within the main geologic units (Ash, 2007). There is clearly more detail in this model than one could ever expect to infer from the surface gravity data set shown in Figure 9. The main purpose of the following inversions is to illustrate the inversion procedure presented here, not to produce the best possible model of the Ovoid from the

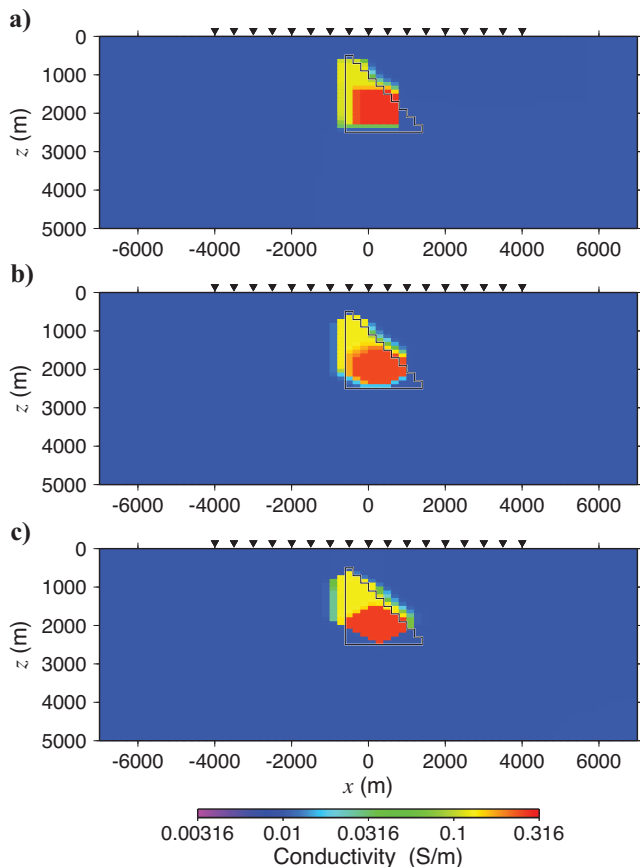


Figure 6. The models produced using an l_1 measure of model structure with three different combinations of spatial differences: (a) just the usual horizontal and vertical differences; (b) equal contributions from the horizontal and vertical differences and the two diagonal differences; (c) the diagonal differences dominant.

surface gravity measurements. Nevertheless, the model in Figure 10 is included as a “true model” with which the following inversion results can be compared and contrasted.

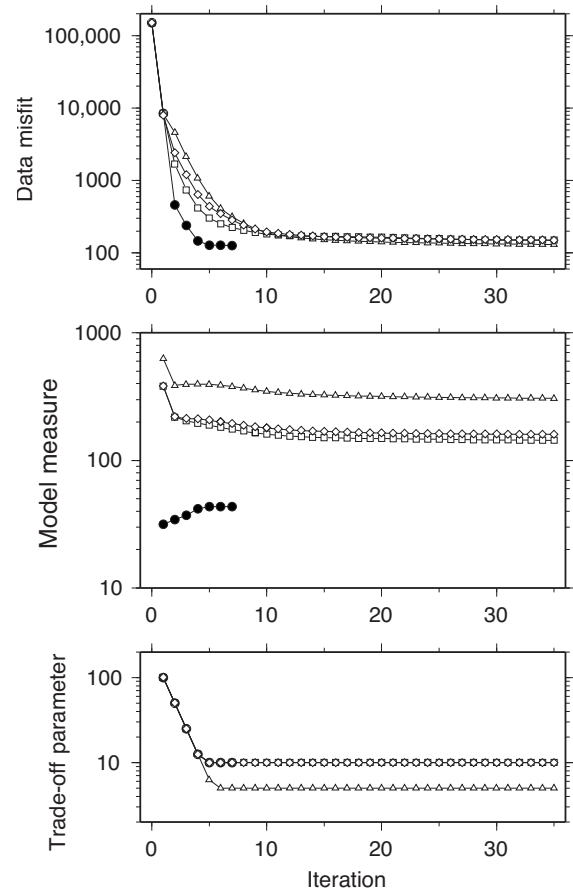


Figure 7. The progression with iteration of data misfit, measure of model structure and trade-off parameter for the four inversions of the 2D MT data set with the usual l_2 -norm measure (disks), the l_1 measure with only horizontal and vertical differences (squares), the l_1 measure with diagonal as well as horizontal and vertical differences (triangles), and the l_1 measure with the diagonal differences dominant (diamonds).




Figure 8. The location of Voisey’s Bay (black star), Labrador, Canada.

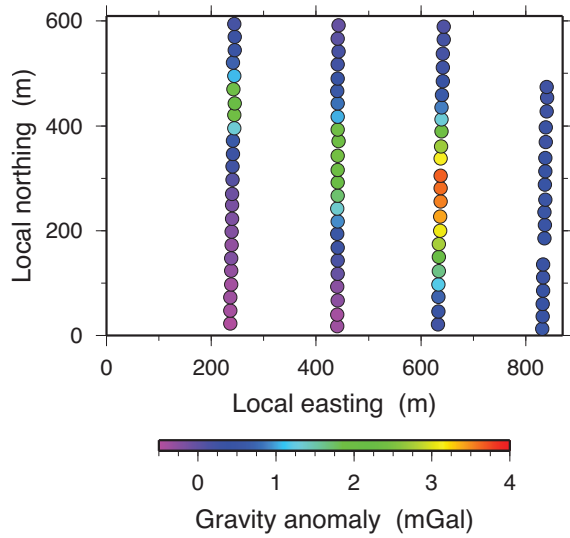


Figure 9. Observed gravity values, in terms of Bouguer anomaly relative to 2.67 g/cm^3 , over the Ovoid deposit at Voisey's Bay. The discs indicate the locations of the observations.

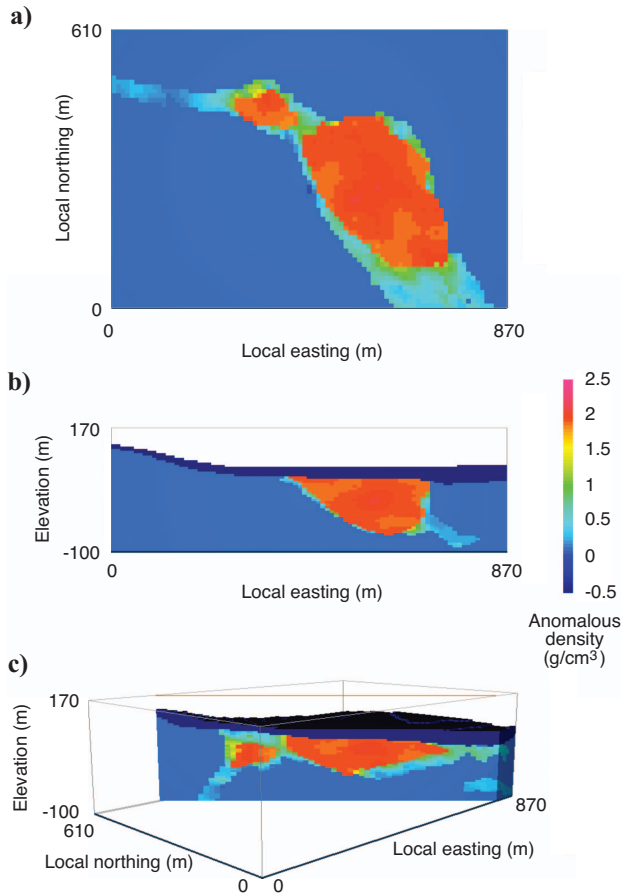


Figure 10. The density model of the Ovoid orebody created by kriging the densities of downhole samples: (a) horizontal slice at 32.5 m elevation, (b) vertical slice at local northing 290 m, (c) a northwest-southeast vertical slice. The anomalous densities are relative to 2.67 g/cm^3 . The mesh is the same as that used for the inversions of the gravity data. The area shown is the same as that in Figure 9. The four main rock types going from most dense to least are massive sulfide, troctolite, gneiss, and overburden.

The same mesh was used for all three inversions described below. It consisted of $87 \times 61 \times 54$ cells in the easting, northing, and vertical directions respectively. All of the cells were of the same dimensions: $10 \times 10 \times 5 \text{ m}$. (There were seven additional planes of padding cells of increasing size in each direction outside this central region. This is for the finite-difference solution used by the forward solver.) The overburden was incorporated by means of the reference model $\mathbf{m}_1^{\text{ref}}$: The anomalous density in the cells in $\mathbf{m}_1^{\text{ref}}$ corresponding to the overburden was set to that for the overburden (-0.75 g/cm^3), the anomalous density in the cells in $\mathbf{m}_1^{\text{ref}}$ below the overburden was set to that for the background gneiss (0.14 g/cm^3), and the elements of the diagonal weighting matrix \mathbf{W}_1 corresponding to the overburden cells were increased by a factor of 100. The other possible reference models for $k = 2, \dots, 14$ were not used. Topography was incorporated in the inversions. The normal l_2 measure of data misfit was used for all three inversions. Measurement uncertainties equal to 0.05 mGal were assumed. All inversions started from a homogeneous half-space of zero anomalous density.

Figure 11 shows the constructed model from an inversion using the usual l_2 measure of model structure. The coefficients in the model measure for the smallest component and for the differences in the easting, northing, and vertical directions (i.e., α_1 to α_4 in equation 4)

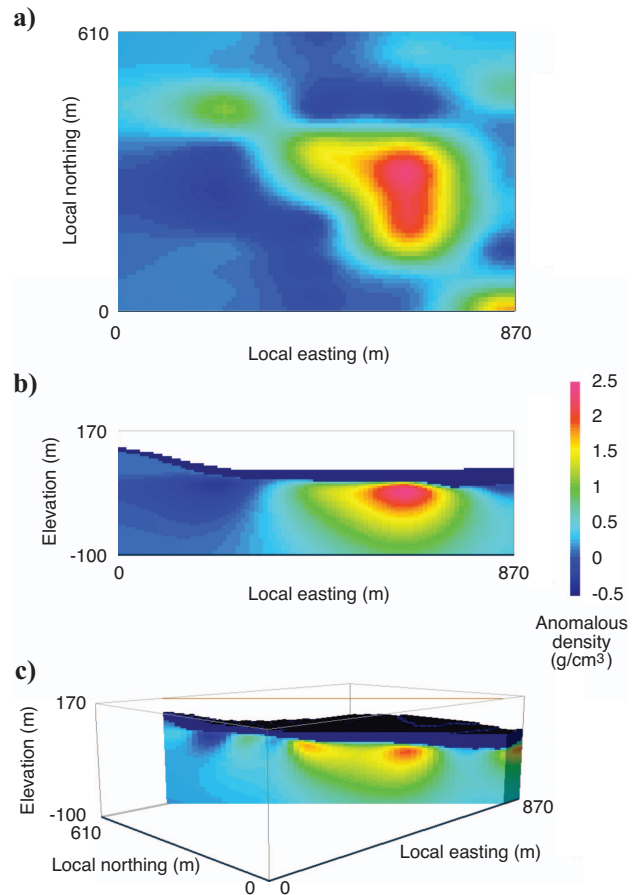


Figure 11. The model produced by a typical inversion of the Ovoid gravity data set (Figure 9) using an l_2 measure of model structure: (a) horizontal slice at 32.5 m elevation, (b) vertical slice at local northing 290 m, (c) a northwest-southeast vertical slice. The anomalous densities are relative to 2.67 g/cm^3 . The area shown is the same as that in Figure 9.

were 0.1, 10, 1, and 0.1, respectively. The choice of 10 and one for the easting and northing directions was to spread out the anomalous density between the survey lines, thus counteracting the effect of the much greater separation between observation locations from one survey line to the next compared to their separation along the lines. The choice of 0.1 for the vertical direction was to reduce the vertical extent of the region of anomalous density to an appropriate amount. The distance weighting of Li and Oldenburg (1998) was also used. The value of the trade-off parameter β was 10^{-7} , and the value of the data misfit for the constructed model was 108. The sections in Figure 11 illustrate the typical fuzzy, smeared-out character of a model created using an l_2 measure of model structure.

Figure 12 shows the model produced using the Eklom measure of equation 9 with $p = 1$ and $\varepsilon = 10^{-4}$ and with only x -, y -, and z -directed finite differences. The coefficients α_1 to α_4 in the model measure were 0.1, 10, 1, and 0.2, respectively. The trade-off parameter was equal to 5×10^{-5} throughout the inversion, and the data misfit for the final model was 103. As seen in Figure 12, the model produced using this l_1 measure comprises regions of uniform density separated by relatively sharp interfaces. However, it is also clear that

the interfaces are all essentially normal to the x -, y -, or z -directions, making rectangular structures.

Figure 13 shows the model produced when diagonal finite differences are included in the l_1 measure of model structure. For this example, the first four coefficients of the model measure were equal to 0.1, 10, 1, and 0.2 as for the preceding example; the remaining 10 coefficients, which correspond to the diagonal difference terms, were all equal to one. The trade-off parameter was equal to 10^{-5} throughout the inversion, and the data misfit for the final model was equal to 100. As the sections in Figure 13 show, the model produced using this measure of model structure also comprises uniform regions separated by relatively sharp interfaces. However, unlike the model in Figure 12, the interfaces are now not restricted to being perpendicular only to the x -, y -, and z -directions; significant dipping and angled interfaces are present.

Figure 14 shows the values of data misfit and measure of model structure as functions of iteration for the three inversions described. The inversion with the l_2 model measure essentially required only one iteration. (The model measure for the constructed model was equal to 2.5×10^{10} , which is not included in Figure 14.) The inversions using the l_1 measure had effectively converged in 20 iterations.

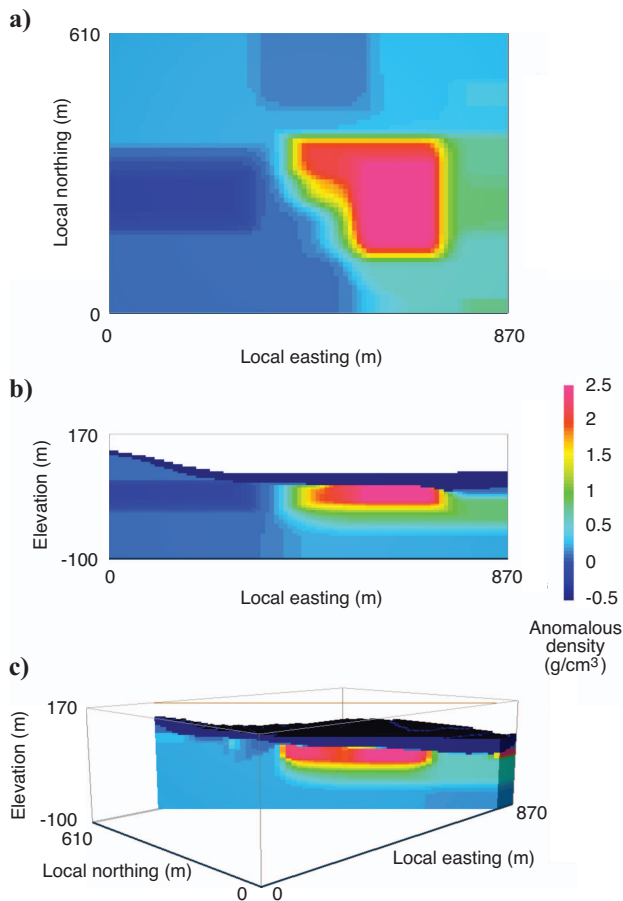


Figure 12. The model produced using an l_1 measure of model structure, with only x -, y -, and z -directed finite differences: (a) horizontal slice at 32.5 m elevation, (b) vertical slice at local northing 290 m, (c) the northwest-southeast vertical slice. (The anomalous densities are relative to 2.67 g/cm^3 .)

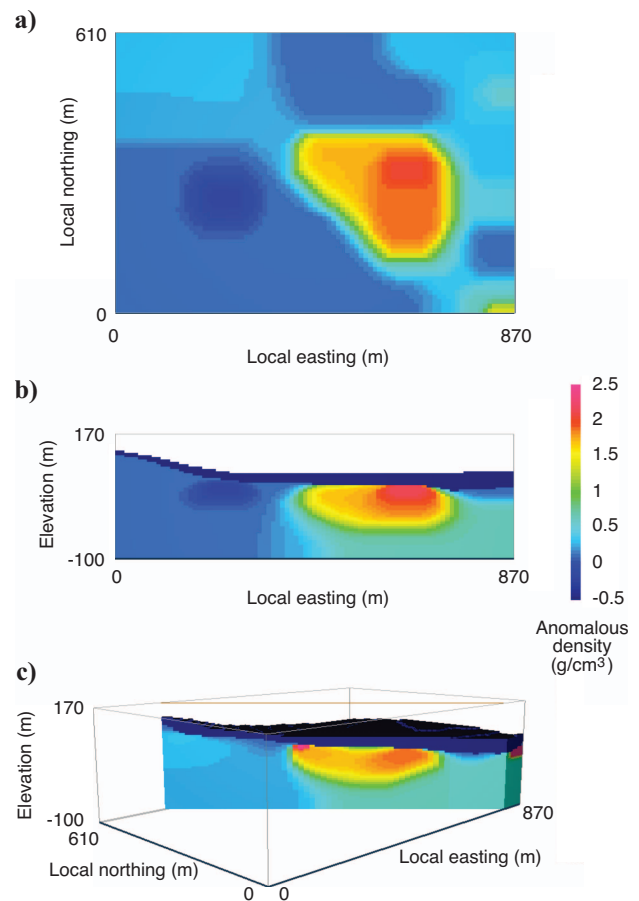


Figure 13. The model produced using an l_1 measure of model structure, with diagonal finite differences as well as the usual x -, y -, and z -directed finite differences: (a) horizontal slice at 32.5 m elevation, (b) vertical slice at local northing 290 m, (c) the northwest-southeast vertical slice. (The anomalous densities are relative to 2.67 g/cm^3 .)




Figure 14. The progression of data misfit and measure of model structure for the three inversions of the Ovoid gravity data set with the usual l_2 -norm measure (disks), the l_1 measure with only horizontal and vertical differences (squares), and the l_1 measure with diagonal as well as horizontal and vertical differences (triangles). (The model measure for the l_2 -norm inversion was equal to 2.5×10^{10} , which is not included here.)

CONCLUSIONS

Minimum-structure inversion algorithms typically are robust and reliable, and they produce models with just enough complexity to reproduce the observations. The iteratively reweighted least-squares procedure can be used to extend the usual implementation to include measures other than the l_2 -norm. This allows the use of robust measures of data misfit and l_1 measures of model structure. The latter results in blocky, piecewise-constant models, which can be more consistent with preconceived notions of the subsurface than the fuzzy, smeared-out models produced using an l_2 measure. As demonstrated here, the explicit inclusion of diagonal finite differences in the l_1 measure of model structure means that dipping and angled interfaces can be constructed, not just interfaces normal to the Cartesian coordinate axes, as was previously the case.

ACKNOWLEDGMENTS

The author thanks Voisey's Bay Nickel Company and CVRD Inco for access to the gravity data from Voisey's Bay; Brian Bengert of VBNC and CVRD Inco, and Hugh Miller and Michael Ash of Memorial University of Newfoundland, for their advice and assistance; and the Atlantic Innovation Fund and CVRD Inco for financial support through the Inco Innovation Centre at Memorial University. Figures 10–13 were generated using the "MeshTools3D" program from the University of British Columbia's Geophysical Inversion Facility; use of this program is greatly appreciated. The author also

thanks Valéria Barbosa and two anonymous reviewers for their constructive and thoughtful comments.

REFERENCES

- Ash, M. R., 2007, Constrained inversion of gravity data over the Ovoid and Mini-Ovoid in the Voisey's Bay Ni–Cu–Co deposit, Labrador: M.S. thesis, Memorial University of Newfoundland.
- Ash, M. R., M. Wheeler, H. Miller, C. G. Farquharson, and A. V. Dyck, 2006, Constrained three-dimensional inversion of potential field data from the Voisey's Bay Ni–Cu–Co deposit, Labrador, Canada: 76th Annual International Meeting, SEG, Expanded Abstracts, 1333–1336.
- Auken, E., and A. V. Christiansen, 2004, Layered and laterally constrained 2D inversion of resistivity data: *Geophysics*, **69**, 752–761.
- Constable, S. C., R. L. Parker, and C. G. Constable, 1987, Occam's inversion: A practical algorithm for generating smooth models from electromagnetic sounding data: *Geophysics*, **52**, 289–300.
- de Groot-Hedlin, C., and S. Constable, 1990, Occam's inversion to generate smooth, two-dimensional models from magnetotelluric data: *Geophysics*, **55**, 1613–1624.
- , 2004, Inversion of magnetotelluric data for 2D structure with sharp resistivity contrasts: *Geophysics*, **69**, 78–86.
- Dozzo, S. E., and D. W. Oldenburg, 1989, Linear and non-linear appraisal using extremal models of bounded variation: *Geophysical Journal International*, **99**, 483–495.
- Eklblom, H., 1987, The L_1 -estimate as limiting case of an L_p - or Huber-estimate, in Y. Dodge, ed., *Statistical data analysis based on the L_1 -norm and related methods*: Elsevier Science Publ. Co., Inc., 109–116.
- Farquharson, C. G., and D. W. Oldenburg, 1998, Nonlinear inversion using general measures of data misfit and model structure: *Geophysical Journal International*, **134**, 213–227.
- , 2003, Constructing piecewise-constant models using general measures in non-linear, minimum-structure inversion algorithms: 6th International Symposium, Society of Exploration Geophysicists of Japan, Expanded Abstracts, 240–243.
- Farquharson, C. G., D. W. Oldenburg, E. Haber, and R. Shekhtman, 2002, An algorithm for the three-dimensional inversion of magnetotelluric data: 72nd Annual International Meeting, SEG, Expanded Abstracts, 649–652.
- Haber, E., and L. Tenorio, 2003, Learning regularization functionals — a supervised training approach: *Inverse Problems*, **19**, 611–626.
- Huber, P. J., 1964, Robust estimation of a location parameter: *Annals of Mathematical Statistics*, **35**, 73–101.
- Last, B. J., and K. Kubik, 1983, Compact gravity inversion: *Geophysics*, **48**, 713–721.
- Li, Y., and D. W. Oldenburg, 1996, 3-D inversion of magnetic data: *Geophysics*, **61**, 394–408.
- , 1998, 3-D inversion of gravity data: *Geophysics*, **63**, 109–119.
- , 2000, Incorporating geological dip information into geophysical inversions: *Geophysics*, **65**, 148–157.
- Loke, M. H., I. Acworth, and T. Dahlin, 2003, A comparison of smooth and blocky inversion methods in 2D electrical imaging surveys: *Exploration Geophysics*, **34**, 182–187.
- Mackie, R. L., and T. R. Madden, 1993, Three-dimensional magnetotelluric inversion using conjugate gradients: *Geophysical Journal International*, **115**, 215–229.
- Naldrett, A. J., 2000, A special issue on the Voisey's Bay Ni–Cu–Co deposit, Labrador, Canada, Introduction: *Economic Geology*, **95**, 675–676.
- Oldenburg, D. W., and R. G. Ellis, 1991, Inversion of geophysical data using an approximate inverse mapping: *Geophysical Journal International*, **105**, 325–353.
- Oldenburg, D. W., and Y. Li, 1994, Inversion of induced polarization data: *Geophysics*, **59**, 1327–1341.
- Oldenburg, D. W., Y. Li, C. G. Farquharson, P. Kowalczyk, T. Aravanis, A. King, P. Zhang, and A. Watts, 1998, Applications of geophysical inversions in mineral exploration: *The Leading Edge*, **17**, 461–465.
- Portniaguine, O., and M. S. Zhdanov, 1999, Focusing geophysical inversion images: *Geophysics*, **64**, 874–887.
- Ramos, F. M., H. F. Campos Velho, J. C. Carvalho, and N. J. Ferreira, 1999, Novel approaches to entropic regularization: *Inverse Problems*, **15**, 1139–1148.
- Rodi, W., and R. L. Mackie, 2001, Nonlinear conjugate gradients algorithm for 2D magnetotelluric inversion: *Geophysics*, **66**, 174–187.
- Smith, J. T., and J. R. Booker, 1988, Magnetotelluric inversion for minimum structure: *Geophysics*, **53**, 1565–1576.
- Smith, T., M. Hoversten, E. Gasperikova, and F. Morrison, 1999, Sharp boundary inversion of 2D magnetotelluric data: *Geophysical Prospecting*, **47**, 469–486.
- van Zon, T., and K. Roy-Chowdhury, 2006, Structural inversion of gravity data using linear programming: *Geophysics*, **71**, no. 3, J41–J50.