Applying geological constraints in UBC-GIF potential field inversions for greenfields or brownfields exploration

Nicholas C. Williams
Mineral Deposit Research Unit
Dept. Earth and Ocean Sciences
The University of British Columbia
6339 Stores Road
Vancouver, BC, V6T 1Z4
Canada
nwilliams@eos.ubc.ca

SUMMARY

The University of British Columbia – Geophysical Inversion Facility GRAV3D and MAG3D programs can be run quickly with default parameter settings to produce physical property distributions that explain observed gravity and magnetic data. In the absence of location-specific geological constraints, this approach gives acceptable first pass models but ones that are unlikely to be geologically realistic.

Even in greenfields mineral exploration there will be some geological information available in addition to the geophysical data. These constraints can be supplied to the inversion software, with adjustable levels of certainty, via a reference model of expected properties, bounds on the expected properties, or smoothness weights based on positions and orientations of the rocks. The constraints can come from mapping, sampling, analogous areas, or neighbouring regions. In the later stages of exploration and development, additional information from drilling, detailed structural interpretation, trenching, and even preliminary mining will also be available.

The UBC-GIF inversion software will produce holistic physical property models consistent with all the supplied information. The models recovered under these circumstances will be more predictive about the geology of a region than the models obtained with default parameter settings because they are consistent with both the geophysical data and the geology constraints. The types of constraints that can be included, and their effectiveness, are demonstrated in this paper using gravity inversion for a simple geologically-based synthetic example. The models that are obtained using different types of constraints are in all cases superior to the model obtained from a default parameter, geologically-unconstrained inversion.

Key words: Inversion, gravity, constraints, geology.

INTRODUCTION

Mineral exploration produces a large amount of diverse geological and geophysical data, yet it has proved difficult to combine all of this information into consistent holistic models of subsurface geology. 2D or 3D geophysical forward modelling can be used to calculate the geophysical response of a property model based on observed or expected geology. Discrepancies between the calculated and observed data are interpreted to indicate differences between the inferred geological scenario used in the forward model and the geology that is present. In recent years advances in computing power have facilitated the forward computation of geophysical responses for very large 2D and 3D models. Methods have also been developed to calculate inverse solutions that predict physical property distributions that give rise to the observed geophysical responses.

Gravity and magnetic data are two of the most common geophysical datasets used in ore deposit exploration. Even in greenfields exploration these datasets may be available from government agencies or from work carried out by previous explorers. The application of inversion methods to obtain estimated models of physical properties within a subsurface region from these gravity and aeromagnetic datasets is a common step in many exploration programs, especially in areas where prospective basement rocks are covered. The recovered inverse models can be used to target regions of anomalous physical properties for further data acquisition or drilling.

Inversion of potential field data is impeded by several numerical difficulties:

1. *Non-existence*: Due to the ubiquitous presence of noise in geophysical and geological observations, there may not exist a single model capable of fitting the measured data.
2. *Instability*: Small changes in the data, such as noise, can result in large changes in the recovered model since the inverse problem is *ill-conditioned*.
3. *Non-uniqueness*: There are two sources of non-uniqueness; the data are limited to a finite number of point observations, and the source distribution for potential fields is intrinsically non-unique.

To recover a physical property model, some type of conditioning or regularisation is required to mitigate these problems. Regularisation typically imposes a set of mathematical constraints that stabilise the problem and recover a model that fulfils certain criteria. The regularisation imposed by the University of British Columbia – Geophysical Inversion Facility (UBC–GIF) inversion approach seeks a model that is *small*, containing as little deviation from a reference model as possible, and *smooth*, with any deviations spread over a number of cells rather than concentrated in individual cells.

While these mathematical constraints can be used to recover physical property models that are visually pleasing and may provide broad geological insight, used in isolation they can never recover an accurate physical property model because there is no direct link between the mathematics and the geology. A holistic model, consistent with all observed information, can only be recovered by including geology-based constraints in addition to the mathematical constraints.

A strength of the UBC–GIF inversion programs is that they allow the inclusion of as much or as little geological information as is available, in the form of a reference model of physical properties or bounds on the range of expected physical properties. The inversion will return a solution that is within the bounds and as close as possible to the imposed reference model while still fitting the geophysical data to within the accepted uncertainty levels.

This paper outlines the basic inversion method used by the UBC–GIF programs GRAV3D and MAG3D, describes the inversion parameters used in these programs, and provides some guidance on parameter settings. The types of geological constraints that can be used are described and illustrated using a synthetic gravity inversion example.

THE UBC-GIF INVERSION METHOD

Linear inversion methods can be used to find a model, **m**, which satisfies:

$$\mathbf{Gm} = \mathbf{d}^{\mathbf{obs}}$$

where **G** is the forward operator, or kernel, that describes the physics of the problem, and $\mathbf{d}^{\mathbf{obs}}$ is the observed data. For potential field data in mineral exploration, we discretise the subsurface into a model (**m**) containing M individual cells, where M is typically greater than the number of data, N, available ($\mathbf{d}^{\mathbf{obs}}$). This results in an $N \times M$ ($M > N$) matrix G that is not square and therefore not invertible. Instead, the problem becomes an optimisation problem, seeking a solution that minimises both a numerical measure of the model and the misfit between the observed and predicted data.

The details of the UBC–GIF inversion approach for potential field data as implemented in the MAG3D and GRAV3D programs are given in Li and Oldenburg (1996, 1998a), and in the software user manuals (UBC–GIF, 2005a, 2005b). A brief summary is included here. In both the inversion programs and the following method and discussion, all distance quantities are measured in metres, all gravity observations and predictions are in mGal, all densities are in g/cm^3, all magnetic field observations are in nT, and all magnetic susceptibilities are induced susceptibilities with units of SI. All other quantities (such as weightings) are effectively unit-less.

Model Objective Function

The UBC–GIF magnetic and gravity inversion codes use a model objective function to quantify various characteristics of the model. The function includes a term that measures the smallness, or difference between the recovered model and a reference model, and terms that measure how the difference between recovered and reference models varies between cells in each of three orthogonal directions. The reference model may be as simple as a uniform (usually zero) half-space, in which case the returned model may be expected to contain the minimum amount of detail necessary to reproduce the observed data. However, the *a priori* geological knowledge portrayed in the reference model may be more substantial, to the point where a full model of the expected physical properties could be used where there is a strong understanding of the subsurface physical property distribution. The model objective function is designed so that minimisation leads to a recovered inverse model that has characteristics that are as close as possible to those in the supplied reference model. The model objective function used is:

$$\begin{aligned}\phi_m(m) = \alpha_s \int_V w_s \left[w_r(z)(m - m_{ref}) \right]^2 dV + \ldots \\ \alpha_x \int_V w_x \left[\frac{\partial}{\partial x} w_r(z)(m - m_{ref}) \right]^2 dV + \ldots \\ \alpha_y \int_V w_y \left[\frac{\partial}{\partial y} w_r(z)(m - m_{ref}) \right]^2 dV + \ldots \\ \alpha_z \int_V w_z \left[\frac{\partial}{\partial z} w_r(z)(m - m_{ref}) \right]^2 dV\end{aligned}$$

The first component measures the smallness. The last three components measure the smoothness of the difference between the recovered model, m, and the reference model, m_{ref}, measured in the direction of each of the three axes. These components ensure that any discrepancies between the recovered model and the reference model are spread over a region rather than concentrated in individual cells. The user-defined parameters α_s, α_x, α_y, and α_z are used to balance the contributions of the smallness and smoothness components. The function w_s may be used to force the physical property of cells to be closer to the supplied reference model at specific cell locations where the physical properties are better understood. The parameters w_x, w_y, and w_z, can be used in a similar fashion to make the model-difference more or less smooth across cell boundaries in the east, north, and vertical directions, for example to reproduce geological continuity or boundaries. The function w_r is a depth weighting function.

Depth Weighting

The depth weighting function is designed to counteract the decay of the potential field response with distance from the source so that all cells have an equal likelihood of containing sources. This is necessary as there is no inherent depth information contained in the observed potential field response and a default solution to the inverse problem would result in a model with sources clustered near the surface. The depth weighting function has the form (Li and Oldenburg, 1996):

$$w_r^2(z) = \frac{1}{(z + z_0)^\beta}$$

where z is the depth to the centre of the cell and z_0 and β are adjustable parameters used to match the weighting function to the kernel's decay with depth. β is usually chosen to reproduce the exponential decay of the gravity or magnetic response of a sphere with distance: $\beta = 2$ for gravity data and

$\beta = 3$ for magnetic data. The parameter z_0 is usually calculated automatically by matching the squared weighting function response with the decay of the kernel.

Data Misfit

For the recovered inverse model to be capable of reproducing the observed data there must also be a measure of how closely the predicted response of the recovered model matches the observed data. Geophysical experiments will obtain measurements:

$$\mathbf{d}^{\mathbf{obs}} = \mathbf{d}^{\mathbf{true}} + \boldsymbol{\varepsilon}$$

where the observed data, $\mathbf{d}^{\mathbf{obs}}$, includes noise, $\boldsymbol{\varepsilon}$, as well as the true response, $\mathbf{d}^{\mathbf{true}}$. As discussed in a later section, the noise fraction includes other contributions besides the errors in the measurements themselves. A model that provided an exact match between the observed noisy data and the predicted data would almost certainly be incorrect because the model response would include both the true signal and the non-zero noise. An approach is used such that the observed data are reproduced to a degree that is commensurate with their perceived accuracy. The data misfit is measured by:

$$\phi_d = \left\| \mathbf{W_d} \left(\mathbf{Gm} - \mathbf{d}^{\mathbf{obs}} \right) \right\|^2$$

$$\mathbf{W_d} = \mathrm{diag}\left(\frac{1}{\sigma_i} \right)$$

where $\mathbf{Gm}$ is the predicted response of the recovered model, and σ_i is the standard deviation associated with the ith data point.

Obtaining a Solution

A suitable model is found by solving the optimisation problem:

$$\text{minimize } \phi = \phi_d + \beta \phi_m$$
$$\text{such that } \phi_d = \phi_d^*$$

where ϕ is the objective function and β is a trade-off parameter to balance the importance of data misfit versus the model objective function, and ϕ_d^* is the target data misfit which, assuming that noise is Gaussian with zero mean and standard deviation σ_i, will be equal to the number of data points, N.

When discretised over a 3D mesh, the problem becomes one of minimising:

$$\phi = \left\| \mathbf{W_d} \left(\mathbf{Gm} - \mathbf{d}^{\mathbf{obs}} \right) \right\|^2 + \beta \left\| \mathbf{W_m} \left(\mathbf{m} - \mathbf{m_{ref}} \right) \right\|^2$$

The minimum occurs at:

$$\nabla_m \phi = \mathbf{0}$$
$$\left(\mathbf{G^T W_d^T W_d G} + \beta \mathbf{W_m^T W_m} \right) \mathbf{m} = \mathbf{G^T W_d^T W_d d^{obs}} + \beta \mathbf{W_m^T W_m m_{ref}}$$

where:

$$\mathbf{W_m^T W_m} = \alpha_s \mathbf{W_r^T W_s^T W_s W_r} + \alpha_x \mathbf{W_r^T W_x^T W_x W_r} + \ldots$$
$$\alpha_y \mathbf{W_r^T W_y^T W_y W_r} + \alpha_z \mathbf{W_r^T W_z^T W_z W_r}$$

and $\mathbf{W_r}$ is a discretised form of the depth weighting function, $\mathbf{W_s}$, $\mathbf{W_x}$, $\mathbf{W_y}$, and $\mathbf{W_z}$ are discretised forms of the weights w_s, w_x, w_y, and w_z,. A solution for $\mathbf{m}$ is found by solving the problem using iterative techniques for a range of β until a solution where $\phi_d = \phi_d^* = N$ is located.

Several additional constraints, the implementation of which is described by Li and Oldenburg (1996, 1998a), are also applied when calculating a solution: positivity and bounds. Logarithmic barrier functions are used in the MAG3D code to ensure that only positive magnetic susceptibilities are obtained. They are also applied in both MAG3D and GRAV3D to ensure that the recovered properties lie between specific bounds. In default inversions, wide bounds are allowed, but when including geological constraints, narrow bounds can be supplied to restrict the properties to some expected range. The bounds can be set for the whole model or for individual cells.

PREPARING INVERSIONS

A number of steps are required to prepare data and a mesh for an inversion. These steps are covered by Li and Oldenburg (1996, 1998a, 1998b) and UBC–GIF (2005a, 2005b). In summary they include:

- Definition of the problem to be addressed
- Definition of the volume of interest (depth, width and length of desired mesh)
- Definition of the data area
- Definition of the cell sizes to match the resolution of the data, the desired resolution of the recovered model, and available computing power (currently about 1.5 million cells is a reasonable upper limit for tenable computation on standard desktop PCs)
- Padding the mesh with a buffer of additional cells to prevent boundary effects where anomalies are located near the edge of the mesh
- Upward continuation of the potential field data to the width of the cells to ensure that high frequency information that could only be reproduced by smaller cell sizes is not included
- Calculation and removal of a regional data trend that accounts for the contribution to the response of all sources located outside the volume of interest.

CHOOSING OPTIMAL PARAMETERS

The implementation of the UBC–GIF inversion method requires selection of appropriate values for a large number of parameters. Each choice can cause large differences in the model, and although use of the default values may be acceptable in some situations, more reliable models will be obtained by tuning the parameters to a particular problem. One exception is the depth weighting which has a basis in the physics of potential fields; the default values are usually best, and if the results are deemed to be inappropriate then a reference model or bounds should be employed.

Noise

It is critical to assign appropriate noise levels to the observed data used in the inversion. If the assigned noise levels are too low, large portions of the noise will be reproduced as artefacts in the model (Figure 1). If the assigned noise levels are too high then too much of the observed data will be treated as noise, and information will be lost in the model. Unfortunately potential field data rarely comes with robust uncertainty estimates, and these estimates need to be revised to account for the effects of data processing, use of a discretised representation of the earth, and numerical inaccuracies.

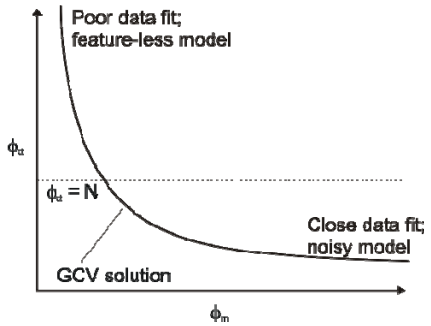

Figure 1. Typical relationship between the model objective function ϕ_m and the data misfit ϕ_d. Only when the noise is Gaussian and the correct standard deviation is used will ϕ_d = *N* provide a balanced solution. The GCV solution will commonly fit the data a little too well and include some noise in the recovered model, but remains a good approximation in the absence of noise estimates.

If an estimate of the noise level is available it can be included in the data file. For inversions gravity will commonly have standard deviations of 1–2 % of the data range, expressed as a constant (i.e. 0.1 mGal). Due to its higher dynamic range, aeromagnetic data may have standard deviations on the order of a couple of percent plus a couple of nT (i.e. 5 % + 5 nT). Older surveys may require higher noise levels than newer surveys, depending on the methods used, and this can be accommodated in the UBC–GIF inversions where data from the two surveys are combined. Likewise, upward continued data will have a lower dynamic range and should have lower standard deviations applied.

One approach to estimating the noise levels in the absence of good estimates is to perform a Generalised Cross Validation (GCV) inversion (e.g., Farquharson and Oldenburg, 2004), an option within the UBC–GIF codes, supplying a standard deviation of 1 (mGal or nT) to all data. This chooses a trade-off parameter β based on how dependent the model is on individual data points. The completed inversion log file reports a final Achieved Misfit. Dividing this misfit by the number of data points gives a rough estimate of the average standard deviation for the dataset. However, GCV inversions commonly fit the data a little too closely so the actual standard deviation chosen should be slightly higher than this estimate, and should be compared to other estimates of the noise in the data.

Alphas

The balance of smoothness versus smallness for the whole model is controlled by the α values: α_s, α_x, α_y, and α_z. Since the values are all subject to scaling by the trade-off parameter, β, it is only the balance between them that is important. A smaller model may contain more structure, manifested as excess detail at depths within the model, whereas a smoother model may exhibit less detail. By simplifying the model objective function to ignore the reference model, and considering the case of only two cells that are adjacent in the x direction, the balance between smoothness and smallness can be evaluated:

$$\phi_m(m) = \alpha_s \int m^2 dA + \alpha_x \int \left(\frac{\partial m}{\partial x} \right)^2 dA$$

If each cell is the same size, with area $A = \Delta x \cdot \Delta z$ (where Δx is the cell width and Δz is the cell height, in metres) and physical properties m_1 and m_2, the integrals transform into summations and the model objective function for the two cells becomes:

$$\phi_m = \alpha_s \left(m_1^2 + m_2^2 \right) A + \alpha_x \left[\frac{m_2 - m_1}{\Delta x} \right]^2 A$$

It is possible to evaluate the α's so as to balance smallness and smoothness by equating the two terms:

$$\alpha_s \left(m_1^2 + m_2^2 \right) A = \alpha_x \left[\frac{m_2 - m_1}{\Delta x} \right]^2 A$$

Equality is reached when their ratio is unity:

$$\alpha_x = \frac{\alpha_s \left(m_1^2 + m_2^2 \right) A}{\left[\frac{m_2 - m_1}{\Delta x} \right]^2 A} = \alpha_s \cdot c_x \left(\Delta x \right)^2$$

where:

$$c_x = \frac{m_1^2 + m_2^2}{\left(m_2 - m_1 \right)^2}$$

The constant, c_x, controls the desired proportionality of the properties of the two cells if all other constraints, such as data misfit, allow. For a full 3D mesh, the constant is more complex and is generalised to an arbitrary model smoothness parameter which can be applied in each direction to also yield:

$$\alpha_y = c_y \cdot \alpha_s \left(\Delta y \right)^2$$
$$\alpha_z = c_z \cdot \alpha_s \left(\Delta z \right)^2$$

If c = 1, the smoothness and smallness should balance. Experimentation shows that better inversion results are usually obtained when c is between 4 and 25, causing a slight bias towards smoother models. In older versions of the UBC–GIF inversion codes, this corresponded to length scales:

$$L_x = \Delta x \cdot \sqrt{c_x}$$

The crucial aspect of the above analysis is that the balance of smoothness versus smallness is controlled by the square of the cell sizes, so it is critical to adjust the α parameters to suit the size of the cells in the centre of the model.

The smoothness parameters, c, can also be used to supply geological information about the continuity of strike in a given area. Higher values can be used to indicate that geology is more continuous in a particular direction, for example, east-west strike can be reproduced by disproportionately higher c_x values. This approach will be more effective when reproducing geological continuity that is sub-parallel to the cardinal axes. An empirical rule of thumb, based on the results of synthetic inversions, takes the form:

$$c_x \geq 2 \cdot (n_x - 7)$$

where n_x is the number of cells in a row in the east-west direction that are influenced by a tight constraint on the property of one cell in the middle of that row, where allowed by the geophysical data. This estimate works for $n \geq 9$ ($c \geq 4$), but the effect on all but those cells closest to the constrained cell is usually minimal. Such tuning is subjective and may be improved by trial and error, and comparison of inversion results with known geology.

REFERENCE MODELS, BOUNDS AND WEIGHTINGS

The UBC–GIF inversion approach is flexible enough to include a wide range of geological information, if available. The formulation provides for several global constraints, i.e., depth weighting and the alphas, together with local constraints in the form of a reference model, smallness and smoothness weights, and model bounds. Without including this information, no inversion method will return a model that is consistent with existing geological knowledge since there are an infinite number of mathematically-feasible, but geologically-unlikely models available. Geologically-unconstrained inversions are unlikely to be consistently predictive about the subsurface physical property distribution.

Types of Geological Constraints

In general the geological information that can be included as constraints falls into five types (Figure 2). Surface data and drill holes may supply actual physical property measurements, possibly even multiple measurements within a single model cell. By assigning a reference model, with appropriate levels of certainty supplied by w_s, or by assigning bounds, these properties can be used to fix the property of cells containing outcrop or drill holes in the recovered inversion model.

All five types can provide geometrical lithological constraints indicating the positions and extents of particular units or types of rocks. With appropriate selection of a physical property or property bounds for each unit, their positions and extents, where known, can be recovered in the inversion model.

In addition to actual observations and measurements, hypotheses can also be tested. Hypothetical bodies or volumes can be created and included in the inversion. The recovered model will show if the hypotheses are possible given the supplied constraints and observed data, or require geologically-unrealistic changes within the recovered model to compensate.

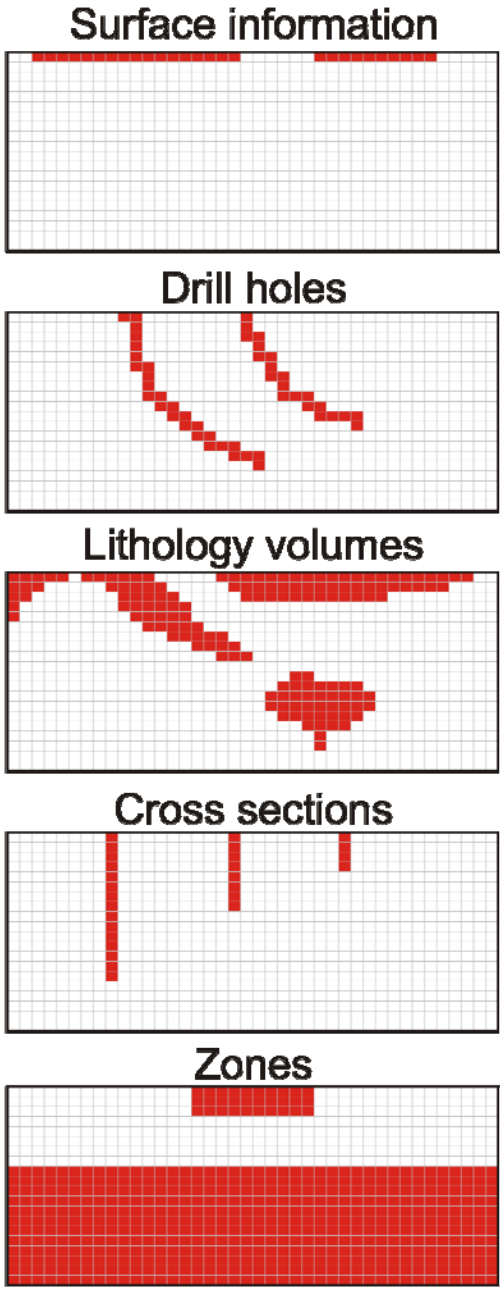

Figure 2. Schematic representations of the styles of geological constraints that can be included in a reference model as viewed in a 2D cross section through a 3D mesh. Red cells are those where information is available to assign either property bounds or a reference property, while white cells remain geologically-unconstrained.

The zone constraint type shown in Figure 2 is a more general constraint based on geological principles or interpretations about the types of rocks expected in different regions of the model. Zones may be used when a detailed 3D physical property model based on a synthesis of all available data (physical property measurements, mapping, drilling, structural interpretation, seismic data, etc.) is available for a region within the inversion model volume. This information can be "painted" into the reference model with reference properties,

high w_s values, and narrow bounds. The inversion can then extrapolate that information outwards into unknown areas based on the geophysical data.

Another application of zones places restrictions on the types of rocks, and therefore the physical properties, that may be present in a region. An example would be where there is weathering at surface, but the depth to basement is uncertain. Based on an understanding of the regolith, drilling, or seismic data, an inference might be made that all rocks below a certain depth (perhaps 250 m) are likely to be unweathered basement rocks and therefore will not have the low densities typical of weathered material. Although the actual geology of the basement may be poorly known, densities less than 2 g/cm^3 would be unlikely, and this can be included as a constraint by applying a narrower range of bounds than the default. In another example, dense carbonate rocks may only be expected in a particular portion of a basin, based on sequence stratigraphic work, and this can be reflected in the inversion by allowing higher bounds in that region, even if the exact location and properties of the rocks are unknown.

Implementing Constraints

As mentioned above, the α values can be used to reproduce strike continuity. But much more information can be included in the definition of the reference model, bounds, and w_s, w_x, w_y, and w_z weightings. They are used in each inversion, and are assigned default values if not explicitly provided by the user. If a reference model file, bound file, or weighting file is to be supplied, it must be defined for every cell in the model, but the user can use appropriate default values anywhere where they have insufficient information.

The reference model consists of a single physical property value for each cell in the model; default values are 0 g/cm^3 and 0 SI. It is used in conjunction with a set of w_s values, also defined for each cell in the model. These smallness values indicate a level of certainty in the physical properties assigned in the reference model. The w_s values are unit-less; the default is unity, but increased certainty can be indicated with higher values.

Bounds provide a powerful means of enforcing a particular range of properties within a region or unit where the physical properties are known to vary, or are difficult to define exactly. They can be supplied with or without a reference model. If a reference model is not supplied, or default values are used for a particular region within the reference model, then bounds can be supplied to restrict the physical properties in that region to some approximate limits based on known, or expected, geology. Where a reference model is supplied, the reference model might be used to define the expected physical property value (perhaps with a low certainty, or w_s value), but the bounds can be used to allow the likely range of values, even if the physical properties are skewed (Figure 3) or bimodal.

The directional smoothness weighting factors, w_x, w_y, and w_z, can be more difficult to apply since they are defined for dual meshes corresponding to the boundaries between cells (Figure 4). Low values can be used to encourage breaks in smoothness of the model across known faults or lithological boundaries; higher values can also be used to define regions where geological strike has different orientations on either side of a contact or fault, in much the same way as the α values except on a local scale.

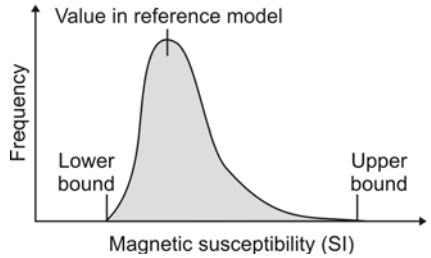

Figure 3. Schematic histogram of magnetic susceptibility measurements for a particular lithologic unit demonstrating how bounds and a reference model could be used to attempt to reproduce the likely physical properties for that unit in a recovered inversion model.

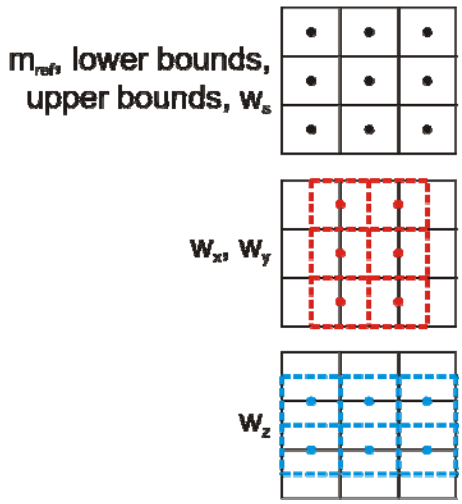

Figure 4. 2D slices through an $n \times n \times n$ mesh (n = 3) indicating how the weighting parameters are set. The reference model properties, bounds, and w_s values are set on the primal grid defined by the black mesh, with black dots indicating cell centres. The smoothness weightings, w_x (as viewed from above or the south), w_y (as viewed from the east), and w_z (as viewed from the south or east), indicate the smoothness to be assigned to each cell boundary in the primal mesh and are defined on three dual grids (red and blue dashed lines) that have their cell centres (red and blue dots) at the boundaries of the primal grid.

Positions of geological boundaries and contacts can also be recovered in the absence of physical property information by applying breaks in smoothness between cells. Assigning low values (< 1) to the w_x, w_y, and w_z parameters along the boundary or contact will allow the overall model smoothness to "break" across the boundary as required to fit the geophysical data and any other geological constraints (Figure 5). This technique is most effective where the bounding surface completely separates two volumes of rock.

An effective way to set reference properties, bounds, or weightings is within the Gocad model building package (Gocad, 2006). Regions are defined based on bounding surfaces, and properties are set for each of those regions or surfaces within a 3D mesh. Sample or drill-hole information

can also be "painted" into the 3D mesh. The mesh can then be translated into UBC–GIF file formats.

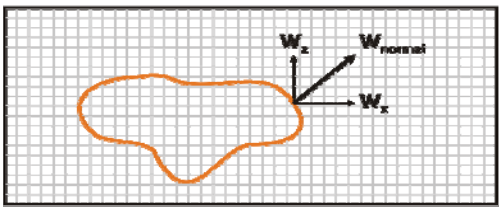

Figure 5. 2D slice through a model where the position of a unit is know or inferred, but its properties are too poorly known to include as constraints. The boundary can be encouraged by assigning a low ($0 \leq w_{normal} << 1$) smoothness to the surface's normal vector (a default property of any surface in Gocad) and separating that into east-west, north-south, and vertical components at each surface vertex (using a trigonometric script within Gocad). These smoothness components can then be "painted" into the appropriate smoothness dual meshes.

Smooth Model or Smooth Model-Difference?

As described above, the model objective function smoothes differences between the recovered model and the reference model over a number of cells. Where the reference model is the default zero model, or is constant through the model, this results in smoothly varying properties. However, if the reference model is defined differently in adjacent cells, the model recovered using smooth model differences can have step discontinuities in the property values (Figure 6). In some situations this may be desirable, but in the majority of cases where the available information is restricted to incomplete surface exposure and a limited number of drill holes, it is generally preferred that the inversion produce a smooth extrapolation of the assigned reference properties out some distance into the model.

One way to achieve this with the existing UBC–GIF inversion software is to only use the default zero reference model, in which case the model objective function effectively becomes:

$$\phi_m(m) = \alpha_s \int_V w_s \left[w_r(z)(m) \right]^2 dV + \ldots$$
$$\alpha_x \int_V w_x \left[\frac{\partial}{\partial x} w_r(z)(m) \right]^2 dV + \ldots$$
$$\alpha_y \int_V w_y \left[\frac{\partial}{\partial y} w_r(z)(m) \right]^2 dV + \ldots$$
$$\alpha_z \int_V w_z \left[\frac{\partial}{\partial z} w_r(z)(m) \right]^2 dV$$

which will always result in a smooth model. This gives a geologically-unconstrained, default, result.

Since the reference model is zero, any available geological constraints must be applied using only upper and lower bounds. The bounds can be made very narrow if the property for a particular cell is well known, or wider if the property is less well known, equivalent to using w_s values to assign certainty to a the property. This also allows the certainty to be more rigorously defined: perhaps 2σ or 3σ, where σ is the standard deviation of the available property measurements appropriately treated to take into account the volume of each model cell. In areas where no constraints are available, appropriate wide bounds are used.

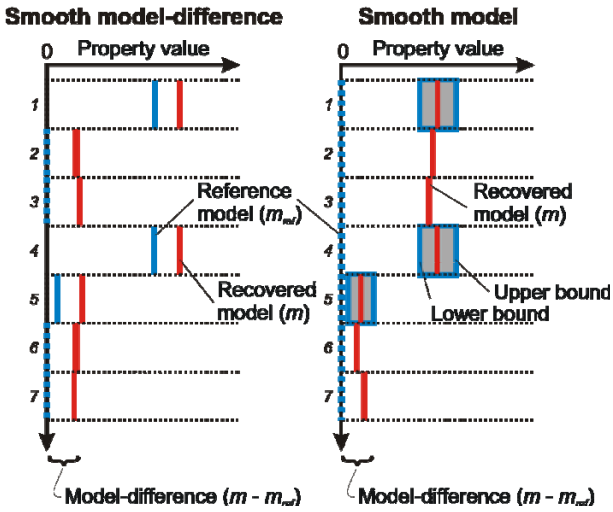

Figure 6. Schematic plots showing recovered models obtained using either a *smooth model-difference* inversion or a *smooth model* inversion within the existing UBC–GIF inversion programs. The seven horizontal numbered bands represent a subset of seven adjacent cells somewhere in an inversion model. The physical properties of cells 1, 4, and 5 are known from drilling. In the left image the reference properties are defined as non-zero for only those three cells (with appropriate high w_s), and default wide bounds are used everywhere. For the right image the reference model is zero everywhere, but tight bounds are used to define the expected properties for the three known cells, with default bounds used elsewhere. In the *smooth model-difference* style the recovered model shows step discontinuities in the physical property wherever the reference model changes. In the *smooth model* style the property is encouraged to vary smoothly away from the constrained cells as required by the geophysical data and other constraints, yet discontinuities can be enforced where known.

A downside of using smooth model inversions is that the inversion is less able to recover sharp boundaries in the model. However, where the positions of boundaries are known, they can be readily recovered by defining different bounds ranges on either side of the boundaries, or by enforcing a break in smoothness with the w_x, w_y, and w_z weighting functions. In other situations, the physical properties are not known well enough to enforce bounds so using a reference model and weightings would provide a better solution. The UBC–GIF software has the flexibility to accommodate these scenarios as well as many others.

In general, the smooth model-difference approach using a non-zero reference model is ideal for regions where a full 3D geology model exists or one can be constructed. In other regions where the raw geological information (mapping, sampling, drill holes, cross sections, etc.) is all that is available, smooth model inversions using only bounds are preferable.

EXAMPLE GEOLOGICALLY-CONSTRAINED INVERSION

A simplified synthetic example that is geologically based on a scenario of nickel exploration in the Eastern Goldfields of Yilgarn Craton in Western Australia is used in this section to demonstrate the benefits of including even a small number of geological constraints in a gravity inversion for mineral exploration. The area has a dipping, north-south-striking granite-greenstone basement, but extensive regolith cover limits basement outcrop. To simplify building the true geology for this example, and to aid visualisation of the results using cross-sections, the north-south-strike is made perfect (similar to a 2.5D model) but full 3D constraints, data, and inversion are used. The topography is flat.

A greenstone belt outcrops above the centre of the volume to be modelled. An ultramafic unit with significant massive sulphide Ni-mineralisation is present in this belt, and additional ultramafic horizons have been identified further to the southwest. It is hoped these continue into the area of interest, and that there may be massive sulphide lenses that have been structural detached from their original ultramafic host rocks.

A surface geological map of the area is shown in Figure 7. The gravity data for the area is on an even 100 m grid, and was upward continued to 100 m, the width of the inversion cells, to remove high frequency information that could only be accommodated by cells smaller than 100-m-wide. The data has a noise level of 0.1 mGal (added to the response of the true model prior to upward continuation). The black line indicates the position of the cross-sections shown in Figure 8.

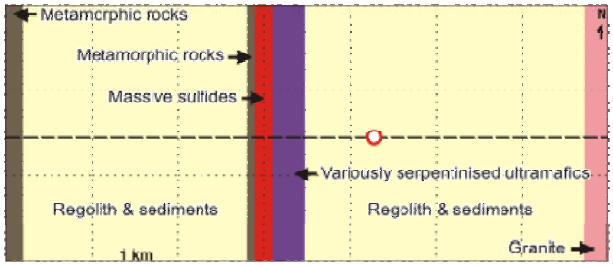

Figure 7. Surface map of the example area, coloured by rock type. The grid cells are 1 km square (total map size is 3 km by 7 km), the dashed black line shows the position of the cross-sections shown in Figure 8, and the red circle is the collar location of a 2 km long stratigraphic drill hole inclined to the west.

Appropriate padding cells were added to the inversion volume. A uniform cell size of 100 m × 100 m × 50 m was used throughout. The smoothness parameters, c, were set to default values in the east-west and vertical directions, $c_x = c_z = 4$, but the known north-south strike suggests a higher value of c_y. It is estimated that rock properties may correlate over a north-south distance of 2 km, so using $n_y = 20$ cells, a smoothness parameter of $c_y = 2 \cdot (n_y - 7) = 26$ is obtained. For the desired cell size appropriate α values are therefore $\alpha_s = 0.0001$, $\alpha_x = 4$, $\alpha_y = 26$, and $\alpha_z = 1$.

A default, geologically-unconstrained GRAV3D inversion was performed on the data and the result is shown in Figure 8(b). Comparison of the result with the known geology shown in Figure 8(c) shows clear differences, especially the lack of low density regolith cover. Six regions of high density that may be considered prospective sulphide targets are marked by labels (1-2). The location labelled (3) shows such high densities over such a large region that it is unlikely to be entirely geological; it probably indicates excess mass that should be distributed elsewhere in the model. There are also two large low density regions (labelled 4) at depth in the model that are also geologically unlikely, indicating further problems in the model. Comparison of the result to the true geology shown in Figure 8(f) reveals that only three of the six potential sulphide targets contain sulphides; these are labelled (1).

Since the recovered model does not reproduce the observed geology, it would be beneficial to include the observed geology into the inversion as constraints to guide the inversion and enhance the result. The constraints shown in Figure 8(c-d) were implemented using the densities shown in Table 1. Densities were converted to density contrasts by subtracting a rough estimate of the mean density (~2.8 g/cm^3) within the inversion volume. The central outcrop of ultramafic, sulphides, and minor metamorphic rock is well sampled for density, allowing narrow density bounds to be used. The surrounding regolith material and the small basement outcrops to the west and east have been sampled, but not in detail, so wider bounds were used. A 1 km^2, 250-m-deep 3D geological model was available for the central area. It was based on outcrop mapping, structural interpretation, and shallow drilling. There is high confidence in this model, and enough density measurements to assign narrow bounds on the expected densities. One 2-km-long stratigraphic drill hole was drilled from east to west through the ultramafic. Although it did not intersect any mineralisation, detailed density measurements were taken.

Table 1. Summary of the densities obtained from measurements of outcrop samples, drill hole samples, and similar rocks in other areas. The densities of sulphidic rocks allow for variable sulphide abundances in cells containing sulphides. The lower table shows the density contrast assigned for each unit based on a loosely estimated mean density for the inversion volume of 2.8 g/cm^3 given the expected rock types.

Rock Type	Density (g/cm^3)	Minimum (g/cm^3)	Maximum (g/cm^3)
Sulphides	3.8	3.4	4.2
Ultramafic	3.1	2.75	3.2
Metamorphics	2.8	2.6	3.0
Granite	2.7	2.65	2.85
Regolith	2.0	1.7	2.4

Rock Type	Density Contrast (g/cm^3)	Minimum (g/cm^3)	Maximum (g/cm^3)
Sulphides	1.0	0.6	1.4
Ultramafic	0.3	-0.05	0.4
Metamorphics	0.0	-0.2	0.2
Granite	-0.1	-0.15	0.05
Regolith	-0.8	-1.1	-0.4

More general constraints, or zones, were also established. The depth of weathering and cover is not known, but is expected to be less than 250 m everywhere. Above a depth of 250 m, there could be either low density regolith material or relatively high

density basement rocks. Below 250 m, only the higher density basement rocks are expected, so the lower bound on density contrasts can be made more restrictive.

Given the sparse constraining information (i.e., the thin layer of outcrop, and single deep drill hole), a smooth model style of inversion was chosen. This was achieved by using a default zero reference model and assigning all the physical property constraints using bounds. As described above and in Figure 8(c-d), varying certainties in the available physical properties were accommodated by using narrower or wider bounds. The results of the geologically-constrained inversion are shown in Figure 8(e) and the true density model is shown in Figure 8(f).

The geologically-constrained inversion fits the gravity data to the same accuracy as the default inversion, but by including the sparse constraints available from the outcrop, a small 3D model, one deep drill hole, and some general expectations about the geology, the geologically-constrained inversion was able to reasonably recover the depth to basement, labelled as 'A' in the Figure 8(e), as well as the position, dip, and extent the central ultramafic unit, labelled 'B'. The spurious features labelled '2', '3' and '4' in Figure 8(b) were also eliminated. However, the inversion still required some anomalous mass on the margins of the ultramafic labelled 'C' and 'D' to fit the data. As there is currently no geological or geophysical explanation for these accumulations, they provide targets for further drilling. In this synthetic example, drilling in the vicinity of 'C' would result in the discovery of new sulphide zones. The anomalous mass accumulation predicted at 'E' in the geologically constrained inversion is enhanced relative to the result obtained in the default inversion and also relative to other shallow anomalies in the default inversion. This might prompt further investigation of the region and further sulphide discoveries.

CONCLUSIONS

In the development of this simple example, the importance of including surface information was clear, even if that surface information doesn't come from outcropping basement rocks. The input of weathered material as a constraint is just as important as delineating basement units due to the low density (and commonly low magnetic susceptibility) of weathered material relative to basement rocks. The addition of deep drill hole data is also key, provided they are sufficiently sampled for physical properties. Drilling enables samples to be obtained and these provide tight constraints at depth. Full or partial 3D models provide tight restrictions on the physical properties of a large volume of rock in comparison with the cell-by-cell restrictions provided by sparse outcrop and drilling data, and consequently restrict the range of possible inverse solutions.

The constraints used in geological zones are derived from carefully assessment of background geological information. A geoscientist may quickly dismiss a recovered inversion model as being "wrong". The question should then be why? What information is available to suggest that the model is wrong? If specific and reliable, this information can and must be included as constraints to derive a realistic model. Often, the criticism of an unconstrained inversion result is based on general knowledge, such as not expecting extremely low density material in basement rocks below 250 m depth. Careful use of bounds, reference models, and weightings can allow this information to be included with the more specific located data. Together, these constraints help recover not only a more geologically-realistic model, but one that is also more accurate.

A specific capacity of the UBC-GIF inversion software is its flexibility to include as much or as little geological information as is available. In early phases of exploration, default inversions may be used to locate possible anomalous regions. As more geological data becomes available during exploration, more constraints can be included to further refine the recovered physical property models, and therefore enhance their potential for targeting. 3D models are clearly a powerful dataset to include if available, but constraints based on the raw geological data can be included without the additional interpretation required to build a full 3D model.

The key to successful development and refinement of geophysical inverse models is to include all available information and to update the constraints as soon as new information becomes available.

ACKNOWLEDGMENTS

The UBC–GIF research group, especially Doug Oldenburg, Roman Shekhtman, Nigel Phillips, and Peter Lelievre have all contributed much to my understanding. Richard Lane at Geoscience Australia is also thanked for helpful reviews, discussions, guidance, and for initially introducing me to potential field inversions.

REFERENCES

Farquharson, C.G., and Oldenburg, D.W., 2004, A comparison of automatic techniques for estimating the regularization parameter in non-linear inverse problems: Geophysics Journal International 156, 411-425.

Gocad, 2006, Welcome to the Gocad Research Group! http://www.gocad.org/, Website accessed 3 January 2007.

Li, Y., and Oldenburg, D.W., 1996, 3-D inversion of magnetic data: Geophysics 61, 394-408.

Li, Y., and Oldenburg, D.W., 1998a, 3-D inversion of gravity data: Geophysics 63, 109-119.

Li, Y., and Oldenburg, D.W., 1998b, Separation of regional and residual magnetic field data: Geophysics 63, 431-439.

UBC–GIF, 2005a, GRAV3D Version 3.0 software manual: A Program Library for Forward Modelling and Inversion of Gravity Data over 3D Structures.

UBC–GIF, 2005b, MAG3D Version 4.0 software manual: A Program Library for Forward Modelling and Inversion of Gravity Data over 3D Structures.

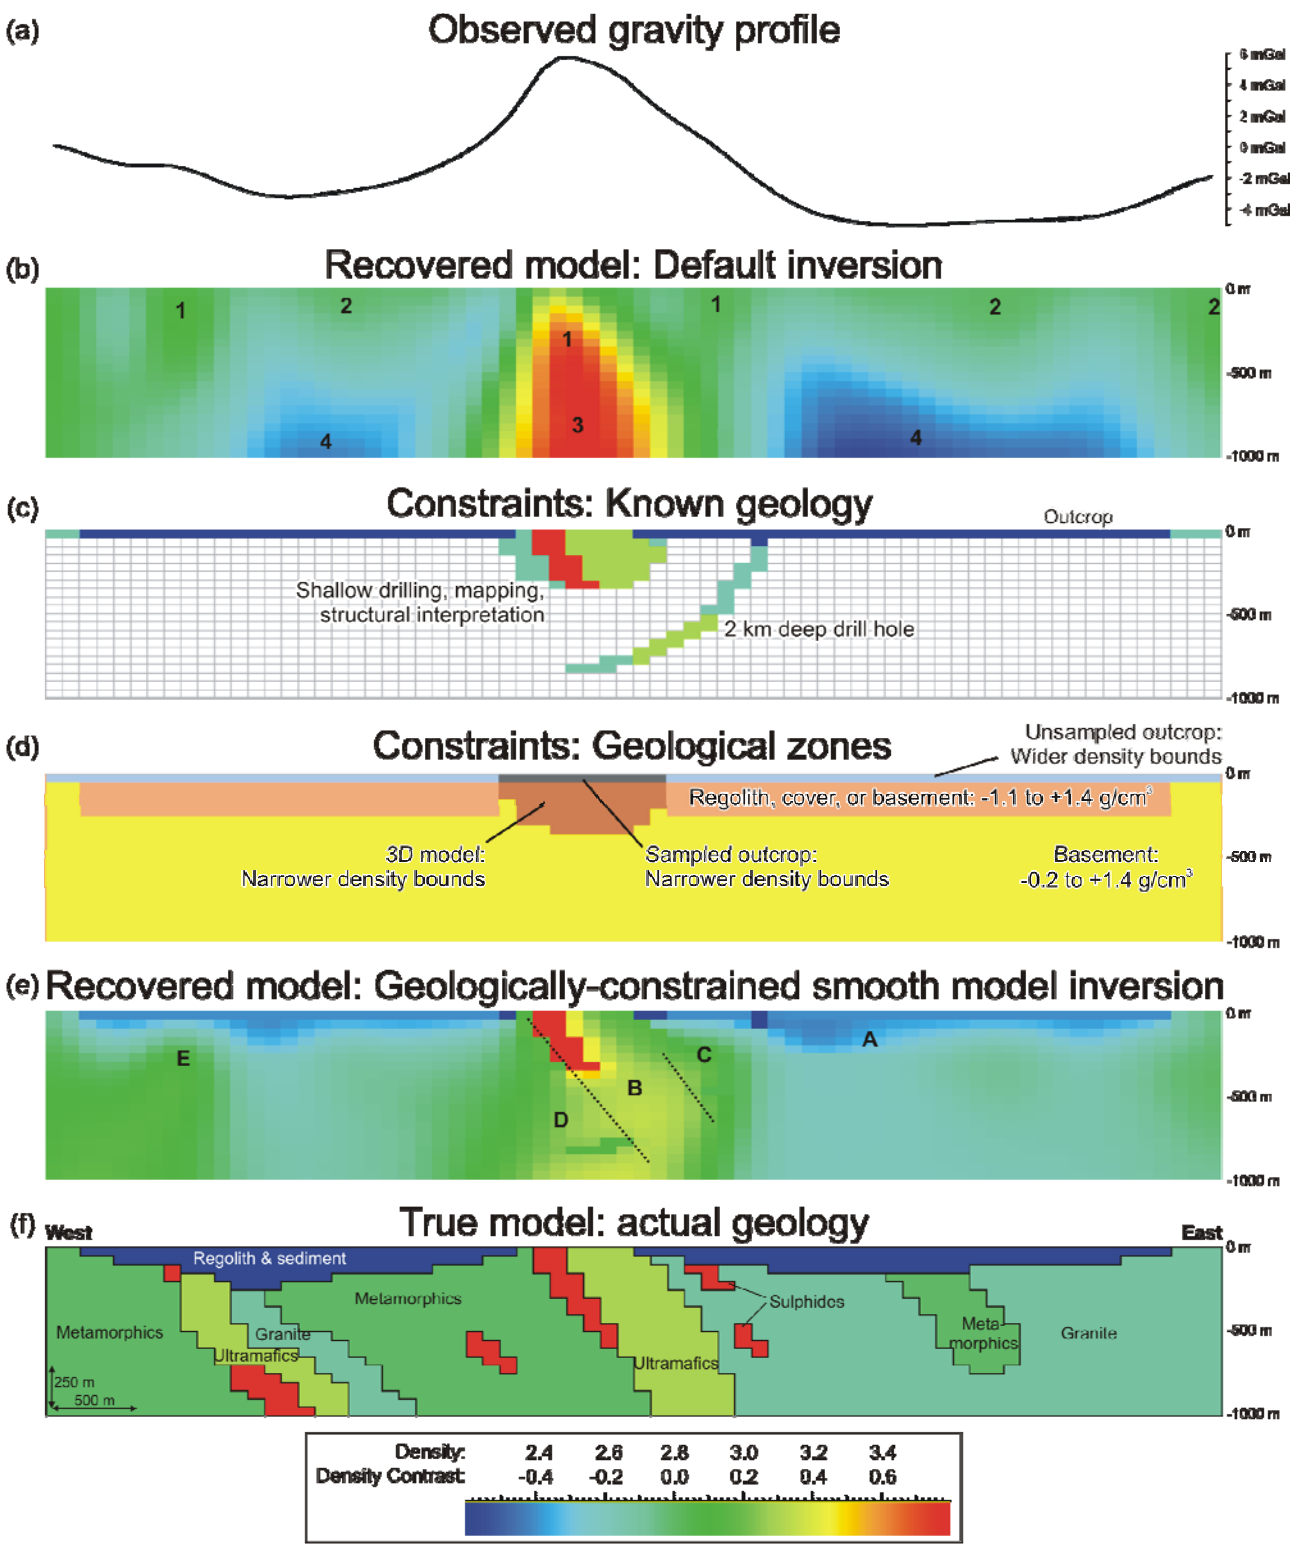

Figure 8. East-west cross sections (along the black line in Figure 7) showing a profile through (a) the gravity data, (b) the default inversion result, (c-d) available geological constraints, (e) the geologically-constrained inversion result, and (f) the true geology. The cross section extends from surface (0 m) to a depth of 1000 m, and is 7 km wide; cells are all 100 m × 100 m and 50 m high. All cells (except for the zones shown in (d)) are coloured by expected density contrast (converted from densities by subtracting an estimated mean density of 2.8 g/cm^3) with the same scale. The default inversion (b) shows: (1) regions with anomalous high densities that correctly identify ore positions, (2-3) regions with anomalous high densities not associated with ore, and (4) regions with anomalously low densities. The geologically constrained inversion result captures several of the key characteristics of the true model such as the depth to basement (A) and easterly ultramafic dip (B), and may suggest the presence of anomalous mass accumulations at (C), (D), and (E). The geologically constrained result has also eliminated the spurious potential targets shown in (b) at locations 2 and 3, as well as the deep low density features (4).