

Large-scale 3D inversion of potential field data

Martin Čuma^{1*}, Glenn A. Wilson² and Michael S. Zhdanov³

¹Center for High Performance Computing, University of Utah, 155 South, 1452 East, Room 405, Salt Lake City, Utah, 84112, USA and TechnoImaging, 4001 South, 700 East, Suite 500, Salt Lake City, Utah, 84107, USA, ²TechnoImaging, 4001 South, 700 East, Suite 500, Salt Lake City, Utah, 84107, USA, ³Department of Geology and Geophysics, University of Utah, 115 South, 1460 East, Room 383, Salt Lake City, Utah, 84112, USA and TechnoImaging, 4001 South, 700 East, Suite 500, Salt Lake City, Utah, 84107, USA.

Received July 2011, revision accepted November 2011

ABSTRACT

Inversion of gravity and/or magnetic data attempts to recover the density and/or magnetic susceptibility distribution in a 3D earth model for subsequent geological interpretation. This is a challenging problem for a number of reasons. First, airborne gravity and magnetic surveys are characterized by very large data volumes. Second, the 3D modelling of data from large-scale surveys is a computationally challenging problem. Third, gravity and magnetic data are finite and noisy and their inversion is ill posed so regularization must be introduced for the recovery of the most geologically plausible solutions from an infinite number of mathematically equivalent solutions. These difficulties and how they can be addressed in terms of large-scale 3D potential field inversion are discussed in this paper. Since potential fields are linear, they lend themselves to full parallelization with near-linear scaling on modern parallel computers. Moreover, we exploit the fact that an instrument's sensitivity (or footprint) is considerably smaller than the survey area. As multiple footprints superimpose themselves over the same 3D earth model, the sensitivity matrix for the entire earth model is constructed. We use the re-weighted regularized conjugate gradient method for minimizing the objective functional and incorporate a wide variety of regularization options. We demonstrate our approach with the 3D inversion of 1743 line km of FALCON gravity gradiometry and magnetic data acquired over the Timmins district in Ontario, Canada. Our results are shown to be in good agreement with independent interpretations of the same data.

Key words: Footprint, Inversion, Parallelization.

1 INTRODUCTION

Reliance on desktop computing has limited the use of 3D potential field inversion, making it impractical to achieve an appropriate level of resolution and detail for geological interpretation of entire airborne surveys that are characterized by very large data volumes. Such surveys typically contain hundreds to thousands of line kilometres of data with measurement locations every few metres, covering thousands of

square kilometres in an area. Regional surveys may be even larger and denser as the result of merging multiple and/or historic surveys. For large surveys, structural interpretations are usually based on a choice or combination of Euler deconvolution (e.g., Thompson 1982), eigenvector analysis (e.g., Beiki and Pedersen 2010), wavelet analysis (e.g., Hornby, Boschetti and Horowitz 1999), analytic signal (e.g., Salem and Ravat 2003; Beiki 2010), or depth-from-extreme-points (DEXP) methods (e.g., Fedi 2007). While such methods may provide information about the sources of the potential field, it is not immediately obvious how this information can be quantified in terms of the physical properties in a 3D earth

*Email: glenn@technoimaging.com

model. For this reason, inversion marks an important step in quantitative interpretation - particularly at deposit scales.

Generalized inversion methods first discretize the 3D earth models into cells of constant density and/or susceptibility. As gravity and magnetic inversion is ill posed, regularization must be introduced so as to recover the most geologically plausible solutions from the infinite number of mathematically equivalent solutions. Regularization effectively selects the class of models from which a unique solution is sought. Over the years, a variety of methods have been developed for 3D inversion of potential field data with both smooth (e.g., Li and Oldenburg 1996, 1998; Li 2001) and focusing (e.g., Portnaguine and Zhdanov 2002; Zhdanov 2002, 2009; Zhdanov, Ellis and Mukherjee 2004) regularization. The computational resources needed for deposit-scale discretization (e.g., <25 m cells) for entire airborne surveys easily exceeds the capacity of high-end desktop computers. From recent case studies, the term *large-scale* can be read to mean 3D earth models with about 3 million cells. This reflects the limited capacity of existing 3D inversion algorithms on desktop computers. In practice, large airborne surveys are usually divided into subsets (circa 1 - 3 million cells each) of reasonably coarse discretization (e.g., 250 m cells) and each of those inverted with the resulting 3D earth models stitched or tiled together post-inversion (e.g., Phillips et al. 2010).

In the present paper, we consider a general method of solving truly large-scale potential field inverse problems with massive parallelization where the modelling domain is discretized into hundreds of millions (even billions) of cells. For such models, computation of the predicted data and direction of steepest descent at every iteration of the inversion is not trivial. We solve these by exploiting the limited sensitivity of the instrument and introducing a moving footprint approach similar to that developed for airborne electromagnetics (AEM) by Cox, Wilson and Zhdanov (2010) and for magnetotellurics (MT) by Zhdanov et al. (2011). Our inversion method is designed to invert any component of the gravity and magnetic fields, including the total, vector, and gradient components. For magnetics, we assume in this paper that the magnetic fields are caused by induced magnetization only. A subsequent paper shall address remanent magnetization. We have developed this as a fully parallelized software that exhibits near-linear scaling, meaning that it can be scaled to thousands of cores.

2 3D MODELING

Gravity (U_g) and magnetic (U_H) potentials are linear with respect to the 3D density (ρ) and magnetic susceptibility (χ)

distributions as per the first-kind Fredholm equations:

$$U_g(\mathbf{r}') = \iiint_V \psi_g(\mathbf{r}', \mathbf{r}) \rho(\mathbf{r}) dv, \quad (1)$$

$$U_H(\mathbf{r}') = \iiint_V \psi_M(\mathbf{r}', \mathbf{r}) \chi(\mathbf{r}) dv, \quad (2)$$

where the volume integration is over the variable, \mathbf{r} and where the kernel functions ψ_g and ψ_H are the respective gravity and magnetic potential Green's functions. All components of the gravity and magnetic fields can be derived from spatial differentiation of $\psi_{g,H}$ in equations (1) and (2), respectively. If the density and susceptibility are constant in each cell, Poisson's relation can be used to transform gravity data and 3D density models to magnetic data and 3D susceptibility models and vice versa (e.g., Li and Oldenburg 1998). Closed-form solutions for the volume integrals in equations (1) and (2) over right rectangular prisms of constant density and/or susceptibility were derived for both gravity (e.g., Okabe 1979) and magnetic (e.g., Bhattacharyya 1980) fields. While exact, these analytic solutions are inefficient to implement. For example, the gravity response due to a right rectangular prism of constant density contains 16 logarithms and 8 arctangents (Li and Chouteau 1998, p. 344). However, the volume integrals in equations (1) and (2) can be evaluated numerically:

$$U_g(\mathbf{r}') = \sum_{i=1}^{N_m} \psi_g(\mathbf{r}', \mathbf{r}_i) \rho_i \Delta x_i \Delta y_i \Delta z_i, \quad (3)$$

$$U_H(\mathbf{r}') = \sum_{i=1}^{N_m} \psi_H(\mathbf{r}', \mathbf{r}_i) \chi_i \Delta x_i \Delta y_i \Delta z_i, \quad (4)$$

where N_m is the total number of cells in the earth model and $\mathbf{r}_i = x_i, y_i, z_i, ; i = 1, \dots, N_m$. Zhdanov (2009) showed that for gravity gradiometry, single-point Gaussian integration with pulse basis functions was as accurate as the analytic solution provided the depth to the centre of the cell exceeded twice the dimension of the cell; i.e., the 3D earth model can be discretized to cubic cells of 40 m dimension or less for an airborne gravity gradiometry survey with 80 m nominal ground clearance. Moreover, numerical integration decreases the run time by more than an order of magnitude compared to the analytic solution. For land or borehole surveys where instruments are located close to or within the cells, higher order Gaussian integration and basis functions can be used to attain the desired level of accuracy.

Since equations (1) and (2) are linear, they can be written in discrete form as the operator equations:

$$\mathbf{d}_g = \mathbf{A}_g \rho, \quad (5)$$

$$\mathbf{d}_H = \mathbf{A}_H \chi, \quad (6)$$

where \mathbf{d}_g and \mathbf{d}_H are the N_d length vectors of observed gravity and magnetic data, \mathbf{A}_g and \mathbf{A}_H are the $N_d \times N_m$ gravity and magnetic operators and ρ and χ are the N_m length vectors of density and susceptibility. Equations (5) and (6) are general in that the data \mathbf{d}_g and \mathbf{d}_H , and operators \mathbf{A}_g and \mathbf{A}_H can be partitioned to contain multiple components. In the following sections we present the explicit forms of the corresponding forward modelling operators.

2.1 Gravity and gravity gradiometry

The gravity field, $\mathbf{g}(\mathbf{r})$, satisfies the following:

$$\nabla \cdot \mathbf{g}(\mathbf{r}) = -4\pi\gamma\rho(\mathbf{r}), \quad \nabla \times \mathbf{g}(\mathbf{r}) = \mathbf{0}, \quad (7)$$

where γ is the universal gravitational constant and $\rho(\mathbf{r})$ is the 3D density distribution within a domain, D . The gravity field can also be expressed by the gravity potential $U_g(\mathbf{r})$:

$$\mathbf{g}(\mathbf{r}) = \nabla U_g(\mathbf{r}). \quad (8)$$

The gravity potential has the following solution:

$$U_g(\mathbf{r}') = \gamma \iiint_V \frac{1}{|\mathbf{r} - \mathbf{r}'|} \rho(\mathbf{r}') dv. \quad (9)$$

As the first spatial derivatives of the gravity potential, the gravity field can be expressed as:

$$\mathbf{g}(\mathbf{r}') = \gamma \iiint_V \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|^3} \rho(\mathbf{r}') dv, \quad (10)$$

for which the components can be written as

$$g_\alpha(\mathbf{r}') = \gamma \iiint_V K_\alpha(\mathbf{r} - \mathbf{r}') \frac{1}{|\mathbf{r} - \mathbf{r}'|^3} \rho(\mathbf{r}') dv, \quad (11)$$

where the kernel $K_\alpha(\mathbf{r} - \mathbf{r}')$ is equal to:

$$K_\alpha(\mathbf{r} - \mathbf{r}') = \alpha - \alpha', \quad \alpha = x, y, z. \quad (12)$$

The second spatial derivatives of the gravity potential,

$$g_{\alpha\beta}(\mathbf{r}) = \frac{\partial^2}{\partial\alpha\partial\beta} U_g(\mathbf{r}), \quad \alpha, \beta = x, y, z, \quad (13)$$

form a symmetric *gravity tensor*:

$$\hat{\mathbf{g}} = \begin{bmatrix} g_{xx} & g_{xy} & g_{xz} \\ g_{yx} & g_{yy} & g_{yz} \\ g_{zx} & g_{zy} & g_{zz} \end{bmatrix}, \quad (14)$$

with zero trace and where:

$$g_{\alpha\beta} = \frac{\partial g_\alpha}{\partial\beta}, \quad \alpha, \beta = x, y, z. \quad (15)$$

The gravity tensor components can be calculated from equation (12):

$$g_{\alpha\beta}(\mathbf{r}') = \gamma \iiint_V K_{\alpha\beta}(\mathbf{r} - \mathbf{r}') \frac{1}{|\mathbf{r} - \mathbf{r}'|^3} \rho(\mathbf{r}') d^3r, \quad (16)$$

where the kernels $K_{\alpha\beta}$ are equal to:

$$K_{\alpha\beta}(\mathbf{r} - \mathbf{r}') = \begin{cases} 3 \frac{(\alpha' - \alpha)(\beta' - \beta)}{|\mathbf{r} - \mathbf{r}'|^2}, & \alpha \neq \beta, \\ 3 \frac{(\alpha' - \alpha)^2}{|\mathbf{r} - \mathbf{r}'|^2} - 1, & \alpha = \beta, \end{cases} \quad \alpha, \beta = x, y, z. \quad (17)$$

A combination of the gravity gradients,

$$g_\Delta = \frac{1}{2} (g_{xx} - g_{yy}), \quad (18)$$

is measured directly by the FALCON airborne gravity gradiometer, and it can be calculated from the following:

$$g_\Delta(\mathbf{r}') = \gamma \iiint_V K_\Delta(\mathbf{r} - \mathbf{r}') \frac{1}{|\mathbf{r} - \mathbf{r}'|^3} \rho(\mathbf{r}') dv, \quad (19)$$

where the kernel K_Δ is given by:

$$K_\Delta(\mathbf{r} - \mathbf{r}') = \frac{3}{2} \frac{(x - x')^2 - (y - y')^2}{|\mathbf{r} - \mathbf{r}'|^2}. \quad (20)$$

Kernels for other combinations of gravity gradients such as those measured by Rio Tinto's VK1 airborne gravity gradiometer can be easily derived. Equations (11), (16) and (19) can be evaluated numerically as per equation (3).

In practice, gravity and gravity gradiometry data require terrain corrections prior to 3D inversion. The terrain correction is generally calculated as the response due to the volume bound by an upper surface of the digital elevation model and a lower surface of a plane that passes through the lowest elevation of the survey area. This volume is discretized into prisms of constant density, for which the responses (11), (16) and (19) are subtracted from the respective free-air data (e.g., Dransfield and Zeng 2009). For mineral exploration, a uniform terrain density between 2.6–3.1 g/cm³ is usually applied, depending on the host geology. For sedimentary basins, the lower surface is generally deeper and the terrain density increases as a function of depth to simulate compaction (G. Jorgensen 2011, pers. comm.).

2.2 Magnetism and magnetic gradiometry

In what follows, we adopt the common assumption that there is no remanent magnetization, that the self-demagnetization effect is negligible and that the magnetic susceptibility is isotropic. Under such assumption, the intensity of magnetization $\mathbf{I}(\mathbf{r})$ is linearly related to the inducing magnetic field,

$\mathbf{H}^0(\mathbf{r})$, through the magnetic susceptibility $\chi(\mathbf{r})$:

$$\mathbf{I}(\mathbf{r}) = \chi(\mathbf{r})\mathbf{H}^0(\mathbf{r}). \quad (21)$$

For any magnetic survey, the inclination (I), declination (D) and azimuth (A) of the inducing magnetic field (in degrees) can be predicted from the International Geomagnetic Reference Field (IGRF) model. It is not the intent of this paper to discuss how the IGRF (or equivalent) is calculated. We do note however that various online calculators exist to compute different generations of the IGRF models (e.g., NOAA) or localized models, such as the Australian Geomagnetic Reference Field (AGRF) model (e.g., Geoscience Australia). Often, airborne contractors provide IGRF reduced magnetic data as a standard deliverable. We assume the inclination is positive below the horizontal, the declination is positive east of true north and the azimuth is positive east of north, so the directional cosines of the inducing magnetic field are as follows:

$$l_x = \cos(I) \cos(D - A), \quad (22)$$

$$l_y = \cos(I) \sin(D - A), \quad (23)$$

$$l_z = \sin(I). \quad (24)$$

The magnetic field, $\mathbf{H}(\mathbf{r})$, has the solution (Zhdanov 1988):

$$\mathbf{H}(\mathbf{r}') = \nabla' \iiint_V \nabla \frac{1}{|\mathbf{r} - \mathbf{r}'|} \cdot \mathbf{H}^0 \chi(\mathbf{r}) dv. \quad (25)$$

It is well-known that the last formula can be rewritten as follows:

$$\mathbf{H}(\mathbf{r}') = H^0 \iiint_V \frac{\partial}{\partial l'} \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|^3} \chi(\mathbf{r}) dv, \quad (26)$$

where H^0 is the magnitude of the inducing field and $\partial/\partial l'$ is a directional derivative in the direction of magnetization, \mathbf{l} , where \mathbf{l} is a unit vector, i.e., $\mathbf{H}^0 = H^0 \mathbf{l}$. These directional derivatives are the directional cosines of equations (22) – (24). Note that

$$\frac{\partial}{\partial l'} \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|^3} = \frac{-1}{|\mathbf{r} - \mathbf{r}'|^3} \left[\mathbf{l} - \frac{3(\mathbf{l} \cdot (\mathbf{r} - \mathbf{r}'))(\mathbf{r} - \mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|^2} \right]. \quad (27)$$

Substituting equation (27) into equation (26), we obtain:

$$\mathbf{H}(\mathbf{r}') = -H^0 \iiint_V \frac{1}{|\mathbf{r} - \mathbf{r}'|^3} \left[\mathbf{l} - \frac{3(\mathbf{l} \cdot (\mathbf{r} - \mathbf{r}'))(\mathbf{r} - \mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|^2} \right] \chi(\mathbf{r}) dv. \quad (28)$$

Now, each component of the magnetic field can be written as

$$H_\alpha(\mathbf{r}') = H^0 \iiint_V K_\alpha(\mathbf{r} - \mathbf{r}') \frac{1}{|\mathbf{r} - \mathbf{r}'|^3} \chi(\mathbf{r}) dv, \quad (29)$$

where the kernel $K_\alpha(\mathbf{r} - \mathbf{r}')$ is equal to

$$K_\alpha(\mathbf{r} - \mathbf{r}') = \frac{3[l_x(x - x') + l_y(y - y') + l_z(z - z')](\alpha - \alpha')}{|\mathbf{r} - \mathbf{r}'|^2} - l_\alpha, \quad \alpha = x, y, z. \quad (30)$$

The amplitude of the magnetic field (or total magnetic intensity) has the following form:

$$H_\alpha(\mathbf{r}') = H^0 \iiint_V K_{TMI}(\mathbf{r} - \mathbf{r}') \frac{1}{|\mathbf{r} - \mathbf{r}'|^3} \chi(\mathbf{r}) dv, \quad (31)$$

where the kernel $K_{TMI}(\mathbf{r} - \mathbf{r}')$ is equal to

$$K_{TMI}(\mathbf{r} - \mathbf{r}') = \frac{3[l_x(x - x') + l_y(y - y') + l_z(z - z')]^2}{|\mathbf{r} - \mathbf{r}'|^2} - 1. \quad (32)$$

Equations (29) and (31) can be evaluated numerically as per equation (4). Derivatives of equation (29) and their respective kernels that form the symmetric magnetic tensor will be detailed in a subsequent paper.

3 3D INVERSION

Potential field data are finite and noisy and their inversion is inherently non-unique, meaning that there is an infinite number of source distributions that can equally satisfy the observed data. In order to solve this inverse problem, regularization must be introduced. The goal of regularization is to recover the most geologically plausible solutions from the infinite number of mathematically equivalent solutions. We note that our term ‘geologically plausible’ is a subjective one and based entirely upon the interpreter’s prejudice. As we have generalized our inversion methodology for both gravity and magnetic data, we describe our model parameters by vector \mathbf{m} , of length N_m , whether they be values of density or susceptibility. Regardless of the iterative scheme used, most regularized inversions seek to minimize the Tikhonov parametric functional, $P^\alpha(\mathbf{m})$:

$$P^\alpha(\mathbf{m}) = \phi(\mathbf{m}) + \alpha s(\mathbf{m}) \rightarrow \min, \quad (33)$$

where $\phi(\mathbf{m})$ is a misfit functional of the observed and predicted potential field data, $s(\mathbf{m})$ is a stabilizing functional. α is the regularization parameter that balances the trade-off between the misfit and stabilizing functionals and is decreased as the inversion progresses (Zhdanov 2002). Data and model weights can be introduced to equation (33) through data and model weighting matrices. We can also re-weight the inverse problem in logarithmic space in order to reduce the dynamic range of both the data and model parameters. For magnetic

inversion, positivity is enforced upon the susceptibility so as to obtain physically meaningful solutions.

In our implementation, all weighting functions are selected based upon their integrated sensitivity (Zhdanov 2002). Our weighting functions provide equal sensitivity of the observed data to cells located at different depths and at different horizontal positions. Thus, our weighting functions automatically introduce appropriate corrections for the vertical and horizontal distribution of the density or susceptibility. This is one of the main differences between our approach and the geometric weighting functions developed by Li and Oldenburg (1996, 1998).

We base our minimization of equation (33) on the re-weighted regularized conjugate gradient method (RRCG). The specifics of this method have been published previously (e.g., Zhdanov 2002) and are not repeated here. This method iteratively updates the vector of model parameters \mathbf{m} so as to minimize the vector of residual errors, \mathbf{r} , akin to:

$$\mathbf{m}_{i+1} = \mathbf{m}_i + k_i \mathbf{A}_{g,H}^T \mathbf{r}_i \text{ subject to } \mathbf{r}_i \rightarrow \min, \quad (34)$$

where k_i is a step length and $\mathbf{A}_{g,H}^T$ is the conjugate transpose of the $N_d \times N_m$ matrix of the gravity or magnetic linear operator. The inversion proceeds to iterate in a manner similar to equation (34) until the residual error reaches a preset threshold, the decrease in error between multiple iterations is less than the preset threshold, or a maximum number of iterations is reached. Upon completion, the quality of the inversion is appraised by the data misfit and visual inspection of the model. Results are also available from any prior iteration, enabling the user to investigate the evolution of the 3D model during inversion, and to re-start the inversion from a prior iteration if necessary.

3.1 Focusing regularization

All geological constraints manifest themselves as regularization that can be quantified through a choice of data weights, model upper and lower bounds, model weights, an a priori model and the type of stabilizing functional. The latter incorporates information about the class of models used in the inversion. The choice of stabilizing functional should be based on the user's geological knowledge and prejudice. In what follows, we will briefly describe different smooth and focusing stabilizers in order to demonstrate the results from the 3D inversion of potential field data produced by each.

A minimum norm (MN) stabilizer will seek to minimize the norm of the difference between the current model and an a priori model:

$$s_{MN}(\mathbf{m}) = \iiint_V (\mathbf{m} - \mathbf{m}_{apr})^2 dv \quad (35)$$

and usually produces a relatively smooth model. Gradient (G) stabilizers implicitly introduce smoothness with the first derivatives of the model parameters:

$$s_G(\mathbf{m}) = \iiint_V (\nabla \mathbf{m} - \nabla \mathbf{m}_{apr})^2 dv, \quad (36)$$

and can result in spurious oscillations and artefacts when the model parameters are discontinuous. In practice, a combination of stabilizers (35) and (36) is often used (e.g., Li and Oldenburg 1996, 1998). However, very little geology exhibits smooth density or susceptibility distributions. Geology is usually characterized by sharp boundaries of contrasting density and/or susceptibility, for example, between an ore deposit and host rock, or across a discontinuity. As such, stabilizers (35) and (36) or their combinations produce results that bear little or no relevance to the actual geology.

Portniaguine and Zhdanov (1999) introduced focusing stabilizers that make it possible to recover models with sharper boundaries and contrasts. We briefly describe these stabilizers here and refer the reader to Zhdanov (2002) for further details. First, we present the minimum support (MS) stabilizer:

$$s_{MS}(\mathbf{m}) = \iiint_V \frac{(\mathbf{m} - \mathbf{m}_{apr})^2}{(\mathbf{m} - \mathbf{m}_{apr})^2 + e^2} dv, \quad (37)$$

where e is a focusing parameter introduced to avoid singularity when $\mathbf{m} = \mathbf{m}_{apr}$. The minimum support stabilizer minimizes the volume with non-zero departures from the a priori model, effectively recovering compact bodies. Thus, a smooth distribution of all model parameters with a small deviation from the a priori model is penalized. A focused distribution of the model parameters is penalized less. The focusing parameter effectively tunes the degree of focusing. Similarly, we present the minimum gradient support (MGS) stabilizer:

$$s_{MGS}(\mathbf{m}) = \iiint_V \frac{\nabla(\mathbf{m} - \mathbf{m}_{apr}) \cdot \nabla(\mathbf{m} - \mathbf{m}_{apr})}{\nabla(\mathbf{m} - \mathbf{m}_{apr}) \cdot \nabla(\mathbf{m} - \mathbf{m}_{apr}) + e^2} dv, \quad (38)$$

which minimizes the volume with non-zero gradients, i.e., sharp transitions in the model parameters are penalized less than smooth transitions.

As an example of the difference between smooth (e.g., minimum norm) and focusing (e.g., minimum support) stabilizers, we used both to invert GETMAG full tensor magnetic gradiometry from a magnetite skarn near Tallawang, New South Wales, Australia (Schmidt et al. 2004). The results are shown in Fig. 1, and clearly demonstrate how the focusing inversion result is more geologically reasonable than the smooth inversion result.




Figure 1 Comparison of focusing (minimum support; panels a-c) and smooth (minimum norm; panels d-f) regularization for 3D inversion of GETMAG full tensor magnetic gradiometry data from Tallawang, New South Wales, Australia (Schmidt et al. 2004). The vertical cross-sections correspond to profile 60 m N. The geology inferred from drilling has been superimposed for further comparison.

3.2 Moving footprint approach

In principle, the regularized inversion methods outlined above can be applied to large-scale problems. Numerically, however, the computational complexity increases linearly with the size of the problem, meaning 3D inversion faces two major obstacles. First is the large amount of computer memory required for storing the kernels of the forward modelling operators, which double as sensitivities for linear problems. Even a small-sized 3D inversion of thousands of data to 3D earth models of hundreds of thousands of cells can exceed memory available for desktop computers. The second obstacle is the amount of CPU time required for the application of the dense matrix of the forward modelling operators to data and model vectors.

To minimize these problems, an alternative formulation has been to exploit the translational invariance of the kernels so as to reduce the matrices to a Toeplitz block structure and use FFTs for matrix-vector multiplication (e.g., Pilkington 1997; Zhdanov et al. 2004). Such strategies minimize memory limitations and reduce the CPU time dramatically. However, these

methods require the data to lie over a regular grid on a flat surface above the topography. This is applicable in some special cases, such as marine gravity. However, for airborne surveys, the data must be upward continued to a flat surface. This is not ideal, particularly for surveys incorporating very rugged terrain.

Other strategies for large-scale 3D inversion involve some form of compression (e.g., Portniaguine and Zhdanov 2002; Li and Oldenburg 2003; Davis and Li 2011). However, the linear operators A_g and A_H can still be too large to store and manipulate on desktop computers for 3D inversion of tens of thousands of data to models with millions of cells. In practice, large surveys are usually divided into subsets and each of these is inverted with the resulting 3D earth models stitched together post-inversion (e.g., Phillips et al. 2010).

In potential fields, sensitivity of the data to the variations of the density or magnetic susceptibility is expressed via the appropriate kernel functions of the forward modelling operators, i.e., via the corresponding Green's functions. For example, Dransfield (2010) used these Green's functions to




Figure 2 Percent of total response of the integrated sensitivity for the vertical gravity component, various gravity gradients and total magnetic intensity as a function of footprint radius. Note that as the footprint is symmetric in the x and y directions, some gravity gradients (e.g., g_{xx} and g_{yy}) overlap.

investigate airborne gravity gradiometry instrument sensitivity. He demonstrated that at some limited distance, which we call the *footprint*, the receiver is no longer sensitive to the 3D earth model. Typically, the size of the footprint is much less than the size of an airborne survey. The footprint size for gravity fields is proportional to $1/r^2$, and for gravity gradiometry and magnetic fields it is proportional to $1/r^3$.

Cox et al. (2010) introduced the concept of a moving footprint for the 3D inversion of AEM data. They showed that, for a single transmitter-receiver pair, there was no need to calculate the responses or sensitivities beyond the AEM system's footprint. The sensitivity matrix for the entire 3D earth model could be constructed as the superposition of footprints from all transmitter-receiver pairs. Zhdanov et al. (2011) also introduced a footprint approach for the large-scale 3D inversion of MT data. The framework of this approach can be described as follows: for a given receiver, compute and store the sensitivities for those cells within the receiver's footprint. The size of the footprint is based on the sensitivity.

As an analogue of this, we introduce a moving footprint for 3D potential field inversion. First, we provide an example. We can assume we have an instrument 60 m above a homogeneous half-space. Figure 2 shows the integrated sensitivities for each of the gravity fields and gravity gradients. The figure shows that the gravity gradients have approximately 95% of the sensitivity within a 15 km footprint. In the same figure, we also show integrated sensitivity for the total magnetic intensity

(TMI). The sensitivity of the TMI with respect to the footprint radius behaves similarly to the gravity gradients with $\sim 95\%$ of the sensitivity being within the 15 km footprint. This is fully expected since kernels (17), (20) and (32) have similar spatial dependencies. Past a 15 km radius, the sensitivity decays very slowly, though increasing the footprint radius beyond 15 km is not practical. Therefore, we conclude that 15 km is an optimal footprint radius for gravity gradiometry and TMI.

With a moving footprint, the linear operators A_g and A_M operate as sparse matrices with reduced memory and computational requirements and minimal loss of accuracy. The number of non-zero elements in each row of the linear operators is just the number of elements within each footprint rather than the total number of elements in the model. In practice, we avoid direct storage of the linear operators by re-computing them as required. On one hand, this adds extra computational burden to each iteration, even though these computations are of simple geometric functions and quickly evaluated. On the other hand, by not explicitly constructing the linear operators, our memory requirements are extremely small, meaning we are effectively unlimited in the size of the inverse problem that we can consider.

In our implementation, we note that the linear operators A_g and A_M can be partitioned to contain an arbitrary combination of components of the gravity or magnetic fields, respectively. As such, we can jointly invert any combination of gravity or magnetic data but not joint gravity *and* magnetic data. Note also that our implementation of the linear operators A_g and A_H includes both topography and variable flight height, obviating the need at least for upward continuation of data. For improved inversion results, terrain corrections to gravity data are still required though, particularly in regions of rough topography. Opportunity does exist however, to perform terrain corrections as part of the inversion process and this is a subject of on-going research.

3.3 Parallel implementation

Our 3D inversion software is implemented as a multilevel parallel application. The 3D earth model is divided in a distributed fashion over the Message Passing Interface (MPI). On a fine-grained level, loops over the observation points (i.e., data points) and a few other auxiliary loops within each MPI process are further parallelized with a shared memory OpenMP standard. This two-level approach has multiple advantages. First, it reduces the number of MPI communicating processes, greatly reducing communication stress on the network. It also tends to save memory since there are data

structures needing to be replicated by each process and most of the data are shared by the OpenMP threads. Finally, it allows for better locality of the processes/threads on the node's boards and sockets, which improves data transfers to/from the main memory. This data locality is critical on current non-uniform memory architecture (NUMA) computers with a growing number of CPU cores.

In a typical cluster configuration, we run one or two MPI processes per cluster node and each of these processes launches a number of OpenMP threads - one thread per processor core. The current generation of clusters ships with two hexa-core CPUs (i.e., 12 cores) per node. We found that it is optimal to run one MPI process per socket (i.e., two per node), with six OpenMP threads per MPI process. This has the advantage of pinning the process to the CPU socket so that it does not move from one socket to another, which improves the memory performance. We found that without pinning, the performance can degrade by up to 20%.

Another important aspect of the program is file management. We need to load the inputs for the inversion (i.e., observed data and a priori model) and save the outputs from 3D inversion (i.e., predicted field data and final model). For large-scale inversion, the sheer size of these data sets requires an intelligent file management design. We implemented several I/O options. Plain text and/or binary I/O is very slow for 3D models with hundreds of millions to billions of cells. For improved I/O performance, we incorporated on-the-fly zip file compression and both serial and parallel HDF5 support. Both the zip files and serial HDF5 files are compressed, but we end up with one file per MPI process (i.e., potentially hundreds of files per 3D model). Parallel HDF5 enables us to write a single file from all MPI processes. However, parallel HDF5 does not support compression. In our tests, we found that the compression is critical for good I/O performance. Furthermore, the zip compression is slightly faster than HDF5, so we used zip compression for all I/O presented in this paper.

In order to be able to invert millions of data to 3D earth models with hundreds of millions to billions of cells, we evaluate the footprint radii and linear operators within each footprint as needed. For large-scale models, even just storing the footprint cell indices for each datum would exceed available system memory. We optimized computations involving the kernels, re-using as many of the variables as possible and re-ordering the loops in order to exploit the vectorization instructions of modern CPUs. This resulted in a several-fold speedup as compared to unoptimized versions. Our 3D inversion is relatively light in MPI communication, largely thanks to the linearity of the forward modelling operators, which makes all

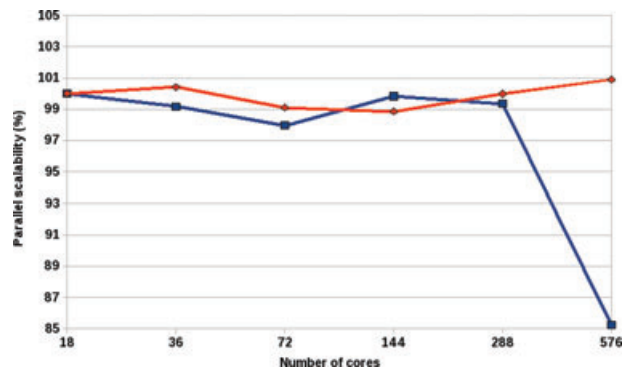


Figure 3 Parallel scaling efficiency for 3D inversion of the Vredefort FALCON airborne gravity gradiometry case study of Wilson et al. (2011). Strong scaling is shown in blue and weak scaling is shown in red.

cells independent of each other. Most MPI communication is located in accumulation of the sensitivities and in regularization as reduction operations. Thanks to this, the program exhibits excellent parallel scaling.

Parallel scaling is usually evaluated with two different metrics. The first one is called *strong scaling*, which measures performance of a fixed problem size with an increasing number of processing elements. If the amount of time to finish the work on one processing element is t_1 and the time to finish work on N processing elements is t_N , then the strong scaling efficiency is expressed as:

$$N \frac{t_1}{t_N} * 100\%. \quad (39)$$

Another parallel scaling evaluation metric is *weak scaling*, which relates the time to complete one unit of work on one processing element, t_1 , to the time to perform N units of work on N processing elements, t_N , as:

$$\frac{t_1}{t_N} * 100\%. \quad (40)$$

In both cases, ideal (linear) scaling is 100%. Any scaling below 100% is sublinear. Any scaling above 100% is super-linear. We note that it is possible to achieve superlinear scaling due to hardware architectural features that multiprocessor programs can exploit.

We evaluated the parallel efficiency of our software using the University of Utah Center for High Performance Computing's Ember cluster, which has 260 nodes, each equipped with two hexa-core Intel Xeon CPUs (12 cores/node) running at 2.8 GHz with 24 GB of RAM and a QDR InfiniBand interconnect. Figure 3 shows the parallel scaling efficiency with the Vredefort FALCON case study by Wilson, Cuma and

Zhdanov (2011). In the case of strong scaling, as depicted by the blue line in Fig. 3, we chose a 3D model with ~ 11 million cells and ~ 600000 data. The scaling efficiency is excellent from 18–288 cores. We see a drop at 576 cores. This is due to running 12 rather than 6 threads per process. The memory load is much more uneven for the single MPI process sharing threads on both CPU sockets within the node, which decreases the efficiency by 15%. The weak scaling, depicted by the red line in Fig. 3, varied the number of inversion cells from ~ 11 million cells on 18 cores to ~ 350 million cells on 576 cores. Again, the scaling is nearly linear with a 1–2 % difference, which can be attributed to system noise. The outcome from the scaling analysis is several-fold. First, our 3D inversion software shows near-linear scaling and it is expected to scale well to massively parallel computers with thousands of cores. Second, we identified how process and thread locality is critical in achieving optimal performance, with MPI processes bound one per socket being the best choice.

4 CASE STUDY - TIMMINS, ONTARIO

We independently applied our large-scale 3D potential field inversion with a moving footprint to FALCON airborne gravity gradiometry (AGG) and TMI data acquired over the Timmins district in Ontario, Canada. Timmins is representative of Archean terranes and is located in the Abitibi Subprovince, including east of the Mattagami River Fault. The area is underlain by Archean (~ 2.7 b.y.) mafic intrusives and metavolcanic rocks in the east and felsic to intermediate metavolcanic rocks in the west. NNW-striking Proterozoic diabase dykes are evident in the TMI data. Copper and lead-zinc vein/replacement and stratabound, volcanogenic massive sulphide (VMS) mineralization occur in the immediate vicinity (Snyder et al. 2008). For example, the Kidd Creek VMS deposit occurs within the survey area.

4.0.1 Survey description

The FALCON AGG system is designed to measure horizontal curvature components of the gravity gradient, i.e., g_{xy} and $g_{\Delta} = \frac{g_{xx} - g_{yy}}{2}$. The full gravity tensor is then derived from these measured components (Lee 2001). In October 2009, Fugro Airborne Surveys flew a FALCON survey of 1743 line km over the Timmins area. The survey was flown east-west with a line spacing of 500 m and with two north-south tie-lines spaced at 5 km, over an area 36 km east-west by 22.5 km north-south. The ground clearance was nominally 140 m flown in a drape over the terrain. The measured gradients were processed by the usual multistep FALCON processing procedures. After the

initial reduction of error due to the residual effects of aircraft motion, the data were demodulated and low-pass filtered with a 6th order Butterworth low-pass filter at a cut-off frequency of 0.18 Hz. The demodulated data were corrected for the self-gradient effects of the aircraft and the tie-lines were levelled. These differential curvature gravity gradient data were further processed to produce terrain-corrected data using a density of 2.67 g/cm^3 and thence the full gravity gradient tensor. In the processing, a low-pass filter with a cut-off wavelength of 750 m was applied to the data. TMI data were also acquired and we inverted upon the anomalous TMI; i.e., the IGRF was removed from the observed TMI.

4.0.2 3D inversion

For both FALCON and TMI data, the $37.8 \text{ km} \times 61.9 \text{ km} \times 2.4 \text{ km}$ inversion domain was discretized to more than 128 million 25 m cubic cells. The inversion domain conformed to topography and contained no a priori density or susceptibility models. Inversion was run for 127 733 AGG stations, each containing all seven gravity gradients (g_{zz} , g_{xx} , g_{yy} , g_{xy} , g_{xz} , g_{yz} , g_{Δ}), giving a total of 894 131 data. Though some of the gravity gradients are redundant, we found that their joint inversion effectively suppresses noise across all channels. Inversion was also run for 596 092 TMI stations. Our previous experience with moving footprint inversion (e.g., Cox et al. 2010) indicated that inversion of redundant data does not aid model recovery.

From an initial weighted misfit of 100% for all FALCON data, the 3D inversion converged to a final weighted misfit of 19% after 50 iterations. Comparisons of the observed and predicted data for all gravity gradient components are shown in Fig. 4, and as one can see, the fit between observed and predicted data is excellent. The 3D FALCON inversion required approximately 29 hours to complete, running on 12 nodes, each equipped with two dual-core (i.e., 8) Intel Xeon CPUs running at 2.4 GHz with 24 GB of RAM and a QDR InfiniBand interconnect.

Similarly, from an initial weighted misfit of 100% for the TMI data, the 3D inversion converged to a final misfit of 10% after 90 iterations. Comparisons of the observed and predicted TMI data are shown in Fig. 5. The 3D FALCON inversion required approximately 39 hours to complete, running on 12 nodes, each equipped with two dual-core (i.e., 8) Intel Xeon CPUs running at 2.4 GHz with 24 GB of RAM and a QDR InfiniBand interconnect.

In Fig. 6, we compare horizontal cross-sections of our 3D FALCON and TMI inversion results with the published

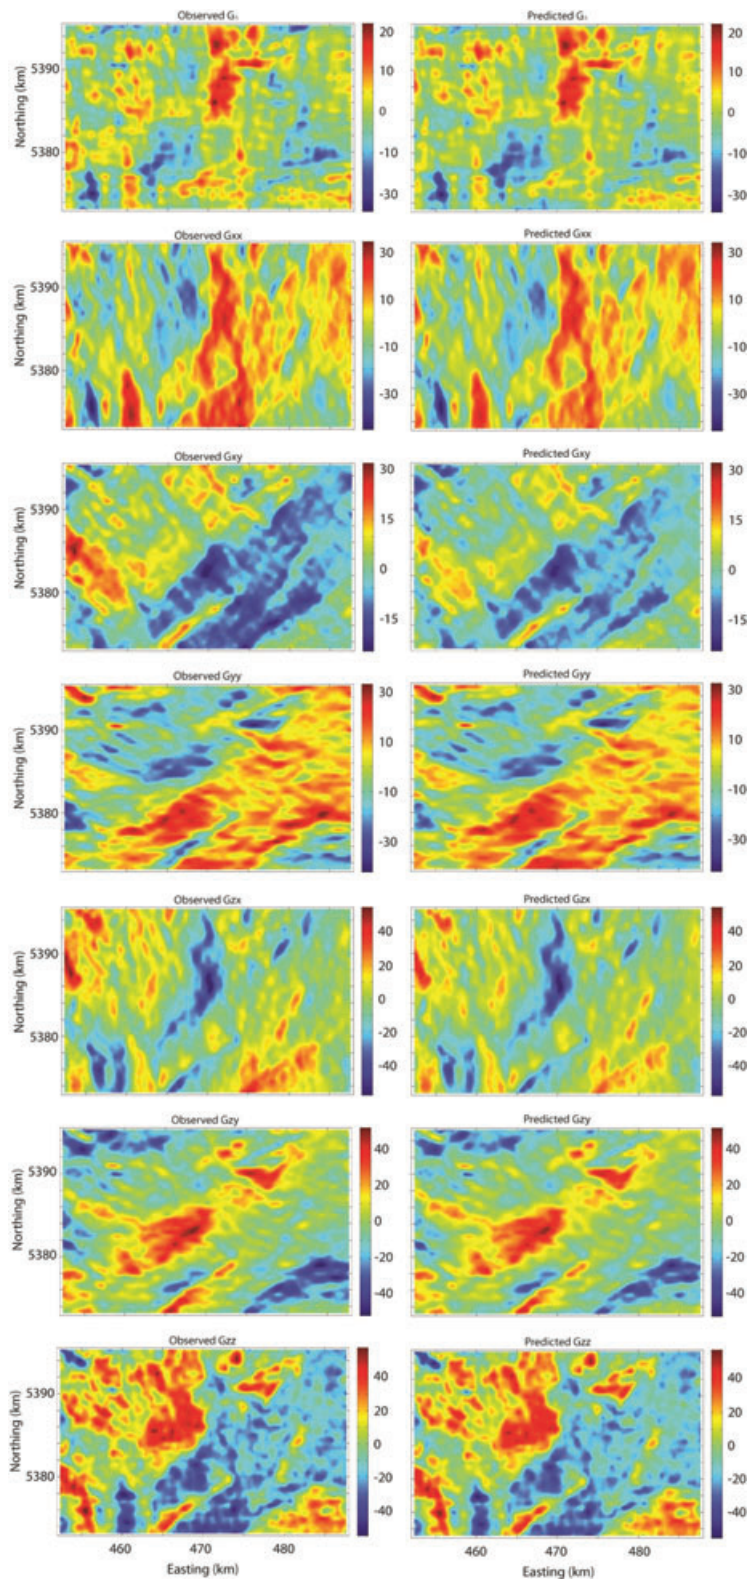


Figure 4 Observed (left) and predicted (right) FALCON data from Timmins.




Figure 5 Observed (left) and predicted (right) anomalous TMI data for Timmins.

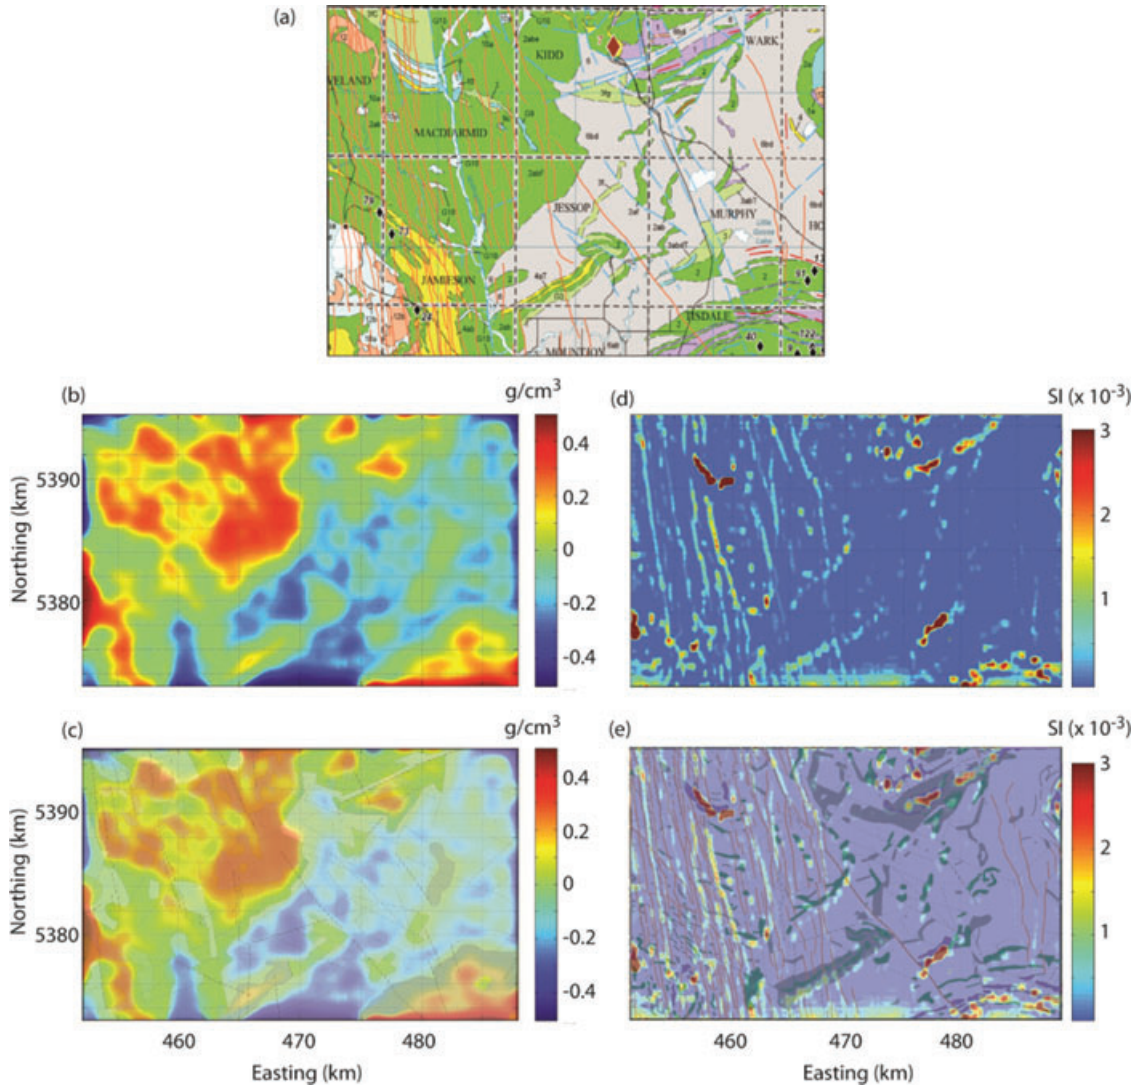


Figure 6 Comparison of (a) geology of the Timmins survey area; (b) a horizontal cross-section through the 3D density model obtained from joint inversion of all FALCON gravity gradiometry data at -1213 m ASL, i.e., ~1500 m depth from the topographic peak; (c) structural interpretation (courtesy of Fugro Airborne Surveys) superimposed on the same horizontal cross-section through the 3D density model; (d) a horizontal cross-section through the 3D susceptibility model obtained from inversion of TMI data at -63 m ASL, i.e., ~300 m depth from the topographic peak; and (e) structural interpretation (courtesy of Fugro Airborne Surveys) superimposed on the same horizontal cross-section through the 3D susceptibility model.

1:250 000 geological map for the Timmins district and with structural interpretations made by Fugro Airborne Surveys. We can observe very good agreement between our inversion results and the inferred geology. In Fig. 7, we present 3D perspectives of the density and susceptibility models. We note that while the diabase dikes are expected to continue to depth, their low amplitudes (circa 250 nT) are sensitive to only several hundred metres depth. Two major intrusive units with high amplitudes (circa 2000 nT) are shown to extend to depth. We note that the FALCON and TMI inversions are recovering different geological structures and this is to be expected. The TMI data are dominated by major intrusives and N-S trending diabase dikes. The FALCON data are dominated by regional tectonics such as high density mafic intrusives and metavolcanics and low density felsic metavolcanics. Together, their interpretations are complimentary of the actual geology.

While the optimal footprint size can be estimated analytically for any desired level of accuracy (e.g., system noise), we practically demonstrate the effectiveness of using a moving footprint with the following table summarizing 3D inversion performance for TMI. We note that all inversions ran to a common misfit. The maximum susceptibility difference is relative to the 3D inversion with no footprint. Use of a footprint that is too small introduces errors in the model while still satisfying the data. We conclude that for 3D inversion with a 15 km footprint, there is negligible difference in the recovered 3D model.

Footprint size (km)	Maximum susceptibility difference (%)
5	30
10	2
15	0.6
20	0.14

5 CONCLUSIONS

We developed a methodology for practical yet truly large-scale 3D regularized inversion of gravity and magnetic data, whether these data be any combination of total field, vector, or gradient components. We implemented a variety of regularization options, including both smooth and focusing stabilizers. We introduced a moving footprint approach, which allows us to reduce memory requirements significantly, as well as operation counts for matrix-vector multiplications. Our parallelization was implemented in such a manner that we have

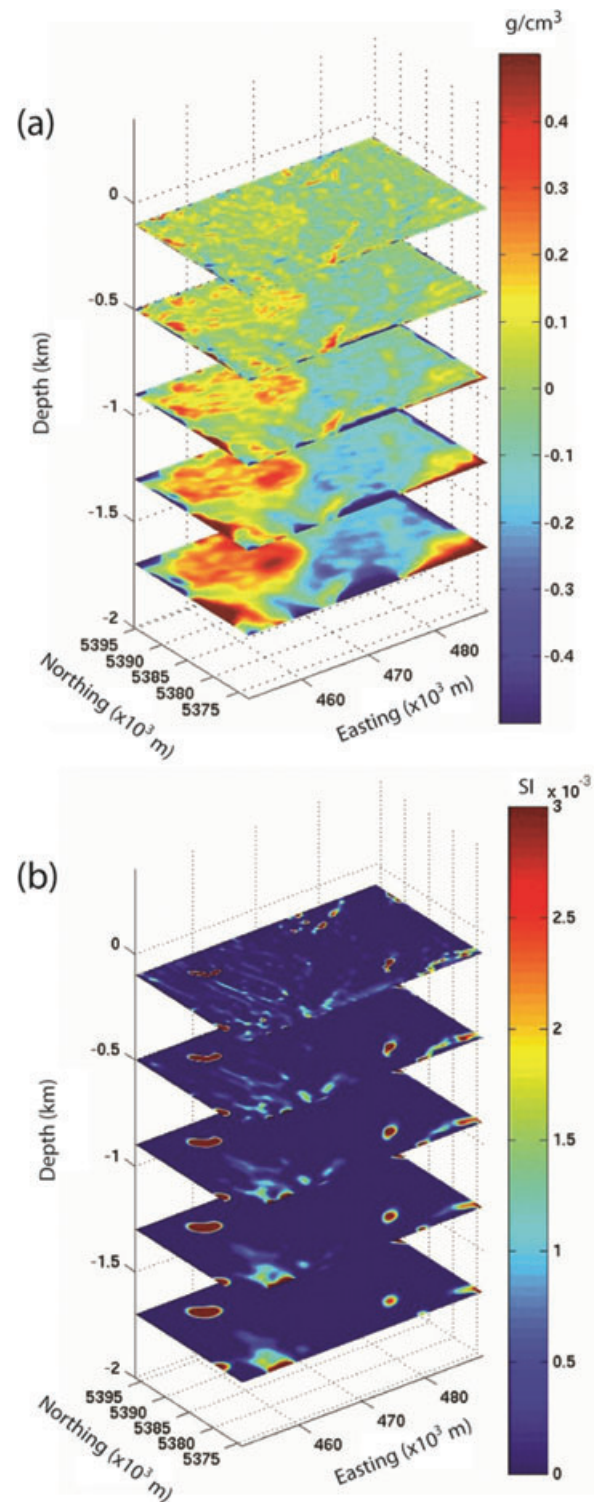


Figure 7 3D perspective view of the Timmins (a) density model obtained from 3D inversion of the FALCON gravity gradiometry data and (b) susceptibility model obtained from 3D inversion of the TMI data.

near-linear scaling. This means we can confidently implement our software on massively parallel architectures and are effectively unlimited in the size of inversion problem we can consider. This has significant implications for airborne surveys, as we are now able to invert large surveys to deposit-scale model discretization. We demonstrated this with an example for both FALCON and TMI from the Timmins district in Ontario, Canada.

In terms of further developments, we implemented 3D magnetic inversion for remanent magnetization and applied this to full tensor magnetic gradiometry (FTMG) data. The theory and related case study will be the subject of a separate publication. We have presently optimized our software to the point that nearly all runtime is dedicated to kernel computations. As these are simple geometric functions, they lend themselves to massive parallelization via GPUs, and this is a subject of on-going development.

6 ACKNOWLEDGEMENTS

The authors acknowledge TechnoImaging for support of this research and permission to publish.

Čuma and Zhdanov acknowledge support of the University of Utah's Consortium for Electromagnetic Modeling and Inversion (CEMI) and Center for High Performance Computing (CHPC).

The authors thank CSIRO Materials Science and Engineering, in particular Drs. Phil Schmidt, Keith Leslie, David Clark and Cathy Foley, for providing the Tallawang GEMTAG data and related material. Hongzhu Cai of CEMI is thanked for producing the 3D GETMAG inversion.

Fugro Airborne Surveys, in particular Dr. Marc A. Vallée, are acknowledged for release of the Timmins FALCON data and related material and permission to publish.

REFERENCES

- Beiki M. 2010. Analytic signals of gravity gradient tensor and their application to estimate source location. *Geophysics* **75**, 159–174, doi: 10.1190/1.3493639.
- Beiki M. and Pedersen L.B. 2010. Eigenvector analysis of gravity gradient tensor to locate geologic bodies. *Geophysics* **75**, 137–149, doi: 10.1190/1.3484098.
- Bhattacharyya B.K. 1980. A generalized multibody model for inversion of magnetic anomalies. *Geophysics* **29**, 517–531, doi: 10.1190/1.1441081.
- Cox L.H., Wilson G.A. and Zhdanov M.S. 2010. 3D inversion of airborne electromagnetic data using a moving footprint. *Exploration Geophysics* **41**, 250–259, doi: 10.1071/EG10003.
- Davis K. and Li Y. 2011. Fast solution of geophysical inversion using adaptive mesh, space-filling curves and wavelet compression. *Geophysical Journal International* **185**, 157–166, doi: 10.1111/j.1365-246X.2011.04929.x.
- Dransfield M. 2010. Potential field searchlights. Presented at the 72nd EAGE Conference and Exhibition, Barcelona.
- Dransfield M. and Zeng Y. 2009. Airborne gravity gradiometry: Terrain corrections and elevation error. *Geophysics* **74**, 137–142, doi: 10.1190/1.3170688.
- Fedi M. 2007. DEXP: A fast method to determine the depth to the sources of potential fields. *Geophysics* **72**, L1–L11, doi: 10.1190/1.239945215.
- Hornby P., Boschetti F. and Horowitz F.G. 2002. Analysis of potential field data in the wavelet domain. *Geophysical Journal International* **137**, 175–196, doi: 10.1046/j.1365-246x.1999.00788.x.
- Lee J.B. 2001. FALCON gravity gradiometer technology. *Exploration Geophysics* **32**, 247–250, doi: 10.1071/EG01247.
- Li Y. 2001. 3-D inversion of gravity gradiometry data. 71st Annual Meeting, SEG, Expanded Abstracts, 1470–1473, doi: 10.1190/1.1816383.
- Li X. and Chouteau M. 1998. Three-dimensional gravity modeling in all space. *Surveys in Geophysics* **19**, 339–368, doi: 10.1023/A:1006554408567.
- Li Y. and Oldenburg D.W. 1996. 3-D inversion of magnetic data. *Geophysics* **61**, 394–408, doi: 10.1190/1.1443968.
- Li Y. and Oldenburg D.W. 1998. 3-D inversion of gravity data. *Geophysics* **63**, 109–119, doi: 10.1190/1.1444302.
- Li Y. and Oldenburg D.W. 2003. Fast inversion of large-scale magnetic data using wavelet transforms and a logarithmic barrier method. *Geophysical Journal International* **152**, 251–265, doi: 10.1046/j.1365-246X.2003.01766.
- Okabe M. 1979. Analytical expressions for gravity anomalies due to homogeneous polyhedral bodies and translations into magnetic anomalies. *Geophysics* **64**, 730–741, doi: 10.1190/1.1440973.
- Phillips N., Nguyen T.H., Thomson V., Oldenburg D. and Kowalczyk P. 2010. 3D inversion modelling, integration, and visualization of airborne gravity, magnetic, and electromagnetic data - The QUEST Project. Presented at the *International Workshop on Electrical, Gravity and Magnetic Methods*, Capri.
- Pilkington M. 1997. 3-D magnetic imaging using conjugate gradients. *Geophysics* **62**, 1132–1142, doi: 10.1190/1.1444214.
- Portniaguine O. and Zhdanov M.S. 1999. Focusing geophysical inversion images. *Geophysics* **64**, 874–887, doi: 10.1190/1.1444596.
- Portniaguine O. and Zhdanov M.S. 2002. 3-D magnetic inversion with data compression and image focusing. *Geophysics* **67**, 1532–1541, doi: 10.1190/1.1512749.
- Salem A. and Ravat D. 2003. A combined analytic signal and Euler method (AN-EUL) for automatic interpretation of magnetic data. *Geophysics* **68**, 1952–1961, doi: 10.1190/1.1635049.
- Schmidt P., Clark D., Leslie K.E., Bick M., Tilbrook D.L. and Foley C.P. 2004. GETMAG - A SQUID magnetic tensor gradiometer for mineral and oil exploration. *Exploration Geophysics* **35**, 297–305, doi: 10.1071/EG04297.
- Snyder D.B., Bleeker W., Reed L.E., Ayer J.A., Houllé M.G. and Bateman R. 2008. Tectonic and metallogenic implications of regional

- seismic profiles in the Timmins Mining Camp. *Economic Geology* **103**, 1135–1150, doi: 10.2113/gsecongeo.103.6.1135.
- Thompson D.T. 1982. EULDPH: A new technique for making computer assisted depth estimates from magnetic data. *Geophysics* **47**, 31–37, doi: 10.1190/1.1441278.
- Wilson G.A., Čuma M. and Zhdanov M. 2011. Massively parallel 3D inversion of gravity and gravity gradiometry data. *Preview* **152**, 29–34, doi: 10.1071/PV2011n152p29.
- Zhdanov M. 1988. *Integral Transforms in Geophysics*. Springer-Verlag.
- Zhdanov M.S. 2002. *Geophysical Inverse Theory and Regularization Problems*. Elsevier.
- Zhdanov M. 2009. New advances in 3D regularized inversion of gravity and electromagnetic data. *Geophysical Prospecting* **57** (4), 463–478, doi: 10.1111/j.1365-2478.2008.00763.x.
- Zhdanov M.S., Ellis R.G. and Mukherjee S. 2004. Regularized focusing inversion of 3-D gravity tensor data. *Geophysics* **69**, 925–937, doi: 10.1190/1.1778236.
- Zhdanov M.S., Smith R.B., Gribenko A., Čuma M. and Green A.M. 2011. Three-dimensional inversion of large-scale EarthScope magnetotelluric data based on the integral equation method – Geoelectrical imaging of the Yellowstone conductive mantle plume. *Geophysical Research Letters* **38**, L08307, doi: 10.1029/2011GL046953.