

1 **Tomofast-x 2.0: an open-source parallel code for inversion of potential field**
2 **data, to recover density, susceptibility and magnetisation vector, with**
3 **topography and wavelet compression**

4
5 Vitaliy Ogarko^{1,2}, Kim Frankcombe³, Taige Liu³, Jeremie Giraud^{1,4}, Roland Martin⁵, Mark Jessell^{1,2}

6
7 1. Centre for Exploration Targeting (School of Earth Sciences), University of Western
8 Australia, 35 Stirling Highway, 6009 Crawley, WA, Australia.

9 2. Mineral Exploration Cooperative Research Centre, The University of Western Australia, 35
10 Stirling Highway, 6009 Crawley, WA, Australia.

11 3. ExploreGeo, PO Box 1191, Wangara DC, WA, 6947, Australia.

12 4. Université de Lorraine, CNRS, GeoRessources, F-54000 Nancy, France.

13 5. Geosciences Environnement Toulouse, CNRS UMR 5563, Toulouse, France.

14
15 Corresponding author: Vitaliy Ogarko (vitaliy.ogarko@uwa.edu.au)

1 **Abstract**

2

3 We present a major release of the Tomofast-x open-source gravity and magnetic inversion code that
4 incorporates several functionalities enhancing its performance and applicability for both industrial
5 and academic studies. The code has been re-designed with a focus on real-world mineral
6 exploration scenarios, while offering flexibility for applications at regional scale or for crustal
7 studies. This new version includes several major improvements: magnetization vector inversion,
8 inversion of multi-component magnetic data, wavelet compression, improved handling of
9 topography with support for non-uniform grids, a new, efficient parallelization scheme, a flexible
10 parameter file, and optimized input-output operations. Extensive testing has been conducted on a
11 large synthetic dataset as well as field data from a prospective area of the Eastern Goldfields
12 (Western Australia) to explore new functionalities with a focus on inversion for magnetisation
13 vectors and magnetic susceptibility, respectively. Results demonstrate the effectiveness of
14 Tomofast-x 2.0 in real-world studies in terms of both the recovery of subsurface features and
15 performances on shared and distributed memory machines. Overall, with its updated features,
16 improved capabilities, and performances, the new version of Tomofast-x provides a free open-
17 source, validated advanced and versatile tool for constrained gravity and magnetic inversion.

18

1 Introduction

In spite of the long-recognised ambiguities in the geological interpretation and inversion of gravity and magnetic datasets (Nettleton, 1942) these two methods remain popular approaches to understanding the continental subsurface. This is at least in part because, unlike theoretically more informative methods such as the array of seismic techniques, there are already widespread gravity and magnetic databases available for most continental areas. In addition, unlike geophysical methods such as gamma ray spectroscopy or Synthetic Aperture Radar, there is significant 3D information content in the data.

One of the challenges faced by geophysical inversion platforms is to make them scalable, so that they not only meet the expectations of research geophysicists trialling new algorithms but can actually be applied to large-scale industrial and governmental applications. As part of their exploration program mineral explorers will often undertake an unconstrained 3D inversion of the available magnetic or gravity data in order to better understand the 3D architecture and to generate prospects. As more data are acquired these unconstrained inversions may be re-run with geological constraints and new data sets. Similarly, government agencies also need to undertake 3D inversions, at a relatively large scale, in order to help their geological mapping teams understand the 3D architecture. While they may be less pressed by time constraints than explorers, they will typically be working with much larger data sets. Tools that can scale from prospect to regional scale are required.

There has been a long history of forward modelling (Jessell, 2001) and inversion of gravity and magnetic datasets (Pears et al., 2017). For much of the past 20 years the tools of choice for most explorers and Government agencies have been the inversion codes developed by the Geophysical Inversion Facility (GIF) at The University of British Columbia (UBC) (Li and Oldenburg, 1996a, 1998). Over the past decade many explorers have switched from the GIF based tools to the Seequent's Voxi product, the interface for which is built into the commercial Windows based, software tools they often use. Voxi is provided as a commercial cloud based SaS running on clusters. There are several other less commonly used commercial programs for 3D potential field regularised inversion, each with unique specialisations (e.g. Geomodeller, VPmg and MGinv3D).

More recently, there has been an increasing focus on developing open-source solutions to solve the gravity and/or magnetic inversion problem, which provide access to inversion solutions to a wider audience (Table 1). When the inversion facility at UBC wound down several of the researchers moved to make the algorithms open source and available in Python in the hope they would be adopted as a teaching and research tool. This project is known as SimPEG and now has a sizeable user group of academic, government and industry users, particularly in Northern America (Cockett et al., 2015). The Python code has been built from scratch using many similar problem formulations to the original Fortran code used in the GIF packages. It uses public libraries (e.g., NumPy) for I/O and for the solver and leverages new advances in high performance computing (dask and zarr libraries) to improve performance. It is however significantly slower than the older Fortran codes and does not offer a practical solution for typical, exploration sized, datasets. Compression is only on file size, not matrix size, and therefore, does not offer rapid solutions to many real exploration

1 problems. To the best of our knowledge, among the open-source codes listed in Table 1, Tomofast-
 2 x is the only one that offers wavelet compression of the sensitivity matrix.

3

Table 1. Some of the more popular open-source 3D gravity and magnetic inversion platforms.

Name	Language	Parallel	Physics	Joint	URL
Tomofast-x	Fortran	X	Grav/Mag	X	https://github.com/TOMOFAST/Tomofast-x
SimPEG	Python	X	Grav/Mag/EM/IP	X	https://github.com/simpeg/simpeg
Escript	Python	X	Grav/Mag	X	https://github.com/esys-escipt
Growth3D	Fortran		Grav		https://github.com/josefern/GROWTH_3.0_software
Potensoft	Matlab		Grav/Mag		Suppl. materials in (Özgü Arısoy and Dikmen, 2011)
IGMAS	Java		Grav/Mag		https://igmas.git-pages.gfz-potsdam.de/igmas-pages/
pyGIMLi	Python		Grav/Mag/EM/IP/Seismic/ NMR	X	https://www.pygimli.org
pyGMI	Python		Grav/Mag		https://patrick-cole.github.io/pygmi/pfmod.html
jif3d	C++		Grav/Mag/Seismic/ERT	X	https://sourceforge.net/projects/jif3d/

4

5 Each of these 3D platforms has strengths and weaknesses in terms of scalability, ease of
 6 installation, ease of generation of input meshes, reliance on commercial platforms, etc. even before
 7 we think of specific functionality and the additional constraints that can be added to the inversion to
 8 try and reduce the aforementioned ambiguity. This paper does not attempt to benchmark the
 9 different codes, interesting as that would be and much as that would be a worthy collaboration
 10 between the different research groups currently maintaining these codes. What has been missing
 11 from this mix is an open-source package able to scale to continental sized problems while offering
 12 rapid answers at prospect scale. Tomofast-x is such a package.

13

14 When Tomofast-x was publicly released in 2021 it strongly reflected its research origins. Any
 15 exploration problem of reasonable size required a large cluster to solve in anything approaching real
 16 time and for magnetic inversions, observations had to be above the highest point in the model space.
 17 It only inverted on susceptibility and density. Through a close collaboration between industry and
 18 academia it now offers equivalent performance and results to the commercial codes. In addition, as
 19 well as susceptibility and density it now also inverts on magnetisation using either single
 20 component Total Magnetic Intensity (TMI) data or three component magnetic data.

21

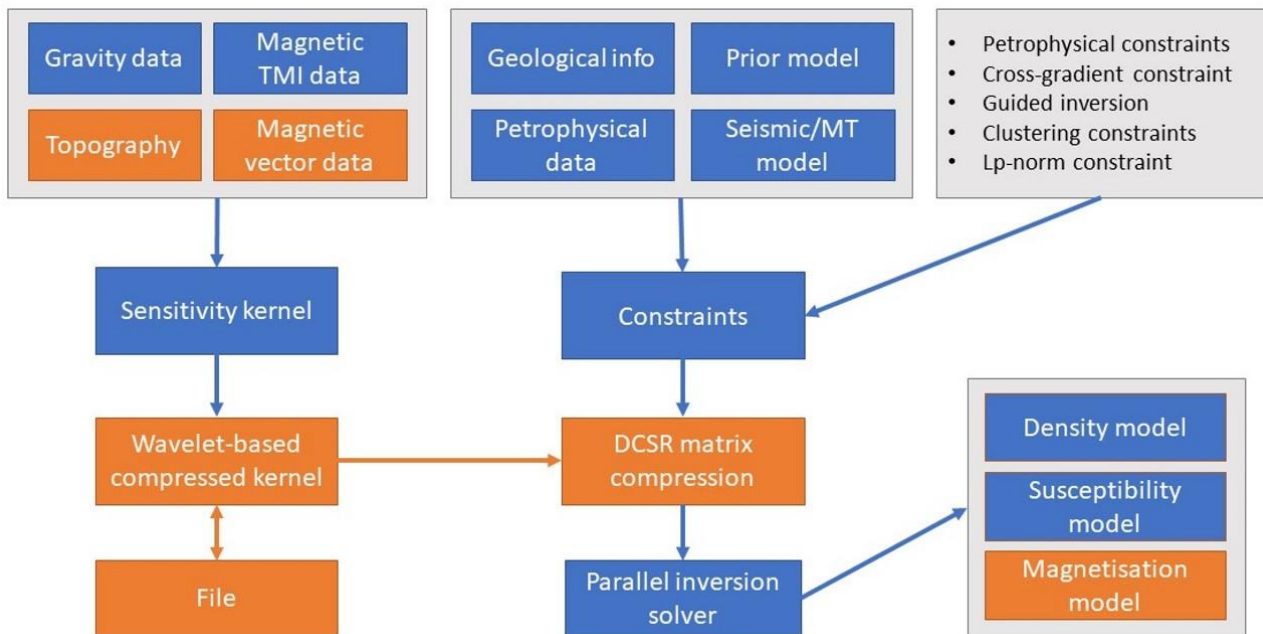
2 Description of code features

2

3 Tomofast-x inversion code is written in Fortran 2008 using classes, which allows flexibility in
4 extending the code such as adding new types of constraints, algorithms, and data I/O. It is fully
5 parallelised using the MPI library which provides both shared and distributed memory parallelism.
6 The code has parallel unit tests covering important parts, which are implemented using "ftnunit"
7 framework (Markus, 2012) extended by us to the parallel case to work with the MPI library.

8 The previous code version Tomofast-x 1.0 included geological and petrophysical constraints
9 implemented, specifically clustering constraints (Giraud et al., 2019b), cross-gradient constraints
10 (Martin et al., 2021), local gradient regularization (Giraud et al., 2019a, 2020), Lp-norm constraints
11 (Martin et al., 2018), and disjoint interval bound constraints (Ogarko et al., 2021; Giraud et al.,
12 2023b), which were summarized in (Giraud et al., 2021). The Figure 1 shows a summary of the
13 Tomofast-x 2.0 inversion workflow, new components added to the version 2.0 are highlighted in
14 orange. In this section we describe new code components added to Tomofast-x 2.0 code version and
15 highlight their importance.

16



17

Figure 1. Tomofast-x 2.0 inversion workflow, new components are highlighted in orange.

18

2.1 Topography

19


20 The impact of topography on gravity data is well known. Less well understood is the effect of
21 topography on magnetic data. However, as clearly demonstrated by (Frankcombe, 2020), it is
22 important to include topography in magnetic modelling. In modern volcanic terrains where
23 moderately magnetic andesitic basalt makes up a large proportion of the subcrop and topography is
24 often severe, the magnetic signal from topography alone can be tens to hundreds of nano-Teslas.
25 This may be greater than what might be a relatively subtle response, from a deeply buried body.
26 Failing to account for topography not only risks missing subtle targets, but it will also lead to
27

1 erroneous interpretations if the magnetic data is interpreted or inverted on the assumption that the
2 ground is flat.

3 There are several approaches used to include topography in the model, the most simple being to
4 ensure that the model space is a rectangular prism extending into air above the highest point in the
5 topography. A more advanced approach is to use a tetrahedral mesh to smoothly model the
6 topography. Intermediate approaches and that used by this code are to drape a rectilinear mesh
7 below the topography. This required that the topography is discretised into steps the size of the
8 mesh cells but does not require that the air be included in the model. As each vertical pillar of cells
9 in the model is independent there is no need for each layer of cells to have a fixed top and bottom.
10 This enables the topography to be discretised into locally thin layers without requiring that all
11 points in the model at that elevation have the same fine discretisation. In areas of severe
12 topography, draping the mesh below the topography and treating each vertical pillar of the mesh
13 independently significantly reduces the size of the mesh, and thus the run time, relative to regular
14 rectangular prism meshes. Figure 2 shows an example of a model grid with topography.

15 The forward magnetic problem in Tomofast-x 1.0 was calculated using Bhattacharyya's formulation
16 (Bhattacharyya, 1964). It has the limitation that all data locations must be above the highest point of
17 the model. Modern aeromagnetic data is acquired at heights of less than 40m above ground. There
18 are very few project areas which have less than 40m difference in elevation across them, limiting
19 the utility of Bhattacharyya's algorithm. In order to allow for arbitrary topography and readings
20 below the Earth's surface, the forward magnetic problem was re-implemented using Sharma's
21 formulation (Vallabh Sharma, 1966). Additionally, distance-based weighting of the sensitivity
22 kernel, which is a generalised version of a simple depth weighting, was included (Li and Oldenburg,
23 1996b). Weighting by distance rather than depth allows for observations beside or even below
24 points in the model.

25



26

Figure 2. A non-uniform model grid with topography for the Synthetic400 model (for more details on this model see Section 3.1).

1

2

3 **2.2 Non-uniform model grid**

4

5 The size of a regularised inversion is $(N_{\text{data}}+1)\times N_{\text{model}}$, where N_{data} is the number of
6 observations being modelled and N_{model} is the number of cells in the model mesh. When stored as
7 8-byte reals this leads to around 7.5 Gb being required for a 100 x 100 x 10 model block with a 100
8 x 100 array of observation points. Such a tiny model size would only be used for test work. For an
9 ordinary exploration sized problem tens of Petabytes might be required while for Government style
10 regional work Exabytes or Zettabytes of storage would be needed. While large problems can be
11 broken up into smaller panels and then re-joined (Goodwin and Lane, 2021), care needs to be taken
12 to avoid edge artefacts and the overall time to run the inversion is considerably longer because of
13 the need to run coarse inversions followed by forward models to minimise these edge effects. In
14 addition, large problems require more computational cycles placing an effective limit on the size of
15 problems that can be handled. In order to reduce the size of the problem, the voxel size can be
16 increased for voxels distant from the observations. In practice it has been found that increasing the
17 size of model voxels in a horizontal sense away from the core area containing the data and in a
18 vertical sense below the ground surface reduces the size of the model array by about 40 times,
19 without degrading from the inversion result.

20 Care needs to be taken not to change the voxel size too quickly as rapid changes in the volume
21 sensitivity will result in material being moved from a smaller, lower sensitivity cell to its larger
22 neighbour introducing artificial stratification to the inversion model. A multiplier of around 1.15
23 with a hyperbolic tangent smoothing into the core of the inversion area has been found to work
24 well. Figure 2 shows an example of such a non-uniform model grid.

25

26 **2.3 Wavelet compression**

27

28 Another tool to reduce the size of the inversion problem is to compress the matrices. (Li and
29 Oldenburg, 2003) proposed using wavelet transforms to compress the sensitivity matrix by nulling
30 all wavelet coefficients with amplitudes below a user defined threshold. This approach can be used
31 to reduce the size of the problem by up to three orders of magnitude without seriously impacting on
32 the solution. Tomofast-x users have the option to choose between the commonly used wavelets -
33 Harr and Daubechies D4 (Martin et al., 2013). The wavelet transform is implemented using
34 multilevel lifting scheme, which reduces the number of arithmetic operations by nearly a factor two.
35 The transform is performed in place, requiring only a constant memory overhead. The wavelet
36 transform is applied to each row of the weighted sensitivity kernel independently, which are then
37 compressed and stored using sparse matrix format (i.e., only non-zero values are stored). The
38 resulting model perturbation vector after each inversion is transformed back from the wavelet
39 domain using the inverse wavelet transform. As discussed in Sect. 2.1 the model grid is
40 implemented by draping the mesh below topography. This allows the model grid to be defined as a
41 3D array of dimensions $N_x\times N_y\times N_z$ for any topography without introducing additional “air” cells.
42 The 3D wavelet transform can therefore be applied efficiently due to the coherency of the

1 sensitivity values along the spatial dimensions. Unit tests were also added to verify that the wavelet
2 transforms preserve a vector norm.

3 The non-compressed sensitivity kernel is a dense matrix of size $N_{data} \times N_{model}$, with
4 $N_{model} = N_x \times N_y \times N_z$. A user selected compression rate R_c defines the fraction of sensitivity values
5 to keep from the original kernel. Thus, the memory to store the compressed sensitivity kernel (using
6 CSR format) nearly linearly scales with the compression rate and can be calculated as:

$$M_{sensit} = (4 + f)R_c N_{model} N_{data} + 8N_{data} \text{ bytes} \quad 1$$

7 where $f=4$ for single precision and $f=8$ for double precision floating values. Comparative tests
8 showed that storing sensitivity values using single precision is sufficient and has no noticeable
9 impact on inversion results, while it is important to use double precision values during the
10 sensitivity calculation, and wavelet transformation. The first term corresponds to two arrays of the
11 CSR format – column indexes (4-byte integers), and matrix values (floats). The second term in Eq.
12 (1) corresponds to the array of indexes in the values-array corresponding to each matrix row. It is
13 stored using long integer type (8 bytes), to avoid integer overflow. An estimation of the total
14 memory requirement is given in Appendix A.

15 An additional memory optimisation for the parallel matrix storage is discussed in section 2.6 and
16 examples of the inverted models using various compression rates are given in section 3.3.

17

18 **2.4 Magnetisation**

19

20 Nearly all magnetic rocks will have a component of remanent magnetisation storing the magnetic
21 field direction from the time of crystallisation of the magnetic mineral. The strength of this
22 remanent field is often very weak and can be ignored, however in a significant number of cases of
23 interest to minerals explorers, the remanent field is strong enough to alter the shape of the observed
24 field. Remanence can affect the host as well as the target. The existence of remanence is often
25 evident from looking at the shape of the anomaly which should alert the interpreter to the potential
26 for error. It may however be cryptic and if not allowed for, will result in a flawed interpretation.

27 In addition, all magnetic bodies are affected by self-demagnetisation. This is a property which
28 produces a magnetic field within the body which is aligned to the shape of the body rather than to
29 the Earth's field or any stored remanent direction. For bodies with susceptibilities less than 0.1 SI
30 the effects of self-demagnetisation are negligible, however, above this, the observed response can
31 be distorted to such a degree that exploration programs fail (Gidley, 1988). A solution to both
32 problems is to invert on magnetisation rather than susceptibility (Ellis et al., 2012). The code was
33 therefore extended to allow for magnetisation inversion.

34 It must be noted that magnetisation inversion needs three sensitivity kernels defined for each
35 magnetic vector component (Liu et al., 2017). Thus, the computational and memory requirements
36 grow nearly by a factor of three. Therefore, using optimisations described previously (i.e., wavelet
37 compression and non-uniform grid) becomes even more important for this type of inversions.
38 Examples of magnetisation inversion are shown in section 3.4.

39

40 **2.5 Multiple data components**

41

42 Historically only a single component of the magnetic field was measured. Initially the local
43 declination or inclination were measured using optical dip meters, subsequently the field strength in
44 the vertical direction, Hz, was measured using a fluxgate. Fluxgates were replaced by proton

1 precession and then optically pumped magnetometers which measured the field strength in the
2 direction of the Earth's field, the Total Magnetic Intensity (TMI). However, fluxgates can be used
3 to also measure the horizontal components, as is routinely done within magnetic down hole survey
4 tools. SQUID magnetometers are now being developed which will enable a tensor measurement.
5 Mutli-component data provide a more diagnostic data set with regard to the body location,
6 particularly when remanence is involved. It is therefore useful to be able to model and invert data
7 acquired in any direction and from any location.

8 Inversion of three magnetic data components was included in the code. As with magnetisation
9 inversion discussed in previous section, this also increases the computation and memory
10 requirements. Particularly, when three data components are used for magnetisation inversion, the
11 number of the sensitivity kernels grows to nine. Thus, these inversions strongly benefit from
12 wavelet compression and non-uniform grid optimisations. These modifications allow for inversion
13 of three-component down hole magnetic data which typically improves drill targeting compared to
14 single or multi component above surface data sets. This improvement is particularly noticeable
15 when magnetic remanence is present.

16

17 **2.6 Efficient Parallelisation scheme**

18

19 The parallelisation of the inversion problem can be approached in several ways: by partitioning the
20 model or partitioning the data. These two approaches correspond to splitting the least-squares
21 system matrix by columns (model) or by rows (data). As both approaches have their advantages, an
22 efficient parallelisation scheme which combines the two approaches was designed. The
23 parallelisation scheme is comprised of the following elements:

24

25 1) The forward problem is parallelised by data: This avoids the need to parallelise the wavelet
26 compression, as the full row of the sensitivity kernel (corresponding to the whole model) is
27 available at each processor. In turn this allows for a saving on the MPI communication cost
28 required to parallelise the wavelet compression. This can be significant as the wavelet
29 transform calculations are non-local and span over the whole model space.

30

31 2) The inverse problem is parallelised by model: In the compressed sparse row (CSR) matrix
32 storage the model vector is accessed in the inner loop in both matrix-vector and transposed-
33 matrix-vector products. This limits the performance due to memory traffic (Chapter 4.3.2 in
34 (Barrett et al., 1994)). In order to minimise the memory traffic in these operations it is better
35 to partition the model vector. Both versions with the inversion problem parallelised by data
36 and by model were implemented and compared, confirming that parallelising by model is
37 faster and leads to better scalability. This was also confirmed by profiling the code using
38 Intel VTune Microarchitecture Exploration tool, which showed an increase in L3 cache
39 misses relative to partitioning by model, resulting in an overall slowdown of the code.
40 Machines with larger L3 caches performed better than those with smaller even though they
41 had older generation CPUs.

42

43 3) Load balancing: When wavelet compression is applied to the sensitivity kernel, the
44 distribution of the non-zero values is not uniform. When the model partitioning is performed
45 using equal size chunks, the number of non-zero kernel elements can vary by a factor of 100

1 on different processors. This leads to non-efficient usage of CPU resources and slows down
2 the calculations. To mitigate this, load balancing was implemented by partitioning the
3 matrix columns into non-equal chunks such that the total number of non-zeros is nearly the
4 same on each CPU.

- 5
- 6 4) Doubly compressed matrix storage: In the CSR matrix storage format, the outer loop in the
7 matrix-vector and transposed-matrix-vector products iterates over the matrix rows. When
8 parallelising the inversion by model (i.e., parallelising the inner loop), the outer loop spans
9 over all matrix rows N_{rows} . This limits the parallelism, defined as the maximum possible
10 speedup on any number of processors, to at most $O(nnz/N_{rows})$ (Buluç et al., 2009), where
11 nnz is the total number of non-zero values. The typical number of least-squares matrix rows
12 in inversions is $N_{row} = N_{data} + N_c N_{model}$, where the N_{data} -rows correspond to sensitivity
13 kernel, and N_{model} -rows correspond to N_c constraint terms (such as model damping term).
14 The typical N_{model} value is much greater than the N_{data} and thus adding constraints
15 significantly limits the scalability. Note that the rows corresponding to constraint terms are
16 usually extremely sparse, for example one non-zero element per row in the model damping
17 term (i.e., a diagonal matrix). Thus, in the parallel runs there will be many empty rows in the
18 local matrices, so the redundant information will be stored. Therefore, double compressed
19 sparse row (DCSR) format which improves the compression efficiency by storing only non-
20 empty rows was implemented (Buluc and Gilbert, 2008). The DCSR format adds one
21 additional array to the CSR format, which stores the indexes of non-empty rows. This
22 improves the scalability and reduces memory traffic. Additionally, this reduces the memory
23 requirements for storing the matrix on several CPUs.
- 24
- 25 5) Efficient memory allocation: In the LSQR inversion solver, the parallel bidiagonalization
26 operations (consisting of matrix-vector and transposed-matrix-vector products) were re-
27 implemented to overwrite the existing vectors thus avoiding memory allocation for the
28 intermediate products, for details see sect. 7.7 in (Paige and Saunders, 1982). The
29 corresponding parallelism was implemented using the same input and output buffer of the
30 collective communication routine `MPI_Allreduce` via `MPI_IN_PLACE`.

31
32 In Section 3.5 the proposed parallelisation scheme is applied to several representative models and
33 the resulting parallelisation efficiency is discussed.

34 **2.7 Sensitivity kernel import/export**

35
36
37 The ability to import and export the sensitivity kernel has been implemented. The compressed
38 sensitivity kernel is stored as a binary file which allows for previously computed sensitivity kernels
39 to be reused. This saves computational time (up to ~95%) when exploring the effect of different
40 types of constraints and different hyperparameter values (i.e., when varying weights in the cost
41 function), as well as different starting and prior models. The sensitivity kernel can also be reused
42 when a different background field needs to be subtracted from the data, as the kernel values depend
43 only on the model grid and data locations. Additionally, the calculated kernel can be imported into
44 other codes to perform sensitivity analysis, null-space navigation, or level-set inversions (Giraud et
45 al., 2023a). Finally, it is possible to import the external kernel corresponding to a different

1 geophysics or another optimisation problem (e.g. geological modelling), and use the Tomofast-x as
2 a high-performance parallel inversion (or optimisation) solver. External to Tomofast-x, a Python
3 script to import/export the sensitivity kernel from/to Python has also been implemented, and it is
4 available in <https://github.com/TOMOFAST/Tomofast-tools>.

3 Inversion examples

The data sets and models used for the inversion and the results of those inversions will now be discussed. Different compression rates were tested, and parallel performance was benchmarked. The Table 2 shows the physical model dimensions and corresponding model grid size and number of data for the test models, which are then described in detail below.

Table 2. Model dimensions. In brackets the dimensions of horizontal padding are specified. Meters are used as units of length.

Name	Δ East	Δ North	Depth	Relief	Nz	Ny	Nz	Ndata
Synthetic400	10000 (+5904)	10000 (+5904)	3290	1457.6	400 (+40)	400 (+40)	28	158006
Callisto	2500 (+3402)	2200 (+3402)	3290	24.6	200 (+42)	176 (+42)	28	35200

In all susceptibility inversions a positivity constraint was applied to reduce the model nonuniqueness (Li and Oldenburg, 1996a). The positivity constraint was added via the ADMM bound constraints (Ogarko et al., 2021). The number of major inversion iterations used was Niter=30 for Callisto model and Niter=50 for Synthetic400 model. The number of minor iterations (in the LSQR inversion solver) used was 100 in all tests.

3.1 Synthetic400 model

While it is important to use real data sets in testing inversion codes it is also important to use data sets where the answer is known. In real exploration data sets the true answer is never known, at least not throughout the model space. A drill hole here or there may constrain the depth to a particular magnetic body and susceptibility measurements on core or chips or from in-hole tools may constrain the likely range of susceptibilities for that body. In general, however, the properties of the vast majority of the model space is unknown. Synthetic models are therefore required in order to calibrate the inversion algorithm and understand its target resolution. Although synthetic models can never achieve the level of complexity that a real data set contains this paper attempts to make the model as realistic as possible.

Topography, extracted from a 12.5 x 12.5m ALOS DEM, for an active exploration prospect in Indonesia was used as the base layer and the ground was ascribed a homogeneous susceptibility of 0.01 SI, which would be at the lower end of expected susceptibilities for this area. The elevation has a range of 1400m over the 10km x 10km extent of the central core of the model space increasing to 1600m if one includes the horizontal padding. The magnetic response for a simulated survey flown with a perfect 60m drape height over this topography was modelled using Tomofast-x and the results are shown in Figure 3 along with the topography that produced it. The field range shown in the key has been clipped to flatten the outer 2% of the actual range which was closer to 12 nT. As shown in (Frankcombe, 2020), typical topographic induced variations in these areas are an order of magnitude higher than this, implying a ground susceptibility closer to 0.1 SI. However, as we are looking at susceptibility contrasts the use of a low background value in the synthetic model does not detract from its results. A parametric potential field modelling package (Potent -

1 <https://geoss.com.au/potent.html>) was then used to compute the magnetic response of several,
 2 geologically plausible, bodies buried in this topography. In order to provide a challenge for the
 3 inversion algorithm, the bodies were placed where their magnetic response might best be masked by
 4 topography. Potent uses line and surface integrals to model the response of bodies. Modelling the
 5 response of the host is only practical in 2D. Topography was used to control the sensor position,
 6 which was defined as a constant drape above topography. The susceptibility of the half-space is
 7 only used by the parametric modelling to compute a contrast to use when computing the response of
 8 the model body. The computed response therefore becomes the response from bodies in air. The
 9 susceptibility of bodies is their contrast relative to the host, which may be negative. These responses
 10 were then added to the computed response of the topography alone from Tomofast-x, to generate a
 11 synthetic data set. The Potent program was used in preference to a generalised voxel-based program
 12 because the regularisation of the mesh would have required very small voxels in order to accurately
 13 reflect the desired shape and thickness of the bodies. We have used a parametric model to better test
 14 the algorithm, by approximating real world structures which are hard to describe with a mesh of
 15 rectangular prisms.

16
 17 It is worthwhile saying a few words about the geological relevance of the synthetic model. In
 18 modern volcanic environments in the western Pacific and elsewhere, the underlying geology often
 19 consists of a series of volcanic flows. These range from relatively nonmagnetic crystal tuffs through
 20 to andesitic basalt and local lahar debris flows. Where the lahar carries andesite boulders it can
 21 become very magnetic and very heterogeneous. The layered package of volcanics may be
 22 interspersed with sediments deposited during periods of quiescence. These sediments are generally
 23 mudstones or limestones but in deeper water environments, may include greywackes. Typically, the
 24 package is only lightly tilted so that dips are relatively flat. For the area to become of interest to
 25 explorers, an intrusion of some sort is required. These are typically pipe like bodies and may have a
 26 positive susceptibility contrast with the host in the case of skarns and potassically altered porphyry
 27 style mineralisation or in the case of epithermal mineralisation or for the alteration halo around a
 28 porphyry system, a negative contrast through magnetite destruction. These elements have been
 29 incorporated into the synthetic model and are labelled in Figure 4. Body 1 is a 100m thick magnetic
 30 disk reflecting a more magnetic basalt flow which has been eroded away, except in the peak near
 31 the centre of the model space. Body 2 is a large, flat 50m thick, magnetic sheet similarly reflecting
 32 an erosional residual from broader magnetic lava flow. Body 3 is a 10m thick dipping slab
 33 reflecting either a magnetically mineralised fault or a dyke like feeder to the lava flows. Body 4 is a
 34 pipe like magnetic body with a radius of 250m and a relatively high susceptibility reflecting a
 35 strongly potassic altered porphyry intrusion or skarn around a finger porphyry. Body 5 is a 20m
 36 thick dipping sheet reflecting a tilted magnetic lava flow. Body 6 is a 500m radius pipe with a
 37 susceptibility lower than the host. This reflects a zone of magnetite destruction around and above a
 38 porphyry system. The bodies are described in Table 3.

39

Table 3. Description of the Synthetic400 model bodies.

Body	Shape	Susceptibility (SI)	Analogue
1	flat disk	0.05	residual lava flow
2	flat slab	0.07	residual lava flow

3	dipping slab	0.03	dyke or feeder
4	vertical cylinder	0.1	magnetic porphyry or skarn
5	dipping slab	0.02	tilted lava flow
6	vertical cylinder	0.0015	magnetite destruction around porphyry

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39


The synthetic data had various levels of noise added to it prior to inversion. Noise was added on the basis of a rescaled percentage of the observed value. Noise levels ranging from 1% up to 10% of the computed synthetic response were rescaled by a random number between -0.5 and +0.5. Where the amplitude of the rescaled noise fell below a noise floor, the noise floor multiplied by the sign of the rescaled noise replaced it. Noise floors of 1 nT and 5 nT were employed. That noise was then added to the computed synthetic response.

As this synthetic model contained no remanence a susceptibility inversion was undertaken. A mesh with a central 400 x 400 voxel core of 25m x 25m voxels with a first layer thickness of 10m was created. This was padded on each side with 20 cells increasing in size away from the core, to add around 3500m to each edge of the horizontal extent of the core 28 cells were used vertically to take the mesh down around 3200m below the ground surface. In total the mesh was 440 x 440 x 28 cells (Nmodel=5,420,800) and there were magnetic observations above the centre of each surface cell in the central core (Ndata=160,000). This would be considered a small to moderate sized exploration problem.

Regardless of the noise level, the inversion does recover the shape and depth to the top of the three most magnetic bodies (Figure 5). The inversion is favouring the top of the vertical cylinder but does focus higher susceptibilities right to the bottom of the model as shown in Figure 6-(a). It also places magnetic material beside the low susceptibility body to compensate for what the inversion perceives as a negative susceptibility (Figure 6-(b)). This is a common artefact in unconstrained susceptibility inversions. For the thin slabs the inversion has overestimated the depth to the bottom of the body. This is also a common artefact and is driven by the smoothness of the L2 norm used to compute convergence.

The bowl-shaped overestimate of the thickness of Body 1 has been compensated for by a "cape" of low susceptibility draped around the body. The noise added to the data is flowing through to the model as variations at the surface. Although the two thin, dipping slabs with small susceptibility contrasts coincide with gradients in the model values, there is not enough contrast to be able to resolve these with any confidence. The bodies have susceptibility contrasts of 0.01 and 0.02 SI and are only 10 and 20m thick. At the depth of the thin plates the model voxels are significantly thicker than the plates and so not being able to resolve them is not particularly surprising.


The eastern body with a negative susceptibility contrast is poorly imaged but as noted above this is typical for unconstrained susceptibility inversions. The inversion sees the body as having a negative susceptibility but the positivity constraints will not allow this so magnetic material is artificially stacked around it to produce the response we see in Figure 6-(b).



1

Figure 3. (a) Synthetic400 model's topography of the whole model, and (b) forward magnetic response of the model's core from the topography alone. The outline of the core is shown in (a). The dashed lines are a 2500m x 2500m graticule.

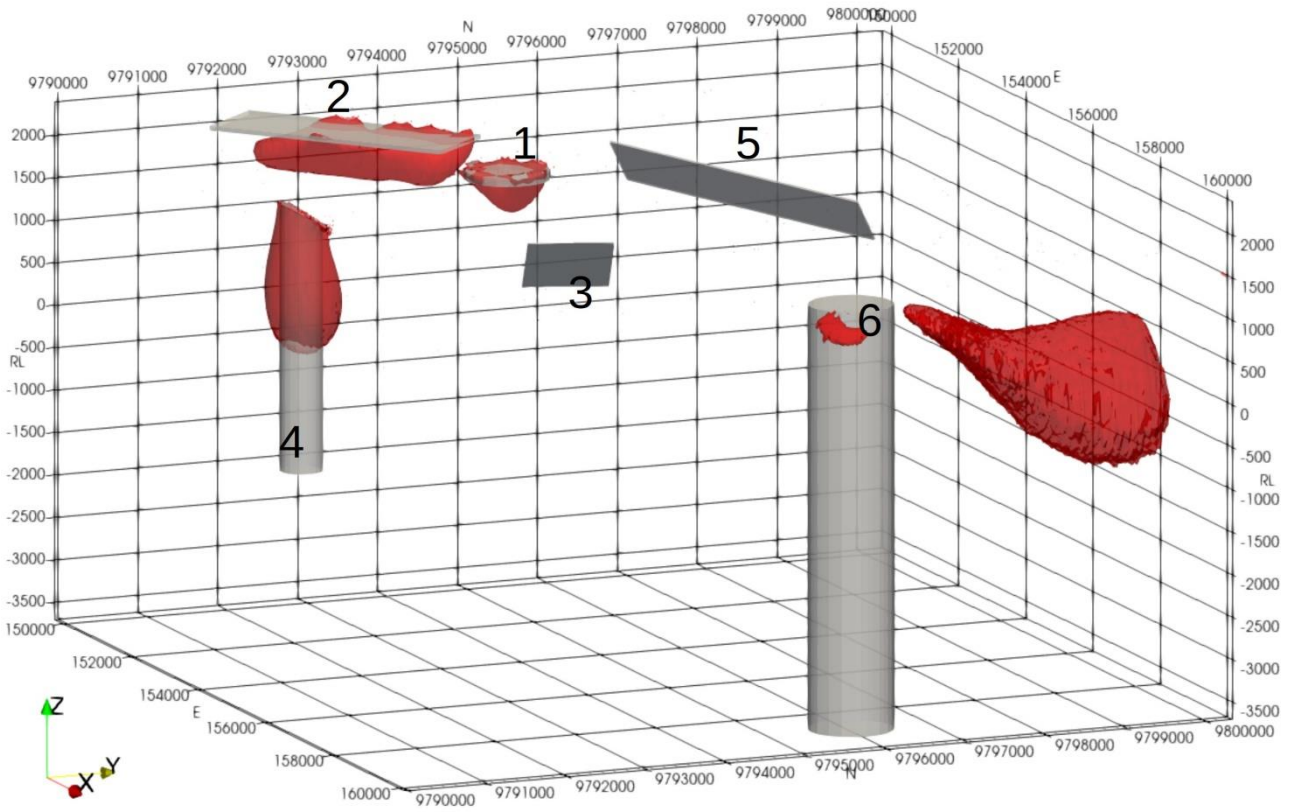
2



3

Figure 4. Synthetic400 model's total magnetic response with magnetic bodies projections. The dashed lines are a 2500m x 2500m graticule.

1

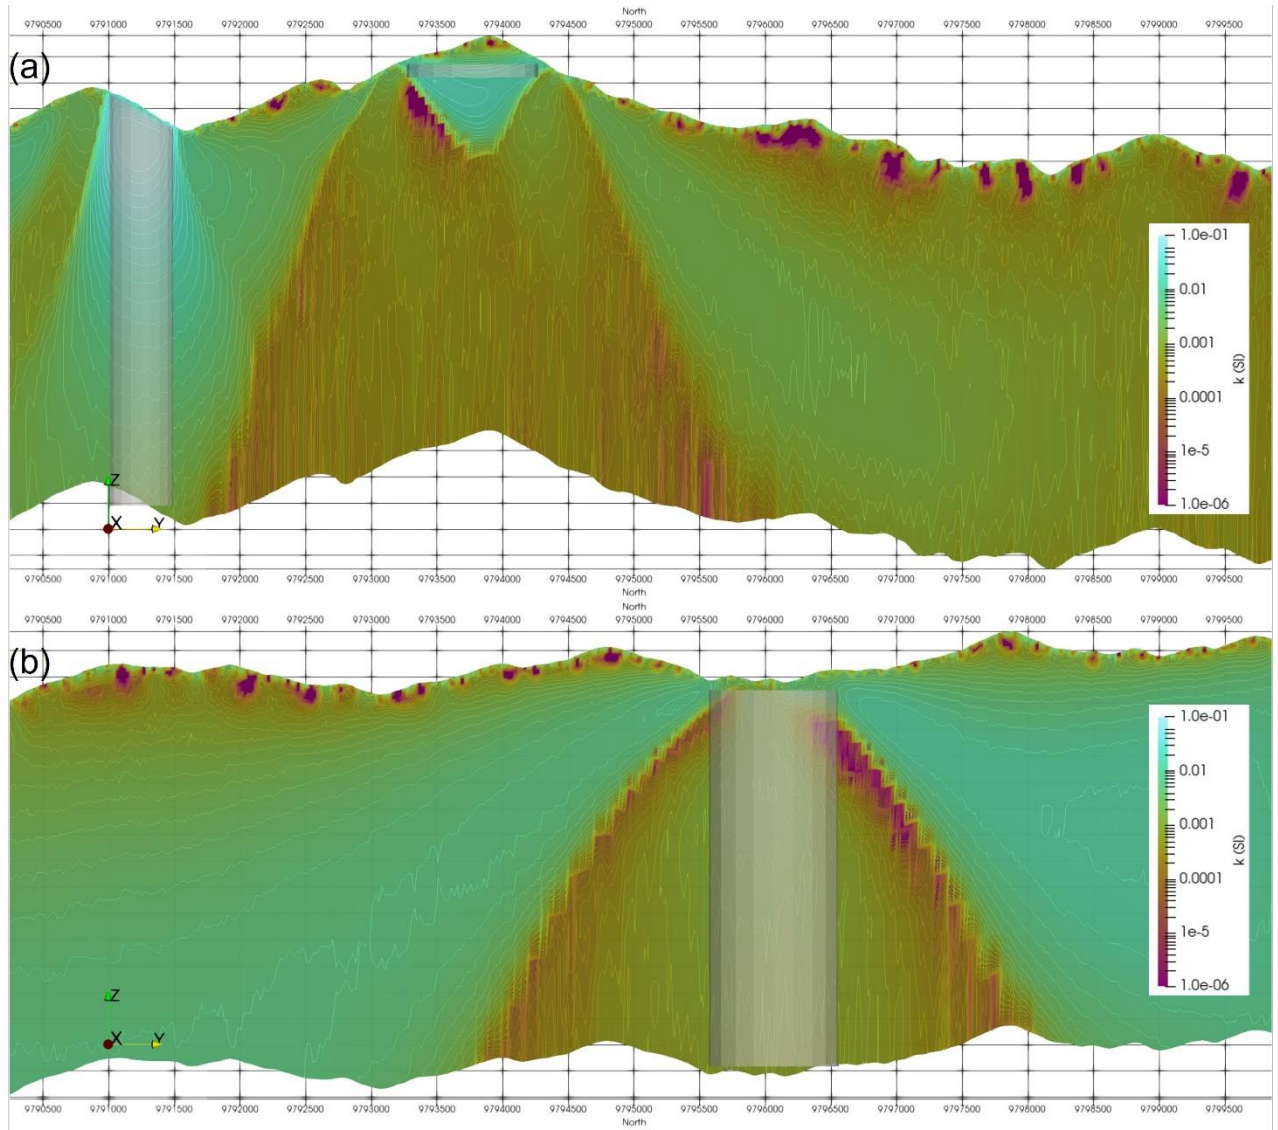


2

Figure 5. Inverted Synthetic400 susceptibility model. The contours are shown at $k=0.01$ SI.

3

1



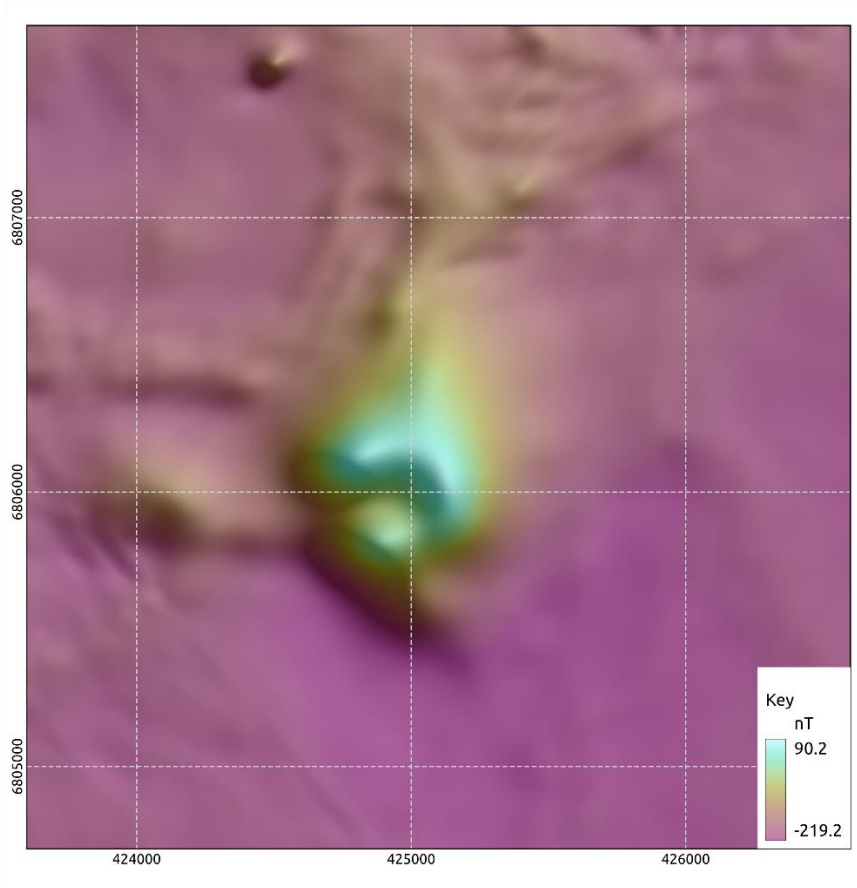
2

Figure 6. Vertical sections of the inverted Synthetic400 model through (a) high susceptibility cylinder and disk, and (b) low susceptibility cylinder.

3.2 Callisto model

The Callisto anomaly is a discrete magnetic anomaly on the edge of Lake Carey in the Eastern Goldfields of Western Australia. Superficially it has many similarities with the neighbouring Wallaby and Jupiter magnetic anomalies which are both magnetically mineralised syenite intrusions carrying significant quantities of gold (Wallaby > 7 MOz, Jupiter > 0.5 MOz). The area was flown in 1998 by its then owners, Homestake Gold of Australia, using 50m line spacing and 40m terrain clearance. The survey data are available from the Western Australian Department of Mines - <https://geodownloads.dmp.wa.gov.au/downloads/geophysics/60251.zip>. The magnetic response of Wallaby has been studied in detail (Coggon, 2003; Banaszczyk et al., 2015) and although good fits to the observed data were obtained from modelling dipping cylinders, these models deviated from the geological model. Using susceptibility data from many drill holes, (Coggon, 2003) showed that the response could be modelled using a hollow pipe which reflected the geological model of a near vertical zone of multiple syenite intrusions which had undergone a strong magnetite enrichment event followed by a strong magnetite destruction event associated with the gold mineralisation. This had created a residual annulus of high susceptibility around the altered, gold rich core. Figure 7 shows an image of the Total Magnetic Intensity for Callisto which suggests that it too may be a hollow pipe like body. A compelling exploration target.


Figure 8 shows the inverted model with drill holes coloured and scaled by susceptibility, measured on core using a handheld meter. There is a good match between increases in measured susceptibility and the edge of the magnetic model but a poor match inside of the inverted body, reinforcing the suggestion that the magnetite distribution has a hollow pipe like shape. Although the shallow parts of the inverted model show fingers around a central low susceptibility core the smoothness in the objective function and a lack of constraining data causes the inversion to fill the pipe at depth. The same shape was recovered from an inversion of these data using the UBC MagInv3D package.



1

Figure 7. Field TMI data for the Callisto model.

2



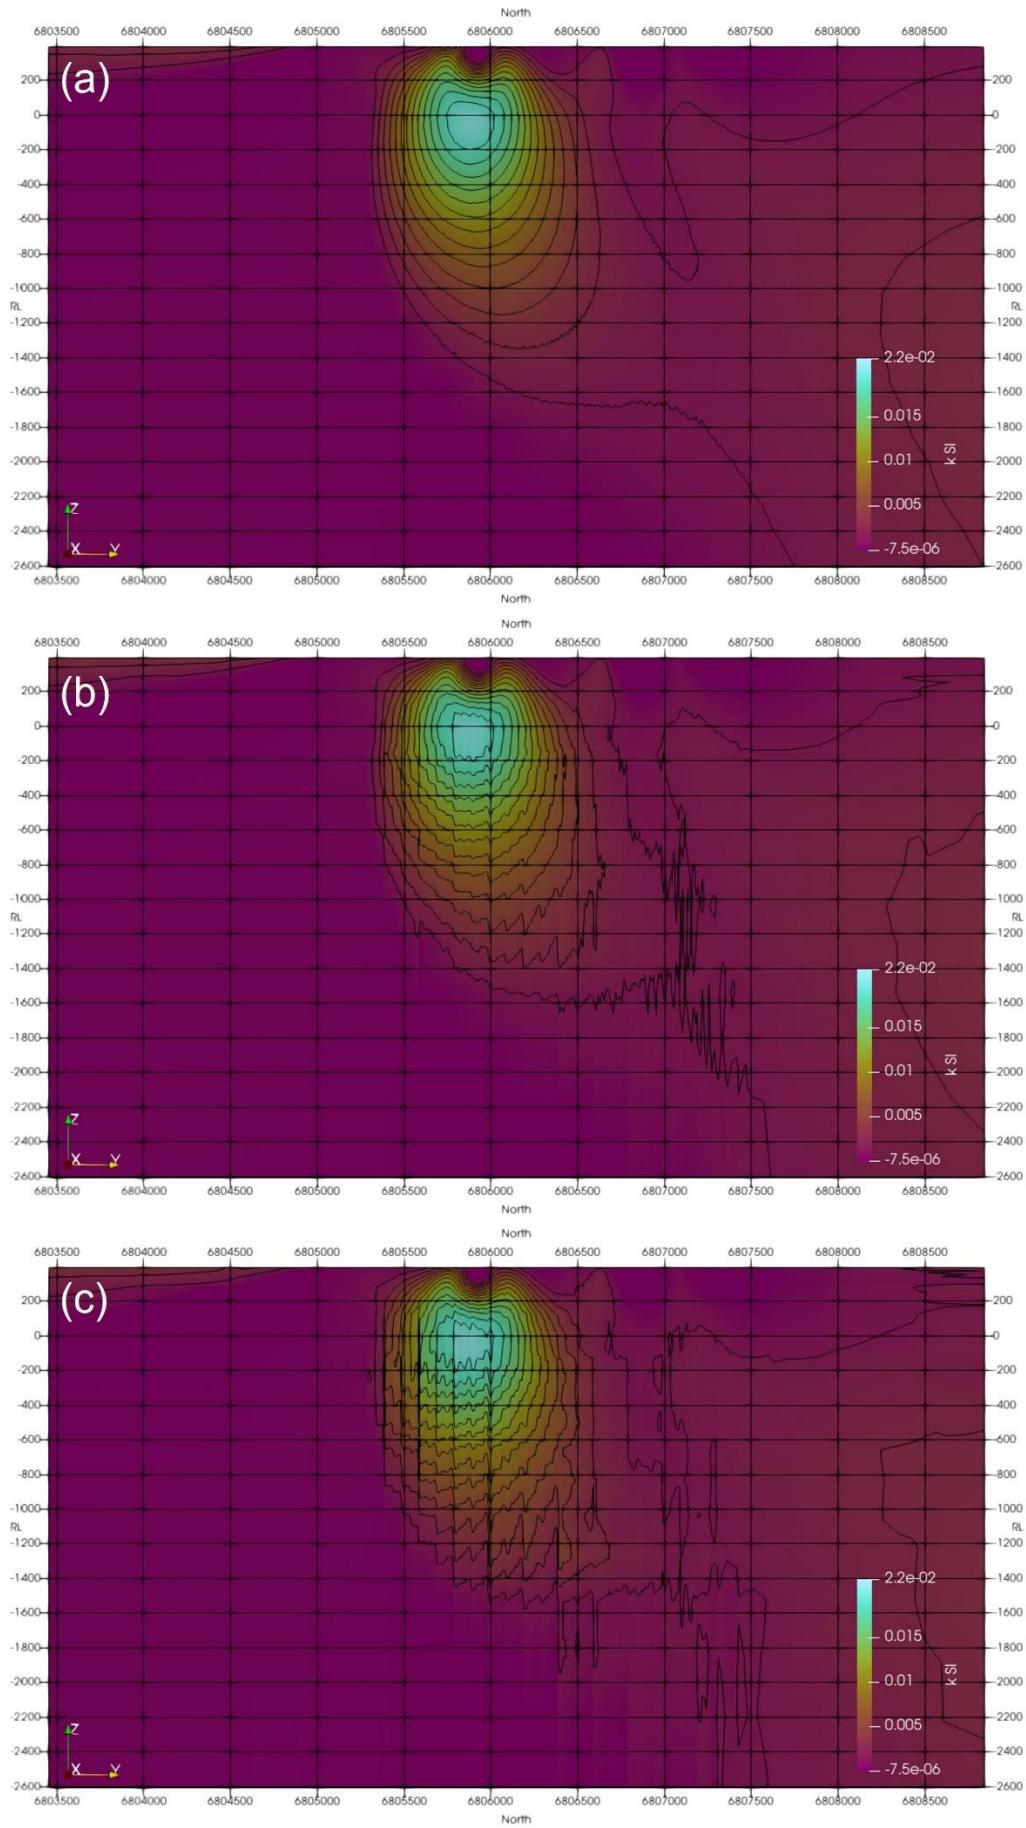
3

Figure 8. The inverted Callisto model with drillholes coloured and scaled by susceptibility. The model contour is shown for susceptibility $k=0.0126$ SI.

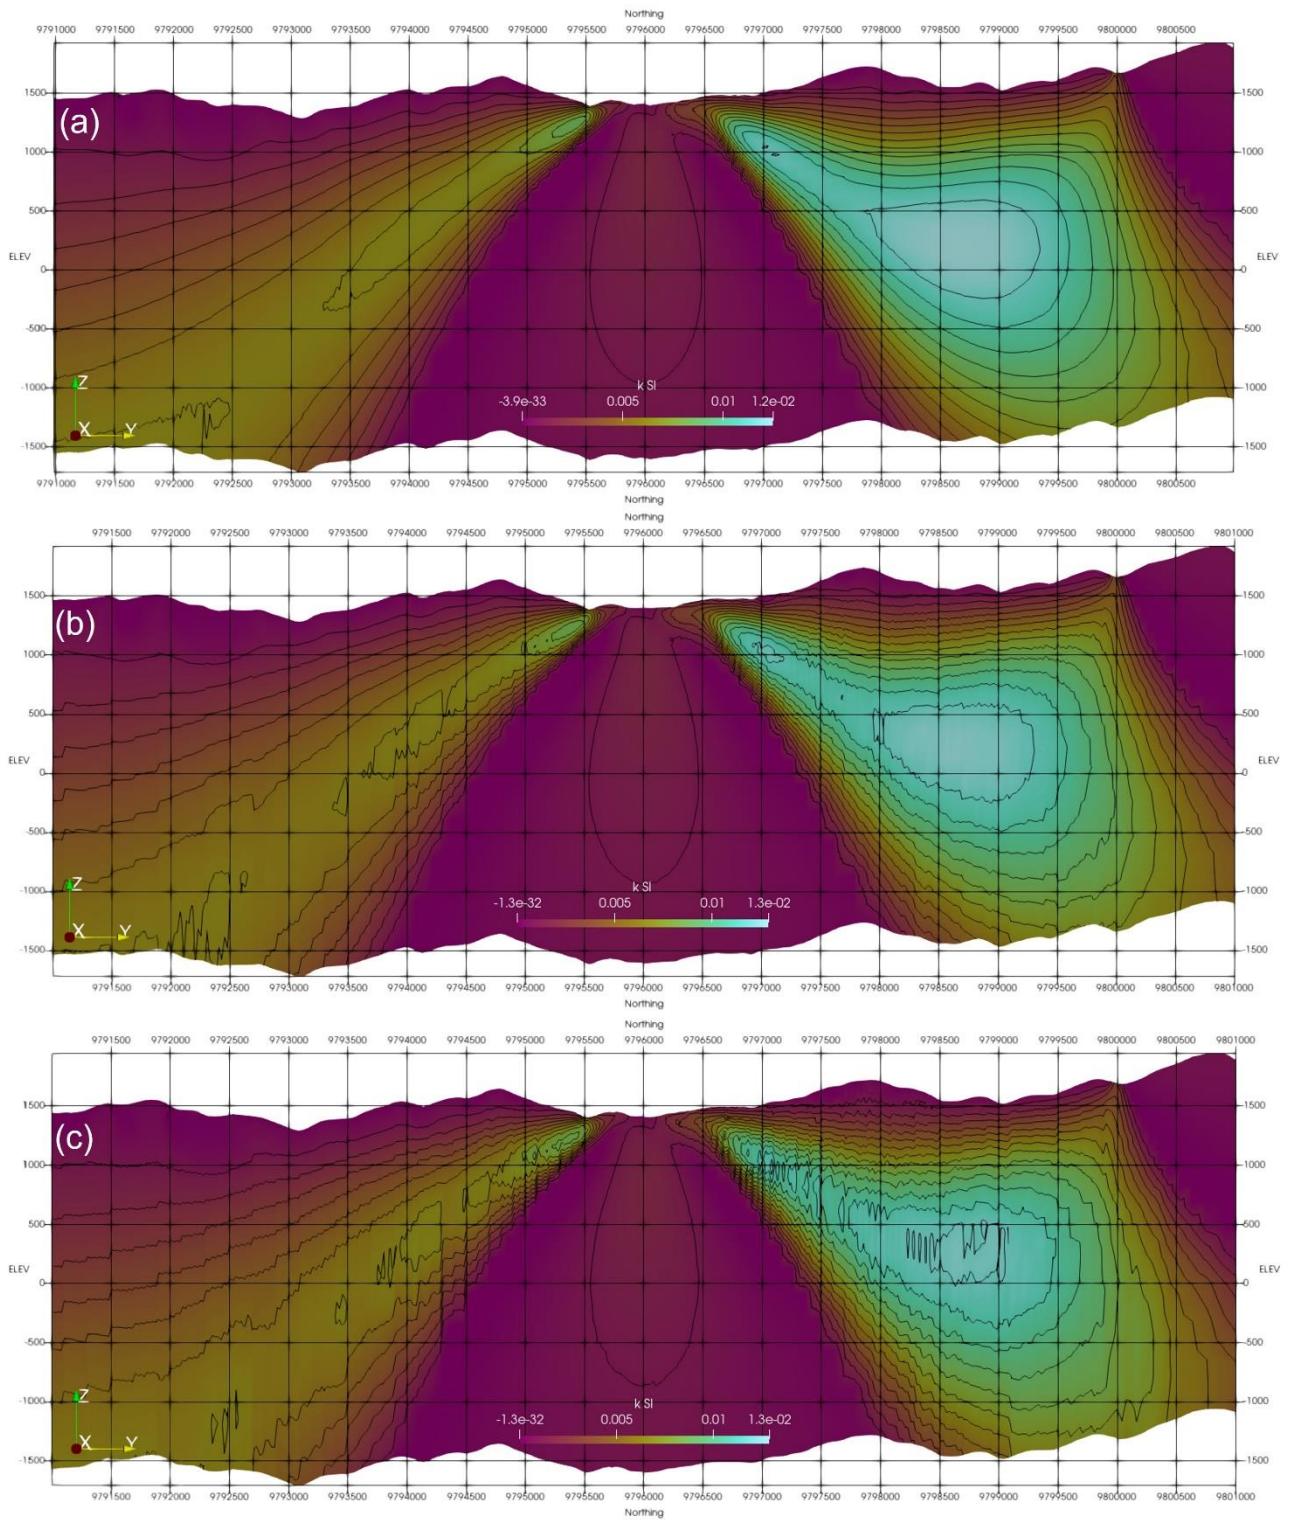
3.3 Different compression rates

To test the effect of differing compression rates, rates of $R_c=5\%$, 1% , and 0.5% , corresponding to a reduction of the sensitivity kernel storage by factors of 20, 100, and 200, respectively were applied. Figure 9 and Figure 10 show a representative section through the inverted models. The effect of increasing degree of compression (i.e., decreasing compression rate) can be seen in the model contours. The weak compression, with $R_c=5\%$, has nearly no influence on the structure of the resulting model. Stronger compressions, with $R_c=1\%$ and 0.5% , result in noise in the contours, while the main model structures remain at the same location with the same shape and model values. As can be seen, the larger Synthetic400 model is less affected by compression than a smaller Callisto model. Generally, the larger the model the greater the compression which can be used without significantly degenerating the model structure. This is due to the usage of multilevel wavelets, where larger models allow more levels to be used. For example, for the large model with dimensions of $N_x=812$, $N_y=848$, $N_z=26$ a compressed rate of $R_c=0.0086$ was successfully used. The optimum compression rate can be either obtained experimentally or calculated from the available memory using Appendix A.

Figure 11 shows the memory usage and CPU time (wall time) for compression rates from $R_c=0.1\%$ to $R_c=50\%$, for the Callisto model, using $N_{cpu}=40$ cores on a machine with shared memory system. The memory was recorded by the value of 'VmHWM' saved in the '/proc/self/status' file. This corresponds to the peak resident set size ("high water mark") memory. The total memory is calculated by summing up this value across all processors. The memory reported in Figure 11 is measured at the end of the run to calculate the maximum RAM memory used. Both the memory used and CPU time scale nearly linearly with the compression rate. The theoretical memory prediction Eq.(A1) has a constant offset from the measured memory. This can be explained by the memory internally allocated by the MPI library. The offset also scales nearly linearly with N_{cpu} and this is explained by the shared memory allocated by the MPI. Thus, the actual amount of the allocated "offset memory" is smaller than reported using this approach, by a factor of $\sim N_{cpu}$.



1
 2 Figure 9. Inversion results for the Callisto model using different compression rates: (a) $R_c = 5\%$, (b)
 3 $R_c = 1\%$, and (c) $R_c = 0.5\%$.



1

Figure 10. Inversion results for the Synthetic400 model using different compression rates: (a) $R_c=5\%$, (b) $R_c=1\%$, and (c) $R_c=0.5\%$.




Figure 11. Memory usage and CPU time for different compression rates, for the Callisto model, using $N_{cpu}=40$.

1

2 **3.4 Magnetisation inversion example**

3

4 In this subsection we show a series of two models to test the magnetisation inversion. In the first
 5 example we demonstrate the capability of the inversion to recover anomalies. In the second
 6 example, we propose a challenging case to evaluate the limits of our code.


7

8 3.4.1 Two-body model

9

10 In order to test the new magnetisation code first a simple two body model was constructed. Each
 11 body was only a single model voxel in size. Two cases are considered with (1) both bodies induced
 12 and (2) both bodies remanently magnetised with a Koenigsberger ratio of 1. The northern (right
 13 hand) body had its remanence directed almost orthogonal to the inducing field ($I_r=60$, $D_r=90$ vs $I_i=-$
 14 60 , $D_i=2$), and it has a susceptibility of 0.1 SI. The southern (left hand) body had a susceptibility of
 15 0.05 SI but its remanent field was parallel to the induced field, which doubled the effective
 16 susceptibility. The host had no remanence and a susceptibility of 10^{-5} SI. The forward response was
 17 computed using both Tomofast-x and UBC's MVIFWD and good agreement found for both the
 18 TMI and three component responses. Three component data were inverted, and the results are
 19 shown in Figure 12. The inverted and true models are in good agreement with respect to anomaly
 20 location, magnetisation vector direction, and the relative anomaly magnitude.

1
2



3

Figure 12. Inverted two-body magnetisation model: (a) both bodies are induced and (b) both bodies are remanent. The true model is shown with white vectors, and inverted model is shown with black vectors. The true and inverted models are not scaled the same, they are both shown for the purpose of gauging fit to direction only. Images have a linear colour stretch, contours are drawn at logarithmic spacing, and vectors are scaled with the log of magnitude.

4

1 3.4.2 Synthetic400 magnetisation model

2

3 In order to test the magnetisation code with a more challenging example the synthetic model
4 described in section 3.1 was revisited and four of the six bodies had remanence applied to them
5 within the parametric model. Table 4 shows the values used.

6


Table 4. Parameters used in the Synthetic400 model for magnetisation tests.

Body	Susceptibility (SI)	Q Koenigsberger ratio	D _r degree	I _r degree
1	0.05	1	0	-10
2	0.07	2	0	-10
3	0.03	0	0	0
4	0.1	3	5	5
5	0.02	1	3	-5
6	0.0015	0	0	0

7

8 As before, the computed response from the bodies was added to the response computed for a
9 magnetic half-space with topography and noise was added. This was then inverted. Despite moving
10 from solving for one unknown (susceptibility) to three unknowns (magnetisation vector) the
11 inversion has done a good job of identifying the bodies also identified previously as shown in
12 Figure 13. The model is smoother than the susceptibility inversion and in the case of the thin sheets
13 and disk the depth is overestimated away from the edges of the sheet or disk. Figure 14-(a) shows a
14 section through the 3D magnetisation model at 154000E which is through the centre of bodies 4 and
15 1. When compared with the same section for the susceptibility inversion (Figure 6-(a)) we note a
16 smoother model and a greater level of sagging below the disk. The low susceptibility cloak below
17 the disk (Body 1) is no longer present and the background is more homogeneous than was the case
18 for the susceptibility inversion. Looking at the response of the low susceptibility body (Body 6) on
19 section 159000E (Figure 14-(b)) we see that the low susceptibility cylinder is resolved to a local
20 magnetisation high and the background is more homogeneous than the susceptibility inversion
21 (Figure 6-(b)).

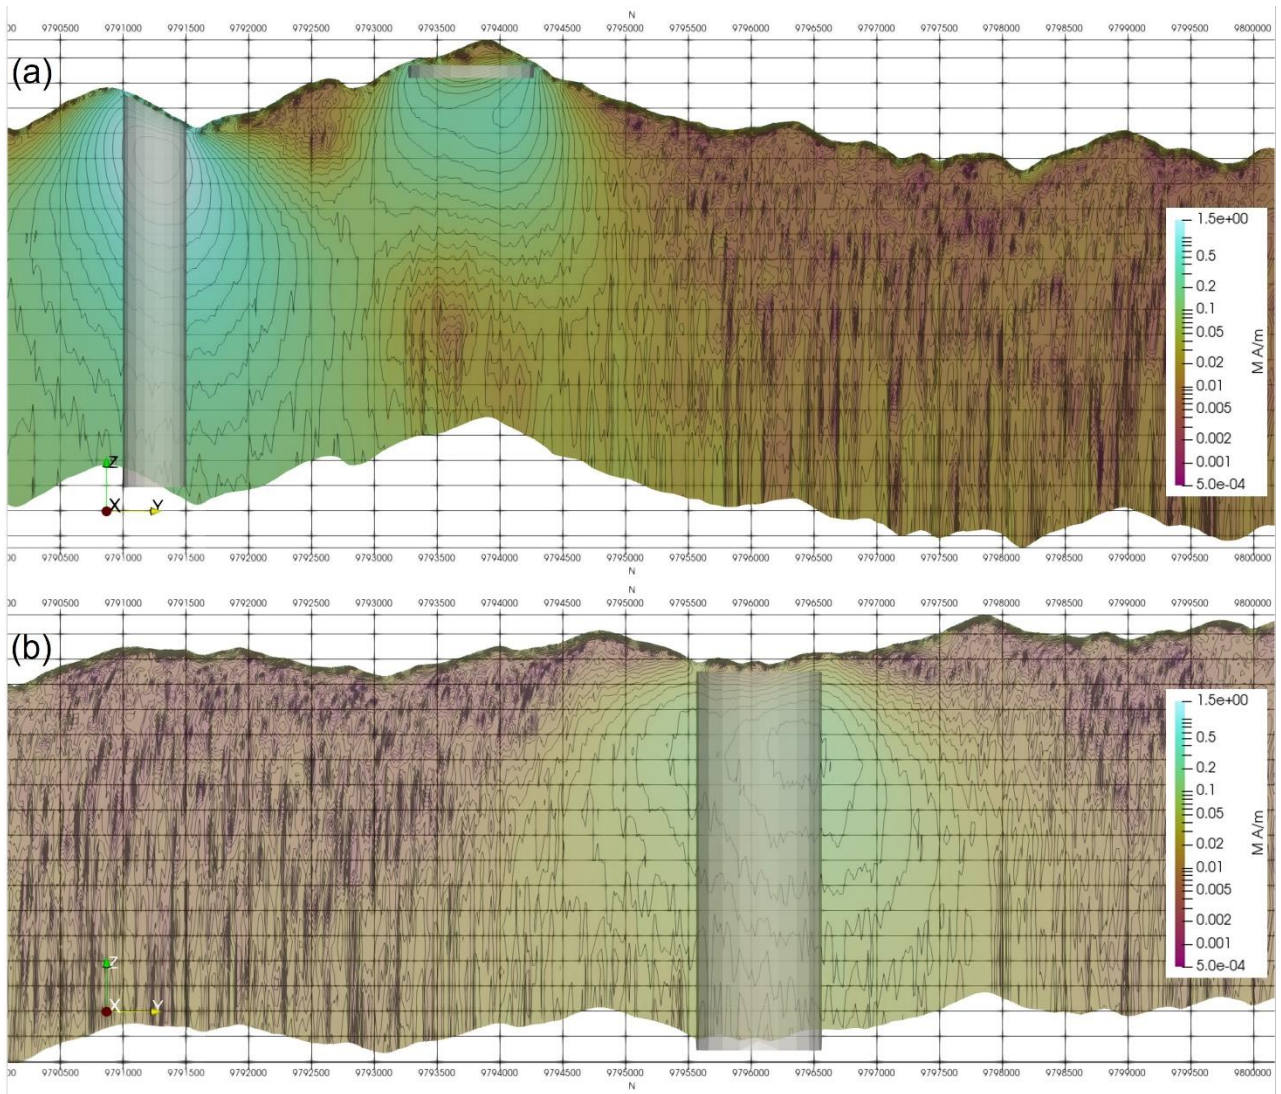
22



1

Figure 13. Inverted Synthetic400 magnetisation model. The contours are shown at $M=0.2$ A/m.

2



1

Figure 14. Vertical sections of the inverted Synthetic400 magnetisation model through (a) high susceptibility cylinder and disk, and (b) low susceptibility cylinder.

2

3 3.5 Parallel performance tests

4

5 We analyse the code parallel performance using two types of machines targeted at different types of
 6 users: (1) a single multi-core machine with shared memory system, and (2) a large supercomputer
 7 with distributed memory across the computational nodes. Most commercial users of the software
 8 are likely to use the first type of machine, while academic users are more likely to use the second
 9 machine type.

10

11 3.5.1 Shared memory tests

12

13 The machine used for shared memory tests is an IBM System x3850 X5 server with four Xeon E7-
 14 4860 decacore processors running at 2.27GHz. The machine has 512 Gb of DDR3 (1066) RAM and
 15 an Ubuntu 20.04 desktop operating system.

16

17 Figure 15 shows the CPU time (wall time) and parallel efficiency for different numbers of MPI
 processors for the Callisto and Synthetic400 models. The parallelisation efficiency is defined as the

1 ratio of the speedup factor and the number of processors. Note that for the Synthetic400 model the
 2 data starting from Ncpu=16 are shown, as it would take very long time to run it on smaller number
 3 of cores; the corresponding parallel efficiency is calculated with respect to Ncpu=16 for this model.
 4 For both models, the CPU time scales nearly linearly with the number of processors up to the
 5 maximum number of cores used Ncpu=40. This is confirmed by high parallel efficiency of greater
 6 than ~90% for all numbers of processors. The nearly linear speed up with number of processors
 7 indicates an efficient use of memory and no sign of bandwidth bottlenecking or I/O limiting
 8 performance barriers. It is expected that this would scale to an even higher number of threads,
 9 particularly on optically linked distributed memory systems.

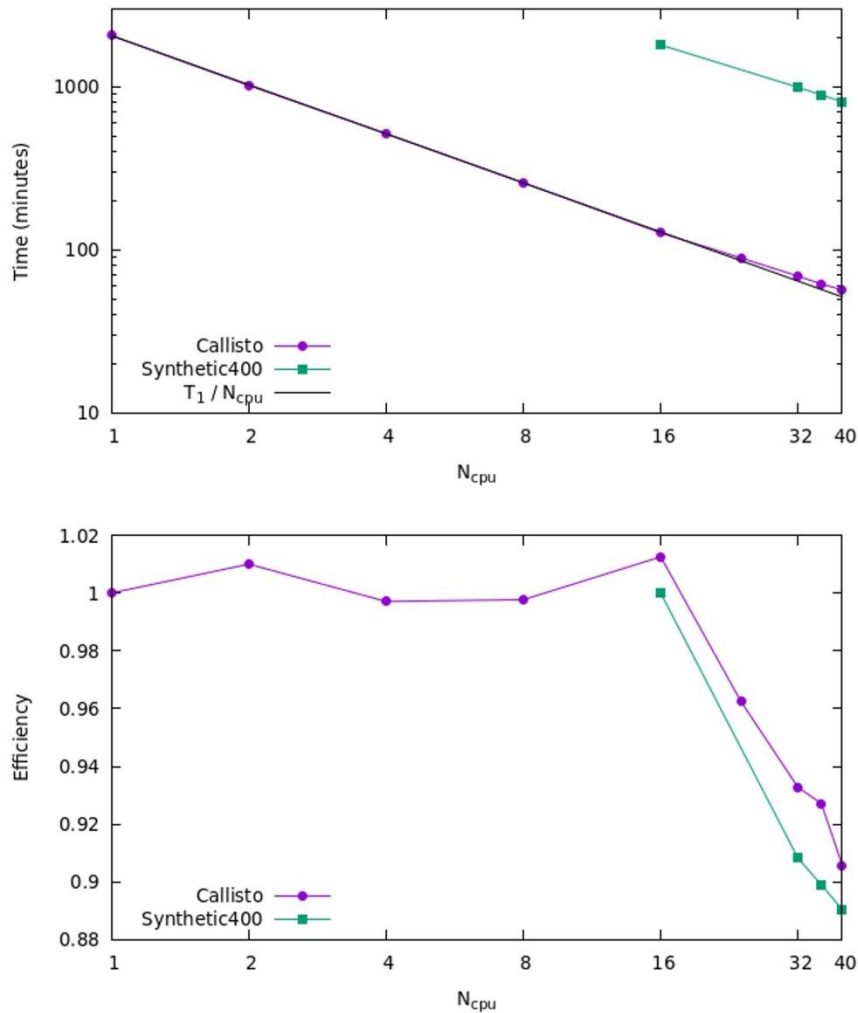


Figure 15. CPU time and parallel efficiency for Callisto and Synthetic400 models, for shared memory tests.

10 To analyse performance, we also measured the MPI time using the Integrated Performance
 11 Monitoring tool (IPM). The IPM has an extremely low overhead and is easy to use, requiring no
 12 source code modification. Various configurations of MPI settings were tested and it was found that
 13 the strongest impact on the performance has the parameter I_MPI_ADJUST_ALLREDUCE. This
 14 controls the algorithm used for the collective MPI communications in the inversion solver. For
 15 example, by setting I_MPI_ADJUST_ALLREDUCE=3, corresponding to “Reduce + Bcast”
 16 algorithm, the MPI time for Synthetic400 model was decreased by 25% on 40 cores. The optimal
 17

1 parameter depends on the machine settings, model size and the number of cores, and must be
2 determined experimentally.


3 3.5.2 Distributed memory tests

4

5 For distributed memory tests we use the Gadi supercomputer of the National Computational
6 Infrastructure (NCI) which is Australia's largest research supercomputing facility. Gadi contains
7 more than 250,000 CPU cores, 930 Terabytes of memory and 640 GPUs. This includes 3,074 nodes
8 each containing two 24-core Intel Xeon Scalable 'Cascade Lake' processors and 192 Gigabytes of
9 memory.

10 Figure 16 shows the CPU time (wall time) and parallel efficiency for different numbers of MPI
11 processors for the Callisto and Synthetic400 models. Note that for the Synthetic400 model, the data
12 starting from Ncpu=8 are shown, and the corresponding parallel efficiency is calculated with
13 respect to Ncpu=8. We observe linear speedup up to Ncpu=48, as confirmed by parallel efficiency
14 greater than ~95%. For a higher number of processors, the efficiency is starting to decline, reaching
15 ~80% at Ncpu=96 and dropping below 50% for Ncpu>624. Overall, through parallelization, we
16 were able to reduce the total computational time by more than two orders of magnitude, achieving
17 maximum speedup of approximately 340 and 420 for Callisto and Synthetic400 models,
18 respectively.

19



1 /
 2 Figure 16. CPU time and parallel efficiency for Callisto and Synthetic400 models, for distributed
 3 memory tests.
 4

1 **Discussion & Conclusions**

2

3 Since its initial release in 2021 the Tomofast-x code base has undergone a significant overhaul in
4 terms of both its versatility and performances. Solving for magnetisation rather than susceptibility
5 alone better reflects the real world, where magnetic remanence is much more common than many
6 would like to admit. In addition, allowing input data to have multiple components improves the
7 inverted model relative to a single component inversion of synthetic data, where the answer is
8 known.

9

10 Meshes with expanding vertical thickness with depth and matrix wavelet compression have allowed
11 larger problems to be solved in real time. Improvements to the parallelisation scheme have allowed
12 for better performance on both shared and distributed memory systems with speedups of ~40 and
13 ~400 times, respectively, for moderate size models. Memory reduction of up to 100 times are
14 achieved due to wavelet compression of the sensitivity kernel.

15

16 Following the changes to the code documented here, Tomofast-x 2.0 produces results very similar
17 to those obtained from the UBC suite of potential field inversion codes. It does so in a similar run
18 time and uses a comparable amount of system resources. However, unlike the commercial version
19 of the UBC suite it runs on Linux platforms which, because of memory and CPU slot limitations
20 imposed on common versions of the Windows operating system, will allow it to be used on more
21 powerful systems than those typically running Windows. In modern Windows systems it can run on
22 the Windows Subsystem for Linux (WSL), rather than as a Windows application. As maintenance is
23 planned for the next foreseeable future, we argue that Tomofast-x can be used for commercial
24 projects and demanding, large scale inversions.

25

26 The optimisation work described in this paper is fundamental to allowing more complete
27 descriptions of the controls on magnetisation, including remanence, which require 3 to 9 times the
28 computer resources of the scalar description.

29

30 Whilst the modifications described above address many of the challenges of industry and
31 government agencies, further developments are planned, including a closer integration with level set
32 inversions (Giraud et al., 2023a). To extend the gravity to plate-scale problems we would need to
33 reframe the calculations with a spherical coordinate scheme. Increasing use of airborne gravity
34 gradiometry data suggests that including this type of data directly in the inversion scheme also
35 warrants some attention. Finally, and in the longer term, the addition of other types of physics to the
36 platform will have obvious benefits.

37

38

1 **Appendix A. Total memory estimation**

2

3 The total memory requirements for an inverse problem with model damping constraints running in
 4 parallel on N_{cpu} cores has been estimated. The matrix is partitioned by columns (i.e., parallelisation
 5 is by model parameters), and single precision floats are used for storing matrix vales. For
 6 comparison, both formats (CSR and DCSR) of sparse matrix storage are shown. Note that DCSR
 7 format requires less memory on multiple cores due to not having to store the indexes of empty rows.
 8 The Table A1 shows memory requirement for different types of data structure used in the code.

9

Table A1. Memory requirements for data structures used in the code.

Data structure	Memory (bytes)
Matrix (CSR format)	$8 (R_c N_{model} N_{data} + N_{model}) + 8 N_{cpu} (N_{data} + N_{model})$
Matrix (DCSR format)	$8 (R_c N_{model} N_{data} + N_{model}) + 12 (N_{cpu} N_{data} + N_{model})$
Right-hand side	$8 N_{cpu} (N_{data} + N_{model})$
Data (obs, calc, residual)	$24 N_{cpu} N_{data}$
Depth weight	$8 N_{model}$
Models (prior, delta, current, full)	$32 N_{model}$
Inversion solver	$16 N_{model}$
Auxiliary model buffers	$24 N_{model}$
Grid (6 doubles, 4 integers)	$64 N_{model}$

10

11 It must be noted that the solver memory is independent of the number of CPUs. To achieve this, the
 12 right-hand side vector is used as a buffer for internal solver calculations, and memory allocation for
 13 the intermediate products is avoided by using existing memory buffers (for details see item 5 in
 14 Section 2.6). The depth-weight and model arrays are split between the processes, so their memory is
 15 independent of the number of cores. Here, the grid memory is allocated only on the master rank to
 16 write the final models. The total memory (for the DCSR matrix format) can be estimated as the sum
 17 of arrays in Table A1 as:

18

$$M_{tot} = 8R_c N_{model} N_{data} + 8N_{cpu} N_{model} + 44N_{cpu} N_{data} + 164N_{model} \quad A1$$

19

20 This equation can be used to estimate the compression rate, R_c , from the total amount of available
 21 memory. It must be noted that the MPI library can allocate additional memory, which is
 22 implementation dependant and can vary between different machines.

23

1 **Code and Data availability**

2

3 The data sets and Tomofast-x input and output files for the Calisto and Synthetic400 models have
4 been packaged and are available for download at Zenodo (Ogarko et al., 2023a). Readers are
5 encouraged to use these when testing Tomofast-x as well as other codes they may be working with.

6 The latest Tomofast-x code version with user manual and some representative tests is available for
7 download in GitHub (<https://github.com/TOMOFAST/Tomofast-x>) and at Zenodo (Ogarko et al.,
8 2023b).

9

1 **Author contribution**

2

3 VO is the main contributor to the writing of this article, Tomofast-x code development, testing, and
4 preparation of the assets. KF and TL worked together with VO on code development, testing and
5 preparation of the datasets for the tests. JG performed distributed memory scaling tests. All authors
6 participated in the discussion and redaction of this paper.

7

8

1 **Competing interests**

2

3 The authors declare that they have no conflict of interest.

4

1 **Acknowledgements**

2

3 This research was supported by the Mineral Exploration Cooperative Research Centre (MinEx
4 CRC), whose activities are funded by the Australian Government's Cooperative Research Centre
5 Programme. This is the MinEx CRC Document number 2023/38. VO acknowledges Rodrigo Tobar
6 for technical discussions. JG acknowledges support from European Union's Horizon 2020 research
7 and innovation programme under the Marie Skłodowska-Curie grant agreement No. 101032994.

8

9

References

- Banaszczyk, S., Dentith, M., and Wallace, Y.: Constrained Magnetic Modelling of the Wallaby Gold Deposit, Western Australia, ASEG Extended Abstracts, 2015, 1–4, <https://doi.org/10.1071/ASEG2015ab290>, 2015.
- Barrett, R., Berry, M., Chan, T. F., Demmel, J., Donato, J., Dongarra, J., Eijkhout, V., Pozo, R., Romine, C., and van der Vorst, H.: Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods, Society for Industrial and Applied Mathematics, <https://doi.org/10.1137/1.9781611971538>, 1994.
- Bhattacharyya, B. K.: MAGNETIC ANOMALIES DUE TO PRISM-SHAPED BODIES WITH ARBITRARY POLARIZATION, GEOPHYSICS, 29, 517–531, <https://doi.org/10.1190/1.1439386>, 1964.
- Buluc, A. and Gilbert, J. R.: On the representation and multiplication of hypersparse matrices, in: 2008 IEEE International Symposium on Parallel and Distributed Processing, 1–11, <https://doi.org/10.1109/IPDPS.2008.4536313>, 2008.
- Buluç, A., Fineman, J. T., Frigo, M., Gilbert, J. R., and Leiserson, C. E.: Parallel sparse matrix-vector and matrix-transpose-vector multiplication using compressed sparse blocks, in: Proceedings of the twenty-first annual symposium on Parallelism in algorithms and architectures, 233–244, <https://doi.org/10.1145/1583991.1584053>, 2009.
- Cockett, R., Kang, S., Heagy, L. J., Pidlisecky, A., and Oldenburg, D. W.: SimPEG: An open source framework for simulation and gradient based parameter estimation in geophysical applications, Comput Geosci, 85, 142–154, <https://doi.org/10.1016/j.cageo.2015.09.015>, 2015.
- Coggon, J.: Magnetism — key to the Wallaby gold deposit, Exploration Geophysics, 34, 125–130, <https://doi.org/10.1071/EG03125>, 2003.
- Ellis, R. G., de Wet, B., and Macleod, I. N.: Inversion of Magnetic Data from Remanent and Induced Sources, ASEG Extended Abstracts, 2012, 1–4, <https://doi.org/10.1071/ASEG2012ab117>, 2012.
- Frankcombe, K.: Magnetics in the mountains: An approximation for the magnetic response from topography, Preview, 2020, 34–44, <https://doi.org/10.1080/14432471.2020.1773235>, 2020.
- Gidley, P. R.: The Geophysics of the Trough Tank Gold-Copper Prospect, Australia, Exploration Geophysics, 19, 76–78, <https://doi.org/10.1071/EG988076>, 1988.
- Giraud, J., Lindsay, M., Ogarko, V., Jessell, M., Martin, R., and Pakyuz-Charrier, E.: Integration of geoscientific uncertainty into geophysical inversion by means of local gradient regularization, Solid Earth, 10, 193–210, <https://doi.org/10.5194/se-10-193-2019>, 2019a.
- Giraud, J., Ogarko, V., Lindsay, M., Pakyuz-Charrier, E., Jessell, M., and Martin, R.: Sensitivity of constrained joint inversions to geological and petrophysical input data uncertainties with posterior geological analysis, Geophys J Int, 218, 666–688, <https://doi.org/10.1093/gji/ggz152>, 2019b.
- Giraud, J., Lindsay, M., Jessell, M., and Ogarko, V.: Towards plausible lithological classification from geophysical inversion: honouring geological principles in subsurface imaging, Solid Earth, 11, 419–436, <https://doi.org/10.5194/se-11-419-2020>, 2020.

1 Giraud, J., Ogarko, V., Martin, R., Jessell, M., and Lindsay, M.: Structural, petrophysical, and
2 geological constraints in potential field inversion using the Tomofast-x v1.0 open-source code,
3 Geosci Model Dev, 14, 6681–6709, <https://doi.org/10.5194/gmd-14-6681-2021>, 2021.

4 Giraud, J., Caumon, G., Grose, L., Ogarko, V., and Cupillard, P.: Integration of automatic
5 implicit geological modelling in deterministic geophysical inversion, EGU sphere, 2023, 1–33,
6 <https://doi.org/10.5194/egusphere-2023-129>, 2023a.

7 Giraud, J., Seillé, H., Lindsay, M. D., Visser, G., Ogarko, V., and Jessell, M. W.: Utilisation of
8 probabilistic magnetotelluric modelling to constrain magnetic data inversion: proof-of-concept
9 and field application, Solid Earth, 14, 43–68, <https://doi.org/10.5194/se-14-43-2023>, 2023b.

10 Goodwin, J. A. and Lane, R. J. L.: The North Australian Craton 3D Gravity and Magnetic
11 Inversion Models: A trial for first pass modelling of the entire Australian continent,
12 <https://doi.org/10.11636/Record.2021.033>, 2021.

13 Jessell, M.: Three-dimensional geological modelling of potential-field data, Comput Geosci,
14 27, 455–465, [https://doi.org/10.1016/S0098-3004\(00\)00142-4](https://doi.org/10.1016/S0098-3004(00)00142-4), 2001.

15 Li, Y. and Oldenburg, D. W.: 3-D inversion of magnetic data, GEOPHYSICS, 61, 394–408,
16 <https://doi.org/10.1190/1.1443968>, 1996a.

17 Li, Y. and Oldenburg, D. W.: Joint inversion of surface and three-component borehole
18 magnetic data, in: SEG Technical Program Expanded Abstracts 1996, 1142–1145,
19 <https://doi.org/10.1190/1.1826293>, 1996b.

20 Li, Y. and Oldenburg, D. W.: 3-D inversion of gravity data, GEOPHYSICS, 63, 109–119,
21 <https://doi.org/10.1190/1.1444302>, 1998.

22 Li, Y. and Oldenburg, D. W.: Fast inversion of large-scale magnetic data using wavelet
23 transforms and a logarithmic barrier method, Geophys J Int, 152, 251–265,
24 <https://doi.org/10.1046/j.1365-246X.2003.01766.x>, 2003.

25 Liu, S., Hu, X., Zhang, H., Geng, M., and Zuo, B.: 3D Magnetization Vector Inversion of
26 Magnetic Data: Improving and Comparing Methods, Pure Appl Geophys, 174, 4421–4444,
27 <https://doi.org/10.1007/s00024-017-1654-3>, 2017.

28 Markus, A.: Modern Fortran in Practice, Cambridge University Press,
29 <https://doi.org/10.1017/CBO9781139084796>, 2012.

30 Martin, R., Monteiller, V., Komatitsch, D., Perrouty, S., Jessell, M., Bonvalot, S., and Lindsay,
31 M.: Gravity inversion using wavelet-based compression on parallel hybrid CPU/GPU systems:
32 application to southwest Ghana, Geophys J Int, 195, 1594–1619,
33 <https://doi.org/10.1093/gji/ggt334>, 2013.

34 Martin, R., Ogarko, V., Komatitsch, D., and Jessell, M.: Parallel three-dimensional electrical
35 capacitance data imaging using a nonlinear inversion algorithm and L_p
36 L_p norm-based model regularization, Measurement, 128,
37 428–445, <https://doi.org/10.1016/j.measurement.2018.05.099>, 2018.

38 Martin, R., Giraud, J., Ogarko, V., Chevrot, S., Beller, S., Gégout, P., and Jessell, M.: Three-
39 dimensional gravity anomaly data inversion in the Pyrenees using compressional seismic
40 velocity model as structural similarity constraints, Geophys J Int, 225, 1063–1085,
41 <https://doi.org/10.1093/gji/ggaa414>, 2021.

42 Nettleton, L. L.: GRAVITY AND MAGNETIC CALCULATIONS, GEOPHYSICS, 7, 293–
43 310, <https://doi.org/10.1190/1.1445015>, 1942.

44

- 1 Ogarko, V., Giraud, J., Martin, R., and Jessell, M.: Disjoint interval bound constraints using
2 the alternating direction method of multipliers for geologically constrained inversion:
3 Application to gravity data, *GEOPHYSICS*, 86, G1–G11, [https://doi.org/10.1190/geo2019-](https://doi.org/10.1190/geo2019-0633.1)
4 0633.1, 2021.
- 5 Ogarko, V., Frankcombe, K., and Liu, T.: Callisto and Synthetic400 models (1.0) [Data set],
6 2023a.
- 7 Ogarko, V., Frankcombe, K., Liu, T., Giraud, J., Martin, R., and Jessell, M.: Tomofast-x 2.0
8 inversion source code, 2023b.
- 9 Ogarko, V., Frankcombe, K., Liu, T., Giraud, J., Martin, R., and Jessell, M.: Tomofast-x v.2.0,
10 2023c.
- 11 Özgü Arısoy, M. and Dikmen, Ü.: Potensoft: MATLAB-based software for potential field data
12 processing, modeling and mapping, *Comput Geosci*, 37, 935–942,
13 <https://doi.org/10.1016/j.cageo.2011.02.008>, 2011.
- 14 Paige, C. C. and Saunders, M. A.: LSQR: An Algorithm for Sparse Linear Equations and
15 Sparse Least Squares, *ACM Transactions on Mathematical Software*, 8, 43–71,
16 <https://doi.org/10.1145/355984.355989>, 1982.
- 17 Pears, G., Reid, J., Chalke, T., Tschirhart, V., and Thomas, M. D.: Advances in geologically
18 constrained modelling and inversion strategies to drive integrated interpretation in mineral
19 exploration, in: *Proceedings of Exploration*, 221–238, 2017.
- 20 Vallabh Sharma, P.: Rapid computation of magnetic anomalies and demagnetization effects
21 caused by bodies of arbitrary shape, *Pure and Applied Geophysics PAGEOPH*, 64, 89–109,
22 <https://doi.org/10.1007/BF00875535>, 1966.

23
24
25