

An efficient 2D inversion scheme for airborne frequency-domain data

Tue Boesen¹, Esben Auken¹, Anders Vest Christiansen¹, Gianluca Fiandaca¹,
Casper Kirkegaard², Andreas Aspmo Pfaffhuber³, and Malte Vöge³

ABSTRACT

In many cases, inversion in 2D gives a better description of the subsurface compared with 1D inversion, but, computationally, 2D inversion is expensive, and it can be hard to use for large-scale surveys. We have developed an efficient hybrid 2D airborne frequency-domain electromagnetic inversion algorithm. Our hybrid scheme combines 1D and 2D inversions in a three-stage process, in which each step is progressively more accurate and computationally more expensive than the previous one. This results in an approximately $2\times$ – $6\times$ speedup compared with full 2D inversions, and with only minor changes to the inversion results. Our inversion structure is based on a regular grid, in which each sounding is discretized individually. The 1D modeling code uses layered

models with derivatives derived through the finite-difference method, whereas our 2D modeling code uses an adaptive finite-element mesh, and it uses the adjoint-state method to calculate the derivatives. By incorporating the inversion grid structure into the 2D finite-element mesh, interpolation between the different meshes becomes trivial. Large surveys are handled by using local meshing to split large surveys into small sections, which retains the 2D information. The algorithm is heavily optimized and parallelized over the frequencies and sections, with good scalability even on nonuniform memory architecture systems, on which it is generally hard to achieve a satisfactory scaling. The algorithm has been tested successfully with various synthetic studies as well as field examples, of which results from two synthetic studies and a field example are shown.

INTRODUCTION

Airborne electromagnetic (AEM) surveys typically contain thousands of line kilometers of data and are routinely flown for mapping of geology, groundwater, saltwater intrusion, etc. Most data are inverted using 1D model algorithms, which have proven to be robust and computationally fast. However, specific targets with a high conductivity contrast between undulating bedrock and sediments, or conductive sheet-like mineralizations, need higher dimensionality in the underlying model to be resolved accurately (Doyle et al., 1999; Wilson et al., 2012; Yang and Oldenburg 2012b). The challenge in moving beyond 1D modeling is that it is prohibitively computationally expensive to invert for a 2D or a 3D model and this limits the usage of these inversion algorithms on a routine basis, even for frequency-domain electromagnetic (EM) data sets of just a few discrete frequencies.

Full 3D EM inversion algorithms have existed for more than a decade (Haber et al., 2007b), and with the concept of the moving

footprint (Cox et al., 2010), several 3D codes using local meshes have been presented (Cox et al., 2012; Yang et al., 2014). Both these algorithms are capable of handling large surveys, in the time- or frequency-domain, by sectioning the survey into smaller parts using a local meshing approach. Local meshing means that the survey is split into smaller parts, in which each part contains a small subset of transmitter-receiver pairs, as well as all the models within their footprint domain (Liu and Becker, 1990; Beamish, 2003; Reid et al., 2006). Here, the footprint is defined as the area of significant lateral sensitivity of the survey system, and its size is thus dependent upon the system itself and the resistivity of the earth. Reid et al. (2006) show that, for a frequency-domain system, the footprint may be as large as 10 times the flight altitude for low induction numbers. Common survey systems operate either in the time domain, or with multiple frequencies spread across the frequency spectrum, and although these latter systems usually have one transmitter frequency

Manuscript received by the Editor 5 May 2017; revised manuscript received 8 January 2018; published ahead of production 22 March 2018; published online 30 May 2018.

¹Aarhus University, Aarhus, Denmark. E-mail: tue@geo.au.dk; esben.auken@geo.au.dk; anders.vest@geo.au.dk; gianluca.fiandaca@geo.au.dk.

²QIAGEN Aarhus, Aarhus, Denmark. E-mail: casper.kirkegaard@gmail.com.

³Norwegian Geotechnical Institute, Oslo, Norway. E-mail: andreas.a.pfaffhuber@ngi.no; malte.voege@ngi.no.

© 2018 Society of Exploration Geophysicists. All rights reserved.

operating within the low induction approximation, most of their transmitter frequencies operate at higher induction numbers, for which the footprint is much smaller. From this, we argue that when a survey is flown with a line separation of 200–500 m, most of any crossline information is lost and the survey results are essentially reduced to only contain inline information. Considering this and the inherent computational burden of doing full 3D inversions, it is clear that there are areas where it is sufficient and even desirable to operate within a 2D formulation. Several 2D inversion algorithms have been presented over the years: Mitsuhashi and Uchida (2002), Wilson et al. (2006), Li et al. (2016), and Key and Oval (2011) develop 2D finite-element algorithms, whereas Abubakar et al. (2008) use a finite-difference approach and Yu and Haber (2012) present a finite-volume approach. In general, finite-difference approaches are considered the most simple and inaccurate of the three approaches, but they can sometimes be justified due to their superior parallel scaling. The finite element is the most accurate of the methods, but it is also the most computationally heavy, and largely, the question of whether the finite element or finite volume is the superior choice remains open (Jahandari et al., 2017).

In this paper, we present a hybrid 1D/2D inversion code for frequency-domain helicopter EM (HEM) data designed for inverting field-scale surveys on desktop computers. Because an efficient 2D modeling algorithm is vital to achieve this goal, and because our 2D algorithm has not previously been published, this paper begins with the construction of our 2D algorithm and the foundation it is built on. The 2D algorithm is based on the 2.5D formulation by Stoyer and Greenfield (1976), along with field separation into primary and secondary fields, where the high-frequency singularity is handled by the introduction of a finite resistivity of the air (Mitsuhashi, 2000). The algorithm has a triangular finite-element mesh for the 2D forward and derivative calculations and a regular grid for the inversion. When having multiple meshes, interpolation schemes are needed to map variables between the meshes. In general, interpolation between meshes is a nontrivial task (Caudillo-Mata et al., 2017), but in our case, the task is made trivial, by using the regular grid as a skeletal structure for building the finite-element mesh. Due to limited memory, as well as performance concerns, we introduce sectioning, which splits large survey lines into smaller sections. Sectioning is only done when carrying out the 2D forward and derivative calculations, during which we enforce sufficient overlap, such that vital 2D information is preserved. Based on the overlap size, we show how section sizes should be chosen to reach optimal performance. The algorithm is written in Fortran and uses OpenMP, Intel MKL libraries, as well as a custom-built block-parallel sparse iterative linear solver. The algorithm is part of AarhusInv, which is provided as freeware for noncommercial academic purposes (Auken et al., 2014).

Following the presentation of our 2D algorithm, we present a hybrid scheme, inspired by the work presented in Christiansen and Auken (2004) and Christiansen et al. (2015). The method starts by performing 1D forward and inverse calculations; later, it switches to 2D forward calculations and 1D derivatives; and, finally, it ends with full 2D calculations. The result of this is a code that produces 2D results, but with a substantially shorter computational time than traditional 2D algorithms. We demonstrate these computational benefits using two synthetic models and a field example. Finally, we discuss the trade-off between computational speed and accuracy, how the algorithm is best parallelized, and we illus-

trate the code's parallel scaling and performance on a multiprocessor system.

METHODOLOGY

The 2.5D forward algorithm is based upon the framework of Stoyer and Greenfield (1976). In a 2.5D formulation, the earth is described in 2D (homogeneous in the cross-section direction), but the source is 3D. The source needs to be 3D, to describe accurately the sources used in AEM because AEM sources do not produce source fields, which are reasonably homogeneous in any direction. Our 2.5D algorithm was originally developed for marine EM measurements, but sources in the air have been added, so it can now be used to invert airborne frequency-domain EM data (Vöge et al., 2015). For hybrid inversion, we include the 1D algorithm of Auken et al. (2014), which is based on the layered 1D model solution presented in Ward and Hohmann (1988).

Governing equations

Starting from Maxwell's equations in the frequency domain,

$$\nabla \times \mathbf{E} + i\omega\boldsymbol{\mu} \cdot \mathbf{H} = 0, \quad (1)$$

$$\nabla \times \mathbf{H} - i\omega\boldsymbol{\epsilon} \cdot \mathbf{E} = \mathbf{J}, \quad (2)$$

where $\mathbf{E}$ and $\mathbf{H}$ are the electric and magnetic fields, $\mathbf{J}$ is the electric source current, i is the imaginary unit, ω is the frequency, $\boldsymbol{\mu}$ is the magnetic permeability, and $\boldsymbol{\epsilon}$ is the complex dielectric permittivity, with $\boldsymbol{\epsilon} = \boldsymbol{\epsilon}_0 - i(\sigma/\omega)$, where σ is the conductivity. To minimize the forward inaccuracy, the fields are split into primary and secondary fields:

$$\mathbf{E} = \mathbf{E}_p + \mathbf{E}_s, \quad \mathbf{H} = \mathbf{H}_p + \mathbf{H}_s. \quad (3)$$

So, equations 1 and 2 can be written as

$$(\nabla \times \mathbf{E}_p + i\omega\boldsymbol{\mu} \cdot \mathbf{H}_p) + (\nabla \times \mathbf{E}_s + i\omega\boldsymbol{\mu} \cdot \mathbf{H}_s) = 0, \quad (4)$$

$$(\nabla \times \mathbf{H}_p - i\omega\boldsymbol{\epsilon} \cdot \mathbf{E}_p) + (\nabla \times \mathbf{H}_s - i\omega\boldsymbol{\epsilon} \cdot \mathbf{E}_s) = \mathbf{J}. \quad (5)$$

We separate the conductivity into a conductivity for the primary field model $\boldsymbol{\sigma}_p$ and a secondary field model $\boldsymbol{\sigma}_s$, where $\boldsymbol{\sigma}_s = \boldsymbol{\sigma} - \boldsymbol{\sigma}_p$. From this, we get $i\omega\boldsymbol{\epsilon} = i\omega\boldsymbol{\epsilon}_0 - (\boldsymbol{\sigma}_p + \boldsymbol{\sigma}_s) = i\omega\boldsymbol{\epsilon}_p - \boldsymbol{\sigma}_s$, which allows us to split equations 4 and 5 into separate equation systems for primary field

$$\nabla \times \mathbf{E}_p + i\omega\boldsymbol{\mu} \cdot \mathbf{H}_p = 0, \quad (6)$$

$$\nabla \times \mathbf{H}_p - i\omega\boldsymbol{\epsilon}_p \cdot \mathbf{E}_p = \mathbf{J}, \quad (7)$$

and the secondary field

$$\nabla \times \mathbf{E}_s + i\omega\boldsymbol{\mu} \cdot \mathbf{H}_s = 0, \quad (8)$$

$$\nabla \times \mathbf{H}_s - i\omega\boldsymbol{\epsilon} \cdot \mathbf{E}_s = \boldsymbol{\sigma}_s \cdot \mathbf{E}_p. \quad (9)$$

The primary field is computed analytically, so the inaccuracy of the finite-element method affects only the secondary field, which is

several orders of magnitude smaller than the total field. In our case, we chose the primary field model to be a uniform full-space air model with magnetic point sources and receivers, which can be easily calculated analytically. Because σ_p is the conductivity of air, we have $\sigma = \sigma_p$ and $\sigma_s = 0$ in the air layer.

The Fourier transform is defined with respect to y (i.e., the strike direction of the survey) as

$$\tilde{F}(x, k_y, z) = \int_{-\infty}^{+\infty} F(x, y, z) e^{ik_y y} dy, \quad (10)$$

where k_y is the wavenumber. A Fourier transformation of the primary field is carried out numerically, following the approach of Streich et al. (2011). The governing equations of the 2D forward response emerge by applying the Fourier transform to equations 8 and 9

$$\tilde{E}_{sx} = \frac{1}{k_y^2 - \omega^2 \mu_z \epsilon_x} \left(ik_y \frac{\partial \tilde{E}_{sy}}{\partial x} - i\omega \mu_z \frac{\partial \tilde{H}_{sy}}{\partial z} - i\omega \mu_z \sigma_{sx} \tilde{E}_{px} \right), \quad (11)$$

$$\tilde{E}_{sz} = \frac{1}{k_y^2 - \omega^2 \mu_x \epsilon_z} \left(ik_y \frac{\partial \tilde{E}_{sy}}{\partial z} + i\omega \mu_x \frac{\partial \tilde{H}_{sy}}{\partial x} - i\omega \mu_x \sigma_{sz} \tilde{E}_{pz} \right), \quad (12)$$

$$\tilde{H}_{sx} = \frac{1}{k_y^2 - \omega^2 \mu_x \epsilon_z} \left(i\omega \epsilon_z \frac{\partial \tilde{E}_{sy}}{\partial z} + ik_y \frac{\partial \tilde{H}_{sy}}{\partial x} - ik_y \sigma_{sz} \tilde{E}_{pz} \right), \quad (13)$$

$$\tilde{H}_{sz} = \frac{1}{k_y^2 - \omega^2 \mu_z \epsilon_x} \left(-i\omega \epsilon_x \frac{\partial \tilde{E}_{sy}}{\partial x} + ik_y \frac{\partial \tilde{H}_{sy}}{\partial z} + ik_y \sigma_{sx} \tilde{E}_{px} \right), \quad (14)$$

$$\begin{aligned} & -i\omega \epsilon_y \tilde{E}_{sy} + \frac{\partial}{\partial x} \left(\frac{i\omega \epsilon_x \partial \tilde{E}_{sy}}{c_{zx}} \right) + \frac{\partial}{\partial z} \left(\frac{i\omega \epsilon_z \partial \tilde{E}_{sy}}{c_{xz}} \right) \\ & - \frac{\partial}{\partial x} \left(\frac{ik_y \partial \tilde{H}_{sy}}{c_{zx}} \right) + \frac{\partial}{\partial z} \left(\frac{ik_y \partial \tilde{H}_{sy}}{c_{xz}} \right) \\ & = -(i\omega \epsilon_{py} - i\omega \epsilon_y) \tilde{E}_{py} + \frac{\partial}{\partial x} \left(\left(\frac{i\omega \epsilon_{px}}{c_{pzx}} - \frac{i\omega \epsilon_x}{c_{zx}} \right) \frac{\partial \tilde{E}_{py}}{\partial x} \right) \\ & + \frac{\partial}{\partial z} \left(\left(\frac{i\omega \epsilon_{pz}}{c_{pzx}} - \frac{i\omega \epsilon_z}{c_{xz}} \right) \frac{\partial \tilde{E}_{py}}{\partial z} \right) - \frac{\partial}{\partial x} \left(\left(\frac{ik_y}{c_{pzx}} - \frac{ik_y}{c_{zx}} \right) \frac{\partial \tilde{H}_{py}}{\partial z} \right) \\ & + \frac{\partial}{\partial z} \left(\left(\frac{ik_y}{c_{pzx}} - \frac{ik_y}{c_{xz}} \right) \frac{\partial \tilde{H}_{py}}{\partial x} \right), \end{aligned} \quad (15)$$

and

$$\begin{aligned} & -i\omega \mu_y \tilde{H}_{sy} + \frac{\partial}{\partial x} \left(\frac{i\omega \mu_x \partial \tilde{H}_{sy}}{c_{xz}} \right) + \frac{\partial}{\partial z} \left(\frac{i\omega \mu_z \partial \tilde{H}_{sy}}{c_{zx}} \right) \\ & + \frac{\partial}{\partial x} \left(\frac{ik_y \partial \tilde{E}_{sy}}{c_{xz}} \right) - \frac{\partial}{\partial z} \left(\frac{ik_y \partial \tilde{E}_{sy}}{c_{zx}} \right) \\ & = -(i\omega \mu_{py} - i\omega \mu_y) \tilde{H}_{py} + \frac{\partial}{\partial x} \left(\left(\frac{i\omega \mu_{px}}{c_{pzx}} - \frac{i\omega \mu_x}{c_{zx}} \right) \frac{\partial \tilde{H}_{py}}{\partial x} \right) \\ & + \frac{\partial}{\partial z} \left(\left(\frac{i\omega \mu_{pz}}{c_{pzx}} - \frac{i\omega \mu_z}{c_{xz}} \right) \frac{\partial \tilde{H}_{py}}{\partial z} \right) + \frac{\partial}{\partial x} \left(\left(\frac{ik_y}{c_{pzx}} - \frac{ik_y}{c_{zx}} \right) \frac{\partial \tilde{E}_{py}}{\partial z} \right) \\ & - \frac{\partial}{\partial z} \left(\left(\frac{ik_y}{c_{pzx}} - \frac{ik_y}{c_{xz}} \right) \frac{\partial \tilde{E}_{py}}{\partial x} \right), \end{aligned} \quad (16)$$

where $c_{ij} = k_y^2 - \omega^2 \mu_i \epsilon_j$. Equations 11–14 only need to be evaluated at the receiver positions, which in our application are in the air. Thus, σ_{sx} , σ_{sy} , and σ_{sz} are always zero, so all primary field terms in equations 11–14 can be ignored. Although the left sides of equations 15 and 16 are identical to the governing equations presented in Mitsuhashi (2000), the right sides are expressed by $\tilde{E}_{py}$ and $\tilde{H}_{py}$ instead of $\tilde{E}_{px}$, $\tilde{E}_{py}$, and $\tilde{E}_{pz}$. However, a simple arithmetic reformulation of the right side of equation 15 using equations 11–14 can show that both forms are equivalent (not shown). Having the same field components and derivatives on both sides of the equations allows us to speed up the assembly of the linear equation system.

With the governing equations defined in equations 11–16, we use the standard finite-element approach to define a set of local equations for each element. By combining all these using the Galerkin method (Zienkiewicz et al., 1977) with second-order nodal elements and Dirichlet boundary conditions, a global set of linear equations is found for the secondary EM fields (see Appendix A for more details). From equations 15 and 16, a linear system of equations for the secondary field is found:

$$\mathbf{A} \tilde{\mathbf{x}} = \mathbf{b}, \quad (17)$$

where $\mathbf{A}$ is the global symmetric stiffness matrix, $\tilde{\mathbf{x}}$ contains the Fourier transformed secondary fields $\tilde{E}_{sy}$ and $\tilde{H}_{sy}$ at the mesh nodes, and $\mathbf{b}$ contains the source terms, where each column represents one source component.

The procedure to solve the system is as follows:

- Assemble the matrix equation resulting from equation 17.
- Solve the linear system using a direct lower-upper (LU)-decomposition solver and interpolate the field values at the receiver positions using the same second-order shape function (which has already been used to assemble the finite-element system) to find $\tilde{\mathbf{x}}$, which contains the Fourier-transformed secondary field components $\tilde{E}_{sy}$ and $\tilde{H}_{sy}$.
- Insert the solution into equations 11–14 to find the remaining components of the Fourier transformed fields of $\tilde{E}_{sx}$, $\tilde{E}_{sz}$, $\tilde{H}_{sx}$, and $\tilde{H}_{sz}$.
- Interpolate the solution to the receiver positions.
- Apply the inverse Fourier transform to obtain the fields E_{sx} , E_{sy} , E_{sz} , H_{sx} , H_{sy} , and H_{sz} at the receiver positions in the frequency domain, which, when all combined, are referred to as the forward response vector $\mathbf{d}$.

Given this procedure, the forward response can formally be written as

$$\mathbf{d} = \mathcal{F}^{-1}(\mathfrak{I}(\tilde{\mathbf{x}})), \quad (18)$$

where $\mathcal{F}^{-1}$ is the inverse Fourier transform operator and $\mathfrak{I}$ is the interpolating operator.

This inverse Fourier transform is done numerically by logarithmic spaced k_y samples, which are splined together over the relevant k_y domain. Tests show that five wavenumbers per decade between 10^{-5} and 10 m^{-1} provide sufficiently accurate results. One important point related to the inverse Fourier transform is that the air conductivity needs to be larger than zero; otherwise, a singularity at $k_y^2 \approx \omega^2 \mu \epsilon$ is encountered (Mitsuhata, 2000). We found that setting the air conductivity to $\sigma = 10^{-6} \text{ Sm}^{-1}$ keeps the air sufficiently resistive, while avoiding the singularity within a frequency range of 0.4–130 kHz. The interpolation to the receiver positions is carried out using the shape functions of the finite elements.

Derivative calculation

The 2D derivatives of the forward response with regards to the model parameters $\mathbf{m}$ can be written as




Figure 1. Inversion grid and forward mesh for a small section of an HEM survey. Transmitters are indicated by a brown circle; receivers are indicated by a blue circle. (a) The inversion grid. In the inversion grid, the column width is determined by the sounding distance, whereas the column depth reflects the thickness of the model layers. (b) The resulting forward modeling finite-element mesh. Note how the inversion grid is still present in the forward modeling mesh because this is used as the skeletal structure for building the finite-element mesh.

$$\frac{d\mathbf{d}}{d\mathbf{m}} = \frac{d(\mathcal{F}^{-1}(\mathfrak{I}(\tilde{\mathbf{x}})))}{d\mathbf{m}}. \quad (19)$$

For model parameters related to source/receiver altitude, the derivatives are calculated by a standard finite-difference approach with small perturbations, as done in the 1D case (Auken et al., 2014). For model parameters related to subsurface resistivities ρ ; the derivatives are calculated as follows.

For inversion parameters related to resistivities, the Fourier transform operator and the interpolation operator are independent of the inversion parameter, so we can write

$$\frac{d\mathbf{d}}{d\rho} = \mathcal{F}^{-1} \left(\mathfrak{I} \left(\frac{d\tilde{\mathbf{x}}}{d\rho} \right) \right). \quad (20)$$

In two dimensions, the derivative of $\tilde{\mathbf{x}}$ is found through the adjoint-state method (McGillivray and Oldenburg, 1990). Equation 17 is differentiated with regard to the model parameters $\mathbf{m}$:

$$\frac{d(\mathbf{A}\tilde{\mathbf{x}})}{d\rho} = \frac{d\mathbf{b}}{d\rho}, \quad (21)$$

which through the product rule gives the Jacobi elements:

$$\frac{d\tilde{\mathbf{x}}}{d\rho} = \mathbf{A}^{-1} \left(\frac{d\mathbf{b}}{d\rho} - \frac{d\mathbf{A}}{d\rho} \tilde{\mathbf{x}} \right). \quad (22)$$

Because only the coefficients of the governing equations are depended on the resistivity, $d\mathbf{A}/d\rho$ can be analytically calculated and assembled. The term $d\mathbf{b}/d\rho$ is zero because all sources are in the air, and thus they are not affected by any of the inverted resistivity cells. Because each inversion cell usually contains only a small number of finite elements, $d\mathbf{A}/d\rho$ is extremely sparse and $d\mathbf{A}/d\rho \tilde{\mathbf{x}}$ can be calculated efficiently, and the result can still be considered sparse. To calculate $\mathbf{A}^{-1}$, however, would be far too expensive. Instead, we use the fact that only the field derivatives at the receiver position are of interest, and thus replacing $\mathbf{A}^{-1}$ with λ^T , where λ^T contains all rows of $\mathbf{A}^{-1}$ that correspond to those column in $d\tilde{\mathbf{x}}/d\rho$, which are necessary to calculate the field values at the receiver positions

$$\frac{d\tilde{\mathbf{x}}_{\text{rec}}}{d\rho} = \lambda^T \left(\frac{d\mathbf{b}}{d\rho} - \frac{d\mathbf{A}}{d\rho} \tilde{\mathbf{x}} \right). \quad (23)$$

Because $\mathbf{A}$ is symmetric, λ can be calculated by solving $\mathbf{A}\mathbf{l} = \mathbf{I}_{\text{rec}}$, with $\mathbf{I}_{\text{rec}}$ being constructed from those columns of the identity matrix necessary to calculate the field values at the receivers. Thus, the same direct LU-decomposition used for the regular forward solutions can be used here. The matrix multiplication in equation 21 then results in $d\tilde{\mathbf{x}}_{\text{rec}}/d\rho$, which is a dense matrix, however, with rather small dimensions, where the number of rows is equal to the number of sources, and the number of columns is equal to 12 times the number of receivers (six nodes per finite element with two field components each). The derivatives at the receiver positions are then calculated from $d\tilde{\mathbf{x}}/d\rho$ using the second-order shape function as interpolator, and the derivative of the forward response are calculated by the interpolated derivatives given in equation 20.

Meshing

The 2D modeling is performed on a triangular finite-element mesh, as shown in Figure 1b, whereas the inversion operates on

a regular grid, as seen in Figure 1a. The column spacing of the inversion grid is determined by the sounding distance, the row spacing by the model layers, and the layer thickness is chosen to be logarithmically increasing, which reflects the decreasing sensitivity of HEM systems with depth. Separating the meshes of forward calculations and inversion has some clear computational benefits because the inversion grid is much coarser than the 2D forward mesh. This decreases the size of the inversion problem, while maintaining the accuracy of the forward modeling.

Interpolation between the inversion grid and the forward mesh is avoided by using the inversion grid as a skeletal structure for the forward mesh (Figure 1). By incorporating the inversion grid into the forward mesh, it is guaranteed that each finite element is fully residing in just one inversion cell, which aligns the forward modeling mesh nodes and edges with the inversion grid.

The mesh density for the forward mesh is adjusted according to the key parameters, such as the frequency and the source/receiver height. The highest mesh density is needed near the surface below the sources, where the primary field is the strongest. Here, the mesh density is selected as a function of the source height and, thus, of the strength of the primary field. For very low altitudes of 1 m and below, a maximum edge length of 0.2 m is required. For altitudes of 20 m and above, a 5 m edge length is sufficient. At deeper locations and at larger horizontal offsets from the source, the mesh density can be reduced without losing accuracy. The mesh density is interpolated between the surface/zero offset mesh density, defined by the altitude, and a background mesh density of 50 m for larger depths and offsets. This interpolation is done by a 2D Gaussian function, with the 2D distance from the closest source/receiver as a parameter and a frequency-dependent standard deviation. Standard deviations in the z -direction are logarithmically interpolated between 200 at 10^{-2} and 20 m at 10^6 Hz. Tests showed, that the mesh density along the surface could not be coarsened as quickly, so the standard deviation in the x -direction is logarithmically interpolated between 800 at 10^{-2} and 80 m at 10^6 Hz. Additionally, the mesh density is increased near the receivers, to calculate the spatial derivatives in equations 11 and 12 accurately.

The mesh of the 2D forward model is appended with large absorbing boundary domains that extend 10 km in each direction. Because the mesh density is coarsening quickly in these boundary domains, the computational overhead is not very high, but tests showed that 10 km boundary domains allow the field to attenuate enough, so that the validity of the Dirichlet boundary conditions is assured.

Sectioning

Even for relatively small surveys, it is computationally inefficient to create and store a sufficiently fine finite-element mesh and do 2D calculations on all soundings at once. Because of this, it is imperative to split large surveys into smaller sections. Sectioning or local meshing, as it is also often called in the literature, can be accomplished in several different ways. Our sectioning method is somewhat similar to the method used in Yang and Oldenburg (2012a). Their method involves a global mesh, and a local mesh for each sounding. Although the forward problem is handled on the local meshes, the inversion is made by subsampling the global mesh. In our case, the local

meshes contain multiple soundings. Because that is more efficient, and it is used for the 2D forward and derivative calculations. In order for these sections to retain the 2D information of the survey, they need to overlap, as shown in Figure 2. Thus, each section L consists of a core section l and one or two overlapping regions Δl . Continuity between different sections is ensured by using sufficient overlap between different sections and by placing lateral constraints on all soundings regardless of the section boundaries. The size of the overlap and sections will be addressed later.

Forward modeling validation

The 2D finite-element forward response was validated against the 1D code of Auken et al. (2014) for a range of frequencies relevant for HEM (0.4–130 kHz) across different half-spaces with resistivities of 10, 100, and 1000 Ωm . A half-space comparison between 1D and 2D is shown in Figure 3. The mesh density is selected such that the resulting responses deviate less than 5% from the 1D responses within the frequency range. The inaccuracy of the 1D response is estimated to be between 0.1% and 0.3% and is insignificant in this context. Note that the deviation between 1D and 2D responses is not just a single number, but it is instead a range because of variations in the mesh density between soundings near the edge of the mesh and those near the center. The overall coarseness surrounding a sounding near the center of a section is will be slightly lower than the overall coarseness surrounding a sounding near the edge of a section. This is reflected in Figure 3b, in which the deviation from 1D is shown for a 300 m long section with 30 soundings equally spaced over the section. The flight height of the system is 30 m and the half-space resistivity is set to 100 Ωm . Similar accuracies are obtained from half-space resistivities at 10 and 1000 Ωm , but for brevity, we only show one representative example. In this case, the inaccuracy is generally less than 2%, while reaching as high as 5% for frequencies beyond the range shown here.

Inversion algorithm

Our inversion technique uses linearized minimization, following the Levenberg-Marquardt adaptive scheme (Menke, 1989). The following is a brief review of our inversion algorithm (for full details, see Auken and Christiansen, 2004; Auken et al., 2014)

The minimized objective function is given as

$$q = q_{\text{obs}} + q_{\text{prior}} + q_{\text{reg}}, \quad (24)$$




Figure 2. Sectioning of survey lines. Each section L , consists of a core section l , and overlap regions Δl . Sections at the end of a survey line only have one overlap region, whereas all other sections have two overlap regions. The dots on top of the ground surface indicate the individual soundings.

with q_{obs} being the observed data (secondary field) misfit, q_{prior} being the prior constraint misfit, and q_{reg} being the regularization misfit. Smooth regularizations are used laterally and vertically. To determine the misfit, we use a standard least-squares solution (L2-norm). With this, the n th iterative update of the model vector $\mathbf{m}$ is given as

$$\mathbf{m}_{n+1} = \mathbf{m}_n + (\hat{\mathbf{G}}_n^T \hat{\mathbf{C}}_n^{-1} \hat{\mathbf{G}}_n + \lambda_n \mathbf{I})^{-1} \cdot (\hat{\mathbf{G}}_n^T \hat{\mathbf{C}}_n^{-1} \delta \hat{\mathbf{d}}_n), \quad (25)$$

where $\mathbf{I}$ is the identity matrix, λ is the damping parameter (Marquardt, 1963), $\delta \hat{\mathbf{d}}$ is the extended perturbed data vector, $\hat{\mathbf{G}}$ is the extended Jacobian, and $\hat{\mathbf{C}}$ is the extended covariance matrix, where the extensions comes from the inclusion of prior information and regularization:

$$\delta \hat{\mathbf{d}} = \begin{bmatrix} \mathbf{d} - \mathbf{d}_{\text{obs}} \\ \mathbf{m} - \mathbf{m}_{\text{prior}} \\ -\mathbf{R}\mathbf{m} \end{bmatrix}, \quad (26)$$

$$\hat{\mathbf{G}} = \begin{bmatrix} \mathbf{G} \\ \mathbf{P} \\ \mathbf{R} \end{bmatrix}, \quad (27)$$

$$\hat{\mathbf{C}} = \begin{bmatrix} \mathbf{C}_{\text{obs}} & 0 & 0 \\ 0 & \mathbf{C}_{\text{prior}} & 0 \\ 0 & 0 & \mathbf{C}_{\text{reg}} \end{bmatrix}, \quad (28)$$




Figure 3. The 1D and 2D forward responses and deviations, on a 100 Ωm half-space at an instrument altitude of 30 m, as a function of frequency. (a) An example of a 1D and 2D forward response for just a single sounding. (b) The relative deviations between a 2D and a 1D forward response. The deviation is expressed as a range because the accuracy of the 2D responses varies between soundings near the edge and near the center due to mesh variations. The 2D responses are from a 300 m section with 30 equally spaced soundings. Deviations at all frequencies and all positions are less than 2%.

where $\mathbf{d}$ is the forward response (see above), $\mathbf{d}_{\text{obs}}$ is the observed data, $\mathbf{m}$ is the model parameters, $\mathbf{m}_{\text{prior}}$ is the a priori model parameters, $\mathbf{R}$ is the roughness matrix, which binds neighboring models/model parameters together, $\mathbf{G}$ is the Jacobian (see above), $\mathbf{P}$ is a matrix containing the a priori information, $\mathbf{C}_{\text{obs}}$ is the covariance of the observed data, $\mathbf{C}_{\text{prior}}$ is the covariance of the a priori information, and $\mathbf{C}_{\text{reg}}$ is the covariance stemming from the roughness matrix.

Calculating the iterative model update as shown in equation 25, requires solving a large linear system. In 1D, this system is sparse, but in 2D, the linear system is, in principle, dense. However, in practice, it can be assumed to be sparse if only the part with the most sensitivity is considered (this will be covered in more detail later). Nevertheless, the 2D linear system will always be considerably less sparse than the 1D case. Solving large sparse linear systems is nontrivial, and the optimal approach is dependent on the system being solved. Our current approach to solving the linear system in the 1D case is thoroughly described in Kirkegaard and Auken (2015), and it starts with a reverse Cuthill and McKee (1969) reordering algorithm, which is used on the ordering of the initial soundings. This results in the matrix being created in such a way that all vital nonzero elements lie relatively close to the diagonal, while retaining the sounding structure in the matrix. The actual matrix is solved in parallel using an iterative sparse solver, which uses CG propagation (Hestenes and Stiefel, 1952; Saad, 2003), along with a preconditioner, which depends on the dimensionality of the inversion problem. For the 1D inversion problem, our method of choice is a block-parallelized version of an incomplete LU factorization with a dual-dropping strategy (Saad, 1994). However, due to the increased bandwidth of the sparse matrix in the 2D case, consistent convergence is not obtained when applying the LU decomposition as a preconditioner. Direct solvers work well for small surveys (up to approximately 5000 soundings), but for larger surveys, direct solvers become inefficient due to memory consumption as well as factorization time. Neither one scales linearly with the size of the survey. Instead, we have found that applying the symmetric-Gauss-Seidel (SGS) preconditioner leads to stable convergence when doing 2D inversions. Furthermore, if applied in cases in which the linear system is sufficiently diagonally dominant, the SGS preconditioner is even more efficient than the incomplete LU factorization, and this can lead to a significant speedup in 1D inversions.

Optimizing section sizes

Previously, we discussed the need for dividing large surveys into smaller sections due to memory concerns. However, even if memory had not been an issue, it still proves computationally advantageous to split a survey into smaller sections. The reason for this is that for unstructured meshes, the computational time scales quadratically with the number of elements as the number of elements becomes large. On the other hand, there is also a size-independent initialization cost associated with each section that needs to be considered. This cost comes from setting up the mesh padding, establishing the equations, spawning the parallel thread pool, and other similar tasks. Even more importantly, it also needs to be assured that each section overlaps its neighboring section by a fixed amount, to retain the 2D information from the survey.

Analysis of our parallel algorithm has led to the identification of an optimal section size. One that is defined by the initial computational cost is the quadratic computational scaling with section size, as well as the overlap size. To find this optimal section size, performance tests for the RESOLVE system (shown in Table 1) were conducted. The experi-

ments were performed over a sweep of section sizes, and the results for forward and derivative computations can be seen in Figure 4a.

As seen in Figure 4a, the computational times present a global minimum, different for forward and derivative calculations. The reason why the derivative calculation favors smaller section sizes over the forward calculation is due to the heavier computational burden associated with derivative calculations. This increased computational burden makes the initialization cost less significant and thus naturally shifts the optimal section size for derivative calculation toward smaller sizes. With the results presented in Figure 4a, the optimal section size as a function of the overlap is determined, as shown in Figure 4b. The optimal section size is determined by using the data in Figure 4a. By subtracting two times, the desired overlap from the section sizes given in Figure 4a, and interpolating the remaining positive core section sizes, an estimate of the computational time for a given section size with a given overlap region can be determined (not shown). To find the optimal section size, the section sizes that fall within 5% of the fastest time for a given overlap are used and shown in Figure 4b. Note that, once again, the optimal section sizes are different for forward and derivative calculations. One caveat to this is that the results in Figure 4b change depending on the 2D finite-element mesh density, which changes slightly between different surveys and systems. Therefore, Figure 4b should not be considered the absolute truth, but it should rather serve as a guide for picking a sensible section size based on overlap size.

2D Jacobian and sensitivity analysis

As mentioned earlier, the structure of the 2D Jacobian is a dense matrix. However, due to the decay in sensitivity as a function of distance, a threshold can be defined, where anything that falls below this threshold is assumed negligible. Thus, in practice, our Jacobian matrix can still be considered sparse even in 2D, even though it has considerably wider nonzero bands around the diagonal than in the 1D case. Our 2D Jacobian matrix resulting from the inversion grid is shown in equations 29 and 30. Where equation 29 shows our 2D Jacobian in block form, where each column/row refers to a single sounding. Note that the 1D Jacobian structure is identical to the one for the 2D Jacobian, but in the 1D case, all the off-diagonal blocks shown in equation 29 would be zero. Equation 30 shows the structure of a single Jacobian block element, which contains several elements equal to the number of perturbable model parameters for this particular model (altitude and resistivities) times the number of data points for the corresponding sounding:

$$\mathbf{G} = \begin{pmatrix} \ddots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \\ \cdots & \frac{\partial \mathbf{D}_{i-2}}{\partial \mathbf{M}_{i-2}} & \frac{\partial \mathbf{D}_{i-2}}{\partial \mathbf{M}_{i-1}} & 0 & 0 & 0 & \cdots \\ \cdots & \frac{\partial \mathbf{D}_{i-1}}{\partial \mathbf{M}_{i-2}} & \frac{\partial \mathbf{D}_{i-1}}{\partial \mathbf{M}_{i-1}} & \frac{\partial \mathbf{D}_{i-1}}{\partial \mathbf{M}_i} & 0 & 0 & \cdots \\ \cdots & 0 & \frac{\partial \mathbf{D}_i}{\partial \mathbf{M}_{i-1}} & \frac{\partial \mathbf{D}_i}{\partial \mathbf{M}_i} & \frac{\partial \mathbf{D}_i}{\partial \mathbf{M}_{i+1}} & 0 & \cdots \\ \cdots & 0 & 0 & \frac{\partial \mathbf{D}_{i+1}}{\partial \mathbf{M}_i} & \frac{\partial \mathbf{D}_{i+1}}{\partial \mathbf{M}_{i+1}} & \frac{\partial \mathbf{D}_{i+1}}{\partial \mathbf{M}_{i+2}} & \cdots \\ \cdots & 0 & 0 & 0 & \frac{\partial \mathbf{D}_{i+2}}{\partial \mathbf{M}_{i+1}} & \frac{\partial \mathbf{D}_{i+2}}{\partial \mathbf{M}_{i+2}} & \cdots \\ \ddots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}. \quad (29)$$

This illustrates our 2D Jacobian in block-matrix form, where the number of off-diagonal bands is equal to the number of surrounding soundings above the sensitivity threshold (here, only the nearest neighbor is above

the threshold). Each entry in the Jacobian block matrix is a dense matrix block, which is given as

$$\frac{\partial \mathbf{D}_j}{\partial \mathbf{M}_k} = \begin{pmatrix} \frac{\partial d_{j,1}}{\partial m_{k,1}} & \frac{\partial d_{j,1}}{\partial m_{k,2}} & \frac{\partial d_{j,1}}{\partial m_{k,3}} & \cdots & \frac{\partial d_{j,1}}{\partial m_{k,N_{m_k}}} \\ \frac{\partial d_{j,2}}{\partial m_{k,1}} & \frac{\partial d_{j,2}}{\partial m_{k,2}} & \frac{\partial d_{j,2}}{\partial m_{k,3}} & \cdots & \frac{\partial d_{j,2}}{\partial m_{k,N_{m_k}}} \\ \frac{\partial d_{j,3}}{\partial m_{k,1}} & \frac{\partial d_{j,3}}{\partial m_{k,2}} & \frac{\partial d_{j,3}}{\partial m_{k,3}} & \cdots & \frac{\partial d_{j,3}}{\partial m_{k,N_{m_k}}} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{\partial d_{j,N_{d_j}}}{\partial m_{k,1}} & \frac{\partial d_{j,N_{d_j}}}{\partial m_{k,2}} & \frac{\partial d_{j,N_{d_j}}}{\partial m_{k,3}} & \cdots & \frac{\partial d_{j,N_{d_j}}}{\partial m_{k,N_{m_k}}} \end{pmatrix}, \quad (30)$$

Table 1. Acquisition parameters mimicking a RESOLVE system used in our computational cost analysis simulation. The flight altitude is 30 m, and the uncertainty on the data is 5%.

Coil #	Orientation	Frequency (kHz)	Separation (m)
1	Z	0.395	7.9
2	Z	1.822	7.9
3	X	5.4	9.06
4	Z	8.199	7.9
5	Z	38.76	7.9
6	Z	128.76	7.9




Figure 4. (a) Time required for an iteration of forward and derivative calculations for the RESOLVE system presented in Table 1 using a 10 m sounding separation, as a function of section size. For small section sizes, the initialization cost becomes dominant and visible as a sharp increase in iteration time per meter, whereas for large section sizes, the quadratic scaling of the finite elements becomes dominant, and it is visible as a linear rise in iteration time per meter (it appears linear in the figure because time is normalized with section size; hence, time is given in seconds per meter). The range bars indicate the variability in computational time between various repetitions, and indicate one standard deviation. (b) Optimal section sizes as a function of the overlap, where the range bars represent section sizes that are within 5% of the optimal compute time.

where j and k represent the individual sounding indices, and $\mathbf{D}_j$ contains N_{dj} forward responses associated with the j th sounding. While $\mathbf{M}_k$ contains N_{mk} model parameters associated with sounding k .

Accurately determining the resulting sensitivity range is important, not just when building the Jacobian, but also when optimizing section size. This is due to the obvious connection between the sensitivity threshold distance and the required overlap distance between adjacent sections. To determine the sensitivity threshold, we follow the convention of Liu and Becker (1990) and define a significant sensitivity range as the distance at which 90% of the full sensitivity is contained. Following this approach, a sensitivity analysis was performed for the coils shown in Table 1, for altitudes ranging between 20 and 50 m. Figure 5a demonstrates the cumulated sensitivity as a function of distance for a 0.4 kHz signal originating at an altitude of 30 m, whereas Figure 5b and 5c shows the correlation between depth and footprint size for a 0.4 kHz signal and a 1.8 kHz signal. Based on the sensitivity analysis as well as performance concerns, we decide to use an overlap of 150 m.




Figure 5. Sensitivity analysis for a frequency-domain system at an altitude of 30 m. (a) Cumulated sensitivity as a function of distance for a 0.4 kHz signal, on a 100 Ωm half-space at a depth of 50–60 m, with the horizontal dashed line indicating the footprint size at a 90% threshold. The two circles indicate the distances where the 0.4 kHz signal reaches the threshold for the real and imaginary response. (b, c) Correlation between footprint size and depth for a 0.4 kHz signal and a 1.8 kHz signal, on a 30, 100, and 300 Ωm half-space. In (b), the two circles from (a) are also plotted.

Although this is less than the footprint size for the real part of the 0.4 kHz signal, more than 75% of the sensitivity for a 100 Ωm half-space is retained, and if considering total sensitivity over all frequencies, then, the total loss of sensitivity is approximately 7%, which we deem an acceptable loss. Based on Figure 4b, we use a section size of 750 m for forward calculations and 550 m for derivative calculations.

1D/2D Hybrid inversion scheme

The conceptual idea behind the hybrid inversion scheme is to use computationally inexpensive approximate forward and derivative computations in the first inversion steps, where accuracy is of little importance. As the iterations start to converge, one can then gradually switch to higher accuracy computations that are more expensive. Within such a scheme, the overall computational time can be greatly reduced without sacrificing the quality of the final model. Such a scheme can be constructed in several ways: Christiansen et al. (2015) create a hybrid scheme using increasingly accurate 1D modeling responses, and a similar approach could be envisioned in 2D by using a coarse mesh in the early iterations and a more refined mesh in the later stages, as is done by Haber et al. (2007a). However, we believe that our hybrid scheme is computationally superior to such a scheme because 1D modeling is so computationally inexpensive compared with 2D modeling, and the number of full 2D iterations used in our scheme is quite low as will be demonstrated later. Our hybrid scheme is a three-stage scheme with

- 1) 1D forward and derivative calculations
- 2) 2D forward, 1D derivative calculations
- 3) 2D forward and 2D derivative calculations.

Each stage is executed with a fixed number of iterations. By running the hybrid scheme over a large number of synthetic models, we have empirically found that the optimal number of iterations is four in the first stage and eight in the second stage. The third stage runs until the algorithm converges.

The inversion is said to have converged if the relative misfit change is less than 1% between two iteration steps. If convergence is reached in stage 1 or 2, then the inversion is advanced to the next stage and the process continues.

RESULTS AND DISCUSSION

Synthetic model

The 2D hybrid inversion algorithm is demonstrated on two synthetic models. A system resembling the RESOLVE system with the parameters shown in Table 1 is modeled. In both examples, the inversion is started from a 100 Ωm half-space with a model discretization of 20 layers, where the thickness of each layer increases logarithmically from 3 to 10 m. Horizontal smoothing constraints are used with a covariance factor of 1.6 and vertical smoothing with a factor of 3.0. There are 51 equidistant soundings distributed more than a 500 m long line. All data have a uniform 5% uncertainty.

Figure 6 shows the results of an inversion of a conductive lens. Figure 6a illustrates the true model, which consists of a 50 Ωm lens in a 500 Ωm half-space, Figure 6b shows a 1D inversion, Figure 6c shows a hybrid inversion, and Figure 6d shows a full 2D inversion. The 1D inversion mostly manages to recover the conductive lens at

the correct depth, but strong pant legs are produced. The 1D inversion is shown with a 1D and 2D residual curve. Both residual curves use the model arrived at through the 1D inversion, but the 1D residual evaluates the 1D forward responses, whereas the 2D residual is relative to 2D forward responses. Variations between the two residuals can therefore be regarded as an indicator of areas where 1D modeling is insufficient. The hybrid and the full 2D inversion reproduce the lens as good as can be expected from an AEM measurement, without any pant-legs effect and with a good determination of the lens boundaries and a misfit well less than 1, which in our synthetic model without noise is a good thing. The speedup gained by using the hybrid inversion was $2.7\times$ compared with the 2D inversion, and it will be discussed in detail in the “Performance” subsection.

Figure 7 shows the results of an inversion of a sharp horizontal conductivity contrast. Figure 7a illustrates the true model, where the left side is $10\ \Omega\text{m}$ and the right side is $200\ \Omega\text{m}$; Figure 7b shows a 1D inversion; Figure 7c shows a hybrid inversion; and Figure 7d shows a full 2D inversion. In this case, the 1D inversion creates a rather wide region around the conductivity contrast where the conductivities are smeared and there are clearly visible pant legs. Once again, the 1D and 2D residuals are shown. The full 2D inversion demonstrates a better determination of the vertical boundary, and although smearing is still observed, the affected region is significantly smaller. The hybrid model again converges to a model that is significantly better than the 1D model, as noticed by the size of the smearing region as well as the residual, which is only slightly higher than for the 2D inversions. The differences between the 2D and the hybrid model are likely a result of the 1D model doing a poor job of accurately modeling the sharp conductivity contrast, combined with model equivalences, as evidenced the similarity of the hybrid/2D residual. The speedup gained by using the hybrid inversion was $6.6\times$ compared with the 2D inversion, and will be discussed in detail in the “Performance” subsection.

Field example

As a final test, the 2D algorithm is used on a field example collected by a RESOLVE system owned by the German Bundesanstalt für Geowissenschaften und Rohstoffe (BGR). The field data were collected on a small island named Langeoog, where the target is a mapping of the freshwater/saltwater boundary (Siemon et al., 2015). Langeoog comprises three geologic features: The base is formed from glaciofluvial sediments from the Pleistocene age. These sediments contain Lauenburg clay, which lies at a depth of 15–35 m below sea level, with a typical thickness of a few meters. Overlaying the Pleistocene layer is a Holocene marine deposit consisting of primarily silt, which lies at 10–20 m depth below sea level. The top layer consists of dunes and beach sand. For more information about the geology of Langeoog, see Costabel et al. (2017).

The field data profile is 1400 m long and consists of 144 soundings. Inversion results are shown in Figure 8. Figure 9a shows a 1D inver-

sion, Figure 9b shows a hybrid inversion, and Figure 9c shows a 2D inversion. Overall, the three different inversions show very consistent results, although there are notable differences in the top layers of the soundings at a distance of approximately 1 km. Although the models deviate in this area, the data residuals for the 2D and hybrid code are only negligibly lower than for the 1D inversion. Upon a closer look at the fit of each individual transmitter frequency, it is revealed that there is excellent correspondence between the measured data and the modeled response for all coils except coil 3. Coil 3 is off by several standard deviations in the high residual area at a distance of 1 km. Figure 9 shows an example of this for sounding 112, which is marked in Figure 8 by a vertical red line. The speedup gained by using the hybrid inversion was $6\times$ compared with the 2D inversion, and this will be discussed in detail in the “Performance” subsection.

Parallelization and scalability

Since the introduction of commodity multicore CPUs in 2005, parallelization has become increasingly important. Although computational speed continues to grow exponentially, it has become a nontrivial issue to fully harness this power. Algorithms often have to be specifically tailored to enable optimal parallelization, and with the shift away from uniform memory access (UMA) systems and toward nonuniform memory access (NUMA) system architecture, this becomes an even harder problem. The architectures are illustrated in Figure 10. The consequences of having a NUMA system are that data placement becomes paramount. If data are not placed in




Figure 6. Inversion results for a synthetic conductive lens. (a) True model, which consists of a $50\ \Omega\text{m}$ lens in a $500\ \Omega\text{m}$ half-space. (b) 1D inversion, where the blue line is the 1D data residual and the red line is the 2D data residual. (c) Hybrid inversion, where the number of iterations are given in each hybrid stage separately. (d) 2D inversion. For all inversions, the number of iterations until convergence is written inside the residual box. Note that the data residuals are normalized by the standard deviation.




Figure 7. Inversion results for a horizontal conductivity contrast. (a) The true model: The left side has a resistivity of $10 \Omega\text{m}$, whereas the right side has a resistivity of $200 \Omega\text{m}$. (b) The result from a 1D inversion, where the blue line is the 1D data residual, and the red line is the 2D data residual. (c) The result from a hybrid inversion. (d) The result from a 2D inversion. The number of iterations until convergence is written inside the residual box. Note that the data residuals are normalized by the standard deviation.




Figure 8. Inversion results from a survey in northern Germany (Siemon et al., 2015). (a) A 1D inversion, where the blue line is the 1D data residual and the red line is the 2D data residual. (b) A hybrid inversion, where the number of iterations are provided for each separate stage. (c) A 2D inversion. The number of iterations until convergence is written inside the residual box. The vertical red line at the top of the figures approximately 1.1 km is sounding 112, which is shown in Figure 9. Note that the data residuals are normalized by the standard deviation.

the local memory associated with the processor working on it, it will need to be accessed over the interconnect by the processor. Not only does this add significant latency, but the interconnect also has limited bandwidth and becomes saturated much before the direct channels to local memory. For this reason, good scaling on the NUMA system is harder to achieve in general than on UMA systems, and if it is not done carefully, it can actually lead to decreased performance, unlike for UMA systems (Dong et al., 2010).

The 2D FEM problem can be parallelized in several ways, but to get the best possible scalability for large surveys, we chose to put our parallelization across the sections. In other words, multiple sections are computed in parallel. This requires more memory than putting the parallelization over the wavenumbers; however, with the sectioning used, the memory requirement during the 2D modeling is less than the 1 GB per thread used (not shown), and thus the total memory requirement is inconsequential on modern hardware. Although parallelization over the sections requires more memory than other approaches, it also gives the best scalability for large surveys because there is practically no intercommunication between the different threads. Section parallelization provides good large-scale scaling, but it does not provide much benefit for small-scale problems. To remedy this, an additional parallelization over the frequencies of each section was implemented using OpenMP's collapse directive. By parallelizing over the sectioning and frequency, good scaling can be achieved for surveys of all sizes. Figure 11 shows the parallel scalability of the code. It can be seen that the scaling is almost linear for low numbers of threads, whereas linear scalability is lost for higher numbers of threads due to memory bandwidth limitations.

Another key concept when doing parallelization is affinity. That is how the parallel threads are bound to the various cores in the system. If thread affinity is not used, it can severely affect performance, especially on NUMA systems. Without affinity, the calculations of a thread are never confined to a single core, but they are rather executed in small portions executed on random cores of the system. This can have dramatic consequences because memory locality cannot be assured and data kept in cache are constantly lost. Figure 11 demonstrates two different affinity schemes that are commonly used: compact affinity and scattered affinity. When using compact affinity, thread spawning tends to cluster together on a NUMA node until all processors on the NUMA are engaged, whereas scattered affinity tends to spread out the thread spawning across all NUMA nodes. The two different

affinity modes can have significantly different performance depending on the problem to which they are used. Because of the low level of intercommunication between the parallel threads, a practically identical scaling for the two affinities can be seen in Figure 11.

As a final comment to scaling, it should be mentioned that due to the way we do sectioning and inversion, our code has linear scaling in compute time as a function of survey size (not shown here).

Performance

Our 2D code is capable of inverting surveys of virtually any size, due to the scalability introduced by sectioning. Thus far, the code has been successfully tested on a 100 line km survey with 10,000 soundings. For such a survey, the code performs a full 2D forward and derivative calculation in approximately 5 h on a NUMA system with two Intel Xeon E5-2650 v3 CPUs, each with 10 cores.

The total inversion time for the 100 line km survey can be seen in Table 2. The hybrid and 2D inversion reach comparable misfits, but the hybrid scheme does so 2.3 times faster than the 2D inversion.




Figure 9. Example of a 2D sounding curve for the inversion results shown in Figure 8c. The sounding curve is marked in Figure 8 by a vertical red line at the top of the figures, approximately 1.1 km distance. The coil # configuration can be found in Table 1.




Figure 10. UMA and NUMA system architectures. P represents physical cores, whereas M represents memory. (a) An UMA system; here, there is only one memory bank, and all the cores have the same bandwidth connection to the memory. (b) A two-node NUMA system, which essentially consists of two connected UMA systems. Thus, each NUMA node has several cores attached to a memory bank, and the different memory banks are then connected together by a relatively small bandwidth.

Although the performance numbers presented are representative, it should be mentioned that the number of iterations needed to reach convergence can vary quite heavily between different surveys. Obviously, this also makes inversion times vary quite heavily. Roughly speaking, a pure 2D inversion can usually be done in approximately 10–20 iterations, whereas a hybrid inversion requires 14–20 iterations. Note that the number of iterations and hence inversion time depend heavily on the stopping criteria, which we have chosen to be a relative misfit change of less than 1%.

In the examples shown previously, the speedup gained from using the hybrid scheme was 2.7x for the synthetic conductive lens, 6.6x for the horizontal conductivity contrast, 6x for the small field example, and 2.3x for a 100 line km. These are all very significant speedups, and they are generally generated without notably worsening the resulting model. The reason why the speedups vary largely depends on the number of iterations spent in stage 3 of the hybrid scheme: If only a few iterations are spent, then the speedup is high in general, whereas if more than a few iterations are spent in stage 3, the speedup will generally be in the low end.




Figure 11. Parallel scaling of the code. The data are generated on a NUMA system with two Intel Xeon E5-2650 v3 CPUs, each with 10 cores. The threads are bound to specific logical processors, following either a compact affinity approach or a scattering affinity approach.

Table 2. Runtime and iteration number for an inversion of a 100 line km survey conducted with the RESOLVE system. For the hybrid system, the iteration numbers and runtimes are given for each of the three stages.

	100 line km	
	Iterations	Runtime (h)
1D	15	0.1 h
Hybrid	4 + 8 + 5	36 s + 4.9 h + 26.0 h = 31 h
2D	15	71 h

CONCLUSION

We have presented an algorithm for hybrid 2D frequency domain forward modeling and inversions. The 2D forward and derivative calculations are done on a triangular finite-element mesh using sectioning, whereas inversions are done on a regular grid. The finite-element mesh is created with the inversion grid as the foundation, which makes interpolation between the meshes significantly easier. By using sectioning and a regular grid for inversion, the code is able to handle large-scale inversions, which are otherwise often problematic for higher dimensional inversion codes. We have demonstrated how section sizes should be chosen to optimize computational times, and we have shown how forward and derivative calculations are optimally performed using different section sizes. Our parallelization goal was to achieve maximum speed; hence, the code is parallelized over the frequencies and the sections. This gives the code high efficiency for large and small surveys, as well as excellent scaling properties even on nonuniform memory architectures. Although the focus was on computational speed, the memory consumption is less than 1 GB per thread, and thus memory consumption for this algorithm was deemed inconsequential. Furthermore, we presented a hybrid 1D/2D scheme, which boosts the computational speed of 2D inversions by approximately 2 – 6x, without significantly reducing the accuracy. The concept of combining lower and higher dimensional algorithms in a hybrid scheme to significantly increase computational speed is a largely unused optimization within the scientific community. We have demonstrated our algorithm with two successful synthetic examples and a field example.

ACKNOWLEDGMENTS

This study was supported by Innovation fund Denmark, through the AirTech4Water project. We furthermore thank B. Siemon and BGR for the field data.

APPENDIX A

DERIVING THE SYSTEM OF LINEAR EQUATIONS

To recast the governing equations 15 and 16 into a system of linear equations, the finite-element method is used. In this regard, we substitute $\tilde{E}_{sy}$ and $\tilde{H}_{sy}$ by interpolated fields combined with second-order shape functions:

$$\begin{aligned}\tilde{E}_{sy}(x, z) &= \sum_{i=1}^n N_i(x, z) \tilde{E}_{sy,i}, \\ \tilde{H}_{sy}(x, z) &= \sum_{i=1}^n N_i(x, z) \tilde{H}_{sy,i},\end{aligned}\quad (\text{A-1})$$

where n is the number of nodes attached to one element. Replacing $\tilde{E}_{sy}$ and $\tilde{H}_{sy}$ in this way leads to an approximation of the Maxwell equations, which bear a residual. We use the weighted residual procedure to minimize the residual averaged over the area of each grid cell. Because our procedure applies to equations 15 and 16 in the same way, we will only focus on equation 15. The weighting function for the residual is the same as the interpolation function, i.e., $N_i(x, z)$, and integration of both sides of the equation and applying the rule for integration by parts combined with Gauss' theorem leads us to

$$\begin{aligned}& -i\omega\epsilon_y \int_{\Omega} N^T N d\Omega \tilde{E}_{sy} + \frac{i\omega\epsilon_x}{c_{zx}} \left[\oint_{\Gamma} N^T \frac{\partial N}{\partial x} d\Gamma + \int_{\Omega} \frac{\partial N^T}{\partial x} \frac{\partial N}{\partial x} d\Omega \right] \tilde{E}_{sy} \\& + i\omega\epsilon_z \left[\oint_{\Gamma} N^T \frac{\partial N}{\partial z} d\Gamma + \int_{\Omega} \frac{\partial N^T}{\partial z} \frac{\partial N}{\partial z} d\Omega \right] \tilde{E}_{sy} \\& - \frac{ik_y}{c_{zx}} \left[\oint_{\Gamma} N^T \frac{\partial N}{\partial x} d\Gamma + \int_{\Omega} \frac{\partial N^T}{\partial x} \frac{\partial N}{\partial z} d\Omega \right] \tilde{H}_{sy} \\& + \frac{ik_y}{c_{xz}} \left[\oint_{\Gamma} N^T \frac{\partial N}{\partial x} d\Gamma + \int_{\Omega} \frac{\partial N^T}{\partial z} \frac{\partial N}{\partial x} d\Omega \right] \tilde{H}_{sy} \\& = -(i\omega\epsilon_{py} - i\omega\epsilon_y) \int_{\Omega} N^T N d\Omega \tilde{E}_{py} \\& + \left(\frac{i\omega\epsilon_{px}}{c_{pzx}} - \frac{i\omega\epsilon_x}{c_{zx}} \right) \left[\oint_{\Gamma} N^T \frac{\partial N}{\partial x} d\Gamma + \int_{\Omega} \frac{\partial N^T}{\partial x} \frac{\partial N}{\partial x} d\Omega \right] \tilde{E}_{py} \\& + \left(\frac{i\omega\epsilon_{pz}}{c_{pxz}} - \frac{i\omega\epsilon_z}{c_{xz}} \right) \left[\oint_{\Gamma} N^T \frac{\partial N}{\partial z} d\Gamma + \int_{\Omega} \frac{\partial N^T}{\partial z} \frac{\partial N}{\partial z} d\Omega \right] \tilde{E}_{py} \\& - \left(\frac{ik_y}{c_{pzx}} - \frac{ik_y}{c_{zx}} \right) \left[\oint_{\Gamma} N^T \frac{\partial N}{\partial x} d\Gamma + \int_{\Omega} \frac{\partial N^T}{\partial x} \frac{\partial N}{\partial z} d\Omega \right] \tilde{H}_{py} \\& + \left(\frac{ik_y}{c_{pxz}} - \frac{ik_y}{c_{xz}} \right) \left[\oint_{\Gamma} N^T \frac{\partial N}{\partial x} d\Gamma + \int_{\Omega} \frac{\partial N^T}{\partial z} \frac{\partial N}{\partial x} d\Omega \right] \tilde{H}_{py}, \quad (\text{A-2})\end{aligned}$$

where Ω is the area of the respective element and Γ is the element's boundary. Within the model domain, the integrals of connected elements cancel each other out, and at the model-domain boundary, we assume Dirichlet boundary conditions, i.e., $\tilde{E}_{sy} = \tilde{H}_{sy} = 0$; hence, the boundary integrals can be ignored. This finally results in the following system of linear equations:

$$\begin{aligned}& -i\omega\epsilon_y \int_{\Omega} N^T N d\Omega \tilde{E}_{sy} + \frac{i\omega\epsilon_x}{c_{zx}} \int_{\Omega} \frac{\partial N^T}{\partial x} \frac{\partial N}{\partial x} d\Omega \tilde{E}_{sy} \\& + i\omega\epsilon_z \int_{\Omega} \frac{\partial N^T}{\partial z} \frac{\partial N}{\partial z} d\Omega \tilde{E}_{sy} - \frac{ik_y}{c_{zx}} \int_{\Omega} \frac{\partial N^T}{\partial x} \frac{\partial N}{\partial z} d\Omega \tilde{H}_{sy} \\& + \frac{ik_y}{c_{xz}} \int_{\Omega} \frac{\partial N^T}{\partial z} \frac{\partial N}{\partial x} d\Omega \tilde{H}_{sy} \\& = -(i\omega\epsilon_{py} - i\omega\epsilon_y) \int_{\Omega} N^T N d\Omega \tilde{E}_{py} \\& + \left(\frac{i\omega\epsilon_{px}}{c_{pzx}} - \frac{i\omega\epsilon_x}{c_{zx}} \right) \int_{\Omega} \frac{\partial N^T}{\partial x} \frac{\partial N}{\partial x} d\Omega \tilde{E}_{py} \\& + \left(\frac{i\omega\epsilon_{pz}}{c_{pxz}} - \frac{i\omega\epsilon_z}{c_{xz}} \right) \int_{\Omega} \frac{\partial N^T}{\partial z} \frac{\partial N}{\partial z} d\Omega \tilde{E}_{py} \\& - \left(\frac{ik_y}{c_{pzx}} - \frac{ik_y}{c_{zx}} \right) \int_{\Omega} \frac{\partial N^T}{\partial x} \frac{\partial N}{\partial z} d\Omega \tilde{H}_{py} \\& + \left(\frac{ik_y}{c_{pxz}} - \frac{ik_y}{c_{xz}} \right) \int_{\Omega} \frac{\partial N^T}{\partial z} \frac{\partial N}{\partial x} d\Omega \tilde{H}_{py}.\end{aligned}\quad (\text{A-3})$$

By creating and combining this system of equations and its counterpart resulting from equation 16, the desired linear system of equations is found:

$$\mathbf{A}\tilde{\mathbf{x}} = \mathbf{b}, \quad (\text{A-4})$$

where $\mathbf{A}$ is the global symmetric stiffness matrix, $\tilde{\mathbf{x}}$ contains the Fourier transformed EM fields, and $\mathbf{b}$ contains the source terms.

REFERENCES

- Abubakar, A., T. Habashy, V. Druskin, L. Knizhnerman, and D. Alumbaugh, 2008, 2.5 D forward and inverse modeling for interpreting low-frequency electromagnetic measurements: *Geophysics*, **73**, no. 4, F165–F177, doi: [10.1190/1.2937466](https://doi.org/10.1190/1.2937466).
- Auken, E., and A. V. Christiansen, 2004, Layered and laterally constrained 2D inversion of resistivity data: *Geophysics*, **69**, 752–761, doi: [10.1190/1.1759461](https://doi.org/10.1190/1.1759461).
- Auken, E., A. V. Christiansen, C. Kirkegaard, G. Fiandaca, C. Schamper, A. A. Behroozmand, A. Binley, E. Nielsen, F. Effersø, N. B. Christensen, K. Sørensen, N. Foged, and G. Vignoli, 2014, An overview of a highly versatile forward and stable inverse algorithm for airborne, ground-based and borehole electromagnetic and electric data: *Exploration Geophysics*, **46**, 223–235, doi: [10.1071/EG13097](https://doi.org/10.1071/EG13097).
- Beamish, D., 2003, Airborne EM footprints: *Geophysical Prospecting*, **51**, 49–60, doi: [10.1046/j.1365-2478.2003.00353.x](https://doi.org/10.1046/j.1365-2478.2003.00353.x).
- Caudillo-Mata, L. A., E. Haber, L. J. Heagy, and C. Schwarzbach, 2017, A framework for the upscaling of the electrical conductivity in the quasi-static Maxwell's equations: *Journal of Computational and Applied Mathematics*, **317**, 388–402, doi: [10.1016/j.cam.2016.11.051](https://doi.org/10.1016/j.cam.2016.11.051).
- Christiansen, A. V., and E. Auken, 2004, Optimizing a layered and laterally constrained 2D inversion of resistivity data using Broyden's update and 1D derivatives: *Journal of Applied Geophysics*, **56**, 247–261, doi: [10.1016/S0926-9851\(04\)00055-2](https://doi.org/10.1016/S0926-9851(04)00055-2).
- Christiansen, A. V., E. Auken, C. Kirkegaard, C. Schamper, and G. Vignoli, 2015, An efficient hybrid scheme for fast and accurate inversion of airborne transient electromagnetic data: *Exploration Geophysics*, **47**, 323–330, doi: [10.1071/EG14121](https://doi.org/10.1071/EG14121).
- Costabel, S., B. Siemon, G. Houben, and T. Günther, 2017, Geophysical investigation of a freshwater lens on the island of Langeoog, Germany: Insights from combined HEM, TEM and MRS data: *Journal of Applied Geophysics*, **136**, 231–245, doi: [10.1016/j.jappgeo.2016.11.007](https://doi.org/10.1016/j.jappgeo.2016.11.007).
- Cox, L. H., G. A. Wilson, and M. S. Zhdanov, 2010, 3D inversion of airborne electromagnetic data using a moving footprint: *Exploration Geophysics*, **41**, 250–259, doi: [10.1071/EG10003](https://doi.org/10.1071/EG10003).
- Cox, L. H., G. A. Wilson, and M. S. Zhdanov, 2012, 3D inversion of airborne electromagnetic data: *Geophysics*, **77**, no. 4, WB59–WB69, doi: [10.1190/geo2011-0370.1](https://doi.org/10.1190/geo2011-0370.1).
- Cuthill, E., and J. McKee, 1969, Reducing the bandwidth of sparse symmetric matrices: *Proceedings of the 24th National Conference*, 157–172.
- Dong, X., G. Cooperman, and J. Apostolakis, 2010, Multithreaded Geant4: Semi-automatic transformation into scalable thread-parallel software: *European Conference on Parallel Processing*, 287–303.
- Doyle, B. J., K. Kivi, and B. H. S. Smith, 1999, The TU Kwi Cho (D027 and D018) Diamondiferous Kimberlite Complex, Northwest Territories, Canada: Presented at the 7th International Kimberlite Conference, 194–204.
- Haber, E., S. Heldmann, and U. Ascher, 2007a, Adaptive finite volume method for distributed non-smooth parameter identification: *Inverse Problems*, **23**, 1659–1676, doi: [10.1088/0266-5611/23/4/017](https://doi.org/10.1088/0266-5611/23/4/017).
- Haber, E., D. W. Oldenburg, and R. Shekhtman, 2007b, Inversion of time domain three-dimensional electromagnetic data: *Geophysical Journal International*, **171**, 550–564, doi: [10.1111/j.1365-246X.2007.03365.x](https://doi.org/10.1111/j.1365-246X.2007.03365.x).
- Hestenes, M. R., and E. Stiefel, 1952, Methods of conjugate gradients for solving linear systems: *Journal of Research of the National Bureau of Standards*, **49**, 409–436, doi: [10.6028/jres.049.044](https://doi.org/10.6028/jres.049.044).
- Jahandari, H., S. Ansari, and C. G. Farquharson, 2017, Comparison between staggered grid finite-volume and edge-based finite-element modeling of geophysical electromagnetic data on unstructured grids: *Journal of Applied Geophysics*, **138**, 185–197, doi: [10.1016/j.jappgeo.2017.01.016](https://doi.org/10.1016/j.jappgeo.2017.01.016).
- Key, K., and J. Owall, 2011, A parallel goal-oriented adaptive finite element method for 2.5-D electromagnetic modeling: *Geophysical Journal International*, **186**, 137–154, doi: [10.1111/j.1365-246X.2011.05025.x](https://doi.org/10.1111/j.1365-246X.2011.05025.x).
- Kirkegaard, C., and E. Auken, 2015, A parallel, scalable and memory efficient inversion code for very large-scale airborne electromagnetics surveys: *Geophysical Prospecting*, **63**, 495–507, doi: [10.1111/1365-2478.12200](https://doi.org/10.1111/1365-2478.12200).
- Li, W.-B., Z.-F. Zeng, J. Li, X. Chen, K. Wang, and Z. Xia, 2016, 2.5 D forward modeling and inversion of frequency-domain airborne electromagnetic data: *Applied Geophysics*, **13**, 37–47, doi: [10.1007/s11770-016-0548-y](https://doi.org/10.1007/s11770-016-0548-y).
- Liu, G., and A. Becker, 1990, Two-dimensional mapping of sea-ice keels with airborne electromagnetics: *Geophysics*, **55**, 239–248, doi: [10.1190/1.1442832](https://doi.org/10.1190/1.1442832).
- Marquardt, D., 1963, An algorithm for least squares estimation of nonlinear parameters: *SIAM Journal of Applied Mathematics*, **11**, 431–441, doi: [10.1137/0111030](https://doi.org/10.1137/0111030).
- McGillivray, P. R., and D. W. Oldenburg, 1990, Methods for calculating Fréchet derivatives and sensitivities for the non-linear inverse problem: A comparative study: *Geophysical Prospecting*, **38**, 499–524.
- Menke, W., 1989, *Geophysical data analysis: Discrete inverse theory*: Academic Press.
- Mitsuhata, Y., 2000, 2D electromagnetic modeling by finite-element method with a dipole source and topography: *Geophysics*, **65**, 465–475, doi: [10.1190/1.1444740](https://doi.org/10.1190/1.1444740).
- Mitsuhata, Y., and T. Uchida, 2002, 2-D inversion of frequency-domain EM data caused by a 3-D source: *Methods in Geochemistry and Geophysics*, **35**, 153–171, doi: [10.1016/S0076-6895\(02\)80091-9](https://doi.org/10.1016/S0076-6895(02)80091-9).
- Reid, J. E., A. Pfaffling, and J. Vrbancich, 2006, Airborne electromagnetic footprints in 1D earths: *Geophysics*, **71**, no. 2, G63–G72, doi: [10.1190/1.2187756](https://doi.org/10.1190/1.2187756).
- Saad, Y., 1994, ILUT: A dual threshold incomplete LU factorization: *Numerical Linear Algebra with Applications*, **1**, 387–402, doi: [10.1002/\(ISSN\)1099-1506](https://doi.org/10.1002/(ISSN)1099-1506).
- Saad, Y., 2003, *Iterative methods for sparse linear systems*: SIAM.
- Siemon, B., S. Costabel, W. Voß, U. Meyer, N. Deus, J. Elbracht, T. Günther, and H. Wiederhold, 2015, Airborne and ground geophysical mapping of coastal clays in Eastern Friesland, Germany: *Geophysics*, **80**, no. 3, WB21–WB34, doi: [10.1190/geo2014-0102.1](https://doi.org/10.1190/geo2014-0102.1).
- Stoyer, C., and R. J. Greenfield, 1976, Numerical solutions of the response of a two-dimensional earth to an oscillating magnetic dipole source: *Geophysics*, **41**, 519–530, doi: [10.1190/1.1440630](https://doi.org/10.1190/1.1440630).
- Streich, R., M. Becken, and O. Ritter, 2011, 2.5D controlled-source EM modeling with general 3D source geometries: *Geophysics*, **76**, no. 6, F387–F393, doi: [10.1190/geo2011-0111.1](https://doi.org/10.1190/geo2011-0111.1).
- Vöge, M., A. Pfaffhuber, E. Auken, C. Kirkegaard, T. Boesen, S. Hendricks, and P. Hunkeler, 2015, 2.5 D inversion of sea ice thickness from helicopter EM data: Presented at the First European Airborne Electromagnetics Conference.
- Ward, S. H., and G. W. Hohmann, 1988, Electromagnetic theory for geophysical applications, in M. N. Nabighian, ed., *Electromagnetic methods in applied geophysics*: SEG, 131–311.
- Wilson, G., A. Raiche, and F. Sugeng, 2006, 2.5 D inversion of airborne electromagnetic data: *Exploration Geophysics*, **37**, 363–371, doi: [10.1071/EG06363](https://doi.org/10.1071/EG06363).
- Wilson, G. A., L. H. Cox, M. Čuma, and M. S. Zhdanov, 2012, Inverting airborne geophysical data for mega-cell and giga-cell 3D earth models: *The Leading Edge*, **31**, 316–321, doi: [10.1190/1.3694899](https://doi.org/10.1190/1.3694899).
- Yang, D., and D. W. Oldenburg, 2012a, Practical 3D inversion of large airborne time domain electromagnetic data sets: 22nd International Geophysical Conference and Exhibition, ASEG, Extended Abstracts, 26–29.
- Yang, D., and D. W. Oldenburg, 2012b, Three-dimensional inversion of airborne time-domain electromagnetic data with applications to a porphyry deposit: *Geophysics*, **77**, no. 2, B23–B34, doi: [10.1190/geo2011-0194.1](https://doi.org/10.1190/geo2011-0194.1).
- Yang, D., D. W. Oldenburg, and E. Haber, 2014, 3D inversion of airborne electromagnetic data parallelized and accelerated by local mesh and adaptive soundings: *Geophysical Journal International*, **196**, 1492–1507, doi: [10.1093/gji/ggt465](https://doi.org/10.1093/gji/ggt465).
- Yu, W. W., and E. Haber, 2012, A 2.5 D inversion of airborne electromagnetic data: 82nd Annual International Meeting, SEG, Expanded Abstracts, doi: [10.1190/segam2012-0592.1](https://doi.org/10.1190/segam2012-0592.1).
- Zienkiewicz, O. C., R. L. Taylor, and R. L. Taylor, 1977, *The finite element method*: McGraw-Hill.