



Accelerating 3D Magnetotelluric Forward Modelling with Domain Decomposition and Order-Reduction Methods

Luis Tao¹, Alba Muixí^{1,2}, Sergio Zlotnik^{1,2}, Fabio Ivan Zyserman^{3,4}, Juan Carlos Afonso^{5,6,7}, and Pedro Diez^{1,2}

¹Laboratori de Càlcul Numèric (LaCàN), ETS de Ingenieria de Caminos, Canales y Puertos, Universitat Politècnica de Catalunya, Barcelona, Spain

²Centre Internacional de Mètodes Numèrics en Enginyeria (CIMNE), Barcelona, Spain

³Centro de Investigaciones Geofísicas, Facultad de Ciencias Astronómicas y Geofísicas, Universidad Nacional de La Plata, La Plata, Argentina

⁴Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET), Argentina

⁵School of Natural Sciences and CODES, University of Tasmania, Australia

⁶Department of Earth and Space Sciences, Southern University of Science and Technology Shenzhen, Guangdong, China

⁷Faculty of Geo-Information and Earth Observation (ITC), University of Twente, Enschede, Netherlands

Correspondence: Luis Tao (luis.tao@upc.edu)

Abstract. Three-dimensional (3D) magnetotelluric (MT) forward modelling is computationally demanding, limiting its use in global uncertainty quantification and sampling-based probabilistic inversion. Here, we introduce a novel forward-modelling framework that combines an iterative domain decomposition (DD) formulation with proper orthogonal decomposition (POD) reduced-order modelling to enable scalable and efficient 3D MT simulations. The DD component partitions the computational domain into subdomains, avoiding the factorization of a single global system, accelerating simulations by over 60% compared to global solvers, and alleviating memory bottlenecks in large problems. The POD component leverages the local DD solutions to construct a reduced-order version of the problem that can deliver accurate and fast solutions to the 3D forward problem during subsequent evaluations. Using the DTM1 benchmark and a real-world conductivity model, we quantify runtime, memory, and accuracy in terms of MT quantities of interest (apparent resistivity and phase). DD–POD achieves speed-ups exceeding 90% relative to full-order solvers and up to 70% relative to existing ROM techniques, while maintaining acceptable accuracy. These results suggest that DD–POD can make higher-resolution 3D MT forward modelling practical within sampling-based workflows by substantially reducing both runtime and memory demands.

1 Introduction

The magnetotelluric (MT) method (Tikhonov, 1950; Cagniard, 1953) is a passive electromagnetic (EM) technique used to estimate the Earth's electrical conductivity structure from near-surface to mantle depths (Vozoff, 1990; Chave and Jones, 2012). Applications range from imaging fluid circulation and magmatic systems (e.g. Heise et al., 2008; Samrock et al., 2018; Piña-Varas et al., 2023) to investigating mantle compositional anomalies and whole-lithospheric structure (e.g. Manassero et al., 2024; Heinson et al., 2018; Jones, 1999; Chase et al., 2023).



Over the past two decades, substantial progress has been achieved in developing 3D MT forward and inversion algorithms (e.g. deGroot Hedlin and Constable, 1990; Mackie et al., 1993; Newman and Alumbaugh, 2000; Zyserman and Santos, 2000; Siripunvaraporn et al., 2005; Avdeev and Avdeeva, 2009; Egbert and Kelbert, 2012; Siripunvaraporn, 2012; Key, 2016; Elías et al., 2022; Grayver, 2015; Wang et al., 2025). Deterministic inversion schemes, typically based on finite-difference, finite-element, finite-volume or integral-equation solvers, remain the most widely adopted methods owing to their efficiency and moderate computational cost. However, such matrix-based approaches generally lead to ill-posed inverse problems and therefore rely on strong regularisation and *a priori* constraints to obtain a single best-fitting model. They also have limited ability to accommodate strong non-linearities, non-uniqueness, and provide limited support for global uncertainty analysis (cf. Aster et al., 2018; Robertson et al., 2020; Trainor-Guitton and Hoversten, 2011; Manassero et al., 2020).

Probabilistic inversion frameworks (cf. Duijndam, 1988; Mosegaard and Tarantola, 1995; Sambridge and Mosegaard, 2002; Tarantola, 2005) overcome these limitations by reformulating the inverse problem as a well-posed, inference problem, where the solution is represented by a multi-dimensional probability density function in data-parameter space. As such, they provide a natural framework for addressing non-uniqueness, complex forward problems, and global uncertainty quantification, while also enabling the coherent integration of multiple data types through joint inversion (eg. Tarantola, 2005; Khan et al., 2007; Afonso et al., 2013a, b, 2016; Manassero et al., 2021; Bissig et al., 2023; Manassero et al., 2024). Nevertheless, their application to realistic 3D MT problems remains computationally limited, as evaluating the posterior probability density function typically requires between 10^5 and 10^7 forward simulations (Manassero et al., 2020) when using Markov Chain Monte Carlo (MCMC) approaches. The overall feasibility of such inversions is therefore strongly controlled by the computational efficiency of the forward solver.

In this context, conventional 3D MT forward solvers are notoriously computationally demanding and scale poorly with increasing mesh resolution. Recent studies have aimed to address the efficiency problem by solving large-scale forward models (1.5×10^9 degrees of freedom) within only a few minutes (Bai et al., 2025; Wang et al., 2025; Qin and Grayver, 2025). Despite these considerable improvements in efficiency, the computational time of the forward MT problem precludes full probabilistic inversions based on sampling algorithms (e.g. MCMC simulations)

Another group of studies have sought to alleviate the computational burden through reduced-order modeling (Manassero et al., 2020) and machine-learning approaches (e.g. Peng et al., 2023; Ling et al., 2024). In particular, Manassero (2019) introduced a reduced-basis (RB) approach to accelerate MT forward modelling within MCMC inversion frameworks, achieving substantial computational speed-ups while maintaining acceptable accuracy and scalability. This methodology was subsequently extended by Manassero et al. (2021, 2024) to develop a joint MT-seismic probabilistic inversion scheme based on adaptive MCMC, and applied to real data from southeastern Australia. While this approach enabled fully probabilistic 3D MT inversion for the first time, the achievable resolution and effective number of degrees of freedom in practical applications remained lower than those typically attainable using deterministic inversion methods.

Building on these developments, the present study proposes a new formulation that integrates Domain Decomposition (DD) with model-order reduction via Proper Orthogonal Decomposition (POD) to produce highly efficient 3D MT forward simulations. The proposed DD-POD framework significantly reduces memory requirements and achieves runtime reductions of up to \sim



2 orders of magnitude compared to conventional implementations. These improvements extend considerably the capabilities of the approach proposed by Manassero (2019), enabling finer spatial discretizations, larger-scale models, and enhanced imaging resolution.

The remainder of this paper is organized as follows. Section 2 reviews the domain-decomposition formulation (Zyserman and Santos, 2000). Section 3 introduces the POD-based reduction strategy. Section 4 outlines implementation details of the DD-POD solver, and Section 5 presents results for a benchmark case and a real-world conductivity model used for validation. Section 6 discusses the results in terms of accuracy and computational performance, and Section 7 summarizes the main findings and perspectives for future work.

2 Forward problem formulation and discretization

2.1 The MT model equations

In the magnetotelluric method, naturally varying external electromagnetic fields induce electric currents within the Earth, whose spatial distribution is controlled by the subsurface electrical conductivity structure (cf. Tikhonov, 1950; Cagniard, 1953). Assuming a time-harmonic dependence of the electric \mathbf{E} and magnetic \mathbf{H} fields, the governing equations in the frequency domain and under the diffusive approximation (Ward and Hohmann, 1988), are:

$$\nabla \times \mathbf{E} = -i\omega\mu_0\mathbf{H}, \quad (1a)$$

$$\nabla \times \mathbf{H} - \sigma\mathbf{E} = 0. \quad (1b)$$

where ω (rad/s) stands for the angular frequency. The electric field $\mathbf{E}(x, y, z, \omega)$ (V/m) and magnetic field $\mathbf{H}(x, y, z, \omega)$ (A/m) describe the Earth's electromagnetic response for a given ω in a region free of sources. The electrical conductivity $\sigma(x, y, z)$ S m^{-1} is assumed isotropic, and μ_0 (H/m) is the magnetic permeability of free space.

Following the scattered-field formulation (Wannamaker et al., 1987), the electrical conductivity distribution is decomposed into a laterally homogeneous background model and three-dimensional perturbations associated with local inhomogeneities. The background conductivity is defined over the primary domain Ω_p and corresponds here to a horizontally layered Earth, such that $\sigma_p(x, y, z) := \sigma_p(z)$. Conductivity contrasts relative to this reference structure are described by $\sigma_i(x, y, z)$ within the subdomain Ω_i containing the 3-D heterogeneities. Under this decomposition, the total conductivity model, $\sigma(x, y, z)$, is expressed as:

$$\sigma(x, y, z) = \begin{cases} \sigma_p(z), & \text{in } \Omega_p \text{ (layered Earth),} \\ \sigma_p(z) + \sigma_i(x, y, z), & \text{in } \Omega_i \text{ (inhomogeneities).} \end{cases} \quad (2)$$

The total electromagnetic $(\mathbf{E}_t, \mathbf{H}_t)$ are the responses induced by a plane monochromatic electromagnetic wave of frequency ω incident upon its top boundary, and are defined as:



$$\mathbf{E}_t = \mathbf{E}_p + \mathbf{E}_s, \quad (3a)$$

$$\mathbf{H}_t = \mathbf{H}_p + \mathbf{H}_s. \quad (3b)$$

where the primary fields $(\mathbf{E}_p, \mathbf{H}_p)$ correspond to the solutions of Eqs. (1a)–(1b) for the horizontally layered background conductivity $\sigma(x, y, z) := \sigma_p(z)$. The analytical expressions for an n -layered Earth are given in Weaver (1994).

The secondary fields $(\mathbf{E}_s, \mathbf{H}_s)$ correspond to the solutions of Eqs. (1a)–(1b) obtained by substituting the total-field decompositions in Eqs. (3a)–(3b) and using the conductivity model defined in Eq. (2), and represent the response of inhomogeneities embedded in a horizontally layered medium.

$$\sigma \mathbf{E}_s - \nabla \times \mathbf{H}_s = (\sigma - \sigma_p) \mathbf{E}_p =: -\mathbf{F} \quad \text{in } \Omega, \quad (4a)$$

$$i\omega\mu_0 \mathbf{H}_s + \nabla \times \mathbf{E}_s = 0 \quad \text{in } \Omega, \quad (4b)$$

where $\Omega \subset \mathbb{R}^3$ denotes the computational domain. To complete the formulation, first-order absorbing boundary conditions (Sheen, 1997) are imposed at the computational domain's boundary, $\partial\Omega$. Reflections at boundaries are therefore suppressed by absorbing the normal component of the EM fields,

$$(1 - i) P_\tau a \mathbf{E}_s + \nu \times \mathbf{H}_s = 0 \quad \text{on } \partial\Omega \equiv \Gamma, \quad (5)$$

with $a = \left(\frac{\sigma}{2\omega\mu_0}\right)^{1/2}$, ν the outward unit normal vector to domain's boundary Γ , and $P_\tau \zeta = -\nu \times (\nu \times \zeta)$ the tangential projection of an arbitrary field ζ onto the boundary, in this case \mathbf{E}_s . This approach removes the need for domain padding (i.e. extended model), thereby reducing computational cost, while maintaining solution accuracy (Manassero et al., 2020).

2.2 Domain decomposition solution

In this section, we introduce the strong and weak formulations and the iterative DD scheme used to solve the problem governed by Eqs. (4a)–(5). A detailed description of these formulations and the DD framework can be found in Zyserman and Santos (2000). Here, we summarize the most relevant aspects.

2.2.1 Strong form

Domain Decomposition methods partition the computational domain into smaller subdomains that can be solved independently, with coupling restricted to shared interfaces. This structure is well suited to parallel computation and can substantially reduce memory usage and runtime when scaling to large problem sizes (e.g. Gauzellino et al., 2009).

The main domain Ω is partitioned into non-overlapping, adjacent subdomains $\Omega_j \subset \mathbb{R}^3$, with $j = 1, \dots, N_s$. In each subdomain Ω_j , Eqs. (4a)–(5) are written as:



$$\sigma \mathbf{E}_j - \nabla \times \mathbf{H}_j = -\mathbf{F}_j \quad \text{in } \Omega_j, \quad (6a)$$

$$i\omega\mu_0 \mathbf{H}_j + \nabla \times \mathbf{E}_j = 0 \quad \text{in } \Omega_j, \quad (6b)$$

$$110 \quad (1-i)P_\tau a \mathbf{E}_j + \nu_j \times \mathbf{H}_j = 0 \quad \text{on } B_j \equiv \Gamma \cap \Gamma_j, \quad (6c)$$

where $(\mathbf{E}_j, \mathbf{H}_j)$ denote the secondary electromagnetic fields in subdomain j , with Γ_j representing the boundaries of Ω_j and Γ the external boundary of Ω . The term $-\mathbf{F}_j$ represents the source, which arises from inhomogeneities and vanishes if none are present in Ω_j . The continuity between subdomains is imposed through Robin-type transmission conditions (Douglas et al., 2000):

$$115 \quad (\nu \times \mathbf{H}_j + \beta_{jk} P_\tau \mathbf{E}_j) = -(\nu \times \mathbf{H}_k + \beta_{kj} P_\tau \mathbf{E}_k) \quad \text{on } \Gamma_{jk} \subset \partial\Omega_j, \quad (7)$$

$$(\nu \times \mathbf{H}_k + \beta_{kj} P_\tau \mathbf{E}_k) = -(\nu \times \mathbf{H}_j + \beta_{jk} P_\tau \mathbf{E}_j) \quad \text{on } \Gamma_{kj} \subset \partial\Omega_k. \quad (8)$$

where Γ_{jk} denotes the interface between the adjacent subdomains Ω_j and Ω_k , and β_{jk} is a complex hyperparameter controlling the convergence rate of the iterative scheme. Following Zyserman and Santos (2000), β_{jk} is defined as the average of a on Γ_{jk} and Γ_{kj} , multiplied by $(1-i)$, in order to mimic the absorbing boundary condition in Eq. (5) on the interior interfaces.

120 2.2.2 Weak form

The variational formulation for each subdomain is obtained by testing Eq. (4a) against the real function $\varphi(x, y, z)$, with the requirement that $\nabla \times \varphi$ is square-integrable, and by testing Eq. (4b) against the real function $\psi(x, y, z)$ (Douglas Jr et al., 1997; Santos and Sheen, 1998). By incorporating the absorbing boundary conditions (Eq. (5)) and the Robin-type transmission conditions (Eqs. (7)–(8)), the following variational problem is obtained:

$$125 \quad (\sigma \mathbf{E}_j, \varphi)_{\Omega_j} - (\mathbf{H}_j, \nabla \times \varphi)_{\Omega_j} + \sum_k \langle \beta_{jk} (P_\tau \mathbf{E}_j - P_\tau \mathbf{E}_k) + \nu_k \times \mathbf{H}_k, P_\tau \varphi \rangle_{\Gamma_{jk}} \\ + (1-i) \langle P_\tau a \mathbf{E}_j, P_\tau \varphi \rangle_{B_j} = -(\mathbf{F}_j, \varphi)_{\Omega_j} \quad (9a)$$

$$i\omega\mu(\mathbf{H}_j, \psi)_{\Omega_j} + (\nabla \times \mathbf{E}_j, \psi)_{\Omega_j} = 0 \quad (9b)$$

2.2.3 Iterative domain decomposition procedure

The idea behind DD is to iteratively solve the problem 4a-4b independently on each subdomain Ω_j , exchange interface solutions on Γ_{jk} between neighboring subdomains, and update the subdomain problems until convergence is achieved on all interfaces. The electromagnetic fields in Ω_j at iteration level $n+1$, $(\mathbf{E}_j^{n+1}, \mathbf{H}_j^{n+1})$, depend on the fields from the previous iteration $(\mathbf{E}_j^n, \mathbf{H}_j^n)$ as well as on the fields from neighboring subdomains Ω_k . Following this methodology, Zyserman and Santos (2000) proposed the following iterative algorithm:



1. Initialize with $(\mathbf{E}_j^0, \mathbf{H}_j^0)$.

135 2. For all Ω_j , compute $(\mathbf{E}_j^{n+1}, \mathbf{H}_j^{n+1})$ as the solution of:

$$\begin{aligned} & (\sigma_j \mathbf{E}_j^{n+1}, \varphi)_{\Omega_j} - (\mathbf{H}_j^{n+1}, \nabla \times \varphi)_{\Omega_j} + \sum_k \langle \beta_{jk} P_\tau \mathbf{E}_j^{n+1}, P_\tau \varphi \rangle_{\Gamma_{jk}} + (1-i) \langle P_\tau a_j \mathbf{E}_j^{n+1}, P_\tau \varphi \rangle_{B_j} \\ & = -(\mathbf{F}_j, \varphi)_{\Omega_j} + \sum_k \langle \beta_{jk} P_\tau \mathbf{E}_k^n - \nu_k \times \mathbf{H}_k^n, P_\tau \varphi \rangle_{\Gamma_{jk}}, \end{aligned} \quad (10a)$$

$$i\omega\mu(\mathbf{H}_j^{n+1}, \psi)_{\Omega_j} + (\nabla \times \mathbf{E}_j^{n+1}, \psi)_{\Omega_j} = 0. \quad (10b)$$

3. Update $\nu_j \times \mathbf{H}_j^{n+1}$ on Γ_{jk} by

$$140 \quad \nu_j \times \mathbf{H}_j^{n+1} = -\nu_k \times \mathbf{H}_k^n - \beta_{jk} (P_\tau \mathbf{E}_j^{n+1} - P_\tau \mathbf{E}_k^n), \quad (11)$$

4. Stop if $\mathbf{E}_j^{n+1} \approx \mathbf{E}_j^n$. If not, return to step 2.

The choice of $(\mathbf{E}_j^0, \mathbf{H}_j^0)$ depends on the model's a priori information. If no such information is available, $(\mathbf{E}_j^0, \mathbf{H}_j^0) = (\mathbf{0}, \mathbf{0})$ is used as the standard initial guess. The complete iterative scheme is summarized in Algorithm 1, whose present implementation is referred to as DD. The global formulation used in Manassero et al. (2020) and Elías et al. (2022) can be recovered by setting

145 $N_s = 1$.

The convergence criterion used to terminate the iterative process evaluates the solution quality across all subdomains and is controlled by the parameter δ . For this purpose, the following metric, denoted by r and based on the L_2 norm, is:

$$r^{(n+1)} = 2 \sqrt{\frac{\sum_{j=1}^{N_s} \|\mathbf{E}_j^{n+1} - \mathbf{E}_j^n\|^2}{\sum_{j=1}^{N_s} \|\mathbf{E}_j^{n+1} + \mathbf{E}_j^n\|^2}}. \quad (12)$$

2.2.4 Finite element discretization

150 For notational simplicity, each subdomain is assumed to consist of a single element; this restriction is relaxed in later sections. The electromagnetic fields within each element are approximated as follows:

$$\mathbf{E}_j^{n+1}(\mathbf{x}, \omega) \approx \sum_{\alpha=1}^{12} \varepsilon_j^{\alpha, n+1}(\omega) \varphi^\alpha(\mathbf{x}), \quad (13)$$

$$\mathbf{H}_j^{n+1}(\mathbf{x}, \omega) \approx \sum_{\eta=1}^9 h_j^{\eta, n+1}(\omega) \psi^\eta(\mathbf{x}), \quad (14)$$



155 where $\boldsymbol{\varepsilon}_j^{n+1} \in \mathbb{C}^{12}$ and $\boldsymbol{h}_j^{n+1} \in \mathbb{C}^9$ denote the vectors of unknown coefficients associated with the electric and magnetic fields, respectively. These vectors represent the local degrees of freedom for element j at iteration $n+1$. Each element has 12 degrees of freedom for the electric field (2 tangential components per face) and 9 for the magnetic field (associated with the element interior), resulting in a total of 21 unknowns per element. However, by construction of the basis functions, and using Eq. (9b), the magnetic field coefficients $h_j^{\eta,n+1}$ can be expressed in terms of the electric field coefficients $\varepsilon_j^{\alpha,n+1}$, thereby
 160 reducing the number of independent unknowns per element to 12. Details of the basis functions associated with each degree of freedom, $\boldsymbol{\varphi}^\alpha$ ($\alpha = 1, \dots, 12$) for \mathbf{E} and $\boldsymbol{\psi}^\eta$ ($\eta = 1, \dots, 9$) for \mathbf{H} , are provided in Zyserman and Santos (2000), following the original formulation in Santos et al. (1999).

To accelerate the domain decomposition (DD) iterative scheme, the formulation is hybridized following Arnold and Brezzi (1985). Rather than enforcing continuity of the tangential components of the electric field over the entire interface faces Γ_{jk} ,
 165 continuity is imposed only at a single representative midpoint $m_{jk} \subset \Gamma_{jk}$. Enforcing the continuity constraint at one point per interface face reduces the number of interface constraints and, consequently, the number of degrees of freedom in each subdomain system as well as the number of globally coupled interface unknowns. The remaining continuity constraints are enforced through Lagrange multipliers defined on the interelement boundaries. The introduction of these Lagrange multipliers $\boldsymbol{\lambda}_{jk}$ yields a block-diagonal algebraic structure that allows the electric and magnetic fields to be computed separately. Together,
 170 the midpoint enforcement and Lagrange multiplier formulation simplify the algebraic problem and reduce the computational cost per iteration of the DD solver. A complex-valued Lagrange multiplier $\boldsymbol{\lambda}_{jk}^{n+1} \in \mathbb{C}^2$ is introduced at each of these midpoints to weakly enforce continuity, with $\boldsymbol{\lambda}_{jk}^{n+1} \approx \nu_j \times \mathbf{H}_j^{n+1}$.

Finally, substituting the approximations of the EM fields from Eqs. (13)-(14) into Eqs. (9a)-(9b), and expressing $h_j^{\eta,n+1}$ in terms of $\varepsilon_j^{\alpha,n+1}$ using Eq. (9b), the left-hand side of Eq. (9a) becomes dependent only on ε_j^{n+1} . The resulting system can
 175 therefore be written in the form of a linear system of equations (LSE):

$$\mathbf{K}_j \boldsymbol{\varepsilon}_j^{n+1} = -\mathbf{f}_j^M + \sum_{k=1}^K (C_{kj} \boldsymbol{\varepsilon}_{kj}^n - \boldsymbol{\lambda}_{jk}^n), \quad (15)$$

where $\mathbf{K}_j \in \mathbb{C}^{12 \times 12}$ is the local system matrix; $\mathbf{f}_j \in \mathbb{C}^{12}$ is the source term associated with the primary fields; $\boldsymbol{\varepsilon}_j^{n+1} \in \mathbb{C}^{12}$ is the vector of electric-field coefficients in element j ; and $\boldsymbol{\varepsilon}_{jk}^n \in \mathbb{C}^2$ denotes the electric-field coefficients on the interface face Γ_{jk} of Ω_j shared with Ω_k . The same notation is used for $\boldsymbol{\lambda}_{jk}$. Here, $C_{jk} = A_{\Gamma_{jk}} \beta_{jk}$, where $A_{\Gamma_{jk}}$ is the interface area,
 180 M denotes the polarization (XY or YX), and K is the number of neighboring subdomains of j . Additionally, the Lagrange multipliers are updated at each iteration by:

$$\boldsymbol{\lambda}_{jk}^{n+1} = -\boldsymbol{\lambda}_{kj}^n - \beta_{jk} \left(\boldsymbol{\varepsilon}_{jk}^{n+1} - \boldsymbol{\varepsilon}_{kj}^n \right). \quad (16)$$



Algorithm 1 Simplified Iterative Domain Decomposition Scheme (based on Zyserman and Santos (2000))

1: **Input:** Initial guesses $(\epsilon_j^0, \lambda_{jk}^0)$; hyperparameters $\{n_{\max}, \beta, \delta\}$; material properties $\{\sigma, \mu_0\}$; frequency ω .
2: **Output:** Approximate solution $\{\mathbf{E}_j(\mathbf{x}, \omega), \mathbf{H}_j(\mathbf{x}, \omega)\}_{j=1}^{N_s}$.
3:
4: **for** $M \in \{XY, YX\}$ **do**
5: **Initialize:** Set $\{\epsilon_j^0\}$ and $\{\lambda_{jk}^0\}$ for all subdomains Ω_j .
6: **for** $n = 0$ to n_{\max} **do**
7: **for** $j = 1$ to N_s **do**
8: **Solve subproblem:** Solve Eq. (15) in Ω_j to obtain ϵ_j^{n+1} .
9: **for** $k = 1$ to K_s **do**
10: **Exchange interface data:** Send $(\epsilon_j^{n+1}, \lambda_{jk}^n)$ to neighboring subdomain Ω_k .
11: **Update Lagrange multipliers:** Update λ_{jk}^{n+1} using Eq. (16).
12: **end for**
13: **end for**
14: **Check convergence:** Compute $r^{(n+1)}$ with Eq. (12).
15: **if** $r^{(n+1)} < \delta$ **then**
16: **break**
17: **end if**
18: **end for**
19: **for** $j = 1$ to N_s **do**
20: **Recover fields:** Compute $(\mathbf{E}_j^{n+1}, \mathbf{H}_j^{n+1})$ using Eqs. (13)–(14).
21: **end for**
22: **end for**
23:
24: **Compute QoIs:** Compute $(\rho_{xy}, \phi_{xy}) - (\rho_{yx}, \phi_{yx})$ using Eqs. (19)–(20).

3 Reduced-order formulation of the forward problem

In this section, a reduced-order formulation of the MT forward problem is incorporated into the DD framework of Section 2.2.

185 We begin by outlining the general dimensionality-reduction strategy based on the Reduced Basis (RB) method (Berkooz et al., 1993; Patera et al., 2007; Nguyen and Peraire, 2008; Rozza et al., 2008; Quarteroni et al., 2015; Muixí et al., 2023, 2025). This reduction framework is then combined with the DD algorithm described in Section 3.3 to obtain a highly efficient approach for solving MT forward problems.

3.1 Quantities of interest

190 In reduced-order modeling, including the RB method, accuracy is typically evaluated with respect to selected quantities of interest (QoIs) that represent the relevant model outputs. Here, the QoIs are not the electromagnetic fields themselves, but the



derived apparent resistivity ρ_a and phase ϕ . These are obtained from the complex impedance tensor $\mathbf{Z} \in \mathbb{C}^{2 \times 2}$ (cf. Simpson and Bahr, 2005), which relates the horizontal components of the electric and magnetic fields through

$$\mathbf{E} = \mathbf{Z} \mathbf{H}, \quad (17)$$

195 or equivalently,

$$\begin{pmatrix} E_x \\ E_y \end{pmatrix} = \begin{pmatrix} Z_{xx} & Z_{xy} \\ Z_{yx} & Z_{yy} \end{pmatrix} \begin{pmatrix} H_x \\ H_y \end{pmatrix}. \quad (18)$$

The apparent resistivity is defined from the magnitude of the impedance tensor components as:

$$\rho_{a,ij}(\omega) = \frac{1}{\mu_0 \omega} |Z_{ij}(\omega)|^2. \quad (19)$$

The apparent phase corresponds to the argument of each impedance component,

$$200 \quad \phi_{ij} = \tan^{-1} \left(\frac{\text{Im}\{Z_{ij}\}}{\text{Re}\{Z_{ij}\}} \right). \quad (20)$$

3.2 POD-RB for a linear system of equations

Consider a parametric LSE where $\boldsymbol{\mu} = (\mu_1, \dots, \mu_p) \in \mathbb{R}^p$ denotes the parameter vector, and let $\mathbf{x}(\boldsymbol{\mu}) \in \mathbb{C}^N$ be the parametric unknown that corresponds to a vector of nodal values, for instance in a finite element (FE) discretization,

$$\mathbf{K}(\boldsymbol{\mu}) \mathbf{x}(\boldsymbol{\mu}) = \mathbf{f}(\boldsymbol{\mu}), \quad (21)$$

205 Furthermore, consider n_s vectors, referred to as *snapshots* \mathbf{x}_i , for $i = 1, \dots, n_s$. Each snapshot represents the solution of a full-order forward problem (21) obtained from an independent forward simulation using a specific set of model parameters $\boldsymbol{\mu}$. The snapshots are collected in the snapshot matrix $\mathbf{X} \in \mathbb{C}^{N \times n_s}$, defined as

$$\mathbf{X} = [\mathbf{x}_1 \quad \mathbf{x}_2 \quad \dots \quad \mathbf{x}_{n_s}] \quad (22)$$

210 The RB approach is then employed to approximate the solution of problem (21) for parameter values not included in the snapshot matrix, by expressing the solution as a weighted linear combination of the precomputed snapshots (cf. Rozza et al., 2008; Muixí et al., 2023). Because snapshots associated with different parameter choices do not necessarily contribute equally to the approximation quality, it is essential to identify and retain only the most informative snapshots in order to minimise redundancy and maintain a compact reduced system. To this end, the POD approach, based on the Singular the Value Decomposition (SVD) algorithm, is employed to extract the dominant modes (cf. Rozza et al., 2008). The SVD of \mathbf{X} is given
215 by

$$\mathbf{X} = \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^H, \quad (23)$$



where $\mathbf{U} \in \mathbb{C}^{N \times N}$ and $\mathbf{V} \in \mathbb{C}^{n_s \times n_s}$ are unitary matrices, and $\Sigma \in \mathbb{R}^{N \times n_s}$ is a diagonal matrix containing the singular values of \mathbf{X} in descending order, $\tau_1 \geq \tau_2 \geq \dots \geq \tau_{n_s} \geq 0$.

The n_s columns of U form an orthonormal basis for the subspace spanned by the snapshots. The most relevant information is captured in the leading modes, which are associated with the largest singular values. By prescribing a tolerance δ_{SVD} , only n_{POD} modes are retained, such that:

$$\sum_{i=1}^{n_{\text{POD}}} \tau_i \geq (1 - \delta_{\text{SVD}}) \sum_{i=1}^{n_s} \tau_i. \quad (24)$$

The linear space spanned from the first n_{POD} columns of \mathbf{U} is used to approximate the solution of Eq. (21) for new parametric points $\boldsymbol{\mu}$. The POD represents the unknown \mathbf{x} using only n_{POD} modes. Typically, $n_{\text{POD}} \ll n_s$, and n_s is expected to be much smaller than the dimension of the full-order system, N (Muixí et al., 2025). This is because \mathbf{X} is constructed from a limited number of representative solutions of the full-order model, and the dominant POD modes capture most of the system's variance. As a result, only a small number of modes is required to accurately approximate the solution space, leading to a reduced basis whose dimension is significantly smaller than that of the original system (Rozza et al., 2008). The full-order approximation $\tilde{\mathbf{x}}$ of the solution vector $\mathbf{x} \in \mathbb{C}^N$ is obtained through the projection:

$$\mathbf{x} \approx \tilde{\mathbf{x}} = \mathbf{U}_r \mathbf{z} = \sum_{i=1}^{n_{\text{POD}}} U_i z_i, \quad (25)$$

where $\mathbf{z} \in \mathbb{C}^{n_{\text{POD}}}$ denotes the vector of reduced coefficients and $\mathbf{U}_r \in \mathbb{C}^{N \times n_{\text{POD}}}$ is the reduced-basis matrix, formed by the first n_{POD} columns of \mathbf{U} . Following a reduced-basis (RB) approach (cf. Rozza et al., 2008; Quarteroni et al., 2015), the coefficients \mathbf{z} are obtained by applying a Galerkin projection to the full-order system (21), leading to

$$[\mathbf{U}_r^H \mathbf{K} \mathbf{U}_r] \mathbf{z} = \mathbf{U}_r^H \mathbf{f}(\boldsymbol{\mu}), \quad (26)$$

which can be compactly denoted as,

$$\mathbf{K}^*(\boldsymbol{\mu}) \mathbf{z} = \mathbf{f}^*(\boldsymbol{\mu}), \quad (27)$$

where \mathbf{U}_r^H denotes the Hermitian transpose of \mathbf{U}_r . The reduced operators satisfy $\mathbf{K}^* \in \mathbb{C}^{n_{\text{POD}} \times n_{\text{POD}}}$ and $\mathbf{f}^* \in \mathbb{C}^{n_{\text{POD}}}$. Consequently, instead of solving the full $N \times N$ linear system in Eq. (21), the reduced $n_{\text{POD}} \times n_{\text{POD}}$ system in Eq. (27) is solved. This yields an approximation $\tilde{\mathbf{x}}$ of the full solution \mathbf{x} , with $n_{\text{POD}} \ll N$.

240 3.3 Joint reduced order-domain decomposition formulation

The POD-RB formulation presented in Section 3.1 is now applied to approximate the solution of Eq.(21). From a general point of view, the parameters $\boldsymbol{\mu}$ correspond to the conductivities $\boldsymbol{\sigma}$ assigned to each element in the FE discretization and $\mathbf{x}(\boldsymbol{\mu})$ corresponds to $\boldsymbol{\varepsilon}_j$.



In this implementation, each subdomain j of the DD is equipped with its own reduced basis $\mathbf{U}_j^M \in \mathbb{C}^{N \times n_{\text{POD}}}$ (the truncation subscript r is omitted for clarity), one for each polarization M , with $M \in \{XY, YX\}$. Consequently, the reduced-order form of Eq. (15) becomes

$$[(\mathbf{U}_j^M)^H \mathbf{K}_j \mathbf{U}_j^M] \mathbf{z}_j^{n+1} = -(\mathbf{U}_j^M)^H [\mathbf{f}_j^M + \sum_{k=1}^K (C_{jk} \tilde{\mathbf{e}}_{jk}^n - \tilde{\boldsymbol{\lambda}}_{jk}^n)], \quad (28)$$

where $\mathbf{z}_j^{n+1} \in \mathbb{C}^{n_{\text{POD}}}$ denotes the reduced solution of the forward problem at iteration $n+1$ in subdomain j . The quantities $\tilde{\mathbf{e}}_{jk}^n$ and $\tilde{\boldsymbol{\lambda}}_{jk}^n$ are the full-order approximations of $\boldsymbol{\varepsilon}_{jk}^n$ and $\boldsymbol{\lambda}_{jk}^n$, respectively. In compact form, the reduced system can be written as:

$$\mathbf{K}_j^{*,M} \mathbf{z}_j^{n+1} = \mathbf{f}_{j,1}^{*,M} - (\mathbf{U}_j^M)^H \sum_{k=1}^K (C_{jk} \tilde{\mathbf{e}}_{jk}^n - \tilde{\boldsymbol{\lambda}}_{jk}^n). \quad (29)$$

The term $\mathbf{f}_{j,1}^{*,M}$ denotes the iteration-independent contribution to the right-hand side. From Eq. (29), $\mathbf{K}_j^{*,M}$ and $\mathbf{f}_{j,1}^{*,M}$ are computed only once, whereas the second term on the right-hand side is updated at each iteration.

The full-order approximation of the electric field at the interfaces is obtained by projection according to Eq. (25),

$$\tilde{\mathbf{e}}_{jk}^{n+1} = \mathbf{U}_j^M \mathbf{z}_j^{n+1}. \quad (30)$$

The Lagrange multiplier approximations are updated through the interface relation in Eq. (16), using only full-order approximations as input:

$$\tilde{\boldsymbol{\lambda}}_{jk}^{n+1} = -\tilde{\boldsymbol{\lambda}}_{kj}^n - \beta_{jk} (\tilde{\mathbf{e}}_{jk}^{n+1} - \tilde{\mathbf{e}}_{kj}^n). \quad (31)$$

To construct the snapshot set, only the solution from the final iteration of each full-order parametric simulation is retained. The algorithm that illustrates this implementation is outlined in Algorithm 2 and is referred to as DD-POD.

4 Implementation aspects

This section discusses the most relevant implementation details of the proposed DD-POD algorithm.

4.1 Domain discretization

The computational domain Ω is discretized into a total of $ne = ngx \times ngy \times ngz$ elements (ngx , ngy and ngz are the number of elements in each direction).



Algorithm 2 DD–POD iterative domain decomposition scheme

```

1: Input: Initial guesses  $(\varepsilon_j^0, \lambda_{jk}^0)$ ; hyperparameters  $\{n_{\max}, \beta, \delta, n_{\text{POD}}\}$ ; material properties  $\{\sigma, \mu_0\}$ ; frequency  $\omega$ .
2:
3: Output: Approximate solution  $\{\tilde{\mathbf{E}}_j(\mathbf{x}, \omega), \tilde{\mathbf{H}}_j(\mathbf{x}, \omega)\}_{j=1}^{N_s}$ .
4:
5: for  $M \in \{XY, YX\}$  do
6:   for  $j = 1$  to  $N_s$  do
7:     Precompute  $\mathbf{K}_j^{*,M}$  and  $\mathbf{f}_{j,1}^{*,M}$  using Eq. (28).
8:   end for
9: end for
10:
11: for  $M \in \{XY, YX\}$  do
12:   Initialize: Set  $(\varepsilon_{jk}^0, \lambda_{jk}^0)$  for all subdomains  $\Omega_j$ .
13:   for  $n = 0$  to  $n_{\max}$  do
14:     for  $j = 1$  to  $N_s$  do
15:       Solve reduced subproblem: Solve Eq. (29) on  $\Omega_j$  for  $\mathbf{z}_j^{n+1}$ .
16:       Reconstruct electric field coefficients: Compute  $\tilde{\varepsilon}_j^{n+1}$  via Eq. (30).
17:       for  $k = 1$  to  $K_s$  do
18:         Exchange interface data: Send  $(\tilde{\varepsilon}_{jk}^{n+1}, \tilde{\lambda}_{jk}^n)$  to neighbor  $\Omega_k$ .
19:         Update multipliers: Update  $\tilde{\lambda}_{jk}^{n+1}$  using Eq. (16).
20:       end for
21:     end for
22:     Check convergence: Compute  $r^{(n+1)}$  using Eq. (12).
23:     if  $r^{(n+1)} < \delta$  then
24:       break
25:     end if
26:   end for
27:   for  $j = 1$  to  $N_s$  do
28:     Recover fields: For each  $j$ , compute  $(\tilde{\mathbf{E}}_j^{n+1}, \tilde{\mathbf{H}}_j^{n+1})$  using Eqs. (13)–(14).
29:   end for
30: end for
31:
32: Compute QoIs: Compute  $(\rho_{xy}, \phi_{xy}) - (\tilde{\rho}_{yx}, \phi_{yx})$  using Eqs. (19)–(20).

```

In the present implementation, the domain Ω is partitioned into N_s horizontal subdomains Ω_j with $ngx \times ngy \times nz$ elements, where $nz = ngz/N_s$ as illustrated in Fig. 1. Increasing the number of subdomains (N_s) reduces the computational burden of the global full-order solver and typically leads to faster solves due to improved parallelism and faster local iterations.




Figure 1. Computational domain Ω partitioned into N_s horizontal subdomains Ω_j , where each subdomain contains $ngx \times ngy \times nz$ finite elements.

In the reduced-order setting, however, the cost of solving each subproblem is relatively small, and the overall efficiency is dominated by the number of iterations required for convergence. As a result, using too many subdomains can be detrimental to the performance of the reduced-order solver. Conversely, using a single subdomain eliminates the need for an iterative coupling scheme, making the reduced-order solver particularly efficient. However, the corresponding full-order solver exhibits poor scalability with increasing problem size, both in terms of computational time and memory requirements. For this work, an optimal number of subdomains was chosen based on tests, providing a balanced trade-off that alleviates the scalability limitations of the global full-order solver (e.g. Manassero, 2019) while avoiding excessive iteration overheads associated with a large number of subdomains (Zyserman and Santos, 2000).

4.2 Parallelization

The algorithm is designed to be parallelized over np processors using the Message Passing Interface (MPI) (Walker and Dongarra, 1996).

In this implementation, a total of $np = 2 \times N_s$ processes are employed to achieve two levels of parallelism (i) at subdomain level, two processes solve the XY and YX polarizations simultaneously; (ii) at the global level, all subdomains are processed in parallel.

Assigning a single processor per subdomain would allow a larger number of subdomains and, consequently, faster local matrix factorizations. However, this configuration requires the XY and YX polarizations to be solved sequentially and typically leads to an increased number of iterations, which is a major drawback in terms of POD-DD computational efficiency. As a result, the adopted configuration (two processors per subdomain) allows the simultaneous solution of both polarizations and yields significantly better overall performance.



Table 1. Summary of the solvers nomenclature and corresponding descriptions.

Nomenclature	Description
G (Zyserman and Santos, 2000)	Full-order global solver ($N_s = 1$)
G-POD (Manassero, 2019)	Reduced-order global solver ($N_s = 1$)
DD (Zyserman and Santos, 2000)	Full-order DD solver
DD-POD (New)	Reduced-order DD solver

Communication between subdomains occurs only across the outer faces of each Ω_j (i.e. along the interfaces Γ_{jk}), which coincide with the interfaces Γ_{kj} of neighboring subdomains Ω_k assigned to other MPI processes. This is because the (RHS) terms updated during the iterative scheme depend exclusively on the solution at these interfaces, no information from the interior of neighboring subdomains is required.

4.2.1 Reduced basis

The reduced-order model is constructed once during an *offline* phase (details are given in Section 5.1.2). A reduced basis, \mathbf{U}_j , is built for each subdomain Ω_j . These bases are constructed by applying the SVD to the snapshot matrix \mathbf{X} and retaining the corresponding matrices \mathbf{U}_j (details in Section 3.2). The snapshot matrix is formed by solving the full-order problem (Eq. (15)) n_s times until convergence, using either the global or DD solver, and extracting only the portion of each solution associated with subdomain Ω_j . Manassero et al. (2020) showed that building independent bases for each polarization and frequency results in smaller basis; here we follow the same approach.

The resulting snapshots are stored locally, ensuring that each processor has access to the data relevant to its subdomain and polarization; a processor associated with a given subdomain Ω_j never requires information from the bases or snapshots of other subdomains Ω_k . During execution, the only quantities exchanged between processors are the full-order approximations of interface variables ϵ_{jk}^{n+1} and λ_{jk}^{n+1} .

4.2.2 Solver variants

Four solver implementations are considered: two for the full-order problem and two for the reduced-order problem. Hereafter, the single-domain full-order solver is denoted by G, while the multi-domain solver based on domain decomposition is denoted by DD. The corresponding reduced-order formulations are referred to as G-POD and DD-POD, respectively.

While the G method requires only a single solve of Eq. (15) ($N_s = 1$), the DD approach requires multiple solves of Eq. (15) until a prescribed convergence criterion is satisfied (Section 2.2.4). Consequently, different parallelization strategies are used for the G and DD solvers. In G, parallelization is achieved through the global linear solver, which distributes the operations required to factorize and solve the full system across MPI processes (see e.g. Manassero et al. (2020)). In contrast, DD employs two levels of MPI-based parallelization, as described in Subsection 4.2. Importantly, when multiple frequencies are considered, an additional level of parallelization can be readily introduced, in which groups of processors are assigned to different



Table 2. Dimensions and resistivities of the bodies present in the DTM1 model.

	x (km)	y (km)	z (km)	ρ (Ω m)
body 1	$[-20, 20]$	$[-2.5, 2.5]$	$[5, 20]$	10
body 2	$[-15, 0]$	$[-2.5, 22.5]$	$[20, 25]$	1
body 3	$[0, 15]$	$[-22.5, 2.5]$	$[20, 50]$	10^4
background				100

frequencies (e.g. Manassero et al., 2020). Table 1 summarizes the different solver implementations and their corresponding descriptions.

315 The resulting LSE are solved using the complex symmetric solver ZMUMPS from the Multi-frontal Massively Parallel Solver (MUMPS) package (version 5.1.2; Amestoy et al., 2001, 2006). In addition, several numerical linear algebra subroutines are employed: METIS (Karypis and Kumar, 1998), LAPACK (Anderson et al., 1999), ScaLAPACK (Blackford et al., 1997), and the Basic Linear Algebra Subprograms (BLAS; Lawson et al., 1979).

5 Application to a benchmark and a real-world case study

320 This section introduces two models employed to evaluate the performance of the proposed method. The first model is a synthetic benchmark, presented to provide a controlled setting in which the solver’s performance can be assessed under progressively refined discretizations. The second model is a real-world subsurface conductivity model proposed by Martí i Castells (2006) to characterize the conductivity structure in southern Spain.

5.1 Benchmark model

325 We use the DTM1 model, originally presented and evaluated during the MT Workshop organised at the Dublin Institute for Advanced Studies (Miensopust et al., 2013). This benchmark model is widely used within the MT community to assess solver performance in the presence of strong resistivity contrasts and is publicly available through `mtnet.info`.

5.1.1 Resistivity structure and problem parameterization

330 The DTM1 comprises three resistivity anomalies embedded in a homogeneous half-space with a resistivity of 100Ω m ($\sigma_b = 0.01 \text{ S m}^{-1}$), as detailed in Table 2 and illustrated in Figs. 2–2. The parametrization of the problem accounts for the resistivity of the three anomalies, therefore, $p = 3$ parameters are considered. Two anomalies are conductive relative to the background, with $\rho_1 = 10 \Omega$ m ($\sigma_1 = 0.1 \text{ S m}^{-1}$) and $\rho_2 = 1 \Omega$ m ($\sigma_2 = 1.0 \text{ S m}^{-1}$), while the third anomaly is resistive, with $\rho_3 = 10^4 \Omega$ m ($\sigma_3 = 1.0 \times 10^{-4} \text{ S m}^{-1}$).

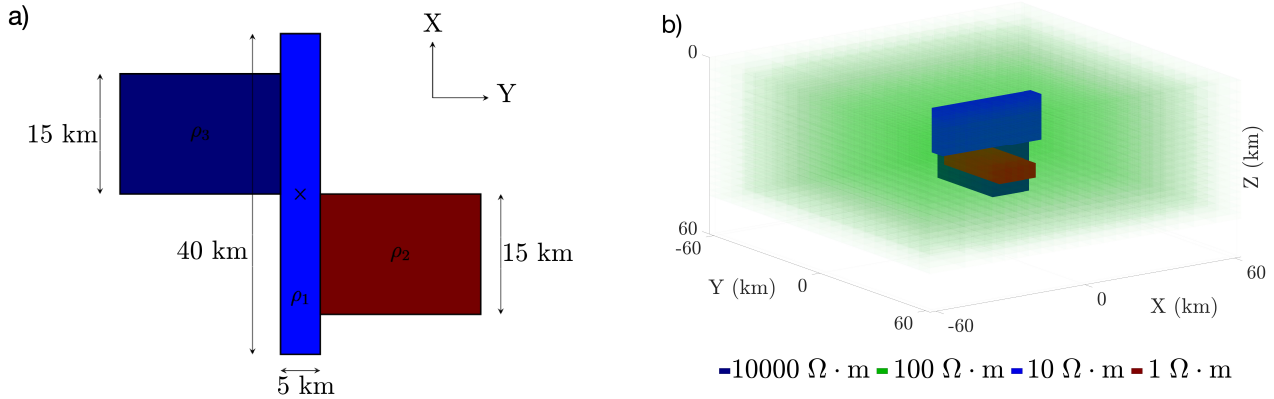


Figure 2. Resistivity model of DTM1. (a) Top view and (b) three-dimensional view of the model, showing three subsurface bodies with contrasting resistivities: a highly resistive dark-blue body ($10\,000\ \Omega\cdot\text{m}$), a moderately resistive blue body ($10\ \Omega\cdot\text{m}$), and a conductive red body ($1\ \Omega\cdot\text{m}$).

5.1.2 Offline basis construction

335 The snapshot matrix \mathbf{X} consists of 125 full-order solutions, generated by considering all possible combinations of the three conductivity parameters ρ_1 , ρ_2 , and ρ_3 , each taking values from the set $\{20000, 200, 100, 2, 0.2\}\ \Omega\cdot\text{m}$.

5.1.3 Frequency selection via skin depth

Each model is solved for a single frequency, chosen to ensure that EM fields penetrate the entire computational domain ($120\ \text{km} \times 120\ \text{km} \times 60\ \text{km}$). The penetration depth is estimated from the skin depth (Simpson and Bahr, 2005), $\sqrt{1/(\pi\mu_0\sigma f)}$,
 340 where μ is the magnetic permeability and σ is the background conductivity, taken as $\sigma = 0.01\ \text{S}\cdot\text{m}^{-1}$.

Given that the largest horizontal dimension of the model is $120\ \text{km}$, we selected a frequency of $f = 10^{-3}\ \text{Hz}$. At this frequency, the skin depth is $\approx 159\ \text{km}$, sufficiently larger than the domain extent to ensure full penetration, while still capturing the effect of smaller, more conductive anomalies within the model.

5.2 Real-world model

345 To validate the proposed method under more realistic conditions, we apply it to a representative case study that reflects the complexity and computational demands encountered in practical MT applications (Martí i Castells, 2006).

The numerical mesh used in this study follows the the model geometry of Martí i Castells (2006) but uses a finer spatial discretization. It consists of 600,000 elements ($100 \times 100 \times 60$), including one layer of air, as illustrated in Fig. 3. The computational domain extends over $240\ \text{km}$ (N-S) \times $220\ \text{km}$ (E-W) \times $91\ \text{km}$ in the north-south, east-west, and vertical directions,
 350 respectively. The top of the mesh is considered flat.

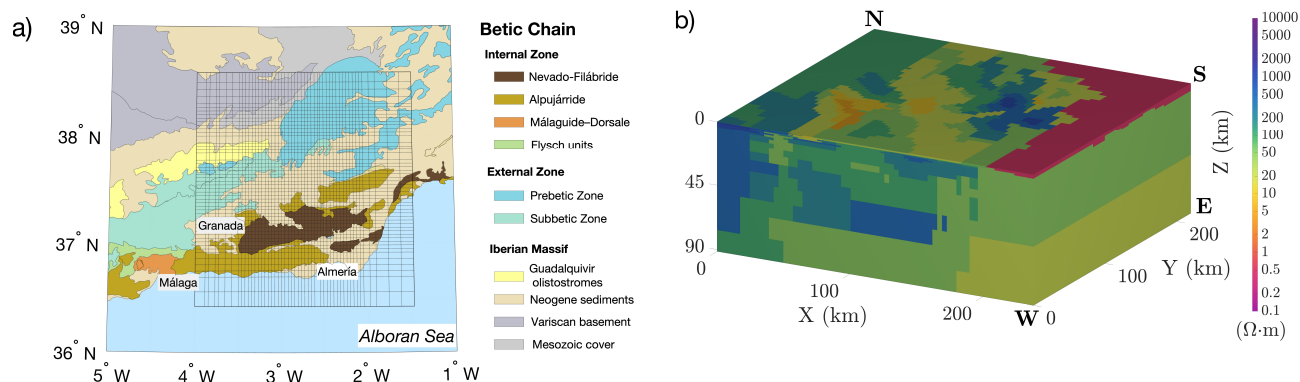


Figure 3. Resistivity model of the real-world case proposed by Martí i Castells (2006). **(a)** Superposition of the finite-element discretization used in Martí i Castells (2006) onto the geological map of the Betic Chain. **(b)** 3D resistivity model proposed by Prof. Anna Martí, obtained by trial and error, with the resistivity scale shown on the right.

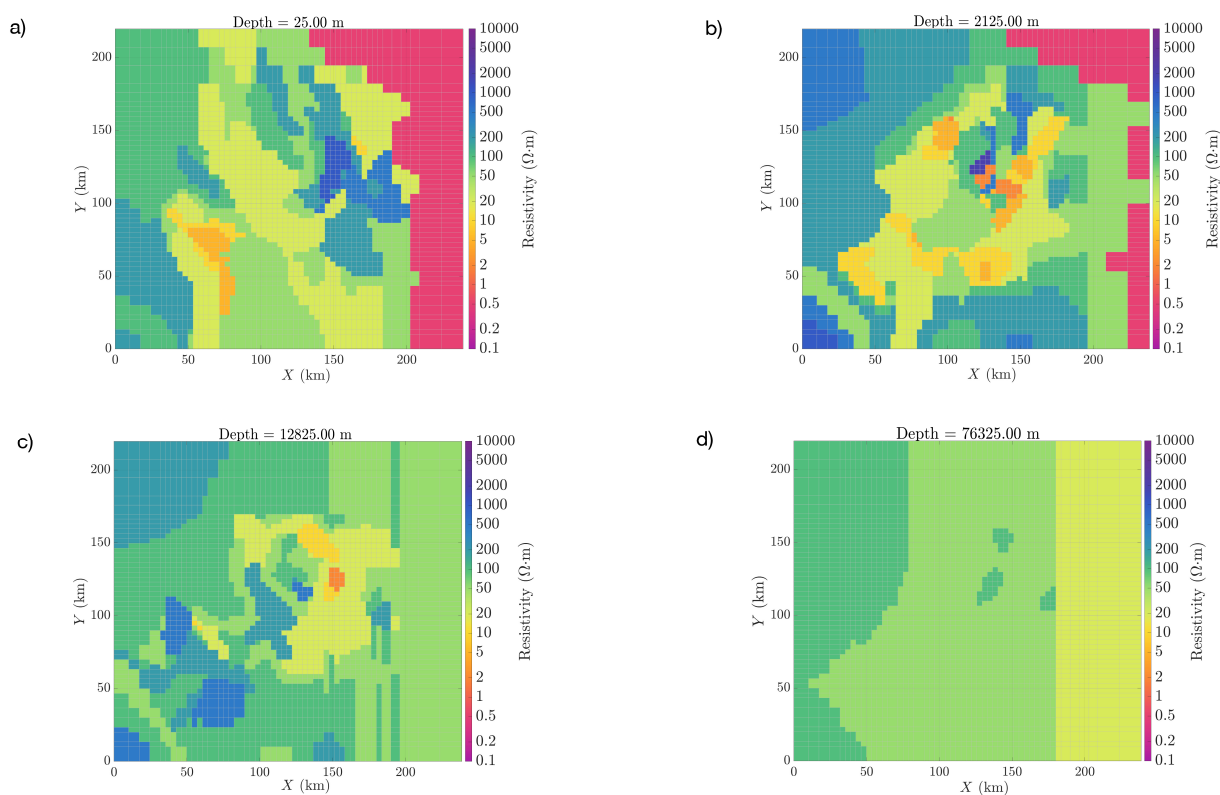


Figure 4. Horizontal slices of the conductivity model proposed by Martí i Castells (2006) for the selected portion of the Betic Chain illustrated in Fig. 3. The panels show slices at depths of **(a)** 25 m, **(b)** 2 125 m, **(c)** 12 825 m, and **(d)** 76 325 m.



5.2.1 Resistivity structure and problem parameterization

The original model was developed through an iterative manual process, starting with an initial structure derived from 1D inversions and geological priors Martí i Castells (2006). This structure was later refined via 3D forward simulations to achieve a balance between data fit and model parsimony.

355 The final resistivity model consists of 12 geological units/bodies with distinct resistivities. The most voluminous body, representing approximately 43% of the total model volume, is considered the background and assigned a constant resistivity of $\rho_b = 50 \Omega, \text{m}$ ($\sigma_b = 0.02 \text{ S, m}^{-1}$). Beneath the model domain, a uniform resistive half-space of $15,000 \Omega, \text{m}$ is assumed.

Variability is introduced only in a subset of the remaining bodies during the construction of the reduced basis. Specifically, variability is confined to the five largest non-background units (i.e. $p = 5$), while smaller bodies are kept fixed at their original
360 values.

This choice is motivated by the dominant contribution of large-volume structures to the electromagnetic response, as they control the principal modes of model variability. Smaller units, by contrast, exert a limited influence on the reduced-order representation and can therefore be held constant without significantly affecting the accuracy of the forward response.

5.2.2 Offline basis construction

365 The snapshot matrix \mathbf{X} consists of $4^5 = 1024$ full-order solutions. These solutions were generated by considering all possible combinations of the five resistivity parameters $\rho_1 - \rho_5$, where $\rho_1, \rho_2, \rho_3, \rho_4 \in \{1000, 216, 47, 10\}$ and $\rho_5 \in \{100, 21.6, 4.7, 1.0\} \Omega \text{m}$.

5.2.3 Frequency selection

Simulations were performed at 0.001 and 10 Hz. The lower frequency ensures sufficient electromagnetic field penetration into the underlying half-space, whereas the higher frequency is used to assess solver accuracy and performance in the high-
370 frequency regime.

6 Results and discussion

6.1 Benchmark results

6.1.1 Computational performance

The computational performance is evaluated in terms of the wall-clock time required to obtain the EM responses. The wall
375 time is reported for the same benchmark model using progressively refined meshes. In all cases, the reported times and errors correspond to solutions obtained using the value of β that maximized computational speed rather than accuracy (i.e., 0.05 of the value proposed in Zyserman and Santos (2000), β_{ref}). Table 3 presents the results obtained.

The size of the LSE is shown for each solver and mesh configuration. In the DD solver, the LSE is smaller because of the subdomain partitioning ($N_s = 9$), with each subdomain containing the same number of degrees of freedom (N_{dof}). Importantly,



Table 3. Computational performance in terms of wall time, relative wall-time reduction, and the size of the LSE for DD, G-POD, and DD-POD solvers compared to the G solver. All simulations were run on two servers, each equipped with 2×8-core Intel Xeon Gold 6134 CPUs (3.20 GHz, 25 MB cache, 2666 MHz FSB) and 192 GB RAM. A total of 18 MPI processes were used (9 per node). The hyperparameters were $N_s = 9$, $\beta = 0.05 \beta_{\text{ref}}$, $\delta = 5 \times 10^{-4}$, and $\delta_{\text{SVD}} = 10^{-3}$. In both DD and DD-POD, the initial guesses were set to $\mathbf{0}$.

Mesh size	Frequency (Hz)	Solver	Size LSE	Wall time (s)
40×40×36 (57 600 elements) ($N_{\text{dof}} = 354,560$)	0.001	G	354,560 × 354,560	8.83
		G-POD	40 × 40	0.93 (↓90%)
		DD	42,240 × 42,240	2.42 (↓73%)
		DD-POD	46 × 46	0.30 (↓97%)
80×80×72 (460 800 elements) ($N_{\text{dof}} = 2,800,640$)	0.001	G	2,800,640 × 2,800,640	193.78
		G-POD	43 × 43	7.38 (↓96%)
		DD	322,560 × 322,560	77.01 (↓60%)
		DD-POD	46 × 46	1.98 (↓99%)

Table 4. Computational performance in terms of wall time, relative wall-time reduction, and the size of the LSE for DD, G-POD, and DD-POD solvers compared to the G solver. All simulations were executed on three servers, each equipped with 2×14-core Intel Xeon Gold 6132 CPUs (2.60 GHz, 19 MB cache, 2666 MHz FSB) and 192 GB RAM. A total of 72 MPI processes were used (24 per node). The hyperparameters were $N_s = 36$, $\beta = 0.05 \beta_{\text{ref}}$, $\delta = 1 \times 10^{-4}$, and $\delta_{\text{SVD}} = 10^{-3}$. In both DD and DD-POD, the initial guesses were set to $\mathbf{0}$.

Mesh size	Frequency (Hz)	Solver	Size LSE	Wall time (s)
160×160×144 (3 686 400 elements) ($N_{\text{dof}} = 22,261,760$)	0.001	G	22,261,760 × 22,261,760	Out-of-Memory
		DD	668,160 × 668,160	209
		DD-POD	46 × 46	15.74 (↓92%)

380 for the reduced-order solvers, the size of the LSE can vary across polarizations (XY, YX) and subdomains. The value reported in Table 3 corresponds to the configuration with the largest number of modes (i.e. the most computationally expensive).

The time reduction reported in the “Wall Time (s)” column is expressed relative to the wall time obtained with the G solver for each model. The results indicate that (i) the DD solver is at least 70% faster than G in terms of wall-clock time; (ii) integrating POD into both implementations yields a time reduction of at least 90% and (iii) the DD-POD implementation is
 385 at least 68% faster than G-POD. It is also worth noting that the number of subdomains (N_s) influences the full-order and reduced-order solvers in different ways. While decreasing N_s (for $N_s > 1$) raises the wall time of the full-order solver (since each iteration becomes more computationally demanding), it reduces the wall time of the reduced-order solver, as each local solve is inexpensive and fewer iterations are required.

The results presented in Table 4 correspond to a large-scale case, where the mesh was refined to 160 × 160 × 144 elements.
 390 In this configuration, $N_s = 36$ subdomains were used with $np = 72$ processors. For this case, the integration of POD within the DD framework also provides a time reduction of at least 90%, consistent with the results obtained in the previous cases.

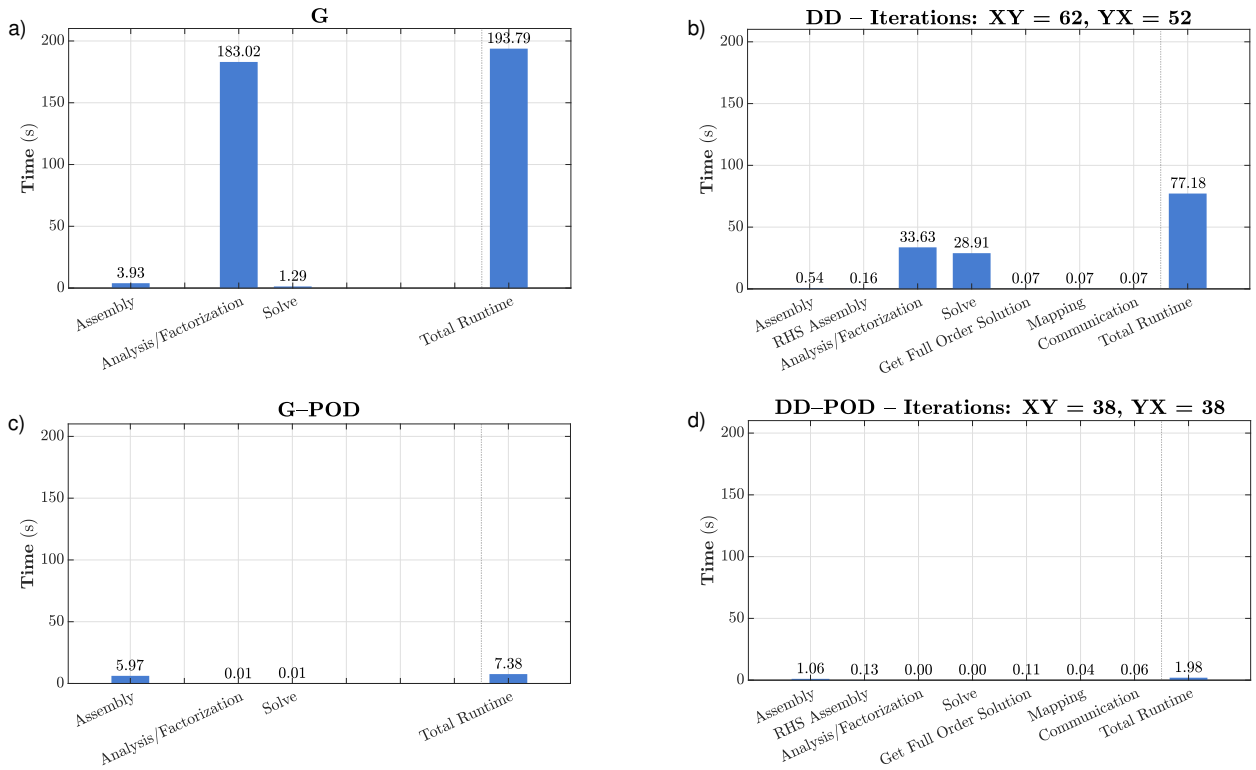


Figure 5. Execution time of the main computational steps for (a) the global full-order solver G, (b) the full-order DD solver, (c) the reduced-order global solver G-POD, and (d) the reduced-order DD solver on an $80 \times 80 \times 72$ mesh using 18 MPI processes. These simulations correspond to the results presented in Table 3. The domain decomposition solvers also display the number of iterations required for each polarization. The procedures along the horizontal axis are described in Section 6.1.1.

The main computational steps involved in the solution of the EM field equations were examined to identify bottlenecks in each solver. These results are shown in Figs. 5 and 6. The steps considered are as follows:

- "Assembly", corresponds to the time required to assemble the LHS. For the full-order cases, this refers to assembling the stiffness matrix \mathbf{K} . For the reduced-order cases, it includes assembling and multiplying matrices (as in Eq. (28)) to obtain the reduced stiffness matrix \mathbf{K}^* .
- "RHS Assembly" denotes the time required to assemble the right-hand side (RHS). For the global G solvers. This is performed only once, whereas in DD, it is repeated at every iteration. In the POD-extended solvers, this step also includes the time needed to project the RHS onto the reduced space.
- "Analysis and Factorization" refer to the time required by ZMUMPS to analyze and factorize the LHS.
- "Solve" represents the time to solve the linear system of equations.

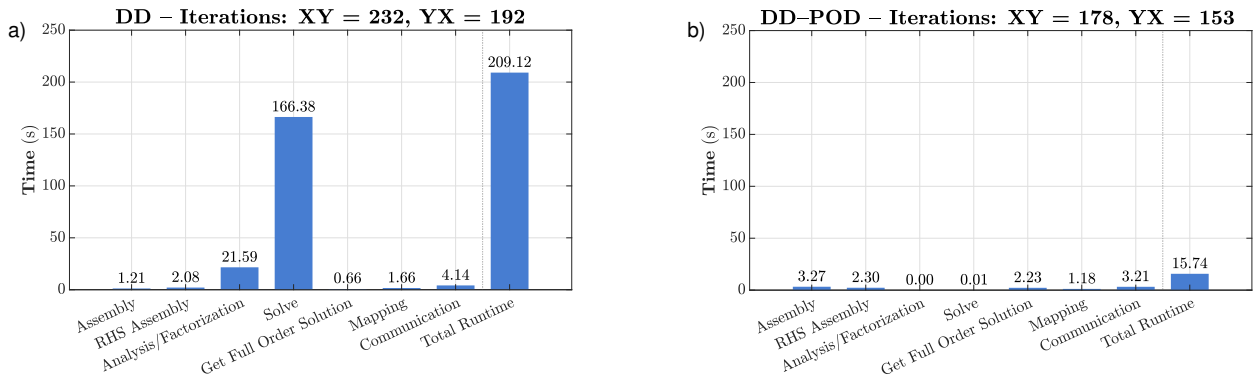


Figure 6. Execution time of the main computational steps for (a) the full-order DD solver and (b) the reduced-order DD solver on a $160 \times 160 \times 144$ mesh using 72 MPI processes. These simulations correspond to the results presented in Table 4, and the reduced-order DD solver shows the number of iterations for each polarization. The procedures along the horizontal axis are described in Section 6.1.1.

- "Get Full-Order Solution" corresponds to the time needed to retrieve the solution in the full space or store it in a new variable.
- "Mapping" accounts for the time required to map the vector solution onto the 3D tensors ε_j and λ_j .
- 405 – "Communication" represents the overhead time spent exchanging information between subdomains.

For the G solver, the dominant cost arises from the analysis and factorization of the stiffness matrix. In the G-POD solver, these steps are entirely bypassed, as the system is projected onto a reduced basis. Consequently, the time associated with the analysis, factorization, and solution of the full system becomes irrelevant. Additionally, allocating all 18 processors to solve the linear system of equations (LSE) in parallel would be inefficient. Instead, the LSE is solved using two MPI processes, one
 410 for each polarization. Within each MPI process, 9 Intel MKL threads are employed exclusively to accelerate the dense linear algebra operations, namely the matrix–matrix and matrix–vector products. Despite this strategy and the drastic reduction in total time, the dominant computational cost in the G–POD approach remains the assembly of \mathbf{K} and its projection onto the reduced space to form \mathbf{K}^* (Eq. (28)).

In the DD solver, the main limitation is the need to solve the linear systems at each iteration. In contrast, the DD-POD solver
 415 exhibits not only a considerable reduction of the total time, but also a more balanced time distribution among computational steps. The most time-consuming tasks in DD-POD are the assembly of the reduced stiffness matrix ("Assembly"), the projection of the right-hand side onto the reduced space ("RHS Assembly"), and the reconstruction of the reduced solution in the full-order space ("Get Full Order Solution").

For the large-scale test case (Fig. 6), the times associated with these steps increase rapidly, and the "Communication"
 420 cost also increases considerably. Also, a more stringent convergence criterion was required to ensure an acceptable solution, resulting in a higher number of iterations and longer computation times. Despite this, DD-POD scheme reduced the total execution time by as much as 92%.



Table 5. Maximum absolute errors, $\max \|\Delta\|$, in apparent resistivity (ρ) and phase (φ) computed by G-POD and DD-POD relative to their reference full-order counterparts, the G and DD solvers, respectively, for different mesh resolutions, at a frequency of 0.001 Hz. These results correspond to the simulations presented in Table 3.

Mesh size	Solver	$\max \ \Delta\rho_{xy}\ $	$\max \ \Delta\phi_{xy}\ $	$\max \ \Delta\rho_{yx}\ $	$\max \ \Delta\phi_{yx}\ $
40×40×36	G-POD	2.93	0.91	0.51	0.07
(57,600 elements)	DD-POD	2.39	0.47	0.67	0.13
80×80×72	G-POD	2.08	0.29	0.39	0.11
(460,800 elements)	DD-POD	1.71	0.50	1.21	0.27

Table 6. Maximum absolute errors, $\max \|\Delta\|$, in apparent resistivity (ρ) and phase (φ) computed by DD-POD relative to the reference solution obtained using the DD solver, at a frequency of 0.001 Hz. These results correspond to the simulations presented in Table 4.

Mesh size	Solver	$\max \ \Delta\rho_{xy}\ $	$\max \ \Delta\varphi_{xy}\ $	$\max \ \Delta\rho_{yx}\ $	$\max \ \Delta\varphi_{yx}\ $
160 × 160 × 144 (3,686,400 elements)	DD-POD	2.70	0.41	0.80	0.13

6.1.2 Validation and accuracy

Having assessed the computational times, we now turn to the assessment of the accuracy of the G-POD, DD, and DD-POD solvers relative to the reference solutions from the G solver.

Table 5 reports the maximum absolute error $\max \|\Delta\|$ of the QoIs (ρ_a and ϕ). Only the components (ρ_{xy}, ρ_{yx}) and (ϕ_{xy}, ϕ_{yx}) are presented. Similarly, Table 6 reports the same metric, but relative to the DD solver for the highest-resolution configuration. From these values, it is observed that the acceleration achieved by DD or DD-POD does not compromise the accuracy of the EM responses when compared with the reference solutions (i.e. G and DD solvers). We note that the reported errors correspond to the value of β that maximized computational efficiency rather than accuracy; smaller errors can be obtained by further adjusting β . Moreover, the maximum relative L_2 -norm errors between the reduced and full-order solutions (Rel L_2), remain below 3 %, which is substantially smaller than the discrepancies typically observed between independent MT solvers (e.g. Manassero, 2019).

6.2 Real-world model results

6.2.1 Computational performance

Table 7 reports the execution wall-time for the same set of simulations across different solvers for f equal to 0.001 Hz and 10 Hz. Although simulations were conducted for various n_{POD} values, the time recorded in Table 7 for the POD solvers corresponds specifically to $n_{\text{POD}} = 50$. For these simulations, the DD solver continues to outperform the G solver, achieving computational time gains of almost 90%. This staggering performance improvement arises from avoiding the solution of a global LSE with 3.644×10^6 degrees of freedom. Instead, the DD solver operates on a sequence of smaller LSEs of size



Table 7. Computational performance in terms of wall time and relative wall-time reduction with respect to the G solver for different solvers at frequencies of 0.001 Hz and 10 Hz, using a $100 \times 100 \times 60$ discretization. All simulations were run on two servers, each equipped with 2×32 -core Xeon Gold 6430 processors (2.1 GHz, 60 MB cache, 4400 MHz FSB) and 256 GB RAM. A total of 40 MPI processes was used. The hyperparameters were $N_s = 20$, $n_{\text{POD}} = 50$, $\beta = 0.05 \beta_{\text{ref}}$, and $\delta = 5 \times 10^{-4}$. The size of the LSE for the reduced-order solvers is set by n_{POD} , and the initial guess was set using the snapshot corresponding to the nearest parametric configuration.

Mesh size	Frequency (Hz)	Solver	Size LSE	Wall time (s)
$100 \times 100 \times 60$ (600 000 elements) ($N_{\text{dof}} = 3\,644\,000$)	0.001	G	$3,644,000 \times 3,644,000$	198.13
		G-POD ($n_{\text{POD}} = 50$)	50×50	7.93 (↓ 96%)
		DD	$201,200 \times 201,200$	23.16 (↓ 88%)
		DD-POD ($n_{\text{POD}} = 50$)	50×50	3.16 (↓ 98%)
	10.0	G	$3,644,000 \times 3,644,000$	190.56
		G-POD ($n_{\text{POD}} = 50$)	50×50	8.06 (↓ 96%)
		DD	$201,200 \times 201,200$	24.50 (↓ 87%)
		DD-POD ($n_{\text{POD}} = 50$)	50×50	2.61 (↓ 99%)




Figure 7. Total execution and assembly times of G-POD and DD-POD for different values of n_{POD} . The definition of assembly time is given in Section 6.1.1. Solid red lines indicate the execution time for G-POD, and dashed red lines indicate the time to assemble \mathbf{K}^* in G-POD. Solid blue lines indicate the execution time for DD-POD, and dashed blue lines indicate the time to assemble \mathbf{K}_j^* . The iterative scheme of the DD-POD solver time is represented by the distance between the solid and dashed blue lines. More details of the simulations are provided in the caption of Table 7. (a) $f = 0.001$ Hz, (b) $f = 10$ Hz.

201,200 \times 201,200, which are solved in parallel. Regarding the POD-based solvers, for an equivalent truncation level n_{POD} , the DD-POD approach is at least 60% faster than the G-POD alternative, for this truncation, demonstrating the benefits of localized reduction.

Fig. 7 illustrates the performance of the POD solvers in terms of execution wall-time across various n_{POD} truncations for the two frequencies, alongside the assembly time as defined in Sect. 6.1.1. It is seen that DD-POD consistently outperforms

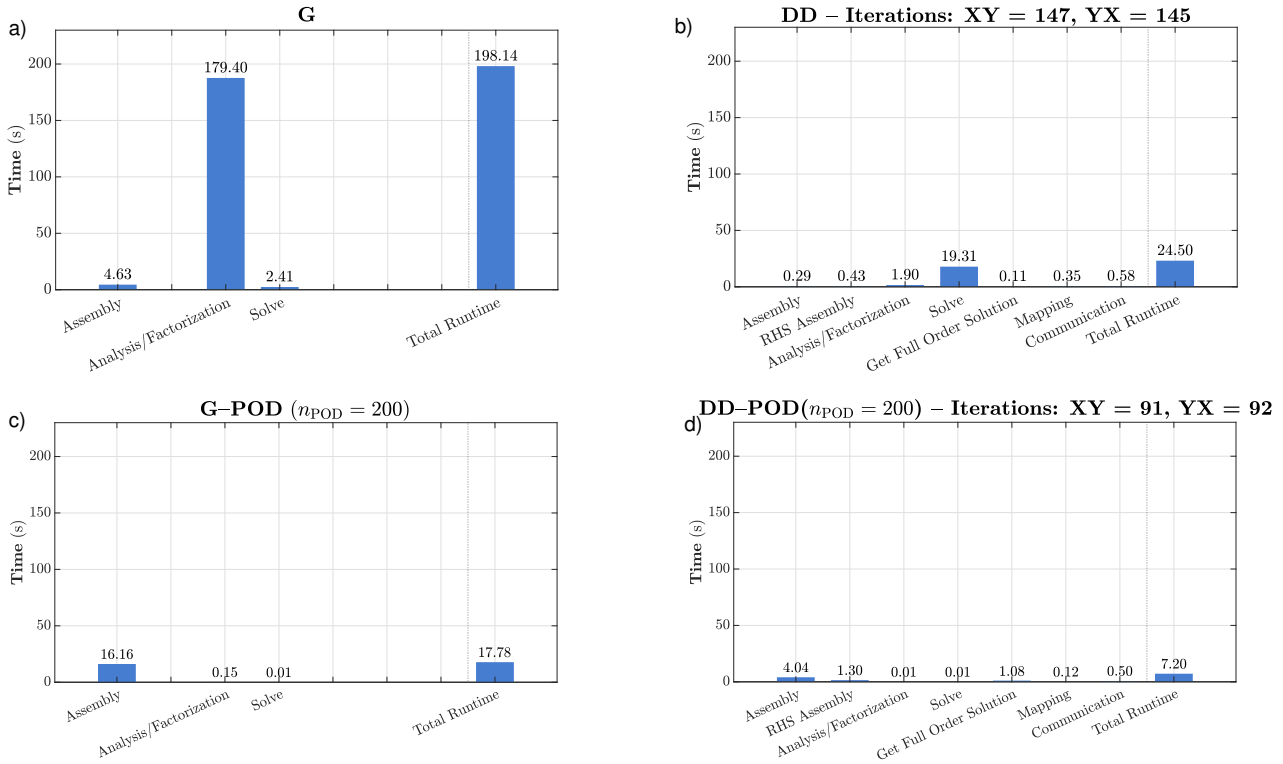


Figure 8. Execution time of the main computational steps for solvers (a) G, (b) DD, (c) G-POD ($n_{\text{POD}} = 200$), and (d) DD-POD ($n_{\text{POD}} = 200$) on a $100 \times 100 \times 60$ mesh using 40 MPI processes. These simulations correspond to the results presented in Table 7 for the full-order solvers and in Fig. 7 for the reduced-order solvers with $n_{\text{POD}} = 200$ at a frequency of 10 Hz. The DD solver plots show the number of iterations for each polarization. The procedures along the horizontal axis are described in Section 6.1.1.

G-POD in computational efficiency for every truncation level tested. In the G-POD framework, the assembly time defined as the time required to assemble \mathbf{K} and perform its projection onto the reduced space accounts for more than 90% of the total simulation time and escalates rapidly as n_{POD} increases. The disparate scaling behaviors observed between the G-POD and DD-POD solvers are primarily attributable to fundamental differences in their respective projection operations.

450 In the G-POD approach, projecting the operator \mathbf{K} onto the reduced basis requires operations that fail to exploit the inherent framework of domain decomposition. Specifically, since \mathbf{K} contains approximately $18^2 = 324$ times more entries than a local subdomain operator \mathbf{K}_j , the computational overhead involved in the projection is significantly higher. As the number of POD modes, n_{POD} , increases, the cost of these projections scales quadratically, eventually dominating the total wall-time. This assembly bottleneck can negate the theoretical advantages of reduced-order modeling, particularly when the overhead of
 455 constructing reduced operators exceeds the savings gained from solving the LSE. For instance, at $n_{\text{POD}} = 200$, the execution time of G-POD becomes comparable to the total execution time of the full DD solver, rendering the reduction less effective.



Conversely, the assembly time for DD-POD scales at a slower pace, maintaining a considerably lower baseline even as the basis size grows. The remainder of the execution time is allocated to the iterative scheme, representing the interval between the two blue lines in the visualization. By reducing the number of iterations through better starting conditions, the execution time of DD-POD can be further decreased, approaching the theoretical minimum of the assembly baseline. For this case, the initial guess for the iterative scheme was set using the snapshot with the closest parametric combination, resulting in a 30% decrease in execution time compared to a null initialization. The choice of the initial guess directly dictates the number of iterations required for convergence, thereby determining the overall speed of the simulation. This bears important implications for the acceleration of MCMC-based inversions (e.g. Manassero et al., 2020, 2021); we discuss this further in Section 6.3.

Finally, Fig. 8 displays the timing of the main computational steps for $f = 10$ Hz. Both POD implementations successfully eliminate the massive computational overhead associated with analyzing and solving the full-scale LSE. In the G-POD solver, the computational cost comprising the assembly of \mathbf{K} and its subsequent projection, rises from 4.63 s to 16 s for $n_{\text{POD}} = 200$. Such a high baseline compromises the viability of the G-POD solver, particularly in high-resolution probabilistic inversion frameworks that require thousands of forward evaluations. In fact, the time required to assemble \mathbf{K} alone in the G-solver exceeds the combined assembly and projection time of the DD-POD solver, proving a markedly more favorable scaling of the latter as the number of retained modes increases. Moreover, the total computational cost of the DD-POD solver remains essentially independent of the frequency considered.

6.2.2 Validation and accuracy

The validation of the full-order approximations and the assessment of their accuracy are carried out using the relative L_2 -norm ($\text{Rel}L_2$) error on the quantities of interest (QoIs), rather than the maximum pointwise difference $\max\|\Delta\|$ used in the benchmark study. This choice of error metric is motivated by the strong heterogeneity of the model, which features material resistivities ranging from 0.1 to 5000 $\Omega \cdot \text{m}$. For the phase components, the denominator of the relative L_2 -norm was shifted to the first quadrant to ensure a fair comparison.

In Figs. 9 and 10, the $\text{Rel}L_2$ is displayed for simulations with different modes (n_{POD}) for both G-POD and DD-POD, with the error measured relative to their full-order counterparts, G and DD, for $f = 0.001$ Hz and $f = 10$ Hz, respectively.

For $f = 0.001$ Hz (Fig. 9), the maximum stabilized error for G-POD across the four QoIs is less than 1%, whereas for DD-POD it is approximately 5%. For $f = 10$ Hz (Fig. 10), G-POD exhibits a maximum $\text{Rel}L_2$ error of 2%, while DD-POD reaches approximately 5%, with the error stabilizing between 50 and 75 modes.

As the number of POD modes increases beyond 75, the incremental improvement in accuracy becomes marginal, suggesting that the POD basis has successfully captured the dominant physical features of the parametric problem. While G-POD provides a slightly better approximation, the substantial gain in computational efficiency (up to 70%) achieved by the DD-POD constitutes a decisive advantage, particularly in high-resolution probabilistic inversion frameworks where overall feasibility is strongly constrained by forward modeling cost. As illustrated in the middle profiles along the middle north-south section in Fig. 11 ($f = 0.001$ Hz) and Fig. 12 ($f = 10$ Hz), the differences between the G, G-POD, DD, and DD-POD solvers are almost imperceptible when plotted using standard MT data visualization scales. In this context, we emphasize that the above




Figure 9. Relative L^2 error ($RelL_2$) of full-order approximations for: (a) ρ_{xy} , (b) ϕ_{xy} , (c) ρ_{yx} , and (d) ϕ_{yx} . Results compare G-POD and DD-POD against full-order G and DD solutions for varying n_{POD} at $f = 0.001$ Hz.

errors, while within typical uncertainty bounds in MT data, are not intrinsic to the DD-POD approach. The hyperparameters ($\beta = 0.05 \times \beta_{ref}, \delta = 5 \times 10^{-4}$) chosen here for DD-POD were selected to maximize computational speed, as in the benchmark problem. For applications requiring higher accuracy, the convergence criteria can be tightened and more iterations can be performed if necessary.

495 6.3 Implications for probabilistic inversions

Manassero et al. (2020, 2021, 2024) demonstrated the application of reduced-order modeling for the inversion of 3D MT data within an MCMC-driven, fully probabilistic framework. However, the monolithic solution of the reduced system, together with the comparatively large number of required basis functions, remained a challenge, constraining the effective dimensionality that could be addressed in practical applications, as seen in Fig. 7.

500 The improvements introduced in this work impact MCMC-driven inversions of MT data (as formulated by Manassero et al., 2020) in two important respects. First, during the initial stages of the inversion (or in previous offline stage), when the reduced basis is constructed on the fly, the use of faster and more memory-efficient DD full-order solver reduces considerably the associated computational cost. This reduction is critical because the initial basis construction accounts for a substantial

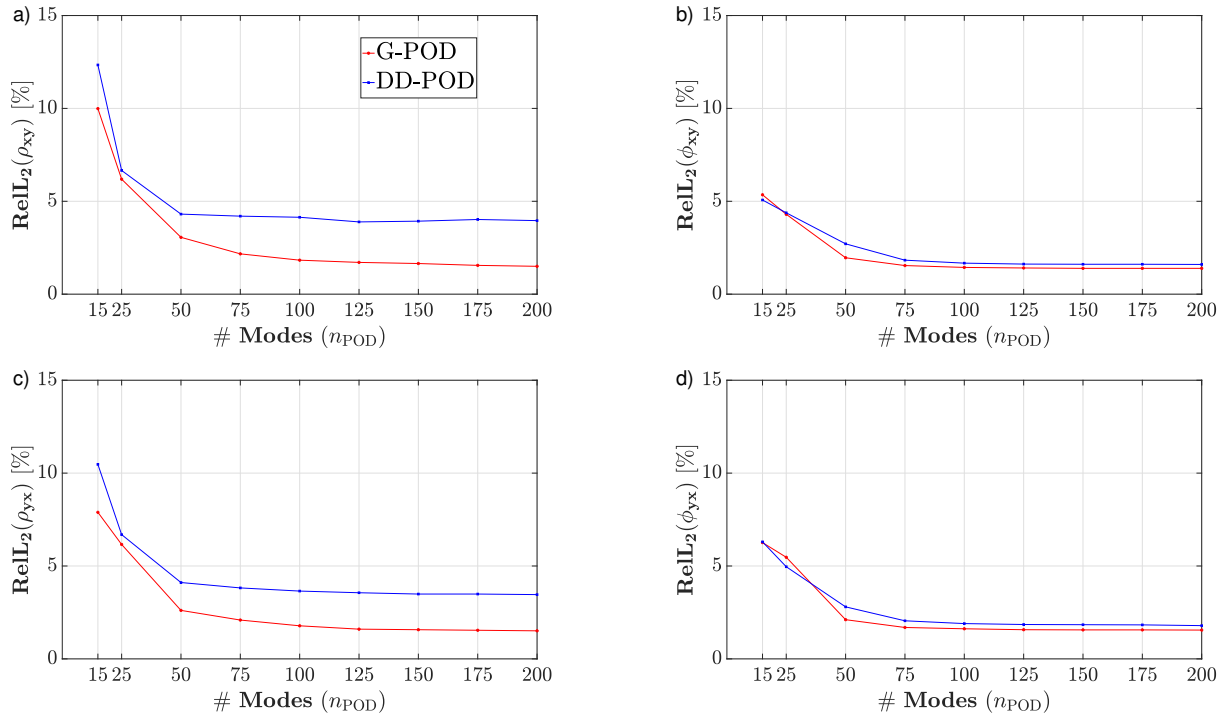


Figure 10. Relative L^2 error ($\text{Rel}L_2$) of full-order approximations for: (a) ρ_{xy} , (b) ϕ_{xy} , (c) ρ_{yx} , and (d) ϕ_{yx} . Results compare G-POD and DD-POD against full-order G and DD solutions for varying n_{POD} at $f = 10$ Hz.

fraction of the total computational cost of the inversion workflow. Second, the combined DD-POD reduced solver exhibits superior memory efficiency and shorter run times than the G-POD approach, thereby accelerating the MCMC inversion and enabling the treatment of larger problems.

Furthermore, the specific MCMC strategy employed by Manassero et al. (2020, 2021), where only a small block of the parameter vector is changed per proposal, creates a particularly favourable setting for the DD-POD approach. As the Markov chain approaches its stationary regime, sampling transitions from an initial exploration of the broad prior support to local exploration of the high-probability region of the posterior distribution. As a result, successive accepted samples tend to be increasingly close in parameter space. In this regime, using the solution from the previous sample as the initial guess for the current iterative solve reduces the number of iterations required, such that the total computational cost of the DD-POD solver approaches that of matrix assembly alone (see Fig. 7).

Based on the results presented here, and considering further potential optimizations of the LSE solvers (e.g. by exploiting affine decomposition), we anticipate that MCMC-driven probabilistic inversions of 3D MT data will become computationally feasible for discretizations on the order of 100^3 elements. This resolution is comparable to, or exceeds, that commonly adopted in contemporary deterministic MT inversions.




Figure 11. QoIs (a) ρ_{xy} , (b) ϕ_{xy} , (c) ρ_{yx} , and (d) ϕ_{yx} along the central N–S section ($100 \times 100 \times 60$ elements). Results compare full-order (G, DD) and reduced-order (G–POD, DD–POD) solvers using the hyperparameters presented in Table 7 at $f = 0.001$ Hz.

Finally, the results presented here also have important implications for uncertainty quantification in deterministic inversion frameworks. For example, once a deterministic optimal model has been obtained, a short and efficient sampling in its vicinity can be used to generate the snapshots necessary to construct a reduced-order surrogate. This surrogate can then be exploited to perform inexpensive Monte Carlo simulations, enabling the estimation of global uncertainty measures and posterior covariance matrices that are typically inaccessible in standard deterministic inversion approaches.

7 Conclusions

This study introduced a new forward modeling framework for 3D MT problems that combines Domain Decomposition (DD) with Reduced-Order Modeling (ROM) based on Proper Orthogonal Decomposition (POD). The method was developed to overcome the memory and scalability limitations of current global full-order and reduced-order solvers.

The results demonstrate that DD significantly reduces the computational and memory costs associated with large global linear systems, enabling efficient forward simulations at resolutions comparable to those used in modern deterministic MT inversions ($\gtrsim 2 \times 10^7$ degrees of freedom). When combined with POD, the DD–POD formulation achieves substantial additional speedups while maintaining errors in apparent resistivity and phase below levels typically observed between independent MT




Figure 12. QoIs (a) ρ_{xy} , (b) ϕ_{xy} , (c) ρ_{yx} , and (d) ϕ_{yx} along the central N–S section ($100 \times 100 \times 60$ elements). Results compare full-order (G, DD) and reduced-order (G–POD, DD–POD) solvers using the hyperparameters presented in 7 at $f = 10$ Hz.

solvers and/or in MT field data. A key outcome is the favourable scaling behaviour of the DD–POD solver with increasing reduced-basis size. Unlike global POD approaches, whose efficiency rapidly deteriorates due to operator assembly and projection costs, the localized structure of DD–POD preserves computational efficiency even for moderately large numbers of modes.

535 The proposed framework has direct implications for probabilistic inversion strategies based on MCMC sampling. In particular, the DD-POD framework can accelerate MCMC-driven probabilistic inversions through two complementary mechanisms. First, the use of a DD-based full-order solver substantially reduces the computational cost associated with reduced-basis construction, which typically represents a major fraction of the total inversion expense. Second, the DD–POD reduced solver exhibits lower memory requirements and shorter run times than its global counterparts, enabling faster forward evaluations
 540 throughout the sampling process. Together, these two effects significantly reduce the overall computational burden of MCMC-based MT inversions and allow larger and more finely discretized models to be considered in practice.

Beyond fully probabilistic inversions, the methodology also opens new perspectives for uncertainty quantification in deterministic workflows. Local sampling around an optimal deterministic solution can be used to construct reduced-order surrogates



at limited cost, enabling Monte Carlo–based uncertainty analyses that are typically inaccessible in standard deterministic in-
545 version approaches.

Future work will focus on further reducing the cost of reduced-operator assembly, improving the convergence of DD–POD
iterations through enhanced initialization strategies, and fully integrating the method within probabilistic inversion frameworks.
Achieving such integration will require addressing additional challenges, including adaptive basis-enrichment strategies, fre-
quency parallelization and more advanced parameterizations of the *a priori* model (e.g. Manassero et al., 2020, 2021, 2024).
550 These developments form the focus of a forthcoming publication.

Code and data availability. The current version of model name is available from the project website <https://github.com/ltao99/DD-POD/tree/v1.0.0> under the MIT licence. The exact version of the model used to produce the results used in this paper is archived on repository under DOI (<https://doi.org/10.5281/zenodo.19567027>), as are input data and scripts to run the model and produce the plots for all the simulations presented in this paper (Tao and Zyserman, 2026)

555 *Author contributions.* SZ was responsible for conceptualization. FZ contributed to software, methodology and provided supervision. AM assisted with conceptualization and validation. LT performed conceptualization, methodology development, software implementation and optimization, validation, visualization, and writing. JC contributed to conceptualization, validation, supervision and writing (review and editing).

Competing interests. The contact author has declared that none of the authors has any competing interests.

560 *Disclaimer.* Publisher’s note: Copernicus Publications remains neutral with regard to jurisdictional claims made in the text, published maps, institutional affiliations, or any other geographical representation in this paper. While Copernicus Publications makes every effort to include appropriate place names, the final responsibility lies with the authors.

Acknowledgements. This work was partially supported by the Marie Skłodowska-Curie Actions (Doctoral Network with Grant agreement No. 101120556), the Spanish Ministry of Science MCIU/AEI/10.13039/501100011033/ FEDER, UE (Grant agreement No. PID2023-148952OB-I00 to AM; PID2023-153082OB-I00 to PD and SZ) and the Generalitat de Catalunya (Grant agreement No. 2021-SGR-01049). AM is Fellow of the Serra Hünter Programme of the Generalitat de Catalunya. The Red Española de Supercomputación has provided computational resources that supported the development of this work through grant AECT-2025-2-0008. The authors express their gratitude to Dr. Constanza Manassero for contributing with two routines used in the stiffness matrix assembly, and to Professor Anna Martí (University of Barcelona) for providing the model and field measurements of the Betic Chain, which enabled testing of the solver in a real-world
570 application.



References

- Afonso, J., Fullea, J., Griffin, W., Yang, Y., Jones, A., D. Connolly, J., and O'Reilly, S.: 3-D multiobservable probabilistic inversion for the compositional and thermal structure of the lithosphere and upper mantle. I: A priori petrological information and geophysical observables, *Journal of Geophysical Research: Solid Earth*, 118, 2586–2617, <https://doi.org/10.1002/jgrb.50124>, 2013a.
- 575 Afonso, J. C., Fullea, J., Yang, Y., Connolly, J., and Jones, A.: 3-D multi-observable probabilistic inversion for the compositional and thermal structure of the lithosphere and upper mantle. II: General methodology and resolution analysis, *Journal of Geophysical Research: Solid Earth*, 118, 1650–1676, <https://doi.org/10.1002/jgrb.50123>, 2013b.
- Afonso, J. C., Rawlinson, N., Yang, Y., Schutt, D. L., Jones, A. G., Fullea, J., and Griffin, W. L.: 3-D multiobservable probabilistic inversion for the compositional and thermal structure of the lithosphere and upper mantle: III. Thermochemical tomography in the Western-Central
580 US, *Journal of Geophysical Research: Solid Earth*, 121, 7337–7370, <https://doi.org/10.1002/2016JB013049>, 2016.
- Amestoy, P. R., Duff, I. S., L'Excellent, J.-Y., and Koster, J.: A fully asynchronous multifrontal solver using distributed dynamic scheduling, *SIAM Journal on Matrix Analysis and Applications*, 23, 15–41, <https://doi.org/10.1137/S089547989935819X>, 2001.
- Amestoy, P. R., Guermouche, A., L'Excellent, J.-Y., and Pralet, S.: Hybrid scheduling for the parallel solution of linear systems, *Parallel computing*, 32, 136–156, <https://doi.org/10.1016/j.parco.2005.07.004>, 2006.
- 585 Anderson, E., Bai, Z., Bischof, C., Blackford, L. S., Demmel, J., Dongarra, J., Du Croz, J., Greenbaum, A., Hammarling, S., McKenney, A., et al.: LAPACK users' guide, SIAM, <https://doi.org/10.1137/1.9780898719604>, 1999.
- Arnold, D. N. and Brezzi, F.: Mixed and nonconforming finite element methods: implementation, postprocessing and error estimates, *ESAIM: Mathematical Modelling and Numerical Analysis*, 19, 7–32, <https://doi.org/10.1051/m2an/1985190100071>, 1985.
- Aster, R. C., Borchers, B., and Thurber, C. H.: Parameter estimation and inverse problems, Elsevier, [https://doi.org/10.1016/C2016-0-04314-](https://doi.org/10.1016/C2016-0-04314-5)
590 5, 2018.
- Avdeev, D. and Avdeeva, A.: 3D magnetotelluric inversion using a limited-memory quasi-Newton optimization, *Geophysics*, 74, F45–F57, <https://doi.org/10.1190/1.3114023>, 2009.
- Bai, N., Zhou, J., Hu, X., Han, B., Zhou, J., Yang, J., and Zhao, H.: Fast three-dimensional inversion of magnetotelluric data based on L-BFGS optimization, *Journal of Applied Geophysics*, 242, 105–119, <https://doi.org/10.1016/j.jappgeo.2021.104385>, 2025.
- 595 Berkooz, G., Holmes, P., and Lumley, J. L.: The proper orthogonal decomposition in the analysis of turbulent flows, *Annual review of fluid mechanics*, 25, 539–575, <https://doi.org/10.1146/annurev.fl.25.010193.002543>, 1993.
- Bissig, F., Khan, A., and Giardini, D.: Joint inversion of PP and SS precursor waveforms and Rayleigh wave phase velocities for global mantle transition zone structure, *Geophysical Journal International*, 233, 316–337, <https://doi.org/10.1093/gji/ggab345>, 2023.
- Blackford, L. S., Choi, J., Cleary, A., D'Azevedo, E., Demmel, J., Dhillon, I., Dongarra, J., Hammarling, S., Henry, G., Petitet, A., et al.:
600 ScaLAPACK users' guide, SIAM, <https://doi.org/10.1137/1.9780898719642>, 1997.
- Cagniard, L.: Basic theory of the magneto-telluric method of geophysical prospecting, *Geophysics*, 18, 605–635, <https://doi.org/10.1190/1.1437915>, 1953.
- Chase, B., Unsworth, M., Atekwana, E., Evans, R., and Zhu, J.: Magnetotelluric imaging of the lithospheric structure of the southern Oklahoma Aulacogen: Evidence for long-term weakening caused by rifting, *Journal of Geophysical Research: Solid Earth*, 128, e2023JB026555, <https://doi.org/10.1029/2023JB027471>, 2023.
- 605 Chave, A. D. and Jones, A. G.: The magnetotelluric method: Theory and practice, Cambridge University Press, <https://doi.org/10.1017/CBO9781139088015>, 2012.



- deGroot Hedlin, C. and Constable, S.: Occam's inversion to generate smooth, two-dimensional models from magnetotelluric data, *Geophysics*, 55, 1613–1624, <https://doi.org/10.1190/1.1442813>, 1990.
- 610 Douglas, J., Santos, J., and Sheen, D.: A Nonconforming Mixed Finite Element Method for Maxwell's Equations, *Mathematical Models and Methods in Applied Sciences*, 10, 593–613, <https://doi.org/10.1142/S021820250000032X>, 2000.
- Douglas Jr, J., Furtado, F., and Pereira, F.: On the numerical simulation of waterflooding of heterogeneous petroleum reservoirs, *Computational Geosciences*, 1, 155–190, <https://doi.org/10.1023/A:1011546114170>, 1997.
- Duijndam, A.: Bayesian estimation in seismic inversion. Part i: Principles 1, *Geophysical Prospecting*, 36, 878–898, <https://doi.org/10.1111/j.1365-2478.1988.tb02198.x>, 1988.
- 615 Egbert, G. D. and Kelbert, A.: Computational recipes for electromagnetic inverse problems, *Geophysical Journal International*, 189, 251–267, <https://doi.org/10.1111/j.1365-246X.2011.05347.x>, 2012.
- Elías, M. W., Zyserman, F. I., Rosas-Carbajal, M., and Manassero, M. C.: Three-dimensional modelling of controlled source electro-magnetic surveys using non-conforming finite element methods, *Geophysical Journal International*, 229, 1133–1151, <https://doi.org/10.1093/gji/ggab157>, 2022.
- 620 Gauzellino, P. M., Zyserman, F. I., and Santos, J. E.: Nonconforming finite element methods for the three-dimensional Helmholtz equation: iterative domain decomposition or global solution?, *Journal of Computational Acoustics*, 17, 159–173, [https://doi.org/10.1016/S0045-7825\(00\)00319-X](https://doi.org/10.1016/S0045-7825(00)00319-X), 2009.
- Grayver, A. V.: Parallel three-dimensional magnetotelluric inversion using adaptive finite-element method. Part I: theory and synthetic study, *Geophysical Journal International*, 202, 584–603, <https://doi.org/10.1093/gji/ggv165>, 2015.
- 625 Heinson, G., Didana, Y., Soeffky, P., Thiel, S., and Wise, T.: The crustal geophysical signature of a world-class magmatic mineral system, *Scientific reports*, 8, 10 608, <https://doi.org/10.1038/s41598-018-29016-2>, 2018.
- Heise, W., Caldwell, T., Bibby, H. M., and Bannister, S.: Three-dimensional modelling of magnetotelluric data from the Rotokawa geothermal field, Taupo Volcanic Zone, New Zealand, *Geophysical Journal International*, 173, 740–750, <https://doi.org/10.1111/j.1365-246X.2008.03737.x>, 2008.
- 630 Jones, A. G.: Imaging the continental upper mantle using electromagnetic methods, *Lithos*, 48, 57–80, <https://doi.org/10.1023/A:1006652328404>, 1999.
- Karypis, G. and Kumar, V.: A fast and high quality multilevel scheme for partitioning irregular graphs, *SIAM Journal on scientific Computing*, 20, 359–392, <https://doi.org/10.1137/S1064827595287997>, 1998.
- 635 Key, K.: MARE2DEM: a 2-D inversion code for controlled-source electromagnetic and magnetotelluric data, *Geophysical Journal International*, 207, 571–588, <https://doi.org/10.1093/gji/ggw054>, 2016.
- Khan, A., Connolly, J., Maclennan, J., and Mosegaard, K.: Joint inversion of seismic and gravity data for lunar composition and thermal state, *Geophysical Journal International*, 168, 243–258, <https://doi.org/10.1111/j.1365-246X.2006.03200.x>, 2007.
- Lawson, C. L., Hanson, R. J., Kincaid, D. R., and Krogh, F. T.: Basic linear algebra subprograms for Fortran usage, *ACM Transactions on Mathematical Software (TOMS)*, 5, 308–323, <https://doi.org/10.1145/355841.355847>, 1979.
- 640 Ling, W., Pan, K., Zhang, J., He, D., Zhong, X., Ren, Z., and Tang, J.: A 3-D magnetotelluric inversion method based on the joint data-driven and physics-driven deep learning technology, *IEEE Transactions on Geoscience and Remote Sensing*, 62, 1–13, <https://doi.org/10.1109/TGRS.2023.3276136>, 2024.
- Mackie, R. L., Madden, T. R., and Wannamaker, P. E.: Three-dimensional magnetotelluric modeling using difference equations; theory and comparisons to integral equation solutions, *Geophysics*, 58, 215–226, <https://doi.org/10.1190/1.1443407>, 1993.
- 645



- Manassero, M. C.: A Reduced Order Approach for Probabilistic Inversions of 3D Magnetotelluric, Phd thesis, Macquarie University, <https://doi.org/10.25949/19441073>, 2019.
- Manassero, M. C., Afonso, J. C., Zyserman, F., Zlotnik, S., and Fomin, I.: A reduced order approach for probabilistic inversions of 3-D magnetotelluric data I: general formulation, *Geophysical Journal International*, 223, 1837–1863, <https://doi.org/10.1093/gji/ggaa415>, 650 2020.
- Manassero, M. C., Afonso, J. C., Zyserman, F. I., Jones, A., Zlotnik, S., and Fomin, I.: A reduced order approach for probabilistic inversions of 3D magnetotelluric data II: Joint inversion of MT and surface-wave data, *Journal of Geophysical Research: Solid Earth*, 126, e2021JB021962, <https://doi.org/10.1029/2019JB019315>, 2021.
- Manassero, M. C., Özaydın, S., Afonso, J. C., Shea, J., Ezad, I., Kirkby, A., Thiel, S., Fomin, I., and Czarnota, K.: Lithospheric structure and melting processes in southeast Australia: new constraints from joint probabilistic inversions of 3D magnetotelluric and seismic data, *Journal of Geophysical Research: Solid Earth*, 129, e2023JB028257, <https://doi.org/10.1029/2023JB027373>, 655 2024.
- Martí i Castells, A.: Magnetotelluric Investigation of Geoelectrical Dimensionality and Study of the Central Betic Crustal Structure, A, Ph.D. thesis, Universitat de Barcelona, 2006.
- Miensopust, M. P., Queralt, P., Jones, A. G., and modellers, D. M.: Magnetotelluric 3-D inversion—A review of two successful workshops on forward and inversion code testing and comparison, *Geophysical Journal International*, 193, 1216–1238, <https://doi.org/10.1093/gji/ggt066>, 660 2013.
- Mosegaard, K. and Tarantola, A.: Monte Carlo sampling of solutions to inverse problems, *Journal of Geophysical Research: Solid Earth*, 100, 12431–12447, <https://doi.org/10.1029/94JB03097>, 1995.
- Muixí, A., García-González, A., Zlotnik, S., and Díez, P.: Linear and nonlinear dimensionality reduction of biomechanical models, *Reduced Order Models for the Biomechanics of Living Organs*, pp. 23–44, https://doi.org/10.1007/978-3-030-43040-5_1, 665 2023.
- Muixí, A., Zlotnik, S., Giacomini, M., and Díez, P.: Data augmentation for the POD formulation of the parametric laminar incompressible navier–stokes equations, *International Journal for Numerical Methods in Engineering*, 126, e7624, <https://doi.org/10.1002/nme.7624>, 2025.
- Newman, G. A. and Alumbaugh, D. L.: Three-dimensional magnetotelluric inversion using non-linear conjugate gradients, *Geophysical journal international*, 140, 410–424, <https://doi.org/10.1046/j.1365-246x.2000.00007.x>, 670 2000.
- Nguyen, N. C. and Peraire, J.: An efficient reduced-order modeling approach for non-linear parametrized partial differential equations, *International Journal for Numerical Methods in Engineering*, 76, 27–55, <https://doi.org/10.1002/nme.2311>, 2008.
- Patera, A. T., Rozza, G., et al.: Reduced basis approximation and a posteriori error estimation for parametrized partial differential equations, 2007.
- Peng, Z., Yang, B., Liu, L., and Xu, Y.: Rapid surrogate modeling of magnetotelluric in the frequency domain using physics-driven deep neural networks, *Computers & Geosciences*, 176, 105360, <https://doi.org/10.1016/j.cageo.2022.105284>, 675 2023.
- Piña-Varas, P., Ledo Fernández, J., Queralt i Capdevila, P., Marcuello Pascual, A., Mitjanas Colls, G., and Martínez Van Dorth, D.: Magnetotelluric applied to deep geothermal exploration: Canary Islands, *Boletín Geológico y Minero*, 2023, vol. 134, num. 3, p. 49–58, <https://doi.org/10.21701/bolgeomin/134.3/004>, 2023.
- Qin, C. and Grayver, A.: Efficient and Scalable Finite-Element Magnetotelluric Modeling on High-Order Meshes, *Geophysical Journal International*, p. ggaf471, <https://doi.org/10.1146/annurev.fl.25.010193.002543>, 680 2025.
- Quarteroni, A., Manzoni, A., and Negri, F.: Reduced basis methods for partial differential equations: an introduction, vol. 92, Springer, <https://doi.org/10.1007/978-3-319-15431-2>, 2015.



- Robertson, K., Thiel, S., and Meqbel, N.: Quality over quantity: on workflow and model space exploration of 3D inversion of MT data, *Earth, Planets and Space*, 72, 2, <https://doi.org/10.1186/s40623-020-01127-z>, 2020.
- 685
- Rozza, G., Huynh, D. B. P., and Patera, A. T.: Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations: application to transport and continuum mechanics, *Archives of Computational Methods in Engineering*, 15, 229–275, <https://doi.org/10.1007/BF03024907>, 2008.
- Sambridge, M. and Mosegaard, K.: Monte Carlo methods in geophysical inverse problems, *Reviews of Geophysics*, 40, 3–1, <https://doi.org/10.1029/2000RG000089>, 2002.
- 690
- Samrock, F., Grayver, A. V., Eysteinnsson, H., and Saar, M. O.: Magnetotelluric image of transcrustal magmatic system beneath the Tulu Moyo geothermal prospect in the Ethiopian Rift, *Geophysical Research Letters*, 45, 12–847, <https://doi.org/10.1016/j.epsl.2017.12.046>, 2018.
- Santos, J. E. and Sheen, D.: Global and parallelizable domain-decomposed mixed FEM for 3d electromagnetic modelling, *Computational and Applied Mathematics*, 17, 265–282, 1998.
- 695
- Santos, J. E., Sheen, D., and Ye, X.: Nonconforming Galerkin methods based on quadrilateral elements for second order elliptic problems, *ESAIM: Mathematical Modelling and Numerical Analysis*, 33, 747–770, <https://doi.org/10.1051/m2an:1999144>, 1999.
- Sheen, D.: Approximation of Electromagnetic Fields: Part I. Continuous Problems, *SIAM Journal on Applied Mathematics*, 57, 1716–1736, <https://doi.org/10.1137/S003613999528656X>, accessed: 2025-05-24 21:07 UTC, 1997.
- 700
- Simpson, F. and Bahr, K.: *Practical Magnetotellurics*, Cambridge University Press, ISBN 9780521817271, <https://doi.org/10.1017/CBO9780511614095>, 2005.
- Siripunvaraporn, W.: Three-dimensional magnetotelluric inversion: an introductory guide for developers and users, *Surveys in geophysics*, 33, 5–27, <https://doi.org/10.1007/s10712-011-9145-4>, 2012.
- Siripunvaraporn, W., Egbert, G., Lenbury, Y., and Uyeshima, M.: Three-dimensional magnetotelluric inversion: data-space method, *Physics of the Earth and planetary interiors*, 150, 3–14, <https://doi.org/10.1016/j.pepi.2004.08.023>, 2005.
- 705
- Tao, L. and Zyserman, F.: ltao99/DD-POD: DD_POD, <https://doi.org/https://doi.org/10.5281/zenodo.19567027>, 2026.
- Tarantola, A.: Inverse problem theory and methods for model parameter estimation, *SIAM*, <https://doi.org/10.1137/1.9780898717921>, 2005.
- Tikhonov, A.: On determining electrical characteristics of the deep layers of the Earth’s crust, in: *Dokl. Akad. Nauk. SSSR*, vol. 73, pp. 295–297, 1950.
- 710
- Trainor-Guitton, W. and Hoversten, G. M.: Stochastic inversion for electromagnetic geophysics: practical challenges and improving convergence efficiency, *Geophysics*, 76, F373–F386, <https://doi.org/10.1190/geo2010-0205.1>, 2011.
- Vozoff, K.: Magnetotellurics: Principles and practice, *Proceedings of the Indian Academy of Sciences - Earth and Planetary Sciences*, 99, 441–471, <https://doi.org/10.1007/BF02840313>, 1990.
- Walker, D. W. and Dongarra, J. J.: MPI: a standard message passing interface, *Supercomputer*, 12, 56–68, <https://www.mpi-forum.org/docs/mpi-1.0/mpi-report.pdf>, 1996.
- 715
- Wang, J., Pan, K., Cai, H., Liu, Z., Han, X., and Ling, W.: Large-scale 3-D magnetotelluric modeling in anisotropic media using extrapolation multigrid method on staggered grids, *Computers & Geosciences*, p. 106019, <https://doi.org/10.1016/j.cageo.2023.105492>, 2025.
- Wannamaker, P. E., Stodt, J. A., and Rijo, L.: A stable finite element solution for two-dimensional magnetotelluric modelling, *Geophysical Journal International*, 88, 277–296, <https://doi.org/10.1111/j.1365-246X.1987.tb01380.x>, 1987.



- 720 Ward, S. H. and Hohmann, G. W.: Electromagnetic theory for geophysical applications, in: *Electromagnetic Methods in Applied Geophysics*, edited by Nabighian, M. N., vol. 1, pp. 131–311, Society of Exploration Geophysicists, Tulsa, <https://doi.org/10.1190/1.9781560802631>, 1988.
- Weaver, J. T.: *Mathematical Methods for Geo-Electromagnetic Induction*, Research Studies Press, Taunton, Somerset, UK, ISBN 978-0863801655, <https://doi.org/10.1111/j.1365-246X.1994.tb00150.x>, 1994.
- 725 Zyserman, F. I. and Santos, J. E.: Parallel finite element algorithm with domain decomposition for three-dimensional magnetotelluric modelling, *Journal of Applied Geophysics*, 44, 337–351, [https://doi.org/https://doi.org/10.1016/S0926-9851\(00\)00012-4](https://doi.org/https://doi.org/10.1016/S0926-9851(00)00012-4), 2000.