Cube2sph-GPU: A GPU accelerated toolkit enabling flexible continental-scale regional and teleseismic full waveform inversion

Nanqiao Du[a], Tianshi Liu[a], Bin He[c], Mijian Xu[b], Ting Lei[a] and Qinya Liu[a,b]

[a]Department of Earth Science, University of Toronto
[b]Department of Physics, University of Toronto
[c]Institute of Geology and Geophysics, Chinese Academy of Sciences

ARTICLE INFO

Keywords:
Seismology
High-performance computing
Full-waveform Inversion
Inverse Theory

ABSTRACT

We present `Cube2sph-GPU`, a GPU-accelerated framework for continental-scale regional and teleseismic full-waveform inversion (FWI). Building upon the capabilities of `SPECFEM3D_Cartesian`, the toolkit introduces: (1) a flexible hybrid simulation scheme for tele-seismic simulations; (2) curvilinear C-PML; (3) spherical PDE-based kernel smoothing; and (4) highly optimized GPU kernels and I/O processing. The framework also includes a modular inversion engine supporting four distinct dataset types and multiple misfit functions. We validate the toolkit through two numerical experiments: an analysis of SKS wave sensitivity kernels and a multi-component ambient noise FWI for a Tilted Transverse Isotropy (TTI) model. These tests demonstrate the framework's ability to resolve complex anisotropic structures with high efficiency. Overall, `Cube2sph-GPU` provides a scalable and robust solution for high-resolution imaging of the Earth's interior.

CRediT authorship contribution statement

Nanqiao Du: Conceptualization,Methodology, Software,Visualization, Writing. **Tianshi Liu:** Methodology. **Bin He:** Software,Validation. **Mijian Xu:** Software,Validation. **Ting Lei:** Software,Validation. **Qinya Liu:** Supervision.

1. Introduction

Accurate numerical simulation of seismic waves is crucial for full-waveform based structural and source inversions. The spectral element method (SEM) is a widely adopted technique for modeling seismic wave propagation, notably implemented in packages like SPECFEM3D_Cartesian and SPECFEM3D_GLOBE (Komatitsch and Tromp, 1999). SPECFEM3D_Cartesian excels at handling complex interfaces, topography, and diverse media. However, its reliance on the Universal Transverse Mercator projection (UTM) projections makes it primarily suitable for small-scale simulations (typically <10°), as significant distortions can arise from Earth's curvature beyond this range. For larger, continental-scale simulations, SPECFEM3D_GLOBE is designed to honor Earth's curvature and ellipticity, and has been widely utilized even when using only a single chunk of its global mesh (Zhu et al., 2015, 2020). Despite its global capabilities, SPECFEM3D_GLOBE currently lacks flexibility in user-defined mesh partitioning and adaptive mesh refinement at user-defined location(e.g., mesh doubling), as the mesh doubling is hard-coded in the package, which can lead to inefficient use of computational resources.

A further challenge with both SPECFEM packages lies in their treatment of absorbing boundary layers. While the Stacey boundary conditions (Clayton and Engquist, 1977) are commonly employed in wave simulations, they

ORCID(s): 0000-0000-0000-0000 (N. Du)

become ineffective for outgoing waves at grazing angles. In contrast, the Perfectly Matched Layer (PML) (Xie et al., 2014) effectively absorbs artificial reflections at the boundary, even at these challenging angles. Although PML is implemented in SPECFEM3D_Cartesian, its requirement for alignment with Cartesian computational axes renders it unsuitable for global simulations that incorporate Earth's true curvature.

To address the aforementioned challenges, Liu et al. (2024) developed the `Cube2sph` toolkit. This package provides solutions to two primary technical hurdles: (1) it generates meshes that accurately incorporate the Earth's spherical curvature, topography, and ellipticity, with the flexibility to terminate at user-defined depths; and (2) it implements the Convolution Perfectly Matching Layers (C-PML) equations within curvilinear grids, thereby enabling accurate and efficient continental-scale simulations using the `SPECFEM3D_Cartesian` solver. Furthermore, tele-seismic data can be useful in tomographic applications which necessitate the simulation of teleseismic waves for high-resolution array imaging (Wang et al., 2022). To this end, Liu et al. (2025) introduced a generalized framework for hybrid simulations, coupling `SPECFEM3D` with external solvers such as `AxiSEM` (Nissen-Meyer et al., 2014). Based on the theory of interface discontinuities, this approach facilitates the modeling of seismic wave propagation across multi-scale domains.

Despite these advancements, several challenges remain. First, integrating these tools into full-waveform inversion (FWI) workflows remains cumbersome due to the fragmented nature of existing FWI facilities and the prevalence of non-standardized in-house codes, which lack compatibility with the `Cube2sph` toolkit. Second, while the development of heterogeneous high-performance computing (HPC) has led to the widespread adoption of Graphics Processing Units (GPUs) in seismology (Gokhberg and Fichtner, 2016), GPU-accelerated anisotropic spherical solvers that support curvilinear C-PML remain limited (e.g., Wang et al., 2023), especially solver within the SEM framework . To address these limitations, we developed the `Cube2sph-GPU` toolkit, which features: (1) GPU-accelerated anisotropic forward and adjoint wavefield simulations; (2) a GPU-based PDE-driven spherical Gaussian smoothing operator for sensitivity kernel regularization; (3) a unified interface for hybrid simulations with external solvers; and (4) a flexible quasi-Newton optimization framework designed for multi-dataset and multi-parameter inversion, including integrated post-processing workflows.

2. Methodology

2.1. Cube2sph Transformation and Curvilinear PML

To generate a mesh covering a spherical domain of angular extent $\Delta\xi \times \Delta\eta$ (radians) down to a maximum depth D, we initiate the process from a Cartesian "cube" domain defined by $[-\frac{R\Delta\xi}{2}, \frac{R\Delta\xi}{2}] \times [-\frac{R\Delta\eta}{2}, \frac{R\Delta\eta}{2}] \times [-D, 0]$, where

R denotes the Earth's radius. The `Cube2sph` mapping is defined as:

$$x = -z\tan\left(\frac{\eta}{R}\right); \quad y = z\tan\left(\frac{\xi}{R}\right),$$
$$z = \frac{R+\zeta}{\sqrt{1+\tan^2\left(\frac{\xi}{R}\right)+\tan^2\left(\frac{\eta}{R}\right)}}, \tag{1}$$

where (ξ, η, ζ) represent the coordinates in the Cartesian cube domain. To ensure a bijective mapping to a specific geographic region, subsequent rotations are applied. A key property of this transformation is that it preserves the relative radial position:

$$r' = \sqrt{x^2+y^2+z^2} = R+\zeta. \tag{2}$$

Consequently, a point with a vertical offset ζ in the cube domain retains the same elevation relative to R in the spherical domain. This property facilitates the inclusion of realistic topography and undulating interfaces, as they can be defined directly in the Cartesian domain without distortion during transformation. To accurately honor the spherical geometry and topographic gradients, we utilize quadratic elements (HEX27; Komatitsch and Tromp (1999)) rather than the linear elements (HEX8) typically employed in `SPECFEM3D_Cartesian`.

Furthermore, since the standard `Cube2sph` mapping does not account for Earth's ellipticity ($\sim 1/298.25$ at the surface), we implement a radial stretching correction to map the spherical radius r to the ellipsoidal radius (Dahlen and Tromp, 2020):

$$r \to r\left[1-\frac{2}{3}\epsilon(r)P_2(\cos\theta)\right], \tag{3}$$

where $\epsilon(r)$ is the ellipticity at radius r, θ is the colatitude, and P_2 is the second-degree Legendre polynomial. The depth-dependent ellipticity $\epsilon(r)$ is determined via Radau's approximation:

$$\epsilon(r) \approx \epsilon(R)\exp\left(-\int_r^R \eta(r')(r')^{-1}dr'\right), \tag{4}$$

where the auxiliary variable $\eta(r)$ is derived from the 1-D density profile (Dahlen and Tromp, 2020). To maintain a self-consistent workflow, this ellipticity correction is applied not only to the mesh coordinates but also to the source/receiver locations including moment tensors, single forces, and seismometer locations. To suppress artificial reflections from the (curved) computational boundaries, the C-PML formulation must be adapted from Cartesian coordinates to the curvilinear system defined by the `Cube2sph` mapping. This adaptation is achieved by utilizing the Jacobian of the

transformation (equation (1)) to reformulate the auxiliary differential equations (ADEs). We derive the corresponding weak form of the momentum equation and its associated ADEs to ensure seamless integration into the SEM framework. This approach maintains the efficiency of the solver while providing robust absorption of outgoing waves in a spherical geometry. The detailed mathematical derivation of the curvilinear PML is provided in Appendix A.

2.2. Teleseismic Simulation via the Method of Interface Discontinuity

One of the key applications in seismic tomography is the utilization of teleseismic data (Wang et al., 2022; He et al., 2024), where the seismic sources are located outside the local study domain. Direct simulation of such wavefields at a global scale while honoring 3-D regional structures is computationally prohibitive. To enhance efficiency, a hybrid approach is typically adopted, assuming that the background medium outside the study region conforms to a reference 1-D Earth model where the wave solution can be computed analytically or semi-analytically (Beller et al., 2018). The wavefield from the global 1-D solver is recorded on the boundaries of the regional domain and subsequently used as a boundary condition. This one-way coupling method is mathematically equivalent to solving an interface discontinuity problem (Pienkowska et al., 2021; Liu et al., 2025):

$$\begin{aligned} &\rho\frac{\partial^2\mathbf{w}}{\partial t^2} - \nabla\cdot(\mathbf{C}:\nabla\mathbf{w}) = 0,\\ &[\mathbf{w}]_\Gamma = \mathbf{w}|_{\Gamma_-} - \mathbf{w}|_{\Gamma_+} = \mathbf{u}^0|_\Gamma,\\ &\left[\hat{\mathbf{n}}\cdot(\mathbf{C}:\nabla\mathbf{w})\right]_\Gamma = \mathbf{T}^0|_\Gamma, \end{aligned} \tag{5}$$

where $\mathbf{u}^0$ and $\mathbf{T}^0$ represent the displacement and traction fields, respectively, from the background solver evaluated on the coupling interface Γ. Here, Γ_- denotes the interior side (within the study region) and Γ_+ the exterior. The total field $\mathbf{w}$ incorporates the interface discontinuity such that $\mathbf{w} = \mathbf{u}^{3D}$ within the study region.

The background solver can be selected based on the required physics, such as QSSP (Wang et al., 2017), AxiSEM (Nissen-Meyer et al., 2014), or AxiSEM3D (Leng et al., 2019). Implementation of equation (5) requires only minor modifications to the SEM framework, specifically by injecting the prescribed tractions and displacements onto the inner side of the coupling interface during each time step.

2.3. GPU acceleration and I/O treatment

The SEM program is designed to allocate replicate arrays on both CPU and GPU memories. And the CPU first allocates all required arrays and initialize them. Then the initialized arrays are copy to GPU for further heavy computations. Data exchange between host and device memories should be as minimal as possible to avoid mitigating the overall performance. To accelerate the solver (with curvilinear C-PML) in heterogeneous HPC facilities using GPU, several engineering side optimization should be applied to the original SPECFEM3D_Cartesian package to enabling

better performance, such as communication between different computation nodes, device data memory layout, and I/O treatment in forward and adjoint runs.

The first optimization is implemented at the level of GPU-to-GPU communication. The traditional workflow in `SPECFEM3D_Cartesian` utilizes host (CPU) memory as an intermediate buffer for data exchange (represented by the black arrows in Figure 1). However, with the advancement of modern interconnect technologies—such as GPUDirect Remote Direct Memory Access (RDMA)—inter-GPU communication bandwidth significantly exceeds that of standard CPU-mediated paths. For instance, on the Trillium cluster at SciNet, the peak GPU-to-GPU communication bandwidth reaches 800 GB/s, whereas the CPU-to-CPU bandwidth is limited to 400 GB/s. To leverage this hardware capability, we introduced a conditional compilation option within the core SEM package to enable direct GPU-to-GPU data transfers (represented by the red arrow in Figure 1), effectively bypassing the host memory bottleneck.

The inclusion of PML necessitates the introduction of auxiliary arrays, significantly increasing the computational overhead per PML element relative to interior elements. This cost is particularly pronounced in the curvilinear C-PML formulation, which requires 36 additional auxiliary arrays per element (see Equation (46)). To maintain load balance across computational nodes, we assign enhanced weights to the PML elements—specifically 5.5 times the weight of regular elements—during the domain decomposition process. We then partition the global mesh into subdomains for each MPI rank using the `SCOTCH` library (Pellegrini, 2012), aiming to minimize both the computational imbalance and inter-partition communication costs.

Furthermore, since PML and regular elements exhibit distinct resource requirements in terms of shared memory and register pressure, we utilize conditional compilation (e.g., C++ `if constexpr`) to branch the GPU kernels. This approach splits the implementation into specialized execution paths, mitigating unnecessary register usage in regular elements and preventing the performance degradation associated with high register pressure in the primary wave-propagation kernels. Guided by performance profiling using NVIDIA Nsight Compute (`ncu`), we further optimized critical kernels through manual vectorization to maximize memory throughput. During each time-marching iteration, data exchange between computational nodes is required. To mask this communication overhead, we implement a double-buffering workflow that prioritizes processing halo elements (both regular and PML elements requiring data transfer). While the asynchronous MPI transfer is active, the solver concurrently processes the interior elements. This integrated approach—combining architecture-aware kernel branching, weighted domain decomposition, and optimized memory layout—ensures that the `Cube2sph-GPU` toolkit achieves near-ideal weak scaling and maximizes hardware throughput in forward simulations, even when handling the complex auxiliary physics required for C-PML simulations.

In full-waveform inversion, the Fréchet kernels of a given misfit function with respect to model parameters are

Table 1
Comparison of computational performance (wall-clock time in seconds) across different hardware configurations and I/O strategies for two mesh sizes.

Resources	DOFs / Elements	Forward	Forward + Async I/O	Forward + Sync I/O
1 × H100 GPU	4.6M / 71,280	2.67	5.65	9.63
1 × CPU (192 cores)	4.6M / 71,280	57.11	58.77	60.32
4 × H100s (NVLink)	21.4M / 327,680	8.05	13.93	20.82
1 × CPU (192 cores)	21.4M / 327,680	105.57	107.53	113.16

derived using the adjoint method (Liu and Tromp, 2006):

$$\delta\chi = \int_\Omega \left[\delta\rho \int_0^T \mathbf{s}(\mathbf{x},t)\cdot\ddot{\mathbf{s}}^\dagger(\mathbf{x},T-t)\,\mathrm{d}t + \int_0^T \nabla\mathbf{s}(\mathbf{x},t) \,:\, \delta\mathbf{C} \,:\, \nabla\mathbf{s}^\dagger(\mathbf{x},T-t)\,\mathrm{d}t \right] \mathrm{d}^3\mathbf{x}, \tag{6}$$

where $\mathbf{s}(\mathbf{x},t)$ and $\mathbf{s}^\dagger(\mathbf{x},t)$ represent the forward and adjoint wavefields, respectively. Efficiently computing these sensitivity kernels requires access to the forward wavefield snapshots during the backward-in-time adjoint simulation. Storing these snapshots often necessitates extensive disk I/O, which introduces performance penalties due to the required data transfers between GPU, CPU memory and the disk. However, for a specific frequency band of interest, accurate kernels can be obtained using a limited number of snapshots—typically around 8-32 samples per dominant period—meaning that kernel updates occur only every 10 to 30 time steps (Fichtner et al., 2009). To mitigate the I/O bottleneck, we developed an asynchronous file-handling routine leveraging non-blocking MPI-IO. This routine performs disk write/read operations in the background, effectively masking the I/O latency by overlapping it with the GPU-accelerated wavefield propagation.

To evaluate the computational speedup provided by the GPU-accelerated implementation, we designed two forward simulation benchmarks with varying mesh densities. The first configuration consists of 71,280 elements (∼4.6 million degrees of freedom, DOFs), while the second represents a larger-scale model with 327,680 elements (∼21.4 million DOFs). In both cases, the source is modeled as a single point force, and the simulation is propagated for 4,500 time steps to produce a seismogram of approximately 400 s in duration. Additionally, C-PML elements constitute roughly 25% of the total mesh. To assess the impact of kernel storage for adjoint-based inversion, wavefield snapshots are written to disk every 7 time steps.

Performance was benchmarked across three scenarios: pure forward simulation, simulation with synchronous I/O, and simulation with our optimized asynchronous I/O. As summarized in Table 1, the NVIDIA H100 GPU implementation provides a significant acceleration over the 192-core CPU configuration (AMD EPYC228 9654 processor (3.7 GHz)), achieving speedups of approximately 21× and 13× for the pure forward component of the small and large models, respectively.

However, the results also demonstrate that I/O operations significantly degrade the overall efficiency of the GPU-

accelerated solver. For the 4-H100 configuration, the introduction of synchronous I/O increases the total execution time from 8.05 s to 20.82 s, representing a nearly 160% overhead. By utilizing our asynchronous I/O routine, we successfully recover a substantial portion of this performance loss, reducing the total time for the large mesh to 13.93 s. This confirms that while the raw computational power of the H100 is transformative, the overall throughput in FWI workflows is heavily dependent on the efficiency of the data-transfer pipeline and the mitigation of the I/O bottleneck through non-blocking communication and storage strategies.

2.4. PDE-Based Gaussian Smoothing

Regularization is essential for stabilizing ill-posed inverse problems by imposing smoothness constraints on the final model, particularly to mitigate strong, non-physical anomalies near sources and receivers (Aster et al., 2018). In FWI, such regularization is typically implemented via Gaussian smoothing in spherical coordinates:

$$K(\mathbf{r}) = \frac{1}{W(\mathbf{r})} \int_V K(\mathbf{r}') \exp\left[-\frac{(r'\Delta)^2}{2\sigma_\Delta^2}\right] \exp\left[-\frac{(r-r')^2}{2\sigma_r^2}\right] \mathrm{d}^3\mathbf{r}', \tag{7}$$

where Δ denotes the angular distance between $\mathbf{r}'$ and $\mathbf{r}$, $r = |\mathbf{r}|$, and σ_r and σ_Δ represent the smoothing lengths in the radial and horizontal directions, respectively. The normalization factor $W(\mathbf{r})$ is defined as:

$$W(\mathbf{r}) = \int_V \exp\left[-\frac{(r'\Delta)^2}{2\sigma_\Delta^2}\right] \exp\left[-\frac{(r-r')^2}{2\sigma_r^2}\right] \mathrm{d}^3\mathbf{r}'. \tag{8}$$

Direct implementation of Equation (7) is computationally expensive and requires extensive MPI data exchange, as the integral is a non-local operation. However, provided that σ_Δ and σ_r are small relative to the Earth's radius, this operation can be approximated by solving a diffusion (heat) equation (see Appendix B for the proof):

$$\begin{aligned} &\frac{\partial u}{\partial t} = \nabla \cdot (\mathbf{C} \cdot \nabla u), \\ &u(\mathbf{x}, t=0) = u_0(\mathbf{x}), \end{aligned} \tag{9}$$

where u_0 is the raw sensitivity kernel. The diffusivity tensor $\mathbf{C}$ is defined as $\mathbf{C} = C_r \hat{\mathbf{e}}_r \hat{\mathbf{e}}_r + C_h(\hat{\mathbf{e}}_\phi \hat{\mathbf{e}}_\phi + \hat{\mathbf{e}}_\theta \hat{\mathbf{e}}_\theta)$, with the following relationship to the smoothing scales:

$$2C_r t = \sigma_r^2, \quad 2C_h t = \sigma_\Delta^2. \tag{10}$$

By leveraging the existing spectral element framework, we can solve this equation efficiently using GPU acceler-

ation. The stability of the explicit time-stepping scheme is governed by the CFL condition:

$$C_{\max}\Delta t = \eta h^2, \tag{11}$$

where h is the minimum spacing between Gauss-Lobatto-Legendre (GLL) points and the stability constant η is empirically set to $1/7$. The required number of time steps N is then:

$$N = \frac{\max(\sigma_r, \sigma_\Delta)^2}{2\eta h^2}. \tag{12}$$

Importantly, since only the product $C_{\max}\Delta t$ determines the total smoothing width, the individual values of $\mathbf{C}$ and Δt do not need to be defined independently for the time-stepping procedure.

As the PDE-based smoothing is integrated directly into the SEM framework, it naturally leverages the same GPU acceleration routines used for wave propagation. To demonstrate this capability, we defined a model centered on the Alaska region with dimensions of $20° \times 20° \times 220$ km (Figure 3). The mesh consists of 327,680 elements, totaling approximately 20 million degrees of freedom. An initial 10% S-wave velocity perturbation, structured as a rectangular box anomaly with dimensions of $400 \times 400 \times 70$ km, was introduced into the background model (Figure 3a). We subsequently applied anisotropic Gaussian smoothing with characteristic lengths of $\sigma_r = 8$ km (radial) and $\sigma_\Delta = 50$ km (horizontal), the results of which are shown in Figure 3b.

To quantify the computational efficiency of our implementation, we compared the execution times of the CPU and GPU versions of the code (Figure 4). The CPU benchmarks were performed using 4 cores of an AMD EPYC 9654 processor (3.7 GHz), while the GPU benchmarks utilized 4 NVIDIA H100 GPUs interconnected via NVLink. Our results indicate that the GPU-accelerated implementation achieves approximately a 40× speedup compared to the multi-core CPU execution.

3. Full-waveform Inversion

3.1. Overview

The FWI framework for applying this toolkit to tomographic problems is illustrated in Figure 5. Each inversion iteration is partitioned into two primary stages: the simulation stage and the optimization stage. In the simulation stage, an automated job scheduler manages source parallelization across available computational resources. This scheduler dynamically allocates simulation tasks to available GPU nodes, ensuring high scalability across different cluster architectures. During this phase, forward and adjoint simulations are executed for each source to compute individual sensitivity kernels.

In the optimization stage, the individual source kernels are weighted and aggregated to construct the global sensitivity kernel. Our package provides two optimization strategies: the steepest descent method and a modified limited-memory Broyden–Fletcher–Goldfarb–Shannon (M-L-BFGS) algorithm. For the steepest descent approach, we utilize a depth-dependent preconditioner $\mathbf{P}$ (or its square root) as an approximation of the inverse Hessian. Alternatively, the M-L-BFGS method employs a two-loop recursion to approximate the inverse Hessian (Algorithm 0), using the aforementioned preconditioner as the initial Hessian approximation (Modrak and Tromp, 2016). The optimal step length is determined via a line search governed by the Wolfe conditions (Nocedal and Wright, 2006). To adapt the L-BFGS framework for seismic functional derivatives, three specific modifications are implemented:

1. The inner product is evaluated variationally, consistent with the definition of the sensitivity kernel as a functional derivative:

$$\langle u, v \rangle = \int_V u(\mathbf{x}) v(\mathbf{x}) \, \mathrm{d}V = \sum_{e=1}^{N_e} \sum_{i,j,k=1}^{N_{GLL}} u^e(\xi_i, \eta_j, \zeta_k) v^e(\xi_i, \eta_j, \zeta_k) \boldsymbol{J}^e(\xi_i, \eta_j, \zeta_k) w_i w_j w_k, \tag{13}$$

   where $\boldsymbol{J}^e$ is the Jacobian of the element transformation and w are the GLL weights.

2. To satisfy the Wolfe conditions efficiently, the gradient of the line search model is cached upon a successful step. This prevents redundant and costly re-evaluations of the misfit function and its derivatives during subsequent iterations.

3. Parameter masking. In many FWI applications, it is necessary to restrict updates to specific physical parameters. We implement a masking operator that is applied to the gradient and model difference vectors during each M-L-BFGS update. This ensures that the Hessian approximation is constructed only from the active (unmasked) degrees of freedom, maintaining the stability of the inversion in the presence of fixed model components.

For enhanced maintainability and extensibility, the general FWI framework is developed in Python using a modular design. High-performance post-processing tasks—including filtering, wavefield rotation, and adjoint source measurement—are accelerated via the `mpi4py` library and the `cupy` library . To leverage established seismic toolkits, such as `measure_adj` (Liu et al., 2012), Python wrappers are provided to facilitate direct integration within the main inversion workflow. For ease of use and standardized configuration management, the framework utilizes:

1. **YAML** for defining simulation parameters and storing L-BFGS optimization histories;
2. The **Seismic Analysis Code (SAC)** format for observed seismograms;
3. **HDF5 (Hierarchical Data Format)** for the efficient storage of sensitivity kernels, model parameters, and synthetic waveforms.

Comprehensive documentation is available via a Sphinx-based user manual hosted on GitHub (see Code Availability).

To ensure compatibility between the FWI framework and the SEM solver—which operates within a global **Rectangular Coordinate System (RCS)** aligned with geocentric coordinates (where the x-axis intersects the Greenwich meridian)—we implemented several coordinate transformation routines. Specifically, source descriptions provided in `CMTSOLUTION` or `FORCESOLUTION` formats are rotated into this global RCS frame. Furthermore, a rotation matrix is dynamically generated to transform the resulting synthetic seismograms from the xyz system in RCS into the local geographic ENZ (East-North-Z/Up) coordinate system. Similarly, anisotropic symmetry axes, which are typically defined relative to the local geographic frame, must be transformed into the global geocentric Cartesian system (Chen and Tromp, 2007). Following the convention of `SPECFEM3D_globe`, the effects of Earth's ellipticity on the orientation of these symmetry axes are currently neglected in the software.

Algorithm 1 Modified L-BFGS Update with Parameter Masking (Single Iteration k)

Input: Current gradient $\mathbf{g}_k$, model $\mathbf{m}_k$, Mask $\mathbf{M}$, Stored pairs $\{\mathbf{m}_i, \mathbf{g}_i\}$ for $i = k-m, \dots, k-1$

1. Initialize $\mathbf{q} \leftarrow \mathbf{g}_k \circ \mathbf{M}$; *(Mask Input Gradient)*
2. **Backward Pass** (Compute α_i):

for $i = k-1$ **down to** $k-m$ **do**
  3. Compute differences and mask:
    $\mathbf{s}_i \leftarrow (\mathbf{m}_{i+1} - \mathbf{m}_i)$;
    $\mathbf{y}_i \leftarrow (\mathbf{g}_{i+1} - \mathbf{g}_i) \circ \mathbf{M}$;
  4. $\rho_i = 1/\langle \mathbf{y}_i, \mathbf{s}_i \rangle$;
  5. $\alpha_i = \rho_i \langle \mathbf{s}_i, \mathbf{q} \rangle$;
  6. $\mathbf{q} \leftarrow \mathbf{q} - \alpha_i \mathbf{y}_i$;
end for

7. **Scaling:** $\mathbf{r} \leftarrow \gamma_k \mathbf{q}$; *(where* $\gamma_k = \frac{\langle \mathbf{s}_{k-1}, \mathbf{y}_{k-1} \rangle}{\langle \mathbf{y}_{k-1}, \mathbf{y}_{k-1} \rangle} \mathbf{P}$*)*

8. **Forward Pass** (Compute β):

for $i = k-m$ **to** $k-1$ **do**
  9. Re-compute masked differences (if not stored):
    $\mathbf{s}_i \leftarrow (\mathbf{m}_{i+1} - \mathbf{m}_i) \circ \mathbf{M}$;
    $\mathbf{y}_i \leftarrow (\mathbf{g}_{i+1} - \mathbf{g}_i) \circ \mathbf{M}$;
  10. $\beta = \rho_i \langle \mathbf{y}_i, \mathbf{r} \rangle$;
  11. $\mathbf{r} \leftarrow \mathbf{r} + \mathbf{s}_i(\alpha_i - \beta)$; *(Crucial:* $\mathbf{s}_i$ *is masked)*
end for

12. Set search direction $\mathbf{p}_k = -\mathbf{r}$;
13. Perform Line Search to find step length η_k, and the next gradient $\mathbf{g}_{k+1}$;
14. Update Model: $\mathbf{m}_{k+1} = \mathbf{m}_k + \eta_k \mathbf{p}_k$;

Output: Updated model $\mathbf{m}_{k+1}$, New pair $\{\mathbf{m}_{k+1}, \mathbf{g}_{k+1}\}$

3.2. Flexible Parameterization and Symbolic Differentiation

The choice of model parameterization is highly dependent on the specific scientific objectives of an FWI study. For example, in subduction zone tomography, inverting for the velocity ratio $\kappa = V_P/V_S$ is often prioritized over independent V_P and V_S updates, as κ is a primary indicator of crustal dehydration and fluid migration (Kan et al., 2023). In anisotropic media, the complexity increases significantly as different regions may require different symmetry assumptions, such as radial anisotropy (VTI), azimuthal anisotropy (HTI) or Tilted Transverse Isotropy (TTI) (Chen and Tromp, 2007).

To ensure the `Cube2sph-GPU` toolkit remains maintainable and extensible, we avoid hard-coding specialized parameter sets within the core SEM solver. Instead, the solver is restricted to computing sensitivity kernels for a fundamental base set: (V_P, V_S, ρ) for isotropic media and the Voigt-notation elastic constants (C_{ij}, ρ) for anisotropic media. To bridge the gap between these base kernels and user-defined parameters, we integrate the `SymPy` symbolic mathematics library into the inversion workflow. This allows for the automatic derivation of the required chain-rule transformations, mitigating the risk of manual derivation errors. The sensitivity kernel for a target parameter M_i is computed from the base kernels m_j as follows:

$$\frac{\partial L}{\partial M_i} = \sum_{j=1}^{N} \frac{\partial L}{\partial m_j} \frac{\partial m_j}{\partial M_i} \tag{14}$$

where L is the misfit function. By evaluating the partial derivatives $\partial m_j / \partial M_i$ symbolically, the framework can flexibly map the physical sensitivities to any arbitrary parameter space defined by the user.

The framework currently supports several pre-configured parameterizations, including (V_S, κ, ρ) for isotropic investigations, as well as VTI, HTI, and TTI for more complex structural studies. A comprehensive guide for defining and integrating custom parameter combinations via this symbolic interface is provided in our online documentation (see Code Availability).

3.3. Inversion Datasets and Misfit Functions

To facilitate both regional and teleseismic tomographic applications, the `Cube2sph-GPU` toolkit supports a diverse array of misfit functions. The framework is currently optimized for four primary data categories: (1) teleseismic body waves, (2) SKS waveforms and splitting intensity, (3) receiver functions, and (4) multi-component ambient noise cross-correlation functions (CCFs).

3.3.1. Teleseismic Body Waves

For teleseismic body-wave inversions, we implement both the L_2-norm waveform difference and the cross-convolution misfit function (Menke and Levin, 2003). In the L_2 case, the source-time function (STF), $s(t)$, is estimated via decon-

volution and PCA analysis (Wang et al., 2022), leading to the adjoint source:

$$u_i^{\dagger}(t) = s(t) \otimes \left[u_i(t) - d_i(t)\right]. \tag{15}$$

Here, $u_i^{\dagger}(t)$ is the adjoint source for the i-th component, $u_i(t)$ and $d_i(t)$ are the synthetic and observed displacement waveforms, respectively, and $\otimes$ denotes the cross-correlation operator.

Alternatively, the cross-convolution method avoids explicit STF estimation by defining the misfit as:

$$\chi = \frac{1}{2} \int \left[V^{obs}(t) * H(t) - V(t) * H^{obs}(t)\right]^2 \, \mathrm{d}t, \tag{16}$$

where V and H represent the vertical and horizontal components of the wavefield, the superscript *obs* denotes observed data, and $*$ represents the convolution operator. The resulting adjoint sources for the horizontal ($H^{\dagger}$) and vertical ($V^{\dagger}$) components are:

$$\begin{aligned} H^{\dagger}(t) &= R(t) * V^{obs}(T - t), \\ V^{\dagger}(t) &= -R(t) * H^{obs}(T - t), \end{aligned} \tag{17}$$

where $R(t) = V^{obs}(t) * H(t) - V(t) * H^{obs}(t)$ is the cross-convolution residual and T is the total duration of the simulation.

3.3.2. SKS Phases and Splitting Intensity

For SKS phases, the cross-convolution method (16) is applied to the radial (R) and transverse (T) components. Additionally, we implement the splitting intensity S (Chevrot, 2006), which is defined as:

$$S = -2 \frac{\int T(t) \dot{R}(t) \, \mathrm{d}t}{\int \dot{R}^2(t) \, \mathrm{d}t}. \tag{18}$$

Here, $\dot{R}(t)$ represents the time derivative of the radial displacement component. To compute the sensitivity kernels, the adjoint sources for the transverse ($T^{\dagger}$) and radial ($R^{\dagger}$) components are derived as:

$$\begin{aligned} T^{\dagger}(t) &= -2 \frac{(S - S^{obs}) \dot{R}(t)}{\int \dot{R}^2(t') \, \mathrm{d}t'} \\ R^{\dagger}(t) &= -2(S - S^{obs}) \left[2 \frac{\ddot{R}(t) \int \dot{R}(t') T(t') \, \mathrm{d}t'}{\left(\int \dot{R}^2(t') \, \mathrm{d}t' \right)^2} - \frac{\dot{T}(t)}{\int \dot{R}^2(t') \, \mathrm{d}t'} \right] \end{aligned} \tag{19}$$

where S^{obs} is the observed splitting intensity and the denominator in the expressions represents the energy of the radial velocity pulse.

3.3.3. Receiver Functions

For receiver functions (RF), we utilize an L_2 misfit formulation derived via time-domain deconvolution ($\circledcirc$) (Ligorría and Ammon, 1999):

$$\chi = \frac{1}{2} \int \left[RF(t) - RF^{obs}(t) \right]^2 \, \mathrm{d}t \tag{20}$$

The corresponding adjoint sources are defined as:

$$\begin{aligned} f_H^{\dagger} &= \Delta RF(t) \circledcirc V(T-t), \\ f_Z^{\dagger} &= -\left[\Delta RF(t) * H(T-t) \right] \circledcirc \left[V(T-t) * V(T-t) \right], \end{aligned} \tag{21}$$

where $\Delta RF(t) = RF(t) - RF^{obs}(t)$ is the receiver function residual, and H and V are the horizontal and vertical components, respectively. The forces $f_H^{\dagger}$ and $f_Z^{\dagger}$ are applied as adjoint point sources at the receiver location to initiate the backward-in-time adjoint simulation.

3.3.4. Multi-Component Ambient Noise

To accurately model non-vertical cross-correlation function (CCF) components—essential for resolving anisotropic structures—the Green's function tensor $\mathbf{G}$ is transformed from the local ENZ (East-North-Z) system into the local Radial-Transverse-Z (RTZ) coordinate system (Wang et al., 2017):

$$G_{IJ}^{rs} = \sum_{\alpha,\beta=1}^{3} R_{I\alpha}^{s} \, G_{\alpha\beta} \, R_{J\beta}^{r}, \tag{22}$$

where $G_{\alpha\beta}$ represents the Green's function components in the local geographic system ($\alpha, \beta \in \{E, N, Z\}$). The rotation matrices for the source (s) and receiver (r) locations are given by:

$$\mathbf{R}^r = \begin{pmatrix} -\sin\phi & -\cos\phi & 0 \\ -\cos\phi & \sin\phi & 0 \\ 0 & 0 & 1 \end{pmatrix}, \qquad \mathbf{R}^s = \begin{pmatrix} \sin\theta & \cos\theta & 0 \\ \cos\theta & -\sin\theta & 0 \\ 0 & 0 & 1 \end{pmatrix}, \tag{23}$$

where θ is the azimuth and ϕ is the back-azimuth angle for the specific source-station pair. To compute the resulting Fréchet kernels, the adjoint forces $\mathbf{F}$ must be projected back into the local ENZ system:

$$\mathbf{F}^n_\alpha(t) = \sum_{I,J,\beta=1}^{3} w^n_{IJ}\, R^s_{I\alpha}\, f^I_J(t) R^r_{J\beta} \hat{e}^r_\beta, \tag{24}$$

where $f^I_J(t)$ is the adjoint source measured in the local RTZ system for the I-J component cross-correlation, w^n_{IJ} is the weighting factor for the n-th pair, and $\hat{e}^r_\beta$ is the unit vector at the receiver in the ENZ frame.

The `Cube2sph-GPU` toolkit currently supports several robust misfit functions for ambient noise analysis, including the exponentiated phase misfit (Yuan et al., 2020), multi-taper measurements (Slepian, 1978), and cross-correlation time shifts (Tromp et al., 2005). Detailed mathematical foundations for these measurements are provided in Appendix C.

3.3.5. Joint Inversion

Joint inversion serves as a powerful methodology for seismic imaging by integrating multiple datasets with complementary sensitivities to different model parameters or depths. The `Cube2sph-GPU` toolkit implements a flexible, weighted joint misfit function defined as:

$$\chi_{total} = \sum_{i=1}^{N} \chi_i w_i \tag{25}$$

where N is the number of distinct datasets and χ_i is the misfit contribution from the i-th dataset. By default, the weighting factors w_i are automatically calculated to normalize the contributions from different data types, ensuring that datasets with different physical units or scales are balanced during the initial iteration:

$$w_i = \frac{\chi^0_1}{\chi^0_i} \lambda_i \tag{26}$$

In this formulation, χ^0_i represents the misfit value of the i-th dataset at the first iteration (the initial model), and χ^0_1 serves as the reference scale (typically the primary dataset). The term λ_i is an additional user-defined scaling parameter, allowing for manual adjustment of the relative importance of specific observables based on data quality or specific research objectives.

4. Numerical Examples

In this section, we provide two numerical examples to demonstrate the accuracy and capabilities of our package.

4.1. SKS Phase Benchmarking and Fréchet Kernels

In this section, we benchmark the synthetic seismograms and evaluate the sensitivity kernels for an SKS dataset. The study area, located in Northeast China, is depicted in Figure 6. The computational domain spans approximately $16° \times 7° \times 400$ km, discretized into roughly 75,000 spectral elements. This configuration results in approximately 5.4 million DOFs, enabling the resolution of seismic waves with a minimum period of $\approx$ 7 s while accounting for the Earth's curvature. To optimize computational resources, a mesh-doubling layer is positioned directly beneath the Moho.

The background velocity structure is based on the 1D `ak135` model. The incident wavefield is generated by a seismic source located at an epicentral distance of approximately 100° from the center of the domain (Figure 6a). We utilize a Gaussian source-time function with a dominant period of 7.0 s. To simulate the teleseismic wave propagation efficiently, the background wavefield is pre-computed using AxiSEM (Nissen-Meyer et al., 2014) and subsequently injected into the simulation domain along the inner boundary of the PML elements.

Figure 7 presents a comparative benchmark between the hybrid SEM seismograms and the AxiSEM reference solutions within the SKS arrival window. The target window is defined as starting 15 s before and ending 25 s after the theoretical SKS arrival time, and all seismograms are filtered between 10-50s. The perfect agreement observed between the two sets of seismograms—specifically the precise alignment of phase and amplitude for the SKS arrivals—validates the accuracy of the wavefield injection scheme and the effectiveness of the curvilinear C-PML in the hybrid simulation framework.

To generate synthetic observations, we manually apply a shift-and-rotate transformation to the radial component of the 1D background wavefield, $s(t)$, following the methodology of Chevrot (2006). For a single station, the radial (R) and transverse (T) components are constructed as follows:

$$\begin{aligned} R(t) &= s(t)\cos^2\phi + s(t-\delta t)\sin^2\phi, \\ T(t) &= [s(t) - s(t-\delta t)]\sin\phi\cos\phi, \end{aligned} \tag{27}$$

where ϕ represents the angle between the polarization direction and the fast axis of anisotropy, and δt denotes the delay time between the fast and slow quasi-shear waves. This setup allows for a test of the splitting intensity misfit and the resulting Fréchet kernels. For the numerical experiments presented here, we prescribe a delay time $\delta t = 1$ s and an orientation angle $\phi = 35°$.

Figure 8 and Figure 9 display the sensitivity kernels for the dimensionless anisotropic parameters G_c' and G_s'. These parameters characterize the strength and orientation of azimuthal anisotropy within the medium. The horizontal and vertical sections reveal the volumetric sensitivity of the SKS phase, exhibiting the characteristic "banana-doughnut"

structure (Tromp et al., 2005). The spatial distribution of the kernels reflects the finite-frequency nature of the wavefield, where the sensitivity is not restricted to the infinitesimal ray path (Figure 9 but extends over the Fresnel zones. The distinct patterns observed in $K_{G_c'}$ and $K_{G_s'}$ result from their specific trigonometric dependence on the fast-axis orientation, demonstrating the toolkit's capability to resolve complex anisotropic structures.

4.2. FWI for Multi-Component Ambient Noise Checkerboard Model

In this example, we perform an anisotropic checkerboard test utilizing multi-component ambient noise cross-correlation functions (CCFs). The study domain covers approximately $1° \times 1° \times 50$ km. The checkerboard pattern consists of sinusoidal positive and negative anomalies with a spatial scale of roughly 50×15 km. While the background structure is based on the isotropic `ak135` model, the target anisotropic model is parameterized as a TTI medium (azimuthal anisotropy + radial anisotropy) defined by $\{V_P, V_S, \kappa_\alpha, \kappa_\beta, G_c', G_s', \rho\}$. Here, V_P and V_S represent the Voigt-average velocities, and the anisotropy ratios κ_α and κ_β are defined as:

$$\begin{aligned} V_P &= \sqrt{\frac{2V_{PH}^2 + V_{PV}^2}{3}}, \quad V_S = \sqrt{\frac{2V_{SH}^2 + V_{SV}^2}{3}} \\ \kappa_\alpha &= \frac{V_{PH} - V_{PV}}{V_{PV}}, \quad \kappa_\beta = \frac{V_{SH} - V_{SV}}{V_{SV}} \end{aligned} \tag{28}$$

Given that ambient noise data are predominantly composed of surface waves, which exhibit preferential sensitivity to shear-wave structures, we apply perturbations of $\pm 10\%$ to V_S and $\pm 5\%$ to κ_β. Additionally, we introduce a 3% anomaly to the azimuthal anisotropy magnitude, $G_0' = \sqrt{G_s'^2 + G_c'^2}$, with corresponding azimuthal angle (orientation of the fast axis) variations of $\pm 30°$ (Figure 10). To generate the synthetic observations, we simulate three orthogonal single forces (E, N, and Z directions) at 8 source stations out of a 49-station array (Figure 10). The source-time function is a Gaussian pulse with a half-duration of 1.0 s. The nine-component synthetic seismograms, initially computed in the ENZ coordinate system, are transformed into the local RTZ (Radial-Transverse-Z) system via Equation 22. From this rotated tensor, we extract two primary components for the inversion: the Z-Z components, which represent the Rayleigh-wave Green's functions (GFs), and the T-T components, which constitute the Love-wave GFs. These specific components are then utilized as the fundamental observation data for the multi-component ambient noise FWI.

During the inversion process, we initiate the optimization from the 1D `ak135` background model. To robustly capture phase information, we utilize an exponentiated phase misfit function to derive the adjoint sources. The total misfit sensitivity $\delta\chi$ is decomposed into contributions from the primary shear-wave and anisotropic parameters as follows:

$$\delta\chi = K_{V_s}\delta \ln V_s + K_{\kappa_\beta}\delta\kappa_\beta + K_{G_c'}\delta G_c' + K_{G_s'}\delta G_s' \tag{29}$$

To regularize the inversion and ensure structural continuity, a Gaussian smoothing operator with a correlation length of 3000×1500 m is applied to the Fréchet kernels at each iteration. The optimization progressed through 50 iterations, achieving a total reduction of 97% in the misfit function relative to the initial `ak135` background model.This multi-component anisotropic FWI was performed on a consumer-grade workstation equipped with an AMD Ryzen 5 5600X CPU and a single NVIDIA RTX 4070 GPU. The entire 50-iteration process, including forward simulations, adjoint calculations for the eight source-station pairs, and L-BFGS optimization was completed in approximately 7.5 hours.

Figures 11 and 12 illustrate the recovery of two representative Green's functions out of the eight simulated source-station pairs. We observe a significant improvement in waveform alignment for both Rayleigh (Z-Z) and Love (T-T) waves. Other comparisons can be found in the supplementary materials.

The results of the multi-component ambient noise FWI checkerboard test are summarized in Figures 13 through 16. Overall, the inversion successfully recovers both the spatial distribution and the magnitude of the target anomalies across the entire simulation domain. For the isotropic component, the shear-wave velocity perturbations ($\Delta \ln V_S$) are reconstructed with high fidelity at both shallow (5 km) and mid-crustal (22.5 km) depths (Figure 13). The vertical profiles (Figure 14) demonstrate that the inversion accurately captures the sinusoidal variations without significant vertical smearing, despite the inherent decrease in surface-wave sensitivity at greater depths. Crucially, the radial anisotropy (κ_β) is also well-resolved, as shown in Figure 15. The retrieval of these $\pm 5\%$ anomalies demonstrates that the multi-component misfit can effectively decouple VTI effects from the isotropic background. Furthermore, the fast-axis orientations—represented by black bars in the horizontal sections—align precisely with the target values. This high-resolution recovery of distinct parameters from multi-component Green's functions underscores the robustness of the `Cube2sph-GPU` toolkit for resolving complex, multi-scale, and anisotropic crustal structures.

5. Discussion and Future Work

The `Cube2sph-GPU` toolkit is developed to facilitate regional and teleseismic full-waveform inversion at continental scales. By integrating curvilinear C-PML with GPU acceleration, the package provides a high-performance solution for both forward and adjoint simulations on heterogeneous HPC architectures. While the current framework offers a comprehensive inversion pipeline for four primary seismic datasets, several critical features are targeted for future development to enhance the physical realism and versatility of the package.

A primary objective for future development is the incorporation of attenuation, particularly within anisotropic media. This can be addressed by introducing standard linear solid (SLS) models (Aki and Richards, 2002) alongside optimized correction methods to integrate the quality factor Q into the forward and adjoint solvers (van Driel and Nissen-Meyer, 2014; Chen et al., 2026). In the anisotropic case, characterizing dissipative behavior requires complex combinations of the elastic tensor C_{ij} (Carcione, 1995), which necessitates additional coordinate transformations

within the Cube2sph-GPU framework to align the attenuation tensor with local coordinates. While adjoint simulations often face numerical instabilities during the time-reversal of attenuated wave equations (Trinh et al., 2019), our strategy of storing the complete forward wavefield provides a stable, unified approach for both elastic and viscoelastic kernel calculations. Although this approach is I/O intensive, the computational overhead is mitigated by our asynchronous I/O routines. Furthermore, because viscoelasticity introduces auxiliary variables that significantly complicate curvilinear C-PML formulations, future versions may explore replacing PMLs with high-attenuation buffer layers in viscoelastic simulations to maintain stability and implementation simplicity.

Beyond attenuation, a further extension of the Cube2sph-GPU toolkit involves implementing elastic-acoustic coupling to facilitate simulations in oceanic or atmospheric environments. While the general FWI framework remains consistent for fluid-solid interactions, the boundary conditions within the C-PML layers require significant modifications (Xie et al., 2016). Extending the formulations of Xie et al. (2016) into curvilinear systems necessitates transforming the geometric normal vector n_i into a stretched normal vector, $n_i' = R_{ip} n_p$, where $\boldsymbol{R}$ is the Jacobian matrix between stretched and un-stretched curvilinear system in frequency domain. This transformation introduces two specific surface integrals at the fluid-solid interface (Σ_{FS}):

$$\begin{aligned} B_{F\to S} &= \int_{\Sigma_{FS}} \sigma_{ij} R_{jk} n_k \psi \, \mathrm{d}\Sigma \\ B_{S\to F} &= \int_{\Sigma_{FS}} R_{ij} u_i n_j \psi \, \mathrm{d}\Sigma \end{aligned} \tag{30}$$

The implementation of these coupling terms requires the introduction of additional memory variables to track the acoustic potential χ and the displacement field u_i. Integrating these curvilinear coupling routines will broaden the toolkit's applicability to global-scale problems involving complex bathymetry and oceanic wave propagation.

The third limitation is that the current framework is not yet optimized for local earthquake tomography. For local-scale applications, automated window selection (Maggi et al., 2009) is essential to isolate specific seismic phases from complex, high-frequency wavefields. Furthermore, to mitigate structural artifacts caused by uncertainties in earthquake locations and origin times, the implementation of double-difference misfit functions (Yuan et al., 2016) and joint source-structure inversion routines is required. While the core spectral-element solver is fully capable of the high resolutions needed for local studies, these specific data-processing and optimization modules are slated for integration in future releases to provide a more versatile multi-scale inversion platform.

6. Conclusions

In this study, we presented the Cube2sph-GPU toolkit, a comprehensive and high-performance framework designed for regional and teleseismic full-waveform inversion in (a)spherical coordinates. The package enables the extension of

inversion domains to continental scales by providing user-defined depth configurations and flexible mesh-refinement options.

A core contribution of this work is the implementation of GPU acceleration across all primary computational bottlenecks, including forward and adjoint wavefield simulations within a curvilinear C-PML framework, PDE-based regularization, and adjoint source measurement. To further optimize throughput, we integrated GPU-specific memory management and an asynchronous I/O pipeline, significantly mitigating the data-transfer overhead inherent in adjoint-based tomography.

Furthermore, `Cube2sph-GPU` offers a modular FWI workflow supporting four distinct seismic datasets—teleseismic body waves, SKS phases, receiver functions, and ambient noise—while leveraging symbolic mathematics to allow for arbitrary, user-defined parameter combinations. Overall, this toolkit provides a robust and scalable solution for high-resolution seismic imaging, bridging the gap between advanced numerical methods and large-scale geophysical applications.

7. Acknowledgments

The authors would like to acknowledge ...

Code availability section

Name of the code/library: Cube2sph-GPU

Contact: nanqiao.du@mail.utoronto.ca

Hardware requirements: CPU or GPU clusters, NVlink is optional but recommened.

Program language: C/C++, Fortran, Python, Bash

Software required: C/C++ and Fortran Compilers, NVCC compilers, NetCDF, OpenMPI, Python

Program size: 125 MB

The source codes are available for downloading at the link: github

A. The Curvilinear PML

Now we want to simulate wave (with PML) in a curvilinear domain with coordinates $\boldsymbol{x}$, the direct PML implementation is tricky. However, if this domain can be related (through a transformation) to a regular domain with coordinates $\boldsymbol{\xi}$, in which we've already know the PML expressions, we seek to find direct relationship between PML equations in $\boldsymbol{x}$ and PML equations in $\boldsymbol{\xi}$. Before we move forward, we define the Jacobian matrix $\boldsymbol{J}$ between two coordinates:

$$\boldsymbol{J}_{ij} = \frac{\partial x_j}{\partial \xi_i} \tag{31}$$

We can write down the action in frequency domain and in the domain $\boldsymbol{x}$ as:

$$S = \frac{1}{2}\int \mathrm{d}^3\boldsymbol{x}\ \left[\rho\omega^2 u_i u_i^* - c_{ijkl}u_{k,l}u_{i,j}^*\right] \tag{32}$$

we define the new wave field $\bar{u}_i(\boldsymbol{\xi}) = u_i(\boldsymbol{x})$ in the $\boldsymbol{\xi}$ domain, and rewrite the action as:

$$S = \frac{1}{2}\int \mathrm{d}^3\boldsymbol{\xi}\ \boldsymbol{J}\left[\rho\omega^2 \bar{u}_i \bar{u}_i^* - c_{ijkl}\boldsymbol{J}_{lp}^{-1}\boldsymbol{J}_{jq}^{-1}\bar{u}_{k,p}\bar{u}_{i,q}^*\right] \tag{33}$$

Suppose we have constructed PML layers by coordinate stretching:

$$\frac{\partial}{\partial \tilde{\xi}_i} = S_{ij}\frac{\partial}{\partial \xi_j} \tag{34}$$

where

$$S_{ij} = \frac{\delta_{ij}}{s_i(\xi_i)} \tag{35}$$

Then we can write down the new action S' in stretched domain $\tilde{\boldsymbol{\xi}}$ as:

$$S' = \frac{1}{2}\int \mathrm{d}^3\tilde{\boldsymbol{\xi}}\ J\left[\rho\omega^2\tilde{\tilde{u}}_i\tilde{\tilde{u}}_i^* - c_{ijkl}J_{lp}^{-1}J_{jq}^{-1}\tilde{\tilde{u}}_{k,p}\tilde{\tilde{u}}_{i,q}^*\right] \tag{36}$$

we can convert the domain to $\boldsymbol{x}$ as:

$$\begin{aligned} S' &= \frac{1}{2}\int \mathrm{d}^3\boldsymbol{\xi}\ |\boldsymbol{S}|J\left[\rho\omega^2\tilde{\tilde{u}}_i\tilde{\tilde{u}}_i^* - c_{ijkl}J_{lp}^{-1}J_{jq}^{-1}S_{pm}S_{qn}\tilde{\tilde{u}}_{k,m}\tilde{\tilde{u}}_{i,n}^*\right] \\ &= \frac{1}{2}\int \mathrm{d}^3\boldsymbol{x}\ |\boldsymbol{S}|J^{-1}J\left[\rho\omega^2\tilde{u}_i\tilde{u}_i^* - c_{ijkl}J_{lp}^{-1}J_{jq}^{-1}S_{pm}S_{qn}J_{mr}J_{ns}\tilde{u}_{k,r}\tilde{u}_{i,s}^*\right] \\ &= \frac{1}{2}\int \mathrm{d}^3\boldsymbol{x}\ |\boldsymbol{S}|\left[\rho\omega^2\tilde{u}_i\tilde{u}_i^* - c_{ijkl}R_{lr}R_{js}\tilde{u}_{k,r}\tilde{u}_{i,s}^*\right] \end{aligned} \tag{37}$$

where $\boldsymbol{R} = \boldsymbol{J}^{-1}\cdot\boldsymbol{S}\cdot\boldsymbol{J}$. So when we transform our original domain to a new region $\boldsymbol{x} = \boldsymbol{x}(\boldsymbol{\xi})$, we can derive the new PML derivative as:

$$\tilde{\nabla} = \boldsymbol{J}^{-1}\cdot\boldsymbol{S}\cdot\boldsymbol{J}\cdot\nabla \tag{38}$$

Then we can rewrite the action as:

$$S' = \frac{1}{2}\int \mathrm{d}^3\tilde{\boldsymbol{x}}\ \left[\rho\omega^2\tilde{u}_i\tilde{u}_i^* - c_{ijkl}\partial_{\tilde{x}_l}\tilde{u}_k\partial_{\tilde{x}_j}\tilde{u}_i^*\right] \tag{39}$$

We set a notation:

$$\begin{aligned} R_{ij}^{\Sigma} &= \sum_k \frac{1}{\kappa_k}J_{ik}^{-1}J_{kj} \\ R_{ij}^{k} &= \frac{\alpha_k - \beta_k}{\kappa_k}J_{ik}^{-1}J_{kj} \end{aligned} \tag{40}$$

We can show that in frequency domain:

$$\frac{\partial f}{\partial\tilde{x}_i} = \sum_{kpj} J_{ik}^{-1}\frac{\delta_{kp}}{s_p(x)}J_{pj}\frac{\partial f}{\partial x_j} = \sum_{kj} J_{ik}^{-1}\frac{1}{s_k}J_{kj}\frac{\partial f}{\partial x_j} \tag{41}$$

By inverse Fourier transform, we have:

$$\begin{aligned} \mathscr{F}^{-1}\left[\frac{\partial f}{\partial\tilde{x}_i}\right] &= \sum_{kj} J_{ik}^{-1}\frac{1}{\kappa_k}J_{kj}\frac{\partial f}{\partial x_j} + \sum_{kj} J_{ik}^{-1}\frac{\alpha_k - \beta_k}{\kappa_k}J_{kj}e^{-\beta_k t}H(t) * \frac{\partial f}{\partial x_j} \\ &= \sum_k R_{ik}^{\Sigma}\frac{\partial f}{\partial x_k} + \sum_{\lambda}\sum_k R_{ik}^{\lambda}e^{-\beta_\lambda t}H(t) * \frac{\partial f}{\partial x_k} \end{aligned} \tag{42}$$

Then the PDE for velocity and stress can be written by index-form:

$$\begin{aligned}
\rho\frac{\partial v_i}{\partial t} &= R^{\Sigma}_{jk}\frac{\partial}{\partial x_k}\sigma_{ij} + \sum_{\lambda} e^{-\beta_\lambda t}H(t) * \left(R^{\lambda}_{jk}\frac{\partial}{\partial x_k}\sigma_{ij}\right) \\
\frac{\partial \sigma_{ij}}{\partial t} &= C_{ijkl}\left\{R^{\Sigma}_{kp}\frac{\partial}{\partial x_p}v_l + \sum_{\lambda} e^{-\beta_\lambda t}H(t) * (R^{\lambda}_{kp}v_{l,p})\right\}
\end{aligned} \tag{43}$$

Then we define the memory variable from the convolution:

$$\begin{aligned}
N^{\lambda}_{ij} &= e^{-\beta_\lambda t}H(t) * u_{i,j} \\
M^{\lambda}_{i} &= e^{-\beta_\lambda t}H(t) * \left(R^{\lambda}_{jk}\frac{\partial}{\partial x_k}\sigma_{ij}\right)
\end{aligned} \tag{44}$$

Then the PML equation for velocity and stress are:

$$\begin{aligned}
\rho\frac{\partial v_i}{\partial t} &= R^{\Sigma}_{jk}\frac{\partial}{\partial x_k}\sigma_{ij} + \sum_{\lambda} M^{\lambda}_{i} \\
\frac{\partial M^{\lambda}_{i}}{\partial t} &= -\beta_\lambda M^{\lambda}_{i} + R^{\lambda}_{jk}\frac{\partial}{\partial x_k}\sigma_{ij} \\
\sigma_{ij} &= C_{ijkl}\left\{R^{\Sigma}_{kp}\frac{\partial}{\partial x_p}u_l + \sum_{\lambda p} R^{\lambda}_{kp}N^{\lambda}_{lp}\right\} \\
\frac{\partial N^{\lambda}_{ij}}{\partial t} &= -\beta_\lambda N^{\lambda}_{ij} + u_{i,j}
\end{aligned} \tag{45}$$

Then we can write it in the weak form:

$$\begin{aligned}
\int_V \rho\frac{\partial v_i}{\partial t}\psi \mathrm{d}V &= -\int_V \sigma_{ij}\frac{\partial}{\partial x_k}\left(R^{\Sigma}_{jk}\psi\right)\mathrm{d}V + \sum_{\lambda}\int_V M^{\lambda}_{i}\psi \mathrm{d}V \\
\int_V \frac{\partial M^{\lambda}_{i}}{\partial t}\psi \mathrm{d}V &= -\int_V \beta_\lambda M^{\lambda}_{i}\psi \mathrm{d}V - \int_V \frac{\partial}{\partial x_k}\left(R^{\lambda}_{jk}\psi\right)\sigma_{ij}\mathrm{d}V + \int_\Sigma \sigma_{ij}R^{\lambda}_{jk}n_k\psi \mathrm{d}\Sigma \\
\sigma_{ij} &= C_{ijkl}\left\{R^{\Sigma}_{kp}\frac{\partial}{\partial x_p}u_l + \sum_{\lambda,p} R^{\lambda}_{kp}N^{\lambda}_{lp}\right\} \\
\frac{\partial N^{\lambda}_{ij}}{\partial t} &= -\beta_\lambda N^{\lambda}_{ij} + u_{i,j}
\end{aligned} \tag{46}$$

Subject to the boundary condition:

$$\sigma_{ij}n_k R^{\Sigma}_{jk} = 0, \boldsymbol{x} \in \partial\Omega; \quad \boldsymbol{u} = 0, \boldsymbol{x} \in \partial\Omega^{PML} \tag{47}$$

where $\partial\Omega$ is the boundary on the free surface and $\partial\Omega^{PML}$ is the PML external boundaries.

B. PDE Smoothing in Spherical Coordinate System

The fundamental solution in spherical coordinates satisfies:

$$\begin{aligned} &\frac{\partial T}{\partial t} = c_r \left[\frac{1}{r^2}\frac{\partial}{\partial r}\left(r^2\frac{\partial T}{\partial r}\right)\right] + \frac{c_h}{r^2\sin\theta}\frac{\partial}{\partial\theta}\left(\sin\theta\frac{\partial T}{\partial\theta}\right) + \frac{c_h}{r^2\sin^2\theta}\frac{\partial^2 T}{\partial\phi^2} \\ &T(t=0) = \frac{\delta(r-r_0)}{r^2}\delta(\Omega-\Omega_0) \end{aligned} \tag{48}$$

We define the angular Laplacian operator Δ_{S^2} as:

$$\Delta_{S^2} = \frac{1}{\sin\theta}\frac{\partial}{\partial\theta}\left(\sin\theta\frac{\partial}{\partial\theta}\right) + \frac{1}{\sin^2\theta}\frac{\partial^2}{\partial\phi^2} \tag{49}$$

Substituting this into Eq. (48):

$$\frac{\partial T}{\partial t} = \frac{c_r}{r^2}\frac{\partial}{\partial r}\left(r^2\frac{\partial T}{\partial r}\right) + \frac{c_h}{r^2}\Delta_{S^2}T \tag{50}$$

Since the angular operator Δ_{S^2} is independent of r and t, we expand the solution using the completeness of Spherical Harmonics $Y_\ell^m(\theta,\phi)$. We recall the eigenvalue property:

$$\Delta_{S^2}Y_\ell^m(\Omega) = -\ell(\ell+1)Y_\ell^m(\Omega) \tag{51}$$

We expand both the temperature T and the angular part of the delta function:

$$T(r,\theta,\phi,t) = \sum_{\ell=0}^{\infty}\sum_{m=-\ell}^{\ell} R_{\ell m}(r,t)Y_\ell^m(\theta,\phi) \tag{52}$$

Using the completeness relation $\delta(\Omega-\Omega_0) = \sum_{\ell,m} Y_\ell^m(\Omega)Y_\ell^{m*}(\Omega_0)$, the initial condition for the radial function becomes:

$$R_{\ell m}(r,0) = \frac{\delta(r-r_0)}{r^2}Y_\ell^{m*}(\theta_0,\phi_0) \tag{53}$$

Substituting the expansion into the PDE and exploiting orthogonality, we obtain the equation for the radial coefficient $R_{\ell m}$:

$$\frac{\partial R_{\ell m}}{\partial t} = c_r\left[\frac{1}{r^2}\frac{\partial}{\partial r}\left(r^2\frac{\partial R_{\ell m}}{\partial r}\right) - \frac{c_h}{c_r}\frac{\ell(\ell+1)}{r^2}R_{\ell m}\right] \tag{54}$$

The radial equation $R_{lm} = g_l Y_l^{m*}(\theta_0, \phi_0)$ simplifies to:

$$\frac{1}{c_r}\frac{\partial g_\ell}{\partial t} = \frac{\partial^2 g_\ell}{\partial r^2} + \frac{2}{r}\frac{\partial g_\ell}{\partial r} - \frac{\lambda^2 \ell(\ell+1)}{r^2} g_\ell \tag{55}$$

where $\lambda^2 = c_h/c_r$ is the anisotropy ratio. We define a **modified order** ν_ℓ to absorb the anisotropy into the standard form of the Bessel differential equation:

$$\nu_\ell(\nu_\ell + 1) = \lambda^2 \ell(\ell+1) = \frac{c_h}{c_r}\ell(\ell+1) \tag{56}$$

The radial equation is now:

$$\frac{\partial g_\ell}{\partial t} = c_r \left[\frac{\partial^2 g_\ell}{\partial r^2} + \frac{2}{r}\frac{\partial g_\ell}{\partial r} - \frac{\nu_\ell(\nu_\ell+1)}{r^2} g_\ell \right] \tag{57}$$

Subject to the initial condition: $g_\ell(r, 0) = \frac{\delta(r-r_0)}{r^2}$.

Note:

$$\frac{\delta(r-r_0)}{r^2} = \frac{2}{\pi}\int_0^\infty \beta^2 j_\nu(\beta r) j_\nu(\beta r_0) \mathrm{d}\beta \tag{58}$$

The solution to this specific radial heat equation can be solved as:

$$g_\ell(r,t) = \frac{2}{\pi}\int_0^\infty \beta^2 e^{-c_r\beta^2 t} j_{\nu_\ell}(\beta r) j_{\nu_\ell}(\beta r_0)\, d\beta \tag{59}$$

Using the relationship between spherical and cylindrical Bessel functions, $j_\nu(z) = \sqrt{\frac{\pi}{2z}} J_{\nu+1/2}(z)$, and the Weber integral identity:

$$\int_0^\infty x e^{-p^2 x^2} J_\mu(ax) J_\mu(bx)\, dx = \frac{1}{2p^2}\exp\left(-\frac{a^2+b^2}{4p^2}\right) I_\mu\left(\frac{ab}{2p^2}\right) \tag{60}$$

we obtain the closed form for the radial part:

$$g_\ell(r,t) = \frac{1}{2c_r t\sqrt{rr_0}}\exp\left(-\frac{r^2+r_0^2}{4c_r t}\right) I_{\nu_\ell+1/2}\left(\frac{rr_0}{2c_r t}\right) \tag{61}$$

Here, $I_{\nu_\ell+1/2}$ is the modified Bessel function of the first kind.

Combining the radial with the angular expansion, the Green's function $G(\mathbf{r}, \mathbf{r}_0, t)$ is:

$$G(\mathbf{r}, \mathbf{r}_0, t) = \frac{e^{-\frac{r^2+r_0^2}{4c_r t}}}{2c_r t\sqrt{rr_0}} \sum_{\ell=0}^{\infty} I_{\nu_\ell+1/2}\left(\frac{rr_0}{2c_r t}\right) \sum_{m=-\ell}^{\ell} Y_\ell^m(\Omega) Y_\ell^{m*}(\Omega_0) \tag{62}$$

Using the Spherical Harmonic Addition Theorem:

$$\sum_{m=-\ell}^{\ell} Y_\ell^m(\Omega) Y_\ell^{m*}(\Omega_0) = \frac{2\ell+1}{4\pi} P_\ell(\cos\Delta) \tag{63}$$

where Δ is the epicenter distance between vectors $\mathbf{r}$ and $\mathbf{r}_0$.

The final expression for the fundamental solution is:

$$\boxed{T(\mathbf{r}, t) = \frac{\exp\left(-\frac{r^2+r_0^2}{4c_r t}\right)}{8\pi c_r t\sqrt{rr_0}} \sum_{\ell=0}^{\infty} (2\ell+1) P_\ell(\cos\Delta)\, I_{\nu_\ell+1/2}\left(\frac{rr_0}{2c_r t}\right)} \tag{64}$$

with the anisotropic order parameter:

$$\nu_\ell = -\frac{1}{2} + \sqrt{\frac{1}{4} + \frac{c_h}{c_r}\ell(\ell+1)} \approx \lambda l \tag{65}$$

when t is a small number, and the two points are close, we can approximate the Modified Bessel function by using the large argument and high order asymptotic expansion:

$$I_\nu(z) \approx \frac{e^z}{\sqrt{2\pi z}}\left(1 - \frac{4\nu^2 - 1}{8z}\right) \approx \frac{\exp\left\{[z - \nu^2/(2z)]\right\}}{\sqrt{2\pi z}} \tag{66}$$

And the Legendre polynomial can be approximated by using the 0-th Bessel function (Mehler-Heine Formula) :

$$P_l(\cos\Delta) = J_0(l\Delta) \tag{67}$$

So we can obtain:

$$(2l+1) P_l(\cos\Delta) I_{\nu_l+1/2}\left(\frac{rr_0}{2c_r t}\right) \approx 2l J_0(l\Delta) \frac{\exp\left\{\left[\frac{rr_0}{2c_r t} - l^2\frac{c_h t}{rr_0}\right]\right\}}{\sqrt{2\pi\frac{rr_0}{2c_r t}}} \tag{68}$$

And the summations can be approximated by integral as:

$$\begin{aligned}\sum_{\ell=0}^{\infty}(2l+1)P_l(\cos\Delta)I_{\nu_l+1/2}(\frac{rr_0}{2c_rt}) &\approx \sqrt{\frac{c_rt}{\pi rr_0}}e^{\frac{rr_0}{2c_rt}}\int_0^{\infty}2lJ_0(l\Delta)e^{-\frac{c_ht}{rr_0}l^2}\,\mathrm{d}l \\ &= \sqrt{\frac{c_rt}{\pi rr_0}}e^{\frac{rr_0}{2c_rt}}\frac{rr_0}{c_ht}e^{-\frac{rr_0\Delta^2}{4\alpha_ht}}\end{aligned} \tag{69}$$

So the fundamental solution can be approximated as ($r \to r_0$):

$$T(\boldsymbol{r},t) \approx \frac{1}{(4\pi t)^{3/2}\sqrt{c_rc_h}}\exp\left\{\left[-\frac{1}{4t}\left(\frac{(r-r_0)^2}{c_r}+\frac{r^2\Delta^2}{c_h}\right)\right]\right\} \tag{70}$$

C. Additional Misfit Formulations

C.1. Exponentiated Phase Misfit

Let $s(t)$ be the synthetic signal and $d(t)$ be the observed data. We can define it's Hilbert transform $\mathcal{H}\{s(t)\}, \mathcal{H}\{d(t)\}$. We can define the instantaneous amplitude as:

$$E_s(t) = \sqrt{s^2(t)+\mathcal{H}^2\{s(t)\}}; \quad E_d(t) = \sqrt{d^2(t)+\mathcal{H}^2\{d(t)\}}; \tag{71}$$

The exponential phase misfit measures the difference between the normalized analytical signals of synthetic and observed data:

$$\chi = \frac{1}{2}\int_0^T\left|\frac{d_a(t)}{E_d(t)}-\frac{s_a(t)}{E_s(t)}\right|^2\mathrm{d}t = \frac{1}{2}\int_0^T\left[\Delta R(t)^2+\Delta I(t)^2\right]\mathrm{d}t \tag{72}$$

where:

$$\begin{aligned}\Delta R(t) &= \frac{d(t)}{E_d(t)}-\frac{s(t)}{E_s(t)} \\ \Delta I(t) &= \frac{\mathcal{H}\{d(t)\}}{E_d(t)}-\frac{\mathcal{H}\{s(t)\}}{E_s(t)}\end{aligned} \tag{73}$$

so the adjoint source is :

$$f^{\dagger}(t) = \left[\Delta I(t)\frac{s(t)\mathcal{H}\{s(t)\}}{E_s^3(t)}-\Delta R(t)\frac{[\mathcal{H}\{s(t)\}]^2}{E_s^3(t)}+\mathcal{H}\left\{\Delta I(t)\frac{s^2(t)}{E_s^3(t)}-\Delta R(t)\frac{[s(t)\mathcal{H}\{s(t)\}}{E_s^3(t)}\right\}\right] \tag{74}$$

C.2. Multi-Taper Measurements

The multi-taper approach (Slepian, 1978) is utilized for frequency-dependent measurements:

$$\Delta\tau(\omega) = \frac{1}{\omega}\mathrm{Im}\left[\ln\left(\frac{\sum_k \tilde{u}_k(\omega)\tilde{d}_k^*(\omega)}{\sum_k |\tilde{d}_k(\omega)|^2}\right)\right], \tag{75}$$

where $\tilde{u}_k(\omega) = \mathcal{F}\left\{h_k(t)u(t)\right\}$ and $\tilde{d}_k(\omega) = \mathcal{F}\left\{h_k(t)d(t)\right\}$, and $h_k(t)$ is the k-th Slepian taper (window), and $\mathcal{F}$ is the Fourier transform. The misfit function is defined as:

$$\chi_P(\mathbf{m}) = \frac{1}{2}\int_0^W W_P(w)\left|\frac{\Delta\tau(\omega)}{\sigma_P(w)}\right|^2 dw \tag{76}$$

where W is the maximum frequency of interest, $W_P(\omega)$ is a frequency-dependent weighting function, and $\sigma_P(\omega)$ represents the uncertainty in the phase measurements, which can be estimated by cross-correlation time shift. The corresponding adjoint source in the time domain is computed via:

$$f^\dagger(t) = \sum_{k=1}^{K} 2\pi h_k(t)\mathcal{F}^{-1}\left[\frac{i}{\omega}\frac{W_P(\omega)}{\sigma_P^2(\omega)}\Delta\tau(\omega)\frac{u_k^*(\omega)}{\sum_{k'} u_{k'}(\omega)u_{k'}^*}\right] \tag{77}$$

where $\mathcal{F}^{-1}$ denotes the inverse Fourier transform, K is the number of tapers, and $\tilde{u}_k(\omega)$ is the Fourier transform of the k-th tapered synthetic waveform.

References

Aki, K., Richards, P., 2002. Quantitative seismology. MIT Press.

Aster, R.C., Borchers, B., Thurber, C.H., 2018. Parameter estimation and inverse problems. Elsevier.

Beller, S., Monteiller, V., Operto, S., Nolet, G., Paul, A., Zhao, L., 2018. Lithospheric architecture of the south-western alps revealed by multiparameter teleseismic full-waveform inversion. Geophysical Journal International 212, 1369–1388.

Carcione, J.M., 1995. Constitutive model and wave equations for linear, viscoelastic, anisotropic media. Geophysics 60, 537–548.

Chen, M., Tromp, J., 2007. Theoretical and numerical investigations of global and regional seismic wave propagation in weakly anisotropic earth models. Geophysical Journal International 168, 1130–1152.

Chen, Y., Lumley, D., Zhu, H., 2026. Viscoelastic wave propagation: spectral-element simulation with a novel gsls-type wave equation. Geophysical Journal International 244, ggaf382.

Chevrot, S., 2006. Finite-frequency vectorial tomography: a new method for high-resolution imaging of upper mantle anisotropy. Geophysical Journal International 165, 641–657.

Clayton, R., Engquist, B., 1977. Absorbing boundary conditions for acoustic and elastic wave equations. Bulletin of the seismological society of America 67, 1529–1540.

Dahlen, F., Tromp, J., 2020. Theoretical global seismology, in: Theoretical global seismology. Princeton university press.

van Driel, M., Nissen-Meyer, T., 2014. Optimized viscoelastic wave propagation for weakly dissipative media. Geophysical Journal International

199, 1078–1093.

Fichtner, A., Kennett, B.L., Igel, H., Bunge, H.P., 2009. Full seismic waveform tomography for upper-mantle structure in the australasian region using adjoint methods. Geophysical Journal International 179, 1703–1725.

Gokhberg, A., Fichtner, A., 2016. Full-waveform inversion on heterogeneous hpc systems. Computers & Geosciences 89, 260–268.

He, B., Wang, K., Liu, T., Lei, T., Du, N., van der Lee, S., Darbyshire, F.A., Frederiksen, A., Zhu, H., Lumley, D., et al., 2024. Crustal and uppermost mantle structures of the north american midcontinent rift revealed by joint full-waveform inversion of ambient-noise data and teleseismic p waves. Earth and Planetary Science Letters 641, 118797.

Kan, L.Y., Chevrot, S., Monteiller, V., 2023. Dehydration of the subducting juan de fuca plate and fluid pathways revealed by full waveform inversion of teleseismic p and sh waves in central oregon. Journal of Geophysical Research: Solid Earth 128, e2022JB025506.

Komatitsch, D., Tromp, J., 1999. Introduction to the spectral element method for three-dimensional seismic wave propagation. Geophysical journal international 139, 806–822.

Leng, K., Nissen-Meyer, T., Van Driel, M., Hosseini, K., Al-Attar, D., 2019. Axisem3d: broad-band seismic wavefields in 3-d global earth models with undulating discontinuities. Geophysical Journal International 217, 2125–2146.

Ligorría, J.P., Ammon, C.J., 1999. Iterative deconvolution and receiver-function estimation. Bulletin of the seismological Society of America 89, 1395–1400.

Liu, Q., Tape, C., Zhou, Y., 2012. measure_adj: A software package for making measurements and computing adjoint sources for seismic tomography .

Liu, Q., Tromp, J., 2006. Finite-frequency kernels based on adjoint methods. Bulletin of the Seismological Society of America 96, 2383–2397.

Liu, T., Du, N., Lei, T., Wang, K., He, B., Tong, P., Grasselli, G., Liu, Q., 2025. Seismic wavefield injection based on interface discontinuity: theory and numerical implementation based upon the spectral-element method. Geophysical Journal International 241, 658–673.

Liu, T., Wang, K., Xie, Y., He, B., Lei, T., Du, N., Tong, P., Yang, Y., Rychert, C.A., Harmon, N., et al., 2024. Cube2sph: a toolkit enabling flexible and accurate continental-scale seismic wave simulations using the specfem3d_cartesian package. Computers & Geosciences 190, 105644.

Maggi, A., Tape, C., Chen, M., Chao, D., Tromp, J., 2009. An automated time-window selection algorithm for seismic tomography. Geophysical Journal International 178, 257–281.

Menke, W., Levin, V., 2003. The cross-convolution method for interpreting sks splitting observations, with application to one and two-layer anisotropic earth models. Geophysical Journal International 154, 379–392.

Modrak, R., Tromp, J., 2016. Seismic waveform inversion best practices: regional, global and exploration test cases. Geophysical Journal International 206, 1864–1889.

Nissen-Meyer, T., van Driel, M., Stähler, S.C., Hosseini, K., Hempel, S., Auer, L., Colombi, A., Fournier, A., 2014. Axisem: broadband 3-d seismic wavefields in axisymmetric media. Solid Earth 5, 425–445.

Nocedal, J., Wright, S.J., 2006. Numerical optimization. Springer.

Pellegrini, F., 2012. Scotch and pt-scotch graph partitioning software: an overview. Combinatorial Scientific Computing , 373–406.

Pienkowska, M., Monteiller, V., Nissen-Meyer, T., 2021. High-frequency global wavefields for local 3-d structures by wavefield injection and extrapolation. Geophysical Journal International 225, 1782–1798.

Slepian, D., 1978. Prolate spheroidal wave functions, fourier analysis, and uncertainty—v: The discrete case. Bell System Technical Journal 57, 1371–1430.

Trinh, P.T., Brossier, R., Métivier, L., Tavard, L., Virieux, J., 2019. Efficient time-domain 3d elastic and viscoelastic full-waveform inversion using a spectral-element method on flexible cartesian-based mesh. Geophysics 84, R61–R83.

Tromp, J., Tape, C., Liu, Q., 2005. Seismic tomography, adjoint methods, time reversal and banana-doughnut kernels. Geophysical Journal International 160, 195–216.

Wang, K., Wang, Y., Song, X., Tong, P., Liu, Q., Yang, Y., 2022. Full-waveform inversion of high-frequency teleseismic body waves based on multiple plane-wave incidence: Methods and practical applications. Bulletin of the Seismological Society of America 112, 118–132.

Wang, R., Heimann, S., Zhang, Y., Wang, H., Dahm, T., 2017. Complete synthetic seismograms based on a spherical self-gravitating earth model with an atmosphere–ocean–mantle–core structure. Geophysical Journal International 210, 1739–1764.

Wang, W., Zhang, Z., Zhang, W., Liu, Q., 2023. Implementation of efficient low-storage techniques for 3-d seismic simulation using the curved grid finite-difference method. Geophysical Journal International 234, 2214–2230.

Xie, Z., Komatitsch, D., Martin, R., Matzen, R., 2014. Improved forward wave propagation and adjoint-based sensitivity kernel calculations using a numerically stable finite-element pml. Geophysical Journal International 198, 1714–1747.

Xie, Z., Matzen, R., Cristini, P., Komatitsch, D., Martin, R., 2016. A perfectly matched layer for fluid-solid problems: Application to ocean-acoustics simulations with solid ocean bottoms. The Journal of the Acoustical Society of America 140, 165–175.

Yuan, Y.O., Bozdağ, E., Ciardelli, C., Gao, F., Simons, F.J., 2020. The exponentiated phase measurement, and objective-function hybridization for adjoint waveform tomography. Geophysical Journal International 221, 1145–1164.

Yuan, Y.O., Simons, F.J., Tromp, J., 2016. Double-difference adjoint seismic tomography. Geophysical Journal International 206, 1599–1618.

Zhu, H., Bozdağ, E., Tromp, J., 2015. Seismic structure of the european upper mantle based on adjoint tomography. Geophysical Journal International 201, 18–52.

Zhu, H., Stern, R.J., Yang, J., 2020. Seismic evidence for subduction-induced mantle flows underneath middle america. Nature Communications 11, 2075.

List of Figures

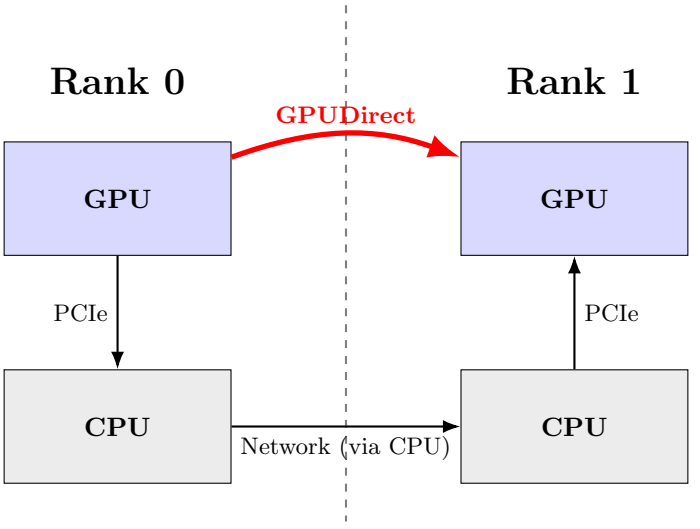

Figure 1: Schematic of data exchange mechanisms between computational nodes. The **black arrows** represent the traditional workflow, which requires staging data through the host CPU memory. In contrast, the **red arrows** highlight the **GPUDirect RDMA** (Remote Direct Memory Access) path, which enables efficient peer-to-peer data transfer between GPUs across different nodes, bypassing the host CPU to significantly reduce latency and overhead.

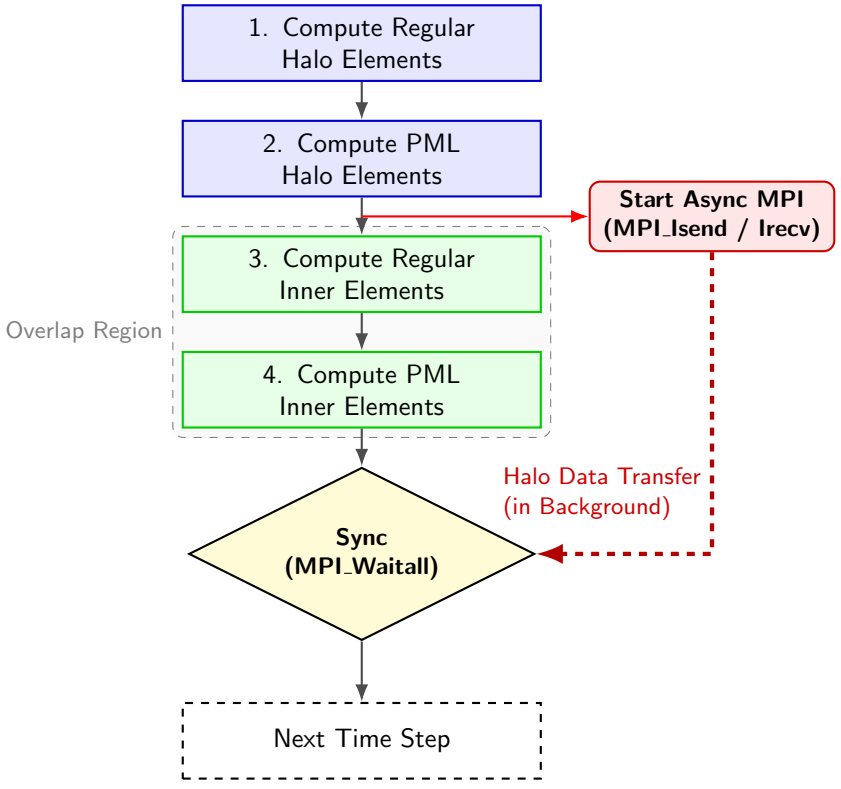

Figure 2: Double-buffering workflow for overlapping communication with computation within a single iteration. Computational elements are partitioned into **halo (boundary) elements (highlighted in purple)**, which are essential for MPI data exchange, and **regular (interior) elements (highlighted in green)**. The solver prioritizes the processing of purple halo elements to facilitate the early packing and initiation of asynchronous MPI routines. While communication occurs in the background, the kernel concurrently processes the green regular elements. This strategy effectively masks data transfer overhead through computational overlap, resulting in substantial performance gains for large-scale simulations.

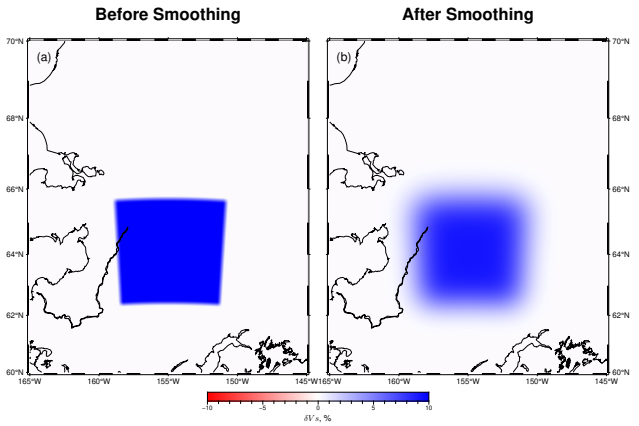

Figure 3: PDE-based Gaussian smoothing with characteristic lengths of $\sigma_r = 8$ km (radial) and $\sigma_\Delta = 50$ km (horizontal). **(a)** Initial 10% velocity perturbation characterized by a rectangular box anomaly with approximate dimensions of $400 \times 400 \times 70$ km. **(b)** Resulting perturbation after the application of the smoothing operator, demonstrating the effective regularization and elimination of high-frequency artifacts while maintaining the large-scale structural features.

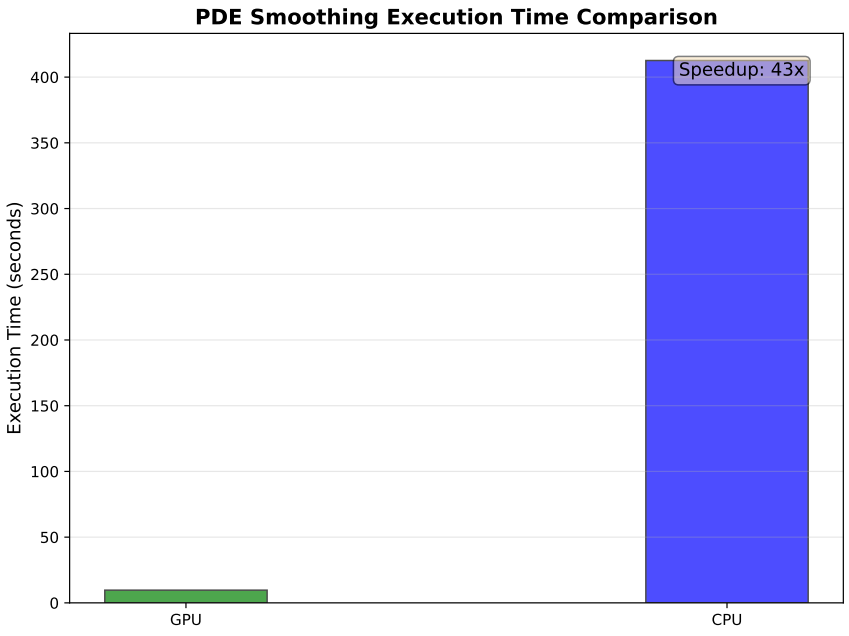

Figure 4: Execution time comparison for PDE-based smoothing implemented on CPU and GPU architectures. The GPU-accelerated version achieves approximately a 43× speedup, demonstrating the efficiency of the spectral element method framework in handling regularization tasks.

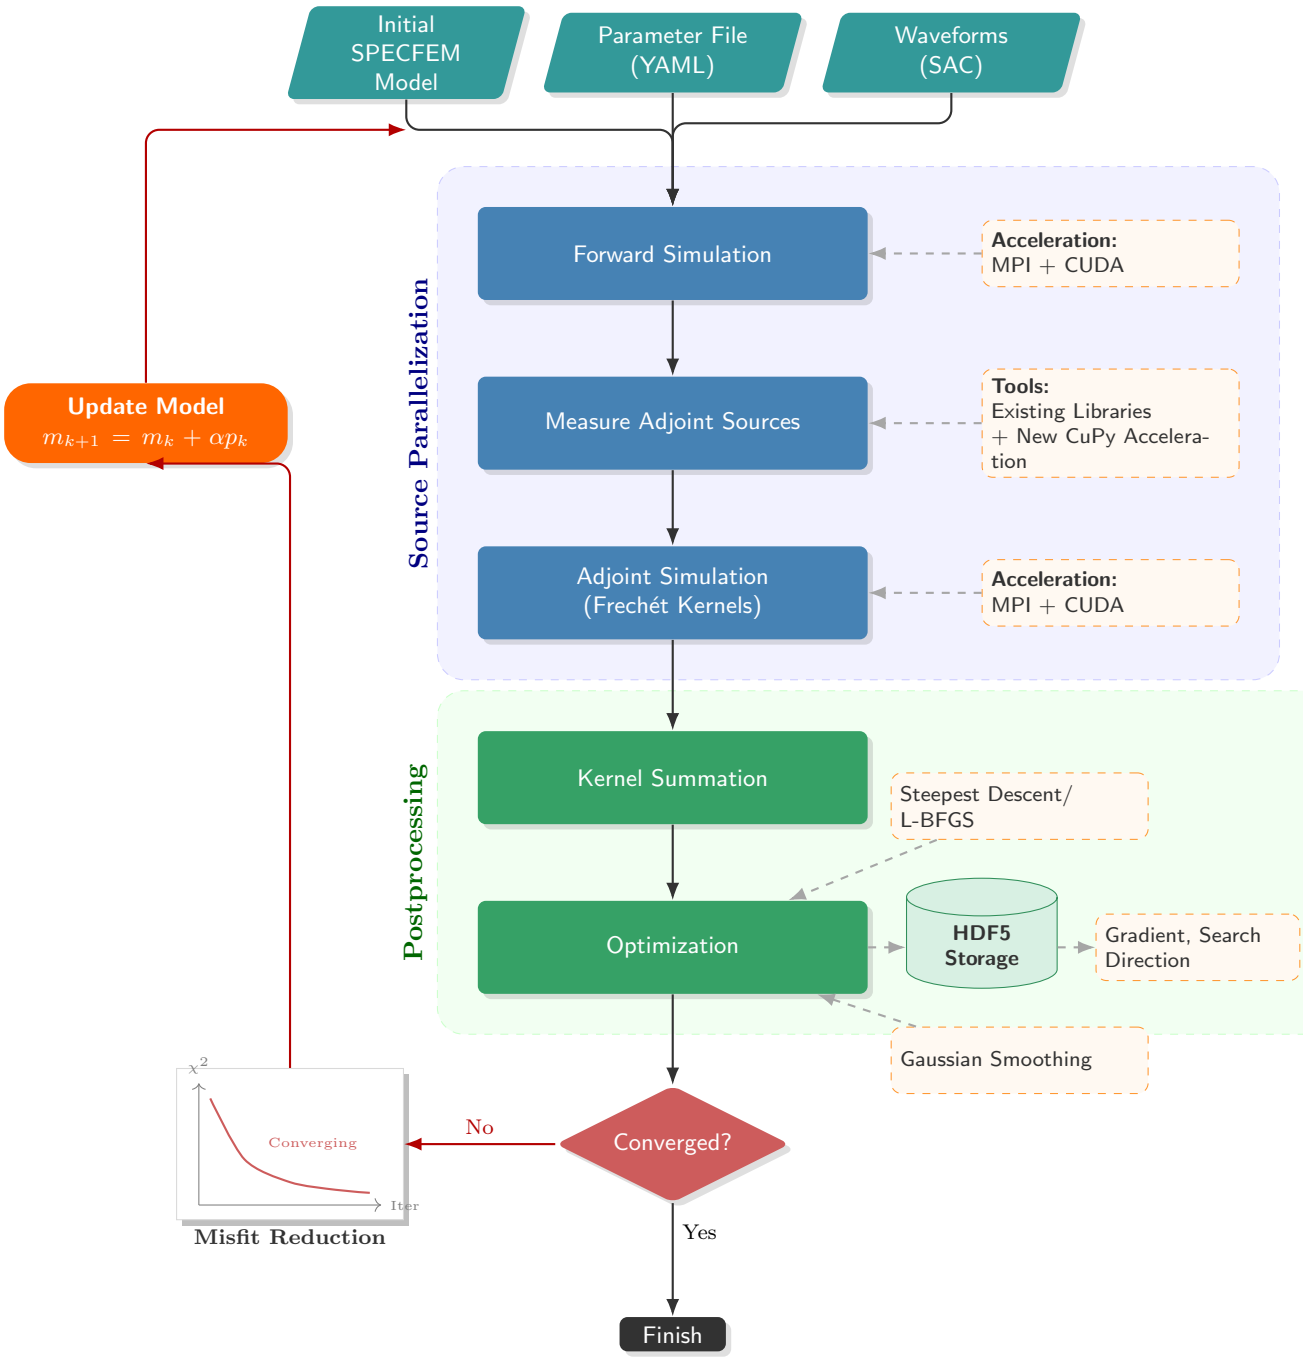

Figure 5: Generate FWI framework.

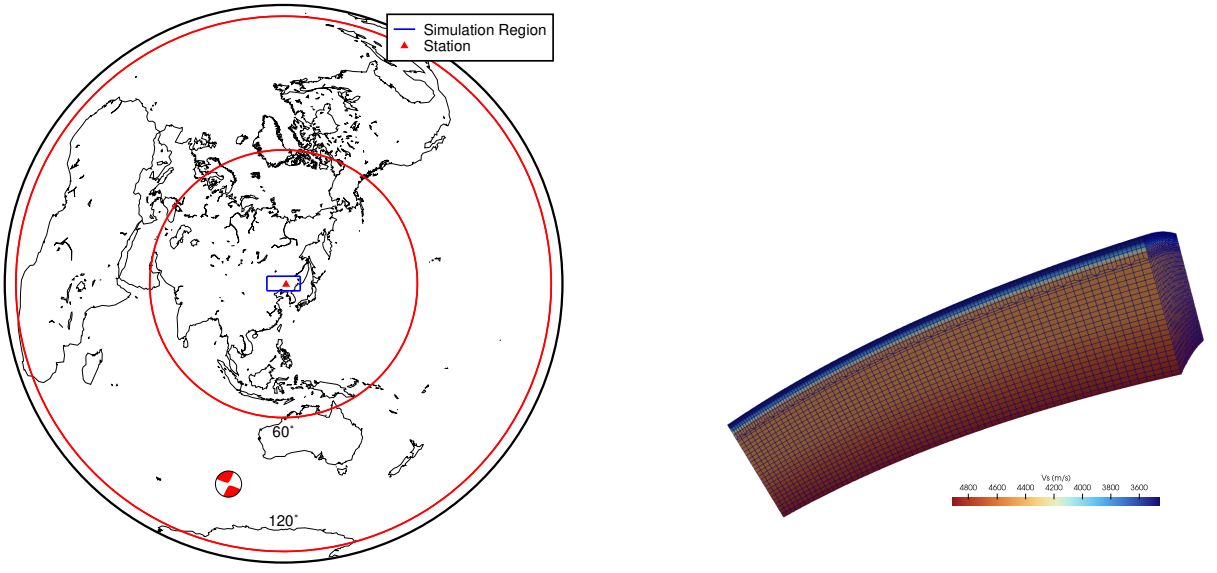

(a) Geographic simulation domain

(b) Computational mesh and velocity model

Figure 6: (a) Geographic extent of the simulation domain (blue box), featuring the seismic source (moment tensor) and the receiver array (red triangles). (b) The unstructured spectral-element mesh used for the simulation. Elements are color-coded by shear-wave velocity (V_S), where warm and cold colors denote low- and high-velocity anomalies, respectively. A mesh-doubling layer is implemented beneath the Moho to optimize computational efficiency.

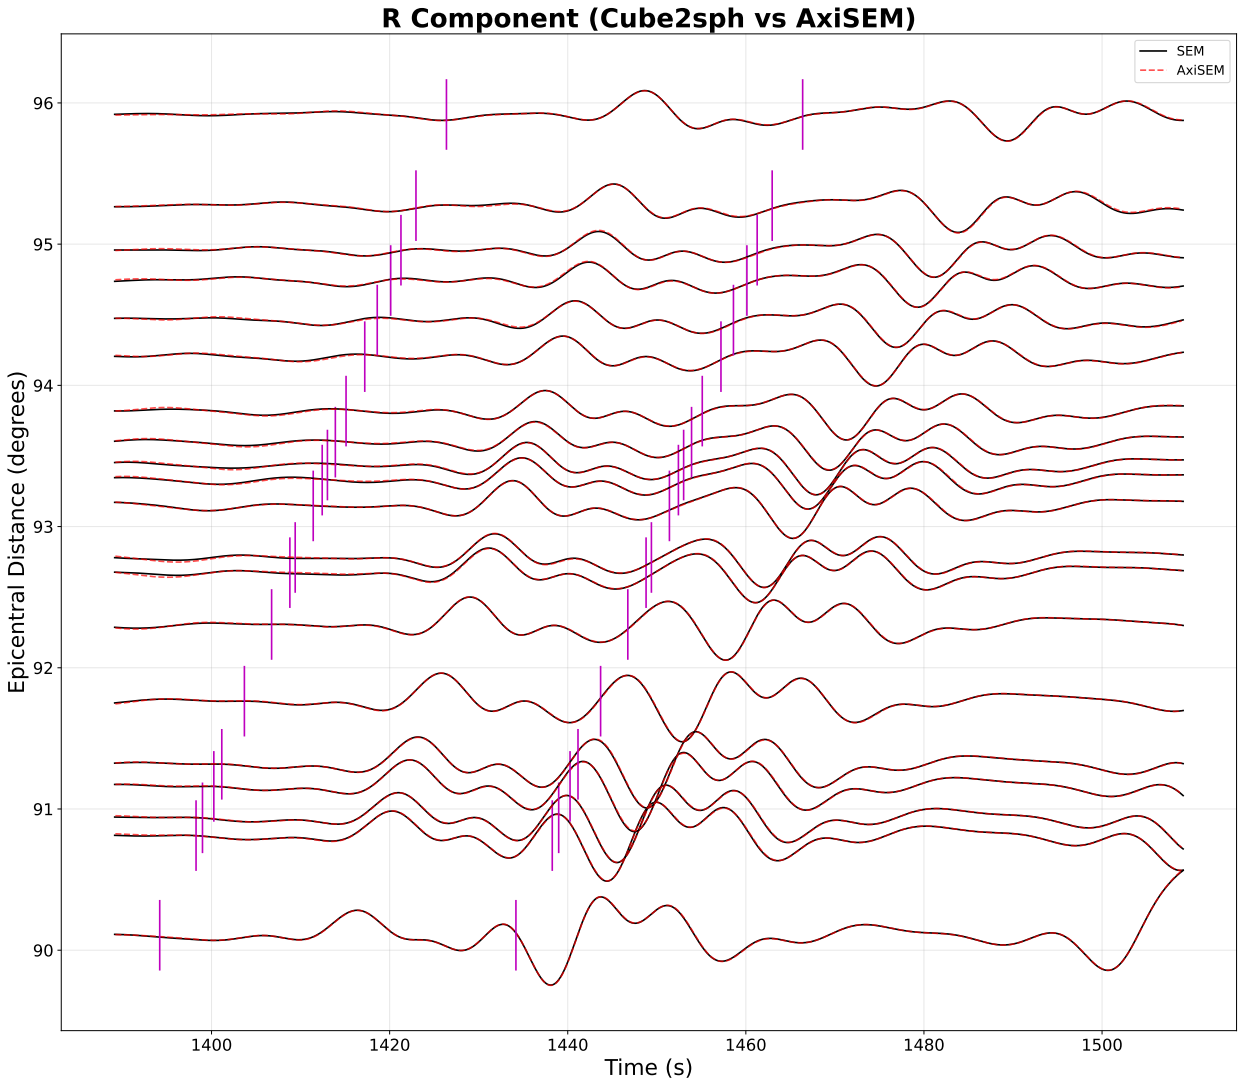

Figure 7: Benchmark of hybrid SEM simulation against AxiSEM reference solutions for the SKS wave window. The magenta bars indicate the specific time windows utilized for analysis ($[t_{SKS} - 15, t_{SKS} + 25]$ s). Black solid lines represent the results from the hybrid simulation, while red dashed lines denote the AxiSEM reference solutions. All seismograms are band-pass filtered between 10–50 s.

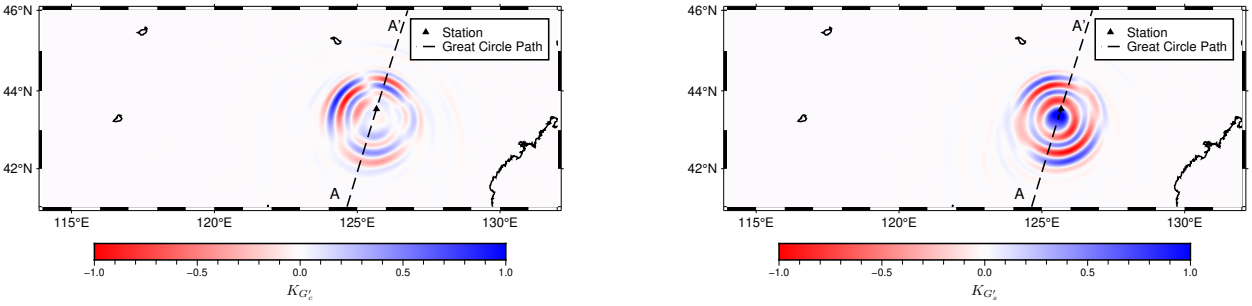

(a) $K_{G'_c}$ (b) $K_{G'_s}$

Figure 8: Horizontal cross-sections of the sensitivity kernels for the anisotropic parameters G'_c and G'_s at a depth of 100 km. Both kernels are normalized to their respective absolute maximum values. Profile A–A′ represents the great-circle path between the source and the target receiver, indicating the location of the vertical slices discussed in Figure 9. (a) Sensitivity kernel $K_{G'_c}$; (b) Sensitivity kernel $K_{G'_s}$.

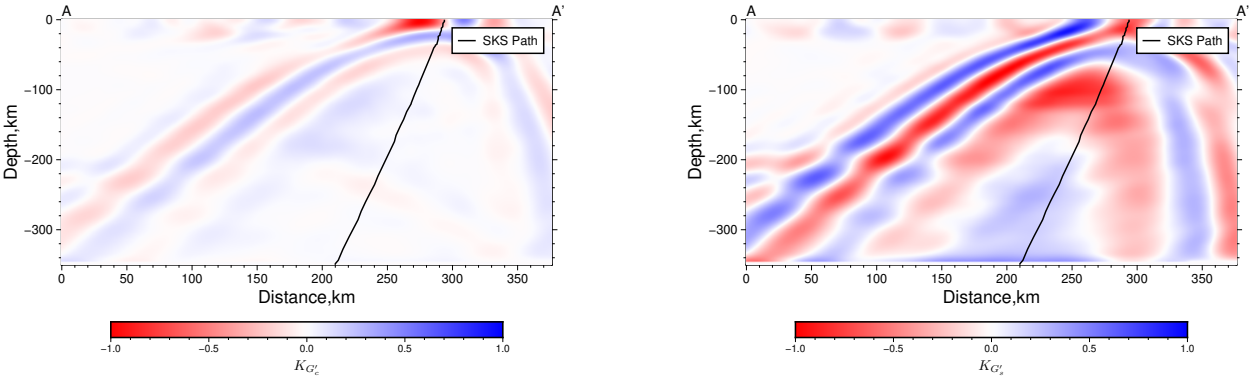

(a) Vertical cross-section of $K_{G'_c}$

(b) Vertical cross-section of $K_{G'_s}$

Figure 9: Vertical cross-sections of the sensitivity kernels $K_{G'_c}$ and $K_{G'_s}$ along the great-circle profile A–A′ defined in Figure 8. The kernels exhibit characteristic Fresnel zone structures associated with finite-frequency SKS phases, centered around the theoretical ray path. All values are normalized to the absolute maximum of the respective kernel. (a) Vertical slice of $K_{G'_c}$; (b) Vertical slice of $K_{G'_s}$.

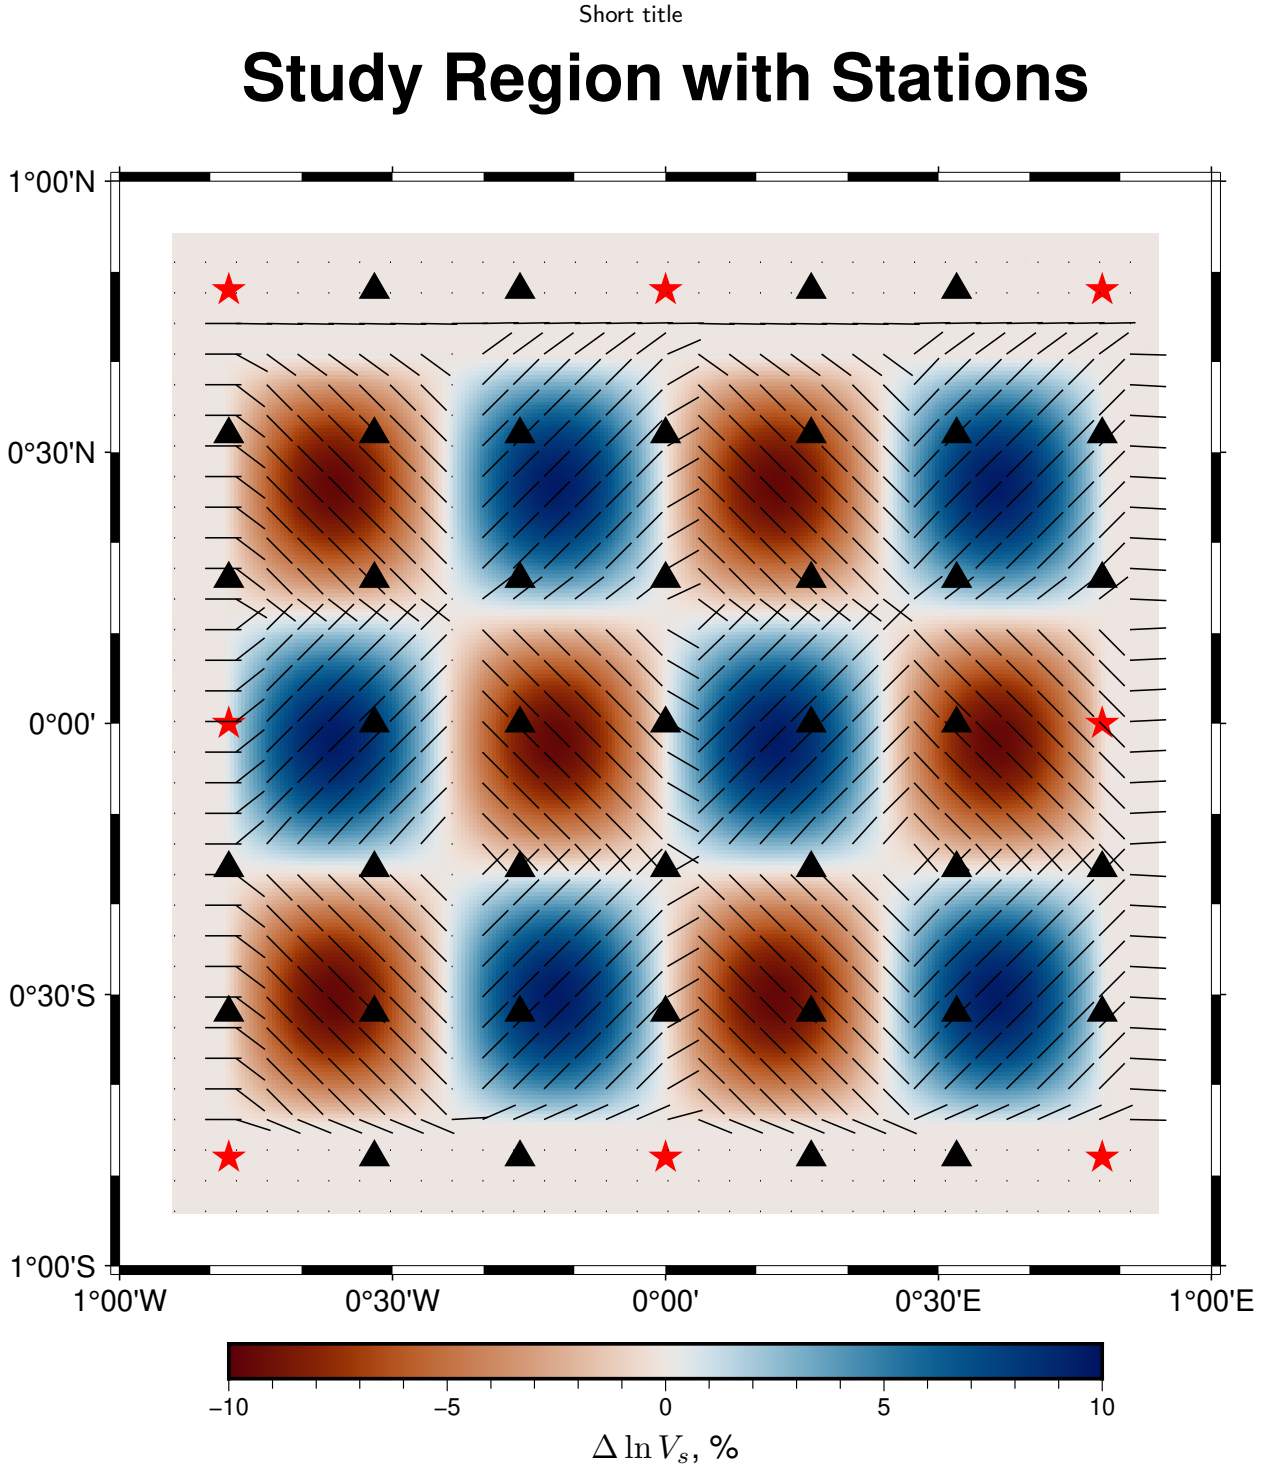

Figure 10: Station configuration for the multi-component ambient noise FWI checkerboard test. Red asterisks denote the master stations where orthogonal (ENZ) single forces are applied to simulate the noise source distribution. The underlying checkerboard patterns for isotropic shear-wave perturbations ($\Delta \ln V_S$) and fast-axis orientations (ϕ) illustrate the target anomalies. Black bars represent the fast-axis orientation, with vectors masked in regions where the anisotropic magnitude G_0' is below a defined significance threshold.

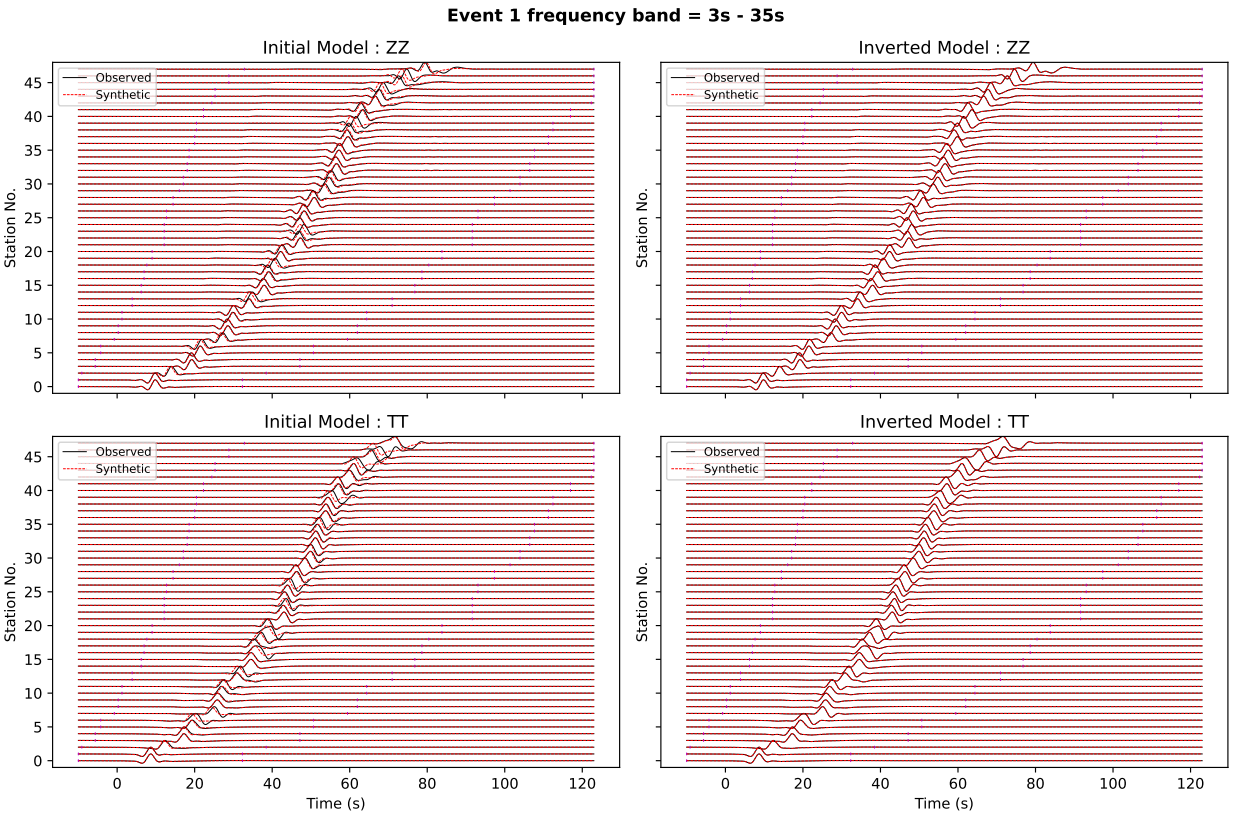

Figure 11: Comparison of synthetic Green's functions in frequency band 3-35s. for the initial and inverted models (Event 1). Top row: Z-Z components (Rayleigh waves) for the (left) initial and (right) inverted models. Bottom row: T-T components (Love waves) for the (left) initial and (right) inverted models.

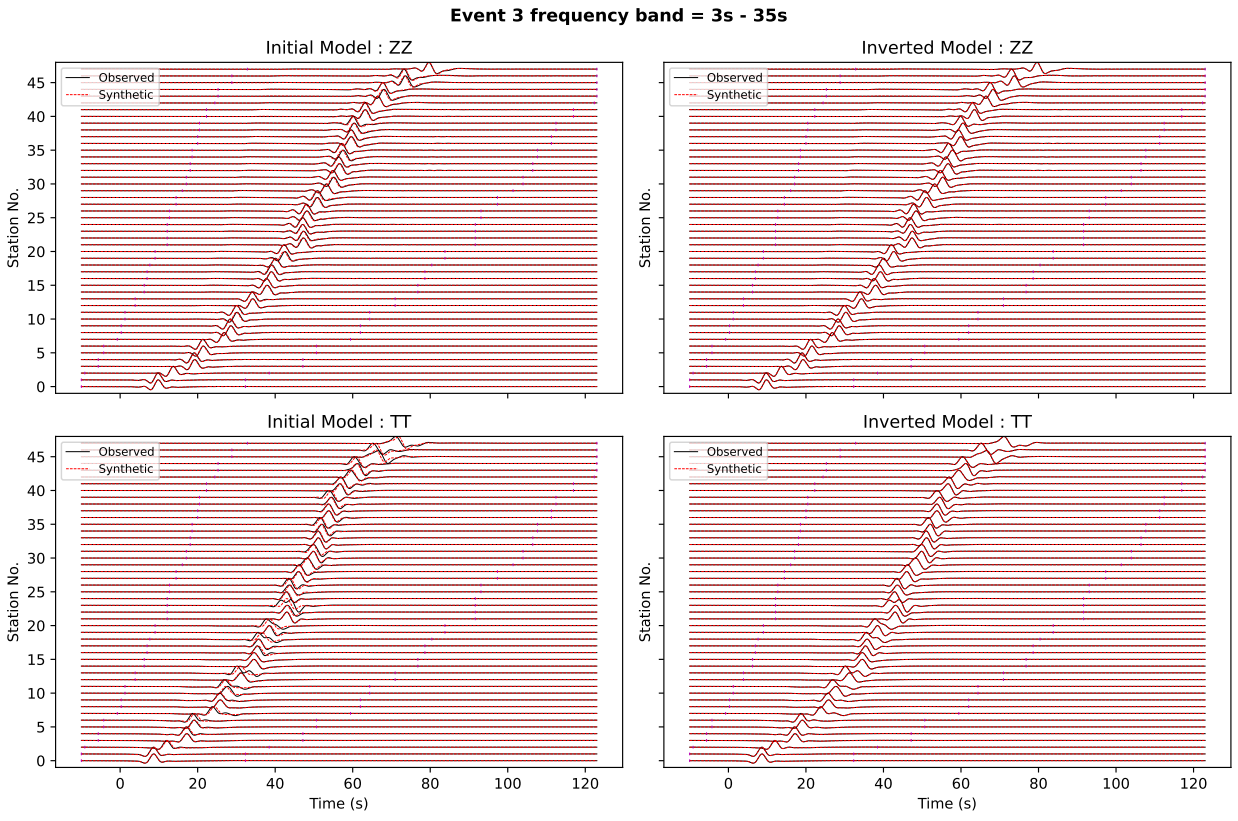

Figure 12: Comparison of synthetic Green's functions for Event 3, following the format of Figure 11.

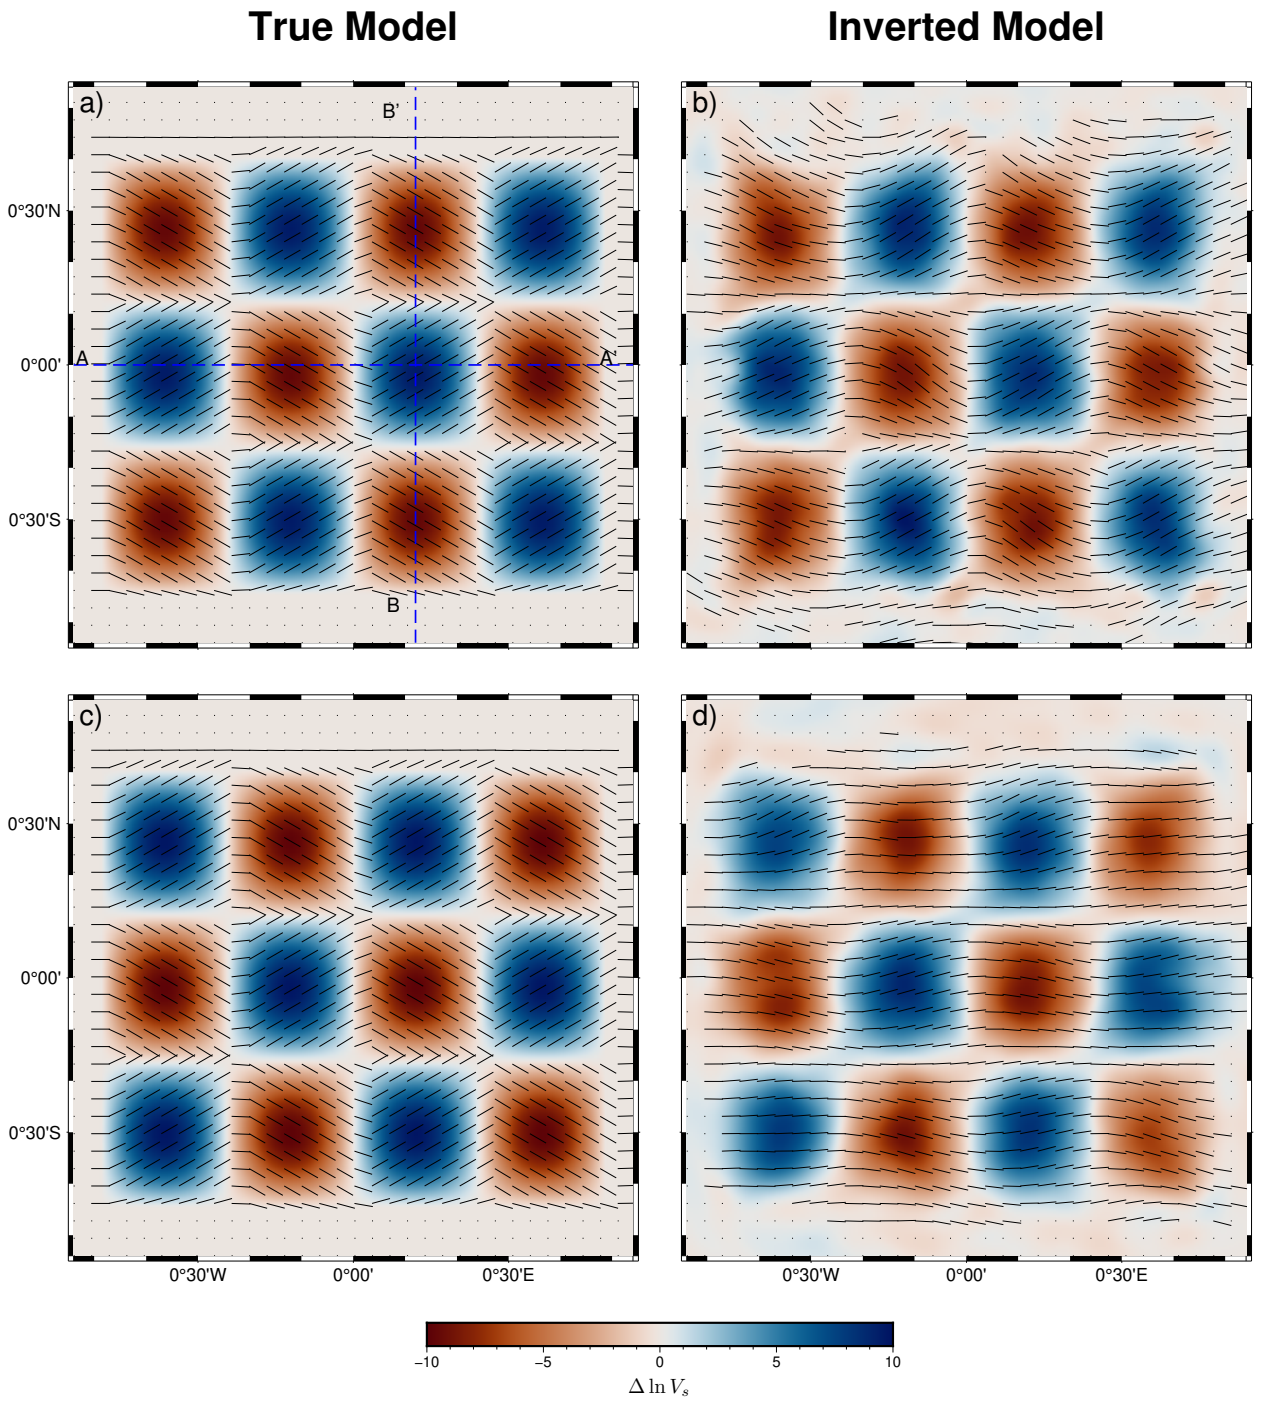

Figure 13: Horizontal cross-sections of the recovered isotropic shear-wave perturbations ($\Delta \ln V_S$) and fast-axis orientations (ϕ). Panels (a) and (b) compare the target and recovered models at 5 km depth, while (c) and (d) show the target and recovered models at 22.5 km depth. The black bars represent the orientation of the azimuthal anisotropy. Blue dashed lines A–A′ and B–B′ indicate the locations of the vertical profiles analyzed in Figure 14.

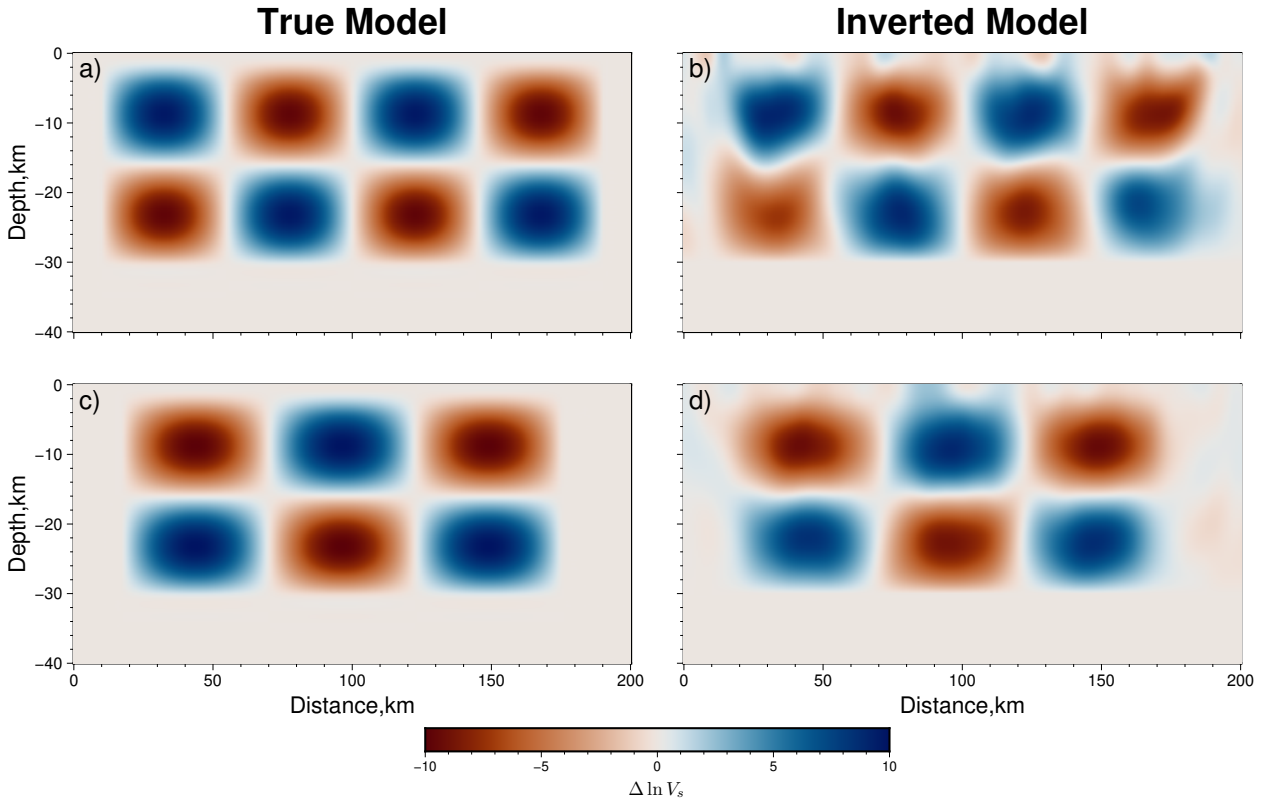

Figure 14: Vertical cross-sections of the isotropic shear-wave recovery ($\Delta \ln V_S$) along profiles A–A′ and B–B′. (a) Target and (b) recovered perturbations for profile A–A′. (c) Target and (d) recovered perturbations for profile B–B′. The vertical axis extends from the surface to a depth of 40 km.

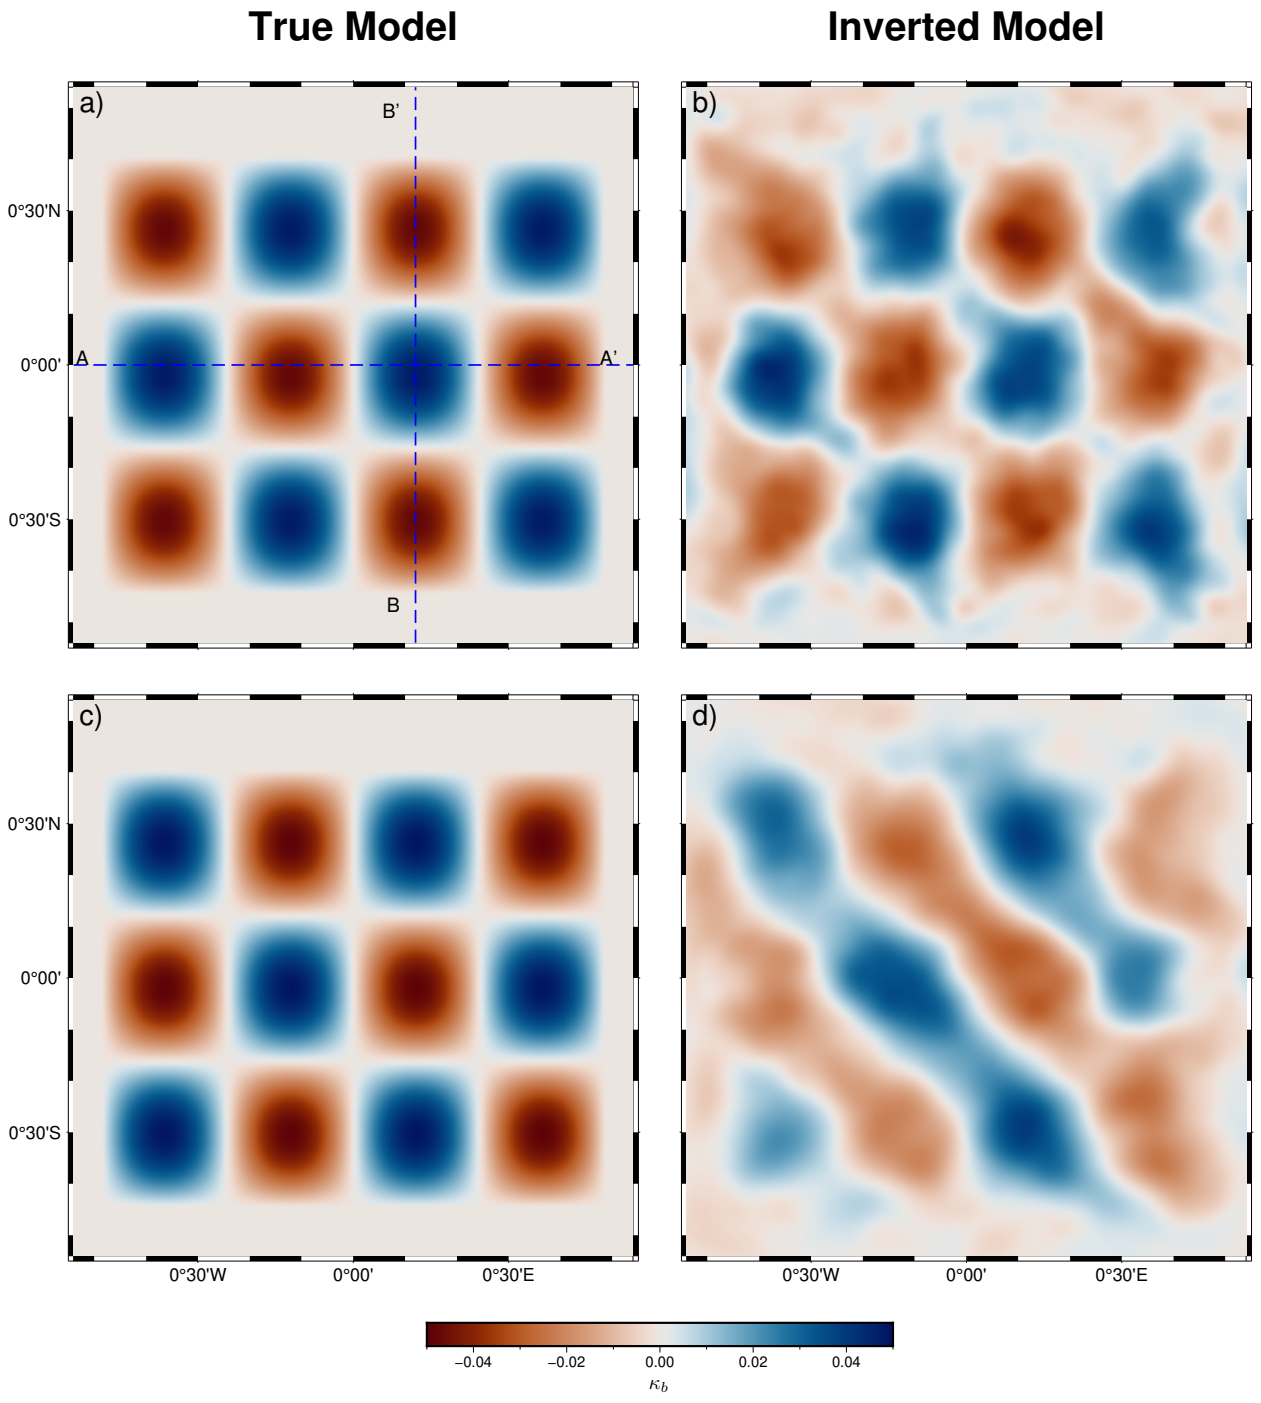

Figure 15: Horizontal cross-sections showing the recovery of the radial anisotropy ratio κ_β. (a) Target and (b) recovered perturbations at 5 km depth; (c) target and (d) recovered perturbations at 22.5 km depth.

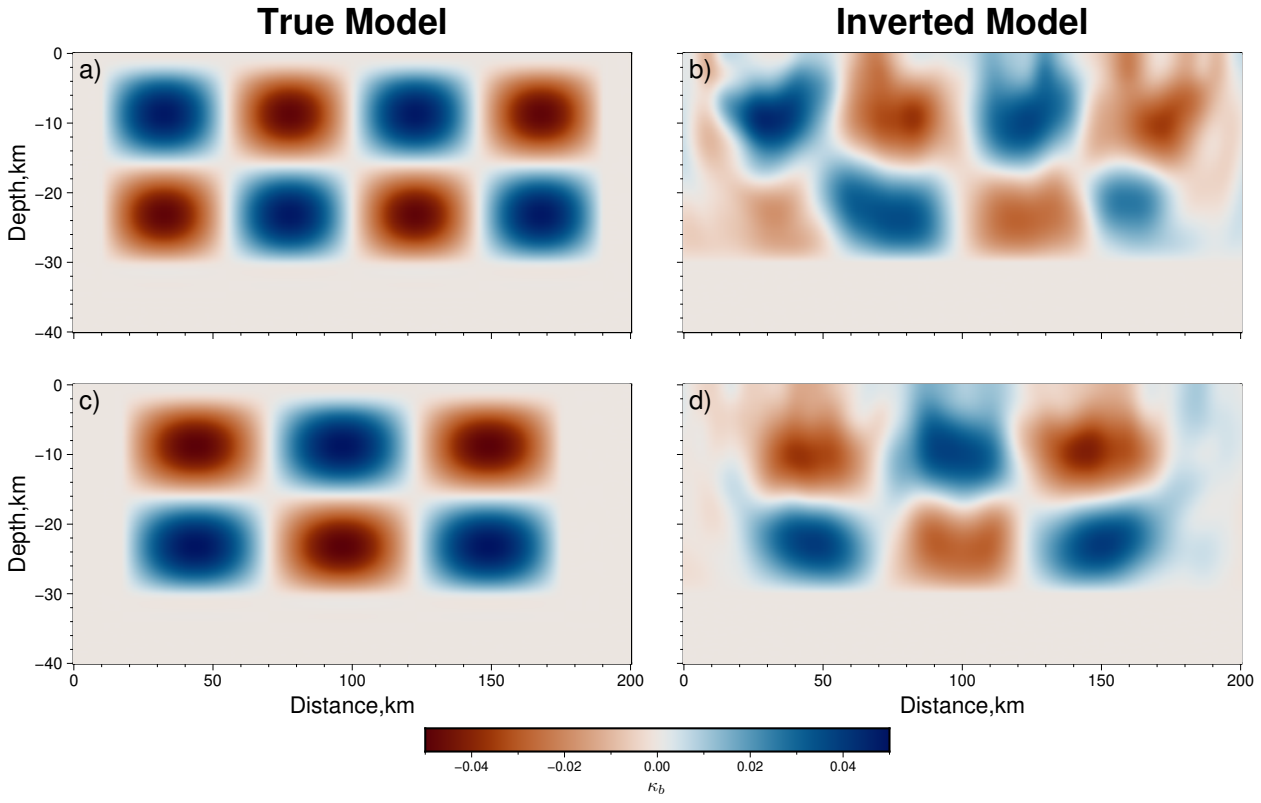

Figure 16: Vertical cross-sections of the radial anisotropy recovery ($\delta\kappa_\beta$) along profiles A–A′ and B–B′. (a) Target and (b) recovered perturbations for profile A–A′. (c) Target and (d) recovered perturbations for profile B–B′. The vertical axis spans the 40 km depth of the simulation domain.