Highlights

SpecFWAT: Specfem3D-based Full-waveform Adjoint Tomography Package for High-resolution Lithospheric Imaging

Mijian Xu,Kai Wang,Nanqiao Du,Tianshi Liu,Bin He,Ting Lei,Jing He,Ping Tong,Qinya Liu

- SpecFWAT is an open-source, high-performance package for lithospheric-scale full-waveform adjoint tomography using multiple datasets.
- GPU-accelerated solvers for PDE-based smoothing and iterative deconvolution achieve significant speedups.
- Implementation of the cross-convolution misfit function significantly mitigates source-time function uncertainties in teleseismic imaging.
- Receiver function adjoint tomography accurately recovers sharp crustal velocity gradients and complex interface geometries.
- A weighted non-dimensionalized sumation strategy effectively balances diverse observables in joint full-waveform adjoint tomography.

SpecFWAT: Specfem3D-based Full-waveform Adjoint Tomography Package for High-resolution Lithospheric Imaging

Mijian Xu[a,b,*], Kai Wang[c,*], Nanqiao Du[d], Tianshi Liu[d], Bin He[a], Ting Lei[d], Jing He[a,e], Ping Tong[f,g] and Qinya Liu[a,d]

[a]Department of Physics, University of Toronto, Toronto, ON, Canada
[b]Faculty of Arts & Science, University of Toronto, Toronto, ON, Canada
[c]School of Earth and Space Sciences, University of Science and Technology of China, Hefei, China
[d]Department of Earth Sciences, University of Toronto, Toronto, ON, Canada
[e]National Institute of Natural Hazards, Ministry of Emergency Management of China, Beijing, China
[f]Division of Mathematical Sciences, School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore, Singapore
[g] Earth Observatory of Singapore, Nanyang Technological University, Singapore, Singapore

ARTICLE INFO

Keywords:
Full-waveform inversion
Adjoint tomography
Crust and upper mantle
Seismic velocity
Spectral element method

ABSTRACT

We present SpecFWAT, a high-performance, open-source software package designed for full-waveform adjoint tomography at the lithospheric scale using multiple data types. Building upon the spectral-element forward solver SPECFEM3D, SpecFWAT introduces a modernized inversion framework featuring an object-oriented Fortran design and a CMake-based build environment. To enhance usability and computational efficiency, the package also integrates a Python-based graphical user interface for interactive data quality control, provides a library of diverse built-in misfit functions, and implements GPU-accelerated utilities for deconvolution and kernel smoothing. Supported by these features, the software facilitates the choice of different individual inversion schemes, including ambient noise, teleseismic waveform inversion, and receiver-function-based teleseismic adjoint tomography. Notably, it incorporates a cross-convolution misfit to mitigate source-time function uncertainties in teleseismic imaging, as well as a receiver-function-based misfit to improve crustal resolution. To further resolve multi-scale lithospheric structures, SpecFWAT offers a flexible joint inversion framework based on a weighted non-dimensionalized gradient summation strategy. A real-world application to the Hangay Dome in Central Mongolia demonstrates the software's efficacy in resolving complex magmatic systems by combining the vertical resolution of teleseismic waveforms with the lateral resolution of ambient noise data. As a scalable and extensible platform, SpecFWAT provides a unified solution for advancing high-resolution lithospheric imaging in diverse tectonic regimes.

CRediT authorship contribution statement

Mijian Xu: Software, Methodology, Formal analysis, Investigation, Writing - Original Draft . **Kai Wang:** Software, Methodology, Writing - Review & Editing. **Nanqiao Du:** Methodology, Software. **Tianshi Liu:** Methodology, Software. **Bin He:** Formal analysis, Writing - Review & Editing. **Ting Lei:** Formal analysis, Writing - Review & Editing. **Jing He:** Writing - Review & Editing. **Ping Tong:** Conceptualization, Funding acquisition, Writing - Review & Editing. **Qinya Liu:** Conceptualization, Supervision, Funding acquisition, Writing - Review & Editing.

1. Introduction

In the past two decades, the development of full-waveform inversion (FWI, also referred to as full-waveform adjoint tomography or FWAT) has significantly advanced quantitative imaging of the Earth's interior (Tromp, 2019; Fichtner

gomijianxu@gmail.com (M. Xu); wangk@ustc.edu.cn (K. Wang)
ORCID(s): 0000-0001-8888-8523 (M. Xu); 0000-0002-6189-7056 (K. Wang); 0000-0003-3623-7686 (N. Du); 0000-0001-9636-3268 (T. Liu); 0000-0002-5994-1652 (B. He); 0000-0001-6732-434X (T. Lei); 0000-0001-6395-6585 (J. He); 0000-0002-1937-3427 (P. Tong); 0000-0002-1071-2314 (Q. Liu)

et al., 2024). By iteratively minimizing the misfit between observed and synthetic waveforms, FWI enables the recovery of detailed seismic velocity structures at multiple spatial scales (e.g., Lei et al., 2020; Tape et al., 2009; Wang et al., 2016; Zhu et al., 2015). With the increase of computational resources and improvements of dense instrumentation, FWI has become a feasible approach for lithospheric-scale imaging. However, conventional FWI applications have primarily relied on regional earthquake data, where limited ray coverage, source uncertainties and uneven source–receiver geometry often restrict spatial resolution, especially in regions with sparse seismicity and complex tectonic settings.

Recent developments have extended the FWI framework to incorporate multiple seismic observations, which provide complementary sensitivity to structures at different depths and scales. Surface waves reconstructed from ambient seismic noise offer complementary path coverage to local earthquakes, improving lateral resolution in the lithosphere (Chen et al., 2014; Sager et al., 2018, 2020; Wang et al., 2018, 2025), although how well the empirical Greens function (EGF) matches with the true Greens function is still being debated (Tromp et al., 2010; Fichtner, 2015). In contrast, teleseismic P-coda waves and scattered phases are more sensitive to velocity gradients and fine-scale heterogeneities, providing stronger vertical resolution (Monteiller et al., 2015; Beller et al., 2018; Wang et al., 2016, 2021a). receiver function (RF), derived from teleseismic waveforms through deconvolution, emphasize converted phases at higher frequencies and thus improve the vertical resolution of lithospheric imaging (Xu et al., 2023, 2025). The joint utilization of these datasets within a consistent adjoint tomography framework exploits their complementary sensitivities, leading to improved model accuracy and stability of the recovered lithospheric models (He et al., 2024; Wang et al., 2021b; Xu et al., 2025). Nevertheless, practical implementation of such combined inversions remains challenging because it requires efficient numerical solvers, flexible parallelization strategies, and automated workflows that can manage diverse data types and computational demands.

To address the need for high-fidelity numerical wavefield simulation, the spectral-element method (SEM) has become a leading approach. By combining the geometric flexibility of finite-element methods with the rapid convergence of spectral methods, SEM can naturally and accurately accommodate complex surface topographies, solid-fluid interfaces, and strong 3D structural heterogeneities (Komatitsch and Vilotte, 1998; Komatitsch and Tromp, 1999). Powered by open-source packages, such as SPECFEM3D and Salvus (Komatitsch and Vilotte, 1998; Afanasiev et al., 2019), SEM has been incorporated in existing earthquake FWI frameworks such as SeisFlows and LASIF (Modrak et al., 2018; Krischer et al., 2015). These tools have been successfully applied to seismic imaging across scales, ranging from global and continental models to regional and local studies (e.g., Gao et al., 2021; Lei et al., 2020; Chow et al., 2022; Rodgers et al., 2022; Van Herwaarden et al., 2023). However, they primarily rely on complex Python scripting to orchestrate the workflow. In high-performance computing (HPC) environments, managing and debugging these multi-layered scripts, alongside handling the I/O overhead between separated solvers and post-processing tools, is often cumbersome. Furthermore, lacking support for teleseismic-based and joint full-waveform inversion, these tools

offer limited resolution in the lithospheric scale.

To bridge this gap, we present SpecFWAT, a high-performance, open-source software package (released under the GPLv3 license) explicitly designed for FWI at the lithospheric scale using multiple data types, including local earthquakes, ambient noise, teleseismic body waves, and receiver functions. SpecFWAT builds directly upon the open-source solver SPECFEM3D (Komatitsch and Vilotte, 1998), inheriting its numerical accuracy and scalability for forward simulation and kernel computation, while introducing a modernized, tightly integrated framework for the inversion process. The software automates the iterative minimization of the misfit between observed and synthetic waveforms: in each iteration, a forward simulation generates synthetic data to compute residuals, which are then back-propagated in an adjoint simulation to construct the raw sensitivity kernel (Tromp et al., 2005; Fichtner et al., 2006). This kernel is subsequently regularized (e.g., Gaussian smoothed) to obtain the gradient (Tape et al., 2010), guiding a line search optimization algorithm to determine the optimal step length for model updates (Zhu et al., 2015). By integrating these procedures, SpecFWAT aims to build a unified and extensible framework for FWI, enabling researchers to efficiently explore the multi-scale structures of the lithosphere in diverse tectonic regimes, including continental rifts, volcanic regions, subduction zones, and orogenic belts.

2. Code Features

2.1. Overview

SpecFWAT is written in modern Fortran, using an object-oriented programming design with a CMake-based build environment, which improves modularity, maintainability, and extensibility. All model and gradient vectors are stored in the HDF5 format, while user configurations are defined through human-readable YAML files for clarity. It supports both Graphics Processing Units (GPU) and CPU parallelization in wavefield simulations, measurements of adjoint source, and regularization of misfit kernel, which allow efficient deployment from local workstations to high-performance computing clusters. In addition, SpecFWAT is complemented by an optional Python module, PyFWAT, which provides tools for data preprocessing, event selection, and quality control through a graphical user interface (GUI), as well as utilities for model visualization and analysis. Comprehensive online documentation for SpecFWAT is provided to support installation, configuration, and demonstrate workflows.

SpecFWAT streamlines the iterative process of FWI into four distinct computational stages (Figure 1). Each stage is managed by a dedicated high-performance executable, allowing the entire inversion loop to be controlled by a concise shell script. A representative workflow for n iterations is illustrated in Algorithm (1).

The workflow begins with `xfwat_mesh_databases`, which projects velocity and density models from input regular grids onto the Gauss-Lobatto-Legendre (GLL) points of hexahedral elements in the mesh required by the SEM. The core computational workload is handled by `xfwat_fwd_measure_adj`, which provides flexible execution strategies

Algorithm 1 Workflow of SpecFWAT

for iteration = 0 ... n **do**

1. Assigning model parameters to spectral-element mesh (`xfwat_mesh_databases`);

2. Forward and adjoint simulations (`xfwat_fwd_measure_adj`);

3. Sum, pre-conditionering, and smoothing of misfit kernel (`xfwat_post_proc`);

4. Line search and optimization (`xfwat_optimize`);

end for

tailored for HPC environments. In its serial mode, the executable employs an internal Fortran-based loop to sequentially conduct forward simulation, adjoint source construction, and kernel simulation for multiple events. Incorporating these coupled tasks into a single executable avoids the instability of complex external scripting and offers excellent fault tolerance, enabling users to seamlessly restart from the point of failure if a cluster job times out. Alternatively, for large-scale inversions, SpecFWAT retains the flexibility of existing tools like Seisflows and LASIF by fully supporting array jobs. This allows users to distribute individual earthquake events across multiple compute nodes, achieving massive task-level parallelism.

Following the kernel calculation, `xfwat_post_proc` performs kernel summation, preconditioning, and smoothing to preparing the gradients for model updating. Finally, `xfwat_optimize` updates model parameters using standard optimization algorithms (e.g., steepest descent (SD), nonlinear conjugate gradient (NLCG), and limited-memory Broyden–Fletcher–Goldfarb–Shanno (L-BFGS)) combined with a line search based on the Armijo condition (Nocedal and Wright, 2006). Additionally, each executable produces standardized outputs and time-stamped log files. This design ensures complete traceability of the computational process, allowing users to focus on model configuration rather than script maintenance. Ultimately, by automating and streamlining iterative workflow, the SpecFWAT enables users to focus on constructing the initial model, ensuring data quality, and fine-tuning inversion configurations.

2.2. Data Preparation

SpecFWAT establishes an integrated inversion framework compatible with the use of local earthquake waveforms, ambient noise EGFs, and teleseismic waveforms. To initialize the solver, users are required to provide specific source parameter files corresponding to the data type: moment-tensor solutions for local earthquakes, point-force solutions for ambient noise virtual sources, and plane wave parameters (defined by back-azimuth and incidence angles) for teleseismic events. For waveform preparation, both local earthquake and teleseismic data must be pre-processed to remove linear trends, mean values, and instrument responses, whereas ambient noise data are processed following standard practices in ambient noise tomography (Bensen et al., 2007).

To streamline quality control across these diverse datasets, a utility module, PyFWAT, offers an interactive GUI built upon PyQt and Matplotlib (Figure 2). For teleseismic data specifically, this interface facilitates waveform visualization and alignment using either P-wave arrival times predicted by the TauP toolkit (Crotwell et al., 1999) or

multi-channel cross-correlation (MCCC) (VanDecar and Crosson, 1990). The aligned P-wave arrival time will be used to associate with predicted teleseismic FK traveltime for adjoint source measurements (Wang et al., 2021a). Across all data types, the GUI facilitates visual inspection and efficient rejection of low-quality traces via simple mouse clicks, ensuring that only reliable data are utilized in the subsequent inversion.

2.3. Window Selection

Following data preparation and quality control, a flexible scheme for time windowing of waveforms is critical for extracting seismic phases for kernel computation. To accommodate diverse imaging targets, SpecFWAT implements four distinct window selection strategies. For local earthquake tomography, the package integrates two established automated algorithms: the windowing method employed in the LASIF package (Krischer et al., 2015) and a re-implementation of the Flexwin package (Maggi et al., 2009). Alternatively, for scenarios targeting specific local or teleseismic body-wave phases (Beller et al., 2018; Chen et al., 2023), windows can be explicitly defined by a fixed time interval centered around reference arrival times. Furthermore, to extract multi-period dispersive surface waves in both ambient noise EGFs and local earthquake waveforms, the software implements a dynamic windowing strategy based on reference group velocities in which time windows are adaptively defined by the travel times corresponding to specific group velocity ranges for different period bands (Wang et al., 2018) .

2.4. Adjoint Sources

In the adjoint tomography framework, sensitivity kernels are computed by interaction between forward and adjoint wavefield. The adjoint wavefield is generated by placing time-reversed adjoint sources at the receiver locations. The adjoint source corresponds to the derivative of the objective function with respect to the synthetic displacement (Tromp et al., 2005; Fichtner et al., 2006). In SpecFWAT, the implementation of adjoint source construction follows the `measure_adj` module in SPECFEM3D, but is fully restructured in modern Fortran with an object-oriented and modular design. In addition to the objective functions inherited from `measure_adj`, SpecFWAT introduces new objective functions proposed by recent studies to enhance flexibility and robustness (Gao et al., 2023). These built-in objective functions include multitaper-based cross-correlation traveltime and amplitude measurements (Tape et al., 2010), waveform L_2 norm misfit, exponentiated phase misfit (Yuan et al., 2020), cross-correlation coefficient (Tao et al., 2018), receiver function based misfit (Xu et al., 2023), and cross-convolution misfit of teleseismic waveforms (Xu et al., 2024). The modular code structure further allows additional misfit functions to be easily incorporated.

Moreover, calculation of both teleseismic and receiver function adjoint source involving time-domain iterative deconvolution (Ligorría and Ammon, 1999) are implemented with GPU parallel support. The deconvolution module is implemented using custom CUDA kernels, featuring memory-optimized element-wise operations and asynchronous CPU–GPU execution to reduce communication overhead during iterative updates. This design follows a hybrid work-

flow in which highly parallelizable vector operations are offloaded to the GPU, while convergence control remains on the CPU. For a receiver function calculation with 13,000 samples, the GPU implementation requires only 0.66 s for 400 iterations on NVIDIA H100 GPU, achieving a 10.5× speedup compared with the CPU version (Figure 3). These CUDA-based optimizations enable efficient deconvolution performance for the source time function (STF), receiver function and related adjoint source calculation.

2.5. Post-processing of Misfit Kernel

During each iteration in the adjoint tomography, the sensitivity kernels from individual events are summarized for a subsequent regularization with a Gaussian smoothing. To improve the efficiency of the smoothing, we employ the PDE smoothing technique used in image processing (Weickert, 1998) and recently in full-waveform inversions (Liu, 2024). The Gaussian smoothing is considered equivalent to solving the following anisotropic diffusion equation

$$\frac{\partial u(\mathbf{x},t)}{\partial t} = \nabla \cdot (\mathbf{D} \cdot \nabla u(\mathbf{x},t)), \tag{1}$$

where $u(\mathbf{x},t)$ represents the kernel field at pseudo-time t, with the initial condition $u(\mathbf{x},0) = K(\mathbf{x})$ being the raw sensitivity kernel computed from adjoint simulations. The D is the diffusion tensor, which allows for anisotropic smoothing lengths. The smoothing operation is controlled by the Gaussian width σ_h in the horizontal direction and σ_v in the vertical direction. The total integration time T relates to the Gaussian width via $max(\sigma_h, \sigma_v) = \sqrt{2DT}$.

We numerically solve Eq. (1) in the spectral element domain to ensure that the mesh discretization of smoothed kernels remains consistent with that used in the forward and adjoint simulations. We derive the weak formulation by multiplying Eq. (1) with a test function $w(\mathbf{x})$ and integrating over the model domain Ω. By applying Green's first identity and assuming homogeneous Neumann boundary conditions ($\nabla u \cdot \mathbf{n} = 0$ on $\partial\Omega$, which implies no energy flow across boundaries), the variational form can be expressed as:

$$\int_\Omega w \frac{\partial u}{\partial t} d\mathbf{x} = -\int_\Omega \nabla w \cdot (\mathbf{D} \cdot \nabla u) d\mathbf{x}. \tag{2}$$

We then discretize the continuous field $u(\mathbf{x},t)$ using a linear combination of Lagrange basis functions $\psi_j(\mathbf{x})$ defined at all the GLL points of the domain:

$$u(\mathbf{x},t) \approx \sum_{j=1}^{N} U_j(t)\psi_j(\mathbf{x}), \tag{3}$$

where N is the total number of grid points and $U_j(t)$ represents the discrete nodal values of the kernel field. Following the standard Galerkin approach, we choose the test functions to be the same as the basis functions (i.e., $w = \psi_i$ for

$i = 1, \ldots, N$). Substituting Eq. (3) into Eq. (2) yields a system of linear equations:

$$\sum_{j=1}^{N} \left(\int_{\Omega} \psi_i \psi_j d\mathbf{x} \right) \frac{dU_j}{dt} = - \sum_{j=1}^{N} \left(\int_{\Omega} \nabla \psi_i \cdot (\mathbf{D} \cdot \nabla \psi_j) d\mathbf{x} \right) U_j. \tag{4}$$

This system can be written in compact matrix notation as:

$$\mathbf{M} \frac{d\mathbf{U}}{dt} = -\mathbf{K}\mathbf{U}, \tag{5}$$

where $\mathbf{U}$ is the vector of nodal values $[U_1, \ldots, U_N]^T$. The elements of the mass matrix $\mathbf{M}$ and the stiffness matrix $\mathbf{K}$ are explicitly defined as $M_{ij} = \int_{\Omega} \psi_i \psi_j d\mathbf{x}$ and $K_{ij} = \int_{\Omega} \nabla \psi_i \cdot (\mathbf{D} \cdot \nabla \psi_j) d\mathbf{x}$, respectively. By utilizing GLL quadrature for numerical integration, the mass matrix $\mathbf{M}$ becomes diagonal (lumped mass matrix). This allows for an efficient explicit time integration scheme:

$$\mathbf{U}^{n+1} = \mathbf{U}^n - \Delta t \mathbf{M}^{-1} \mathbf{K} \mathbf{U}^n. \tag{6}$$

where Δt is the pseudo-time step. The time integration continues until the total simulation time reaches T, at which point the field $\mathbf{U}$ represents the smoothed sensitivity kernel with the target Gaussian widths.

To computationally expedite the iterative smoothing process, we develop a GPU-accelerated solver that leverages the massive parallelism of GPU architectures by mapping each spectral element directly to a GPU thread block. The implementation utilizes on-chip shared memory to cache nodal values and differentiation matrices, thereby minimizing global memory transactions. Furthermore, on multi-GPU clusters, the code employs an efficient parallel strategy that performs data exchange between boundaries simultaneously with internal calculations, effectively eliminating communication bottlenecks and maximizing overall efficiency. To quantify the performance gains, we benchmarked the solver using the mesh of the size of $136 \times 46 \times 35$ elements. As illustrated in Figure 3, the elapsed time for the PDE smoothing on 4 AMD EPYC 9654 CPU cores is 62.16 s. In contrast, the optimized solver on 4 NVIDIA H100 GPUs completes the same task in merely 0.37 s. This corresponds to a dramatic speedup of 168.8×, confirming that the GPU implementation effectively resolves the computational bottleneck associated with model regularization.

3. Synthetic Examples

3.1. Ambient Noise Adjoint Tomography

The ambient noise adjoint tomography module in SpecFWAT follows the implementation of Wang et al. (2018), which optimizes the S-wave velocity structure beneath a seismic array by matching the empirical and synthetic Green's functions. Because ambient noise interferometry retrieve surface-wave EGFs between station pairs, the resulting ray

paths densely sample the region beneath the array. This configuration provides improved lateral resolution and robust sensitivity to shallow crustal structure. In SpecFWAT, surface-wave windows of ambient noise cross-correlations are automatically selected for different periods based on the approximated group velocity (He et al., 2024; Wang et al., 2018). The default misfit function is the frequency-dependent travel-time misfit derived using the multitaper method (Zhou et al., 2004; Tape et al., 2010). Several alternative and well-established objective functions are also available (see Section 2.4).

To demonstrate the workflow, we conduct a synthetic test using a dense array of 49 stations, 24 of which are selected as virtual sources evenly distributed along the array edges (Figure 4a). The model domain covers an area of $1° \times 1°$ with a depth of 32 km. The target model consists of a $2 \times 2 \times 2$ checkerboard pattern superimposed on a linearly increasing background velocity model (Figure 4b – 4d). The inversion successfully recovers the main features of all eight checkers (Figure 4e – 4g), though minor artifacts appear due to the limited data coverage. The entire synthetic test is performed on a laptop equipped with an AMD 7950HX CPU and an NVIDIA 4060 GPU, requiring 14 iterations and a total runtime of approximately 1 hour 1 minute.

3.2. Teleseismic Full-waveform Adjoint Tomography

Teleseismic full-waveform adjoint tomography (Tele-FWAT) utilizes scattered and converted phases in teleseismic P coda waves to provide enhanced vertical resolution for crustal and upper-mantle structures compared to ambient noise and local earthquake adjoint tomography. Such approach has been widely applied to subduction zones and intracontinental orogenic regions (Wang et al., 2016, 2021b; Kan et al., 2023).

In SpecFWAT, teleseismic wavefields are simulated using a spectral-element and frequency-wavenumber (SEM–FK) hybrid method (Tong et al., 2014a,b; Monteiller et al., 2020), which efficiently solves the wave equation for a plane-wave injection into a local 3-D model with a global 1D background model. The misfit function is defined as

$$\chi_{Tele} = \frac{1}{2} \| s_i(t) * W(t) - d_i(t) \|^2 dt, \tag{7}$$

where $s_i(t)$ is the synthetic teleseismic waveform in the i-th component. $d_i(t)$ is the observed teleseismic waveform of the corresponding component. W(t) denotes the estimated average source wavelet (i.e., source time function, STF). We calculate individual STF at each receiver via time-domain iterative deconvolution of the observed waveforms with the synthetics (Ligorría and Ammon, 1999). The final STF is then extracted by applying Principal Component Analysis (PCA) to this set of STFs (Wang et al., 2021b).

To evaluate the influence of the STF on inversion performance, we design a comparative synthetic test. The model domain spans $1.8° \times 1.0°$ laterally and extends to a depth of 80 km, containing 33 receivers distributed at the surface (Figure 5a). 8 teleseismic events are simulated, each with a ray-parameter of 0.06 s/km and azimuths uniformly dis-

tributed from 0° to 360° at 45° intervals. Each event is assigned a distinct STF (Figure 5b) to prepare for subsequent tests. The target model consists of alternating $5 \times 3 \times 4$ checkerboard patterns ($\sim$ 20 km for each checker) with 12% velocity perturbations, while the initial model includes two homogeneous layers (Figure 5a).

In the first test, synthetic data are generated with convolution by a Gaussian impulsive source. The inversion successfully recovered the checkerboard pattern and the magnitude of perturbations across all depths (Figure 6a – 6c). In the second test, synthetic data are convolved with STFs before inversion. While the shallow anomalies are well-recovered, the deeper features suffered significant distortion in both amplitude and pattern (Figure 6d – 6f). Estimating complex STFs is prone to errors because the deconvolution relies on imperfect synthetic Green's functions derived from the reference model. Furthermore, complex STFs inherently act as a low-pass filter that degrades effective resolution. These tests suggest strict selection criteria to retain only events with short half-durations and Gaussian-like STF. However, the application of such strict selection criteria may substantially reduce the number of usable events, which remains a practical challenge for teleseismic adjoint tomography (Wang et al., 2021b).

3.3. Teleseismic Cross-convolution Full-waveform Adjoint Tomography

The cross-convolution misfit function has been used in seismology to minimize the effect of STF uncertainties. It has been successfully applied in estimating shear-wave splitting parameters (Menke and Levin, 2003; Menke, 2017) and in 1-D inversion for crustal and uppermost mantle velocity model (Bodin et al., 2013; Li et al., 2019). Motivated by these applications and a recent application in the Southern Alaska (Xu et al., 2024), SpecFWAT implements a cross-convolution-based objective function for teleseismic waveform inversion to improve its robustness against STF-related uncertainties. The misfit function is defined as

$$\chi_{TeleCC} = \frac{1}{2} \| s_R(t) * d_Z(t) - s_Z(t) * d_R(t) \|^2 dt, \tag{8}$$

where s_R and s_Z denote the synthetic radial and vertical components at each iteration, and d_R and d_Z represent the corresponding observed teleseismic waveforms. Unlike conventional Tele-FWAT, the cross-convolution approach (TeleCC-FWAT) does not require explicit deconvolution to estimate the STF, thereby avoiding the uncertainties associated with that process. To assess the influence of the STF and to compare the cross-convolution approach with the conventional Tele-FWAT, we conduct the same checkerboard tests as in the Section 3.2. Figures 7(a) – 7(c) show the inversion results for data generated with Gaussian sources. All checkers are successfully recovered, although some smearing artifacts appear in the third (40 - 60 km) and fourth layers (60 - 80 km). When the synthetic data are generated by convolution with complex STFs, the recovered model shows only minor differences from the impulsive-sources case, indicating that the cross-convolution objective function is insensitive to STF complexity.

A further comparison shows that for data generated by convolving with complex STFs (Figures 7d – 7f) , TeleCC-

FWAT yields overall better results than Tele-FWAT (Figure 7d – 7f), particularly at depths of 40–80 km. These results demonstrate that the cross-convolution formulation provides a stable and reliable alternative when teleseismic sources have complex or poorly constrained STFs. By contrast, when the STF is simple and nearly Gaussian function, the conventional Tele-FWAT achieves higher resolution (Figure 6a – 6c). Both objective functions are implemented in SpecFWAT, allowing users to select the appropriate approach based on source complexity and data quality.

3.4. Receiver Function Adjoint Tomography

While the theoretical resolution of FWAT is dictated by the half-wavelength limit, in practice, adjoint tomography based on teleseismic waveforms (Tele-FWAT or TeleCC-FWAT) often faces challenges in resolving fine-scale crustal structures. This limitation arises primarily from the spectral complexity of teleseismic STFs, which typically exhibit cutoff periods greater than 5s (Beller et al., 2018). Consequently, even with dense seismic arrays, the obtainable resolution within the crust remains constrained. To address this issue, receiver functions can be utilized by deconvolving the vertical component from the radial component of teleseismic records. This process effectively eliminates source signatures and near-source structural responses, thereby significantly enhancing the dominant frequency of the signal (e.g., a Gaussian width factor of 1.0 in deconvolution corresponds to a dominant frequency of approximately 0.58 Hz). Thus, the waveform difference of receiver functions has been adopted as the objective function which was developed into the receiver function adjoint tomography (RFAT) approach (Xu et al., 2023). This approach is particularly advantageous for imaging regions characterized by strong lateral heterogeneities in crust and uppermost mantle, such as continental orogenic belts (Xu et al., 2023) and intraplate volcanic systems (Xu et al., 2025).

To demonstrate the capability of RFAT in resolving strong lateral heterogeneity by using the SpecFWAT package , we conduct a synthetic test with a target model (Figure 8a) constructed based on prior receiver function images of dense short-period nodal array (Wu et al., 2024), representing the structural deformation features of the Altyn-Tagh thrust zone. The model exhibits strong lateral heterogeneity in both S-wave velocity (Figure 8c) and its gradient (Figure 8d). Key structural features include thrust of crustal layers between 150 – 300km, as well as significant undulation and offsets along the Conrad and Moho discontinuities. Additionally, an underthrusting wedge is incorporated near the Moho offset to simulate the partial eclogitization (Wu et al., 2024).

For the RFAT setup, the initial model is generated by smoothing the 1D average of the target velocity model (Figure 8e). The computational domain is discretized into a mesh of $136 \times 46 \times 35$ spectral elements to ensure accurate wavefield simulations with maximum frequency of 1.5 Hz. We utilize synthetic data from a total of 24 teleseismic events (Figure 8b) to ensure sufficient illumination. The source geometry consists of 8 back-azimuthal directions (0° to 360° with 45° intervals) combined with 3 distinct ray parameters (0.044, 0.060, and 0.076 s/km). The synthetic receiver functions are computed with Gaussian factors of 1.0 and 1.5. This event count is comparable to the number

of high-quality teleseismic records typically available in practical applications. The data are recorded by a linear array with a station spacing of 2 km. We performed 11 iterations until the misfit reduction became stable, resulting in the final model M10 (Figure 8e).

To evaluate the computational efficiency, the inversion is conducted on a HPC system equipped with 4 NVIDIA H100 GPUs. The total elapsed time for the 11 iterations is approximately 6 hours and 8 minutes. Given the comparable number of sources and elements utilized in practical applications, this test suggests that SpecFWAT achieves excellent computational performance in real-data tomographic inversions. The capability to obtain high-resolution crustal imaging within a manageable time frame highlights the practical efficiency of the RFAT workflow.

The tomographic results demonstrate the superior capability of RFAT in recovering detailed crustal structures. As shown in the final model M10 (Figure 8f) and its corresponding gradient norm (Figure 8g), the geometry of the undulations and fault-induced offsets of the Conrad and Moho discontinuities are recovered with remarkable accuracy compared to the target model. Furthermore, the method successfully retrieves the sharp velocity gradients associated with the intracrustal thrust structures. The successful recovery of these detailed crustal features validates the efficacy of using high-frequency receiver function waveforms in adjoint tomography for crustal imaging. However, the strong velocity contrasts in the shallow crust, specifically the low velocities within sedimentary layers and the high velocities beneath the orogeny, are not fully recovered. We attribute this limitation to the complex behavior of crustal multiples (Xu et al., 2023), whose high sensitivity to S-wave velocity structure in the upper crust exacerbates the non-linearity of the optimization problem.

4. Joint Inversion over Multiple Datasets

While teleseismic data provide excellent vertical resolution and are sensitive to velocity discontinuities down to 100 km depth (Figure 8), their lateral resolution is often constrained by the density of the seismic array. Conversely, ambient noise surface waves excel at resolving lateral heterogeneities in the shallow crust but lose sensitivity at greater depths. To overcome the inherent resolution limitations of using a single dataset, Joint Full-waveform Adjoint Tomography (Joint-FWAT) that integrates teleseismic (e.g., teleseismic waveform or receiver functions) and ambient noise data are often performed (Wang et al., 2021b; He et al., 2024). This method leverages the complementary sensitivities of the two datasets and facilitates the construction of high-resolution 3-D velocity models of the lithosphere.

In the SpecFWAT, the Joint-FWAT framework is designed as a flexible platform capable of integrating diverse seismic datasets (e.g., local earthquake, teleseismic waveform, receiver function and ambient noise data) into a unified inversion scheme. Since different seismic observations utilize distinct misfit functions (e.g., waveform differences, cross-correlation phase shift, or exponentiated phase misfit), their respective sensitivity kernels possess different physical units and orders of magnitude. To balance the contributions of different datasets on the gradient, we employ a

weighted non-dimensionalized summation strategy to compute the final model update direction (He et al., 2024; Xu et al., 2025).

Assuming the individual misfit gradients for two distinct datasets are computed as $G_1(\mathbf{m}_n)$ and $G_2(\mathbf{m}_n)$ at the n-th iteration (after necessary pre-conditioning and smoothing), the unified descent direction $\mathbf{d}(\mathbf{m}_n)$ is constructed by normalizing each gradient by its respective norm before summation:

$$\mathbf{d}(\mathbf{m}_n) = w_1 \frac{G_1(\mathbf{m}_n)}{||G_1(\mathbf{m}_0)||} + w_2 \frac{G_2(\mathbf{m}_n)}{||G_2(\mathbf{m}_0)||}, \tag{9}$$

where $||G(\mathbf{m}_0)||$ represents normalization factor of the gradient at the initial iteration. This factor can be defined as either the L_∞ norm or L_2 norm of the gradient vector. The coefficients w_1 and w_2 represent the relative weights assigned to each dataset to balance their contributions. Although the amplitude of the gradients naturally diminishes as the data misfit decreases over iterations, this weighted non-dimensionalized approach is crucial for unbalanced the decay, thereby ensuring the balanced contribution of all datasets to the optimization direction.

To demonstrate the practical utility and robustness of the developed software in processing real-world observational data, we select the application in Hangay Dome, Central Mongolia as a benchmark example (Figure 9a). In this experiment, we perform Joint-FWAT using teleseismic cross-convolution and ambient noise EGFs, which aim to investigate the complex intraplate volcanism in this region. The computational domain, spanning from 95.8°E to 103.8°E and 44.8°N to 49.25°N, is discretized into a mesh of $112 \times 96 \times 24$ spectral elements. The average element size of this high-resolution mesh is approximately 5 km, ensuring the accurate wavefield simulation with period up to 2s required for joint inversion.

The inversion integrate two distinct datasets to constrain the velocity structure. We adopt the fully processed dataset of ambient noise (Xu et al., 2025) and newly selected dataset of teleseismic waveforms. The input comprises 28 sources for TeleCC-FWAT and 65 virtual sources derived from ambient noise EGFs (Figure 9a). By employing the weighted non-dimensionalized summation strategy, the optimization process achieved stable convergence after 14 iterations, with a total elapsed time of approximately 16 hours 14 minutes on 4 NVIDIA H100 GPUs. Specifically, the normalized data misfit for teleseismic cross-convolution is reduced to 64.5%, while the misfit for ambient noise data decrease to 32.2%. The result validates the efficacy of the joint inversion algorithm in balancing constraints from multiple datasets.

The resulting 3D S-wave velocity model demonstrates the software's capability to integrate Teleseismic cross-convolution and ambient noise EGFs to resolve complex heterogeneities. Overall, the velocity structure is highly consistent with the benchmark results from the receiver function and ambient noise joint inversion (Xu et al., 2025). The multi-level magmatic system remains clearly identifiable, revealing significant low-velocity zones in the mid-lower

crust (3.4–3.6 km/s) beneath the volcanoes (Figure 10a – 10c), interpreted as the crustal magma reservoir (Figure 10d), as well as distinct low-velocity anomalies in the uppermost mantle (70–100 km) indicating partial melting (Figure 10e and 10f). However, due to the lower frequency content inherent in the teleseismic waveforms, the effective resolution is slightly reduced, yielding a comparatively smoother velocity model. Furthermore, while Teleseismic cross-convolution provides weaker constraints on sharp subsurface interfaces, such as Moho and lithosphere-asthenosphere boundary (LAB), than receiver functions (Xu et al., 2025). Although the lateral variations of these boundaries appear less drastic than those constrained by receiver functions, the overarching structural features, such as the crustal thickening and Moho depression beneath the Orkhon volcano, are still robustly resolved (Figure 10e and 10f). These findings confirm that SpecFWAT can reliably handle diverse data combinations to recover consistent 3D geological structures.

5. Conclusions and Future Works

In this study, we introduce SpecFWAT, a modernized and high-performance open-source framework designed for full-waveform adjoint tomography at the lithospheric scale. Built upon the spectral-element method and the popular SPECFEM3D solver, the software employs a modular object-oriented architecture and a streamlined workflow to enhance reproducibility and user accessibility. A distinguishing feature of SpecFWAT is its robust heterogeneous computing capability. By offloading computationally intensive tasks to GPU accelerators (e.g., forward and adjoint simulations, iterative time-domain deconvolution, and PDE-based model regularization), the software significantly mitigates the computational burden of high-frequency 3-D wavefield simulations, rendering lithospheric imaging computationally tractable on moderate-scale hardware clusters.

The methodological versatility of SpecFWAT is illustrated by its ability to support various seismic observations and a robust joint inversion framework that employs weighted, non-dimensional gradient summation. Systematic validation via synthetic benchmarks and a practical application to the Hangay Dome in Central Mongolia confirm the software's efficacy in resolving multi-scale lithospheric heterogeneity. By synergizing the vertical constraints of receiver functions with the lateral sensitivity of ambient noise, the framework successfully reconstructs complex subsurface structures that remain ill-constrained by single-observable approaches.

Future advancements of the SpecFWAT ecosystem will prioritize the integration of anisotropic parameters, specifically incorporating Transverse Isotropy with a Vertical (VTI) and Horizontal (HTI) axis of symmetry, to facilitate more detailed characterizations of lithospheric deformation and dynamics (Wang et al., 2025). Additionally, recent studies have demonstrated that station-pair double-difference adjoint tomography can effectively mitigate the influence of source parameter uncertainties in regional earthquake imaging (Chen et al., 2023; Yuan et al., 2016). Building on these developments, we plan to incorporate the double-difference misfit function into SpecFWAT, utilizing station-pair formulations for both travel-time shifts and cross-convolution measurements. Furthermore, we intend to implement

full compatibility with the Cube2sph meshing package (Liu et al., 2024), enabling forward and adjoint simulations in spherical coordinates to support increasingly complex regional-scale models.

6. Acknowledgments

This work is supported by the National Natural Science Foundation of China (Grant No. 42574106). M.X is supported by Arts & Science Postdoctoral Fellowship Program at University of Toronto. K.W is supported by startup funding from University of Science and Technology of China (Grant No. KY2080000115). Q. Liu is supported by NSERC Discovery Grant. The high-performance computing used for this research is provided by the Digital Research Alliance of Canada (`http://alliance.can.ca`). We thank Prof. Tao Xu and Prof. Chenglong Wu to share their synthetic model of Altyn-Tagh thrust zone. The figures are made by PyGMT (Tian et al., 2025). We appreciate Specfem community to develop and maintain the open-source Specfem3D_Cartesian (Komatitsch and Vilotte, 1998; Komatitsch and Tromp, 1999) (`https://specfem.org`). We also acknowledge the developers of open-source packages yaml-cpp for sharing their codes.

Code availability section

Name of the code/library: SpecFWAT.

License: GNU General Public License v3.0.

Contact: gomijianxu@gmail.com, +1 437-340-6344.

Hardware requirements: PC or HPC

Program language: Fortran, C++, CUDA, Python

Software required: Fortran compiler, C++ compiler, CUDA, MPI compiler, CMake, yaml-cpp, HDF5

Program size: 19.8 MB

The source codes are available for downloading at the link: https://specfwat.xumijian.me

References

Afanasiev, M., Boehm, C., van Driel, M., Krischer, L., Rietmann, M., May, D.A., Knepley, M.G., Fichtner, A., 2019. Modular and flexible spectral-element waveform modelling in two and three dimensions. Geophysical Journal International 216, 1675–1692. doi:10.1093/gji/ggy469.

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. doi:10.1093/gji/ggx216.

Bensen, G.D., Ritzwoller, M.H., Barmin, M.P., Levshin, A.L., Lin, F., Moschetti, M.P., Shapiro, N.M., Yang, Y., 2007. Processing seismic ambient noise data to obtain reliable broad-band surface wave dispersion measurements. Geophysical Journal International 169, 1239–1260. doi:10.1111/j.1365-246X.2007.03374.x.

Bodin, T., Yuan, H., Romanowicz, B., 2013. Inversion of receiver functions without deconvolution–application to the Indian craton. Geophysical Journal International 196, 1025–1033. doi:10.1093/gji/ggt431.

Chen, G., Chen, J., Tape, C., Wu, H., Tong, P., 2023. Double-Difference Adjoint Tomography of the Crust and Uppermost Mantle Beneath Alaska. Journal of Geophysical Research: Solid Earth 128. doi:10.1029/2022JB025168.

Chen, M., Huang, H., Yao, H., van der Hilst, R., Niu, F., 2014. Low wave speed zones in the crust beneath SE Tibet revealed by ambient noise adjoint tomography. Geophysical Research Letters 41, 2013GL058476. doi:10.1002/2013GL058476.

Chow, B., Kaneko, Y., Tape, C., Modrak, R., Mortimer, N., Bannister, S., Townend, J., 2022. Strong Upper-Plate Heterogeneity at the Hikurangi Subduction Margin (North Island, New Zealand) Imaged by Adjoint Tomography. Journal of Geophysical Research: Solid Earth 127. doi:10.1029/2021JB022865.

Crotwell, H.P., Owens, T.J., Ritsema, J., 1999. The TauP Toolkit: Flexible Seismic Travel-time and Ray-path Utilities. Seismological Research Letters 70, 154–160. doi:10.1785/gssrl.70.2.154.

Fichtner, A., 2015. Source-structure trade-offs in ambient noise correlations. Geophysical Journal International 202, 678–694. doi:10.1093/gji/ggv182.

Fichtner, A., Bunge, H.p., Igel, H., 2006. The adjoint method in seismology: I. Theory. Physics of the Earth and Planetary Interiors 157, 86–104. doi:10.1016/j.pepi.2006.03.016.

Fichtner, A., Kennett, B.L.N., Tsai, V.C., Thurber, C.H., Rodgers, A.J., Tape, C., Rawlinson, N., Borcherdt, R.D., Lebedev, S., Priestley, K., Morency, C., Bozdağ, E., Tromp, J., Ritsema, J., Romanowicz, B., Liu, Q., Golos, E., Lin, F.C., 2024. Seismic Tomography 2024. Bulletin of the Seismological Society of America 114, 1185–1213. doi:10.1785/0120230229.

Gao, Y., Tilmann, F., van Herwaarden, D.P., Thrastarson, S., Fichtner, A., Heit, B., Yuan, X., Schurr, B., 2021. Full Waveform Inversion Beneath the Central Andes: Insight Into the Dehydration of the Nazca Slab and Delamination of the Back-Arc Lithosphere. Journal of Geophysical Research: Solid Earth 126, e2021JB021984. doi:10.1029/2021JB021984.

Gao, Y., Tilmann, F., Rietbrock, A., 2023. A review of misfit functions for adjoint full waveform inversion in seismology. Geophysical Journal International 235, 2794–2827. doi:10.1093/gji/ggad372.

He, B., Wang, K., Liu, T., Lei, T., Du, N., Van Der Lee, S., Darbyshire, F.A., Frederiksen, A., Zhu, H., Lumley, D., Halls, H., Liu, Q., 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. doi:10.1016/j.epsl.2024.118797.

Kan, L., 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. doi:10.1029/2022JB025506.

Komatitsch, D., Tromp, J., 1999. Introduction to the spectral element method for three-dimensional seismic wave propagation. Geophysical Journal International 139, 806–822. doi:10.1046/j.1365-246x.1999.00967.x.

Komatitsch, D., Vilotte, J.P., 1998. The spectral element method: An efficient tool to simulate the seismic response of 2D and 3D geological structures. Bulletin of the Seismological Society of America 88, 368–392. doi:10.1785/BSSA0880020368.

Krischer, L., Fichtner, A., Zukauskaite, S., Igel, H., 2015. Large-Scale Seismic Inversion Framework. Seismological Research Letters 86, 1198–1207. doi:10.1785/0220140248.

Lei, W., Ruan, Y., Bozdağ, E., Peter, D., Lefebvre, M., Komatitsch, D., Tromp, J., Hill, J., Podhorszki, N., Pugmire, D., 2020. Global adjoint tomography—model GLAD-M25. Geophysical Journal International 223, 1–21. doi:10.1093/gji/ggaa253.

Li, G., Niu, F., Yang, Y., Tao, K., 2019. Joint Inversion of Rayleigh Wave Phase Velocity, Particle Motion, and Teleseismic Body Wave Data for Sedimentary Structures. Geophysical Research Letters 46, 6469–6478. doi:10.1029/2019GL082746.

Ligorría, J.P., Ammon, C.J., 1999. Iterative deconvolution and receiver-function estimation. Bulletin of the Seismological Society of America 89, 1395–1400. doi:10.1785/BSSA0890051395.

Liu, T., 2024. Full-Waveform Seismic Tomography Using Earthquake and Ambient Noise Data: With Application to the Alaskan Lithosphere. Ph.D.. University of Toronto (Canada). Canada – Ontario, CA. Publication Title: ProQuest Dissertations and Theses 31490390.

Liu, T., Wang, K., Xie, Y., He, B., Lei, T., Du, N., Tong, P., Yang, Y., Rychert, C.A., Harmon, N., Grasselli, G., Liu, Q., 2024. Cube2sph : A toolkit enabling flexible and accurate continental-scale seismic wave simulations using the SPECFEM3D_cartesian package. Computers & Geosciences 190, 105644. doi:10.1016/j.cageo.2024.105644.

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. doi:10.1111/j.1365-246X.2009.04099.x.

Menke, W., 2017. Sensitivity Kernels for the Cross-Convolution Measure. Bulletin of the Seismological Society of America 107, 2213–2224. doi:10.1785/0120170045.

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. doi:10.1046/j.1365-246X.2003.01937.x.

Modrak, R.T., Borisov, D., Lefebvre, M., Tromp, J., 2018. SeisFlows—Flexible waveform inversion software. Computers & Geosciences 115, 88–95. doi:10.1016/j.cageo.2018.02.004.

Monteiller, V., Beller, S., Plazolles, B., Chevrot, S., 2020. On the validity of the planar wave approximation to compute synthetic seismograms of teleseismic body waves in a 3-D regional model. Geophysical Journal International 224, 2060–2076. doi:10.1093/gji/ggaa570.

Monteiller, V., Chevrot, S., Komatitsch, D., Wang, Y., 2015. Three-dimensional full waveform inversion of short-period teleseismic wavefields based upon the SEM–DSM hybrid method. Geophysical Journal International 202, 811–827. doi:10.1093/gji/ggv189.

Nocedal, J., Wright, S.J., 2006. Numerical optimization. Springer series in operations research. 2nd ed ed., Springer, New York. OCLC: ocm68629100.

Rodgers, A., Krischer, L., Afanasiev, M., Boehm, C., Doody, C., Chiang, A., Simmons, N., 2022. WUS256: An Adjoint Waveform Tomography Model of the Crust and Upper Mantle of the Western United States for Improved Waveform Simulations. Journal of Geophysical Research: Solid Earth 127. doi:10.1029/2022JB024549.

Sager, K., Boehm, C., Ermert, L., Krischer, L., Fichtner, A., 2020. Global-Scale Full-Waveform Ambient Noise Inversion. Journal of Geophysical Research: Solid Earth 125, e2019JB018644. doi:10.1029/2019JB018644.

Sager, K., Ermert, L., Boehm, C., Fichtner, A., 2018. Towards full waveform ambient noise inversion. Geophysical Journal International 212, 566–590. doi:10.1093/gji/ggx429.

Tao, K., Grand, S.P., Niu, F., 2018. Seismic Structure of the Upper Mantle Beneath Eastern Asia From Full Waveform Seismic Tomography. Geochemistry, Geophysics, Geosystems 19, 2732–2763. doi:10.1029/2018GC007460.

Tape, C., Liu, Q., Maggi, A., Tromp, J., 2009. Adjoint Tomography of the Southern California Crust. Science 325, 988–992. doi:10.1126/science.1175298.

Tape, C., Liu, Q., Maggi, A., Tromp, J., 2010. Seismic tomography of the southern California crust based on spectral-element and adjoint methods. Geophysical Journal International 180, 433–462. doi:10.1111/j.1365-246X.2009.04429.x.

Tian, D., Leong, W.J., Fröhlich, Y., Grund, M., Schlitzer, W., Jones, M., Toney, L., Yao, J., Tong, J.H., Magen, Y., Materna, K., Belem, A., Newton, T., Anant, A., Ziebarth, M., Quinn, J., Uieda, L., Wessel, P., 2025. PyGMT: A Python interface for the Generic Mapping Tools. doi:10.5281/ZENODO.17156962.

Tong, P., Chen, C.w., Komatitsch, D., Basini, P., Liu, Q., 2014a. High-resolution seismic array imaging based on an SEM-FK hybrid method. Geophysical Journal International 197, 369–395. doi:10.1093/gji/ggt508.

Tong, P., Komatitsch, D., Tseng, T.L., Hung, S.H., Chen, C.W., Basini, P., Liu, Q., 2014b. A 3-D spectral-element and frequency-wave number hybrid method for high-resolution seismic array imaging. Geophysical Research Letters 41, 7025–7034. doi:10.1002/2014GL061644.

Tromp, J., 2019. Seismic wavefield imaging of Earth's interior across scales. Nature Reviews Earth & Environment 1, 40–53. doi:10.1038/s43017-019-0003-8.

Tromp, J., Luo, Y., Hanasoge, S., Peter, D., 2010. Noise cross-correlation sensitivity kernels: Noise cross-correlation sensitivity kernels. Geophysical Journal International 183, 791–819. doi:10.1111/j.1365-246X.2010.04721.x.

Tromp, J., Tape, C., Liu, Q., 2005. Seismic tomography, adjoint methods, time reversal and banana-doughnut kernels: Seismic tomography, adjoint methods, time reversal and banana-doughnut kernels. Geophysical Journal International 160, 195–216. doi:10.1111/j.1365-246X.2004.02453.x.

Van Herwaarden, D., Thrastarson, S., Hapla, V., Afanasiev, M., Trampert, J., Fichtner, A., 2023. Full-Waveform Tomography of the African Plate Using Dynamic Mini-Batches. Journal of Geophysical Research: Solid Earth 128, e2022JB026023. doi:10.1029/2022JB026023.

VanDecar, J.C., Crosson, R.S., 1990. Determination of teleseismic relative phase arrival times using multi-channel cross-correlation and least squares. Bulletin of the Seismological Society of America 80, 150–169. doi:10.1785/BSSA0800010150.

Wang, K., Tong, P., Zhang, Z., Su, L., 2025. Mapping the crustal magmatic system and regional deformation of the toba region by multiparameter ambient noise adjoint tomography. Journal of Geophysical Research: Solid Earth 130, e2024JB030550. doi:10.1029/2024JB030550.

Wang, K., Wang, Y., Song, X., Tong, P., Liu, Q., Yang, Y., 2021a. 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. doi:10.1785/0120210094.

Wang, K., Yang, Y., Basini, P., Tong, P., Tape, C., Liu, Q., 2018. Refined crustal and uppermost mantle structure of southern California by ambient noise adjoint tomography. Geophysical Journal International 215, 844–863. doi:10.1093/gji/ggy312.

Wang, K., Yang, Y., Jiang, C., Wang, Y., Tong, P., Liu, T., Liu, Q., 2021b. Adjoint Tomography of Ambient Noise Data and Teleseismic P Waves: Methodology and Applications to Central California. Journal of Geophysical Research: Solid Earth 126. doi:10.1029/2021JB021648.

Wang, Y., Chevrot, S., Monteiller, V., Komatitsch, D., Mouthereau, F., Manatschal, G., Sylvander, M., Diaz, J., Ruiz, M., Grimaud, F., Benahmed, S., Pauchet, H., Martin, R., 2016. The deep roots of the western Pyrenees revealed by full waveform inversion of teleseismic P waves. Geology 44, 475–478. doi:10.1130/G37812.1.

Weickert, J., 1998. Anisotropic Diffusion in Image Processing. volume 1. Teubner Stuttgart.

Wu, C., Xu, T., Tian, X., Mitchell, R.N., Lin, J., Yang, J., Wang, X., Lu, Z., 2024. Underthrusting of Tarim Lower Crust Beneath the Tibetan Plateau Revealed by Receiver Function Imaging. Geophysical Research Letters 51, e2024GL108220. doi:10.1029/2024GL108220.

Xu, M., Wang, K., Chen, J., He, J., Liu, Q., Liu, Y., Huang, Z., Tong, P., 2025. Multilevel Mechanisms Driving Intraplate Volcanism in Central Mongolia Revealed by Adjoint Waveform Tomography of Receiver Function and Ambient Noise Data. Earth and Planetary Science Letters 650, 119137. doi:10.1016/j.epsl.2024.119137.

Xu, M., Wang, K., Chen, J., Yu, D., Tong, P., 2023. Receiver Function Adjoint Tomography for Three-Dimensional High-Resolution Seismic Array Imaging: Methodology and Applications in Southeastern Tibet. Geophysical Research Letters 50, e2023GL104077. doi:10.1029/2023GL104077.

Xu, M., Wang, K., Liu, T., Du, N., Tong, P., Liu, Q., 2024. Full-waveform Joint Inversion of Teleseismic Cross-convolution, Receiver Function, and Ambient Noise Data for High-resolution Lithospheric Seismic Imaging, in: AGU Fall Meeting Abstracts, pp. S43A–342.

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. doi:10.1093/gji/ggaa063.

Yuan, Y.O., Simons, F.J., Tromp, J., 2016. Double-difference adjoint seismic tomography. Geophysical Journal International 206, 1599–1618. doi:10.1093/gji/ggw233.

Zhou, Y., Dahlen, F.A., Nolet, G., 2004. Three-dimensional sensitivity kernels for surface wave observables. Geophysical Journal International 158, 142–168. doi:10.1111/j.1365-246X.2004.02324.x.

Zhu, H., Bozdağ, E., Tromp, J., 2015. Seismic structure of the European upper mantle based on adjoint tomography. Geophysical Journal International 201, 18–52. doi:10.1093/gji/ggu492.

List of Figures

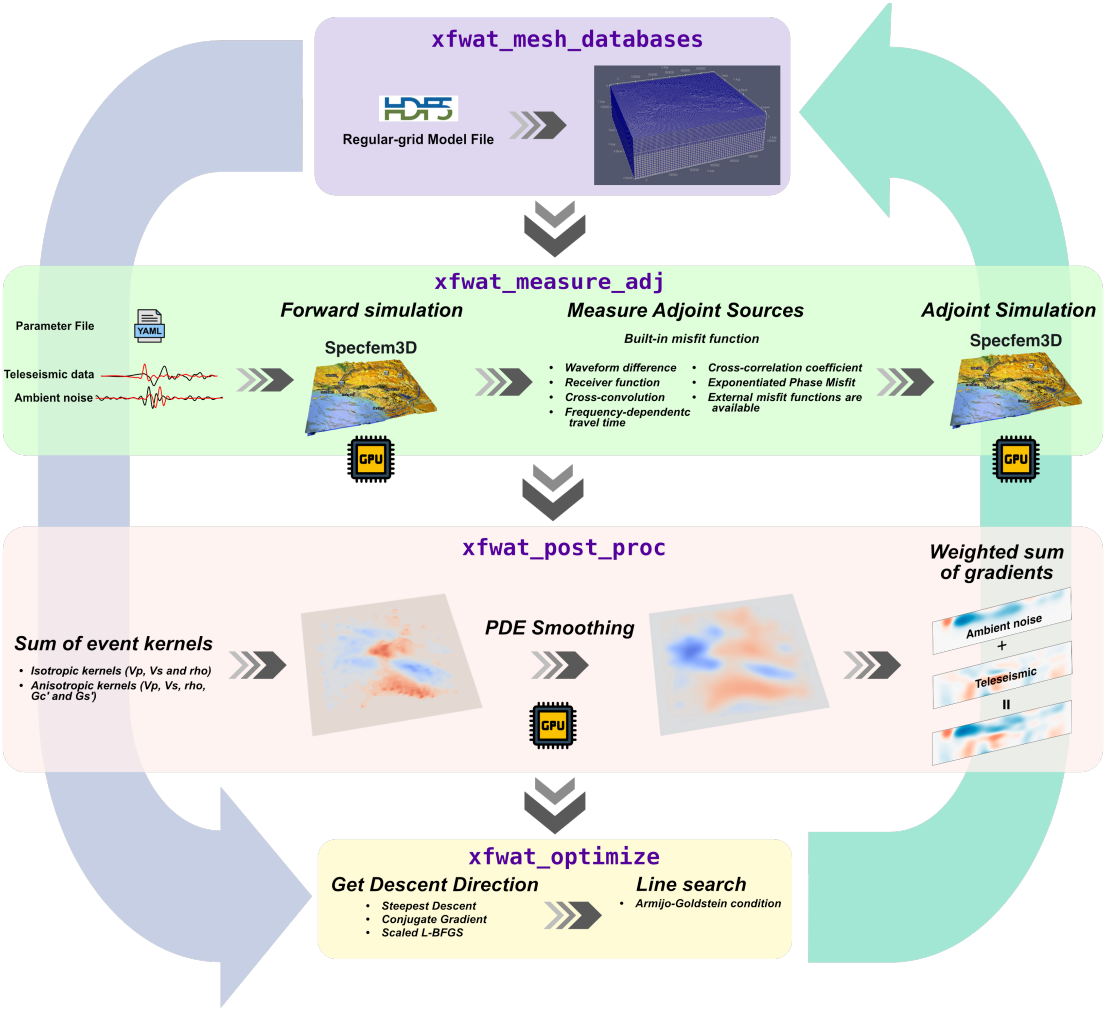

Figure 1: The general framework of SpecFWAT.

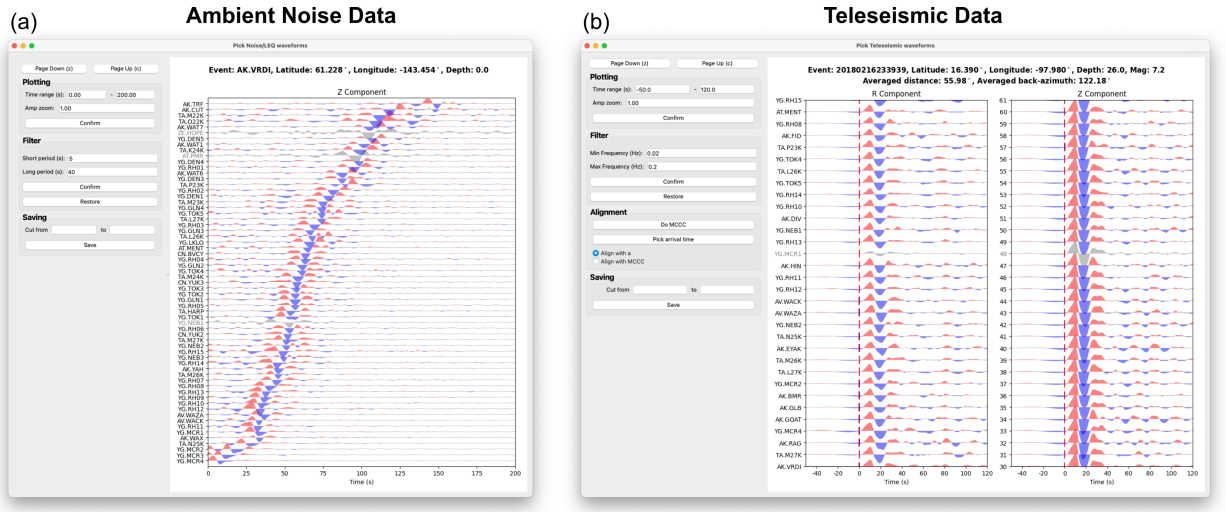

Figure 2: (a) The graphic user interface (GUI) designed for conducting quality control on ambient noise data. In the graphic window, the waveforms of each virtual source are displayed, with gray wiggles indicating data earmarked for rejection. The left panel provides control options, including figure adjustment, filtering, and saving functionalities. (b) The GUI for quality control on teleseismic data. Radial and vertical components are shown in the graphic window. Blue bars denote the P arrival times estimated using the Taup toolkit (Crotwell et al., 1999), while red bars represent the P arrival times corrected through the multi-channel cross-correlation method (VanDecar and Crosson, 1990). The "Alignment" group in the left panel provides options for waveform alignment.

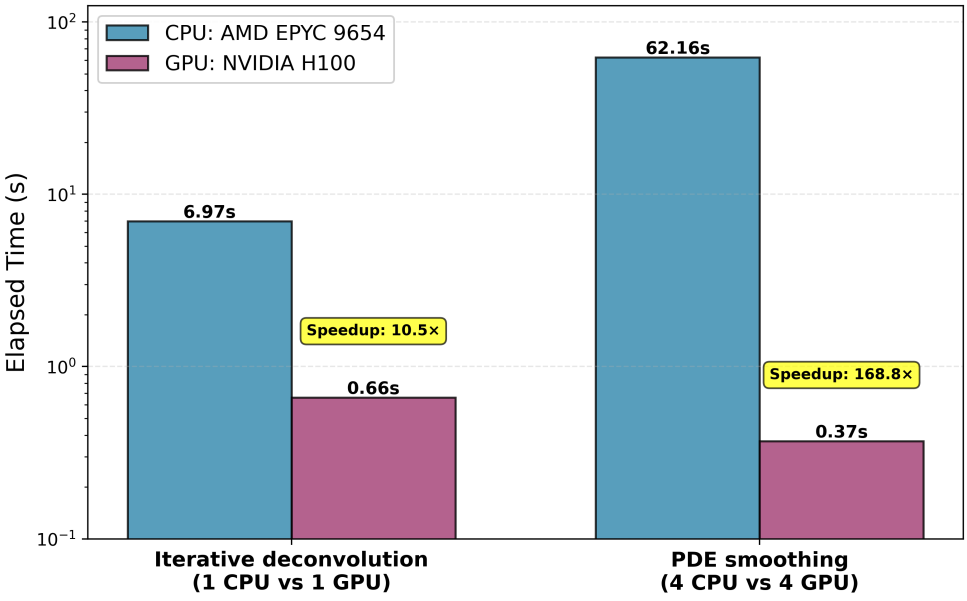

Figure 3: Elapsed time comparison between CPU and GPU architectures for two schemes in the FWI workflow: iterative deconvolution and PDE smoothing. Note the logarithmic scale on the y-axis. The GPU implementation (NVIDIA H100) achieves speedups of 10.5× and 168.8×, respectively, compared to the CPU baseline (AMD EPYC 9654).

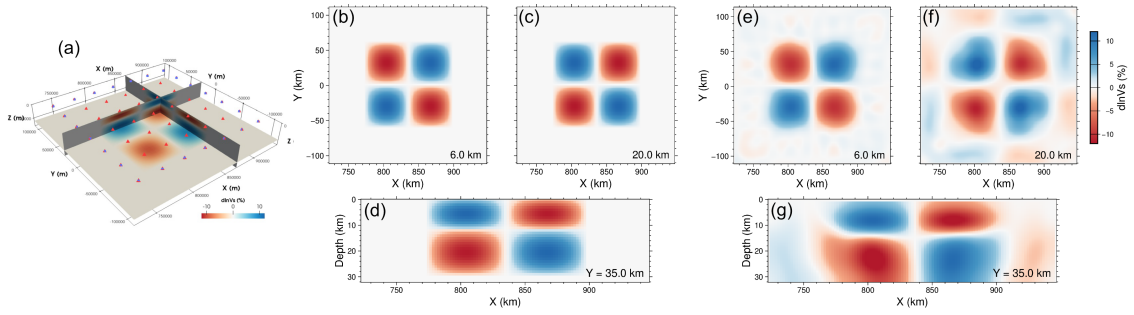

Figure 4: Example of ambient noise adjoint tomography. (a) Slices view of computational domain. Blue dots represent virtual sources. Red triangles represent receivers. (b) – (d) Vs anomalies of the target model relative to initial model. (e) – (g) Vs anomalies of the final model relative to initial model.

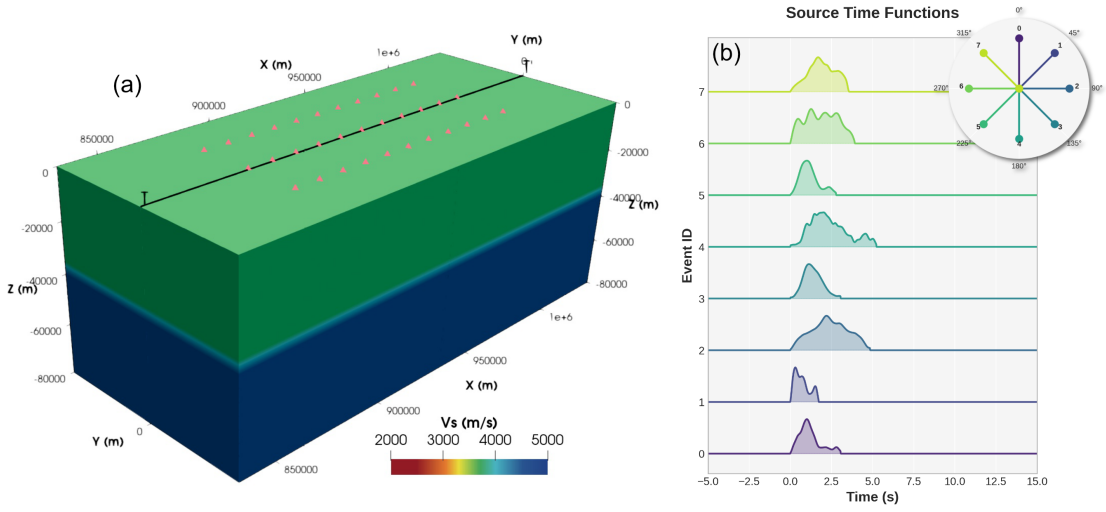

Figure 5: (a) Computational domain and initial model for teleseismic synthetic tests. The red triangles represent receivers. The black lines represent the location of profile TT'. (b) STFs that convolved with synthetic teleseismic waveforms to generate synthetic data of the target model. The top right insert shows back-azimuth of the teleseismic sources.

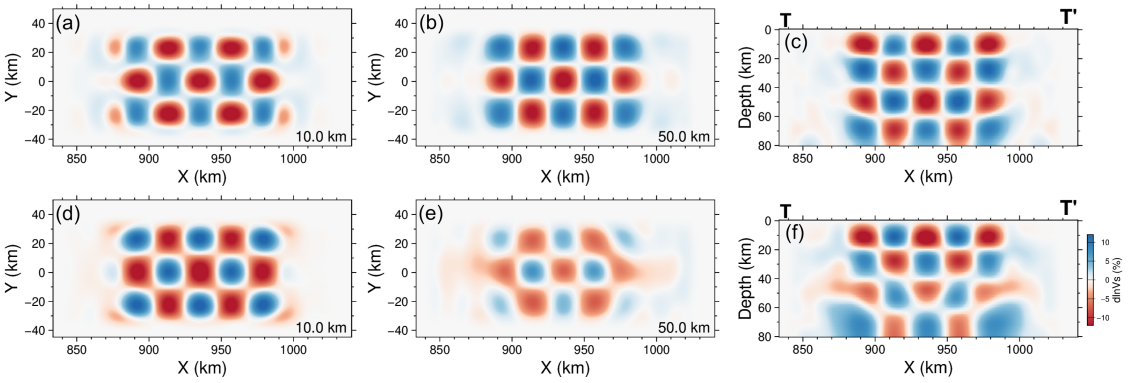

Figure 6: Synthetic example of Tele-FWAT (a–c): Inversion results using data from an impulsive source (no STF complexity). (d–f): Inversion results using data convolved with the complex STF (Figure 5b).

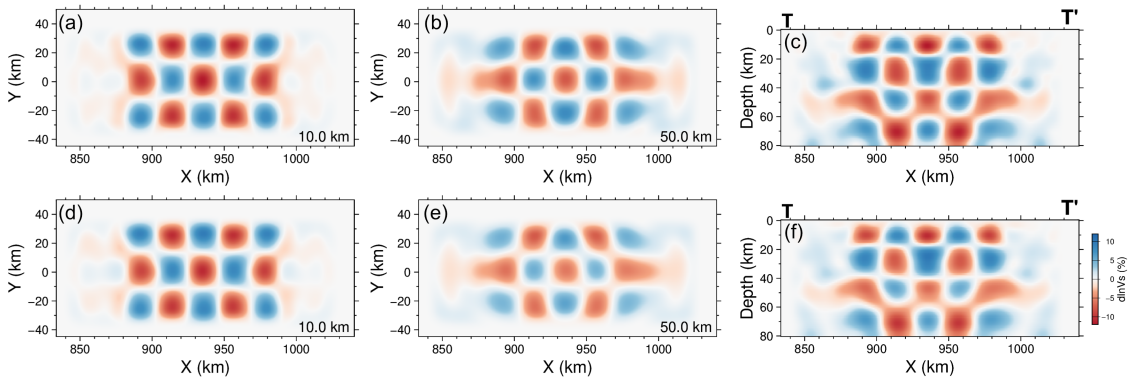

Figure 7: Synthetic example of TeleCC-FWAT with the same content as Figure 6.

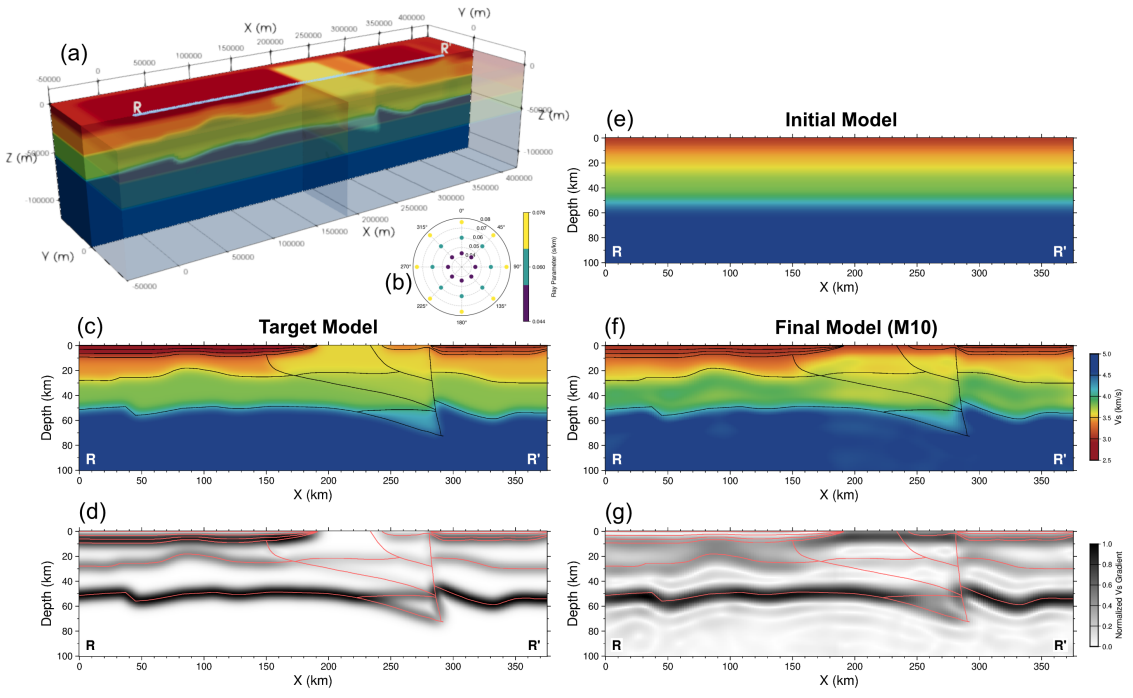

Figure 8: Synthetic example of receiver function adjoint tomography (RFAT). (a) Computational domain with target S-wave velocity model. Blue dots represent receivers with 2 km spacing. (b) Teleseismic events for this test. Colored dots represent ray-parameter and back-azimuth. (c) and (d) S-wave velocity and its gradient in the target view along the cross-section RR' (shown in (a)). (e) S-wave velocity of initial model along the cross-section RR'. (f) and (g) Same content as (c) and (d), but the S-wave velocity derived by RFAT with 10 iterations.

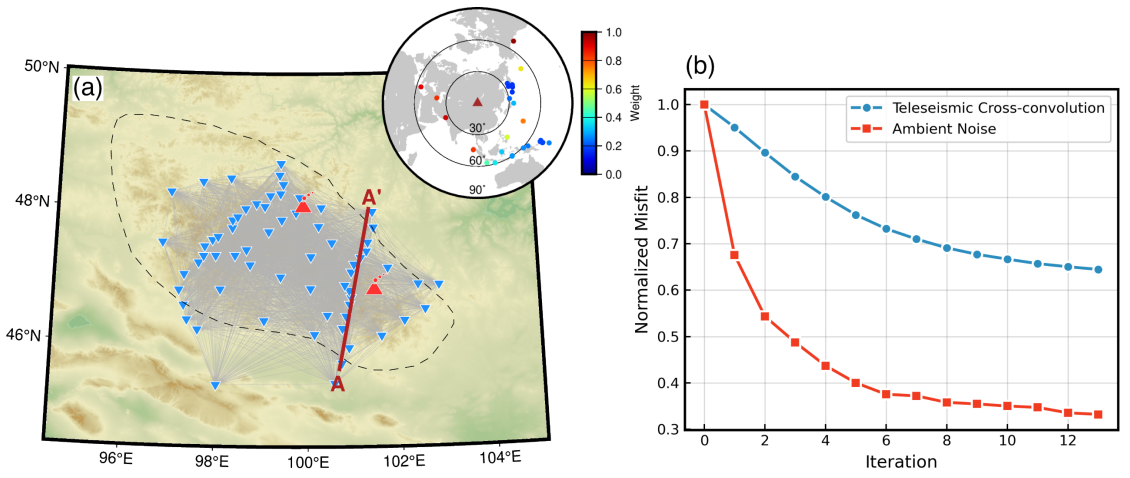

Figure 9: (a) Stations (blue triangles) for the example of Joint-FWAT, Gray lines denote the ray path coverage of the ambient noise data. The red line denotes profile AA'. The gray dashed line denotes the trace of the Hangay dome. The teleseismic events for cross-convolution measurements with weights are shown in the top right insert. (b) Normalized misfit of receiver function (blue line) and ambient noise data (red line) over iterations.

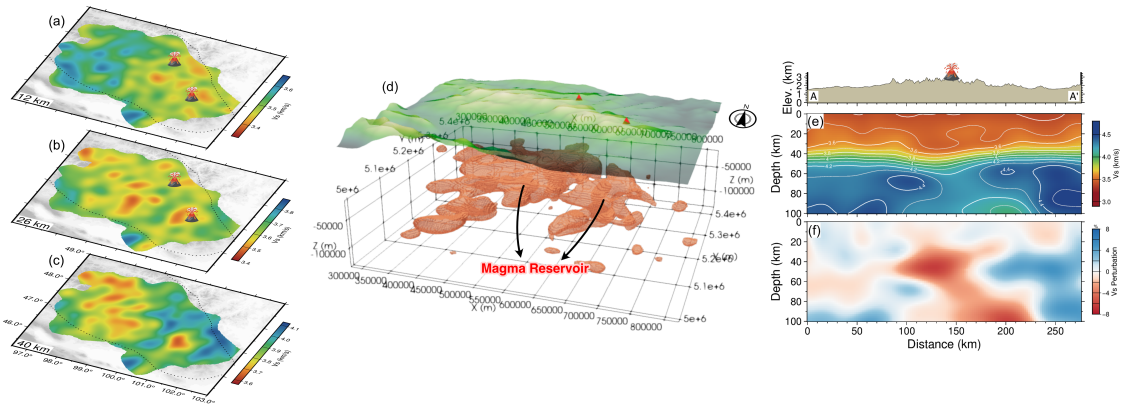

Figure 10: (a – c) Map view of S-wave velocity at 12, 26, and 40 km depth beneath Hangay Dome. (d) Schematic view of S-wave velocity model. Red blocks represent low-velocity with perturbation less than -3%. Blue blocks represent high-velocity with perturbation greater than 3%. (e – f) Cross section of absolute S-wave velocity, S-wave velocity perturbations along the profile AA' (shown in Figure 9a).