Preconditioned BFGS-based Uncertainty Quantification in Elastic Full Waveform Inversion

Qiancheng Liu¹, Stephen Beller¹, Wenjie Lei¹, Daniel Peter², and Jeroen Tromp^{1,3}

¹Department of Geosciences, Princeton University, Princeton, NJ 08544, USA

²Division of Physical Sciences and Engineering, King Abdullah University of Science and Technology (KAUST), Thuwal, Saudi Arabia

³Program in Applied and Computational Mathematics, Princeton University, Princeton, NJ 08544, USA

2020

SUMMARY

Full Waveform Inversion (FWI) has become an essential technique for mapping geophysical subsurface structures. However, proper uncertainty quantification is often lacking in current applications. In theory, uncertainty quantification is related to the inverse Hessian (or the posterior covariance matrix). Even for common geophysical inverse problems its calculation is beyond the computational and storage capacities of the largest high-performance computing systems. In this study, we amend the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm to perform uncertainty quantification for large-scale applications. For seismic inverse problems, the limited-memory BFGS (L-BFGS) method prevails as the most efficient quasi-Newton method. We aim to augment it further to obtain an approximate inverse Hessian for uncertainty quantification in FWI. To facilitate retrieval of the inverse Hessian, we combine BFGS (essentially a full-history L-BFGS) with randomized singular value decomposition to determine a low-rank approximation of the inverse Hessian. Setting the rank number equal to the number of iterations makes this solution efficient and memory-affordable even for large-scale problems. Furthermore, based on the Gauss-Newton method, we formulate different initial,

diagonal Hessian matrices as preconditioners for the inverse scheme and compare their performances in elastic FWI applications. We highlight our approach with the elastic Marmousi benchmark model, demonstrating the applicability of preconditioned BFGS for large-scale FWI and uncertainty quantification.

Key words: computational seismology, seismic tomography, waveform inversion

1 INTRODUCTION

Seismic full-waveform inversion (FWI) is a compelling method for characterizing subsurface properties. FWI seeks to estimate the “optimal” model by minimizing a measure of data misfit between simulated and observed seismograms, most commonly via an iterative inversion procedure in a least-squares sense (Lailly 1983; Tarantola 1984). Its ultimate goal is to resolve geophysical properties from all available information in observed seismic measurements (Virieux & Operto 2009b; Tromp 2019). Thanks to advances in data acquisition, high-performance computing, and numerical simulation methods (Komatitsch & Tromp 1999; Peter et al. 2011; Lefebvre et al. 2017; Polychronopoulou et al. 2018), FWI can constrain seismic models with increasing resolution. Successful applications of FWI across scales have been reported in global (French & Romanowicz 2014; Bozdağ et al. 2016; Fichtner et al. 2018; Lei et al. 2020), regional (Tape et al. 2010; Zhu et al. 2012; Krischer et al. 2018), exploration (Warner & Guasch 2016; Métivier et al. 2016), and medical (Bachmann & Tromp 2020; Guasch et al. 2020) imaging. However, due to its ill-posed nature, FWI suffers from non-unique solutions owing to limited data coverage and uncertainties in measurements and theories. Therefore, uncertainty quantification in FWI is essential, but to date only a few solutions have been proposed for larger applications (Fichtner & Trampert 2011a,b; Zhu et al. 2016; Fichtner & Simutè 2018; Gebraad et al. 2019; Liu & Peter 2019, 2020; Thurin et al. 2019).

Optimization methods in FWI can be categorized into two families: deterministic and statistical. Deterministic methods, mainly referring to gradient-based optimization (Pratt 1999; Virieux & Operto 2009a), have been well developed to tackle challenges such as cycle-skipping (Warner & Guasch 2016; Métivier et al. 2016) and source-encoding (Tromp & Bachmann 2019). Statistical methods rely on random sampling, based on Markov chain Monte Carlo or

Hamiltonian Monte Carlo algorithms (Duane et al. 1987; Betancourt 2017). Such methods are theoretically preferable in FWI, because they provide not only the maximum *a posteriori* (MAP) model, but also statistical metrics for uncertainty quantification (Biswas & Sen 2017; Fichtner & Simutè 2018; Gebraad et al. 2019). Unfortunately, statistical approaches are often prohibitive for large-scale applications due to their computational expense (Tarantola 2005). For this reason, deterministic methods are popular in FWI due to their low cost and fast convergence rate. Tarantola (2005) expressed hope for FWI and uncertainty quantification by deterministic optimization if the posterior covariance matrix can be constructed for model appraisal (Tarantola & Valette 1982).

Posterior model covariance is closely related to the inverse data-misfit Hessian under the assumption of linearizable forward modeling and Gaussian model priors (Tarantola 2005). For large-scale applications involving millions of parameters, it becomes unfeasible to store, assemble, and analyze such huge matrices. To tackle this issue, Zhang & McMechan (1995) compress the data volume of classic inversion algorithms using least-squares QR factorization. Trampert et al. (2012) random-probe tomographic models to estimate the resolution length of waveform tomography. Rawlinson et al. (2014) provide a review of uncertainty estimation in waveform inversion. Fichtner & van Leeuwen (2015) analyze direction-dependent resolution lengths from the randomly sampled Hessian via auto-correlation. Bui-Thanh et al. (2013) approximate the posterior covariance matrix by eigen-decomposing the data-misfit Hessian for its inversion with randomized singular-value decomposition (RSVD) (Liberty et al. 2007; Halko et al. 2011). Zhu et al. (2016) exploit the point-spread function (PSF) test to improve Hessian-computation efficiency. Luo (2012) and Liu et al. (2019) introduced a memory-affordable vector-version square-root variable metric (SRVM) algorithm. Subsequently, Liu & Peter (2019) used SRVM and RSVD to efficiently probe the inverse Hessian for uncertainty quantification in FWI and characterize non-uniqueness based on the SRVM-based null-space shuttle (Thurin et al. 2019; Liu 2019; Liu & Peter 2020).

In this paper, we explore the feasibility of a classic quasi-Newton method, the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm (Broyden 1970; Liu & Nocedal 1989), for elastic FWI with uncertainty quantification. Although the limited-memory BFGS (L-BFGS) algorithm has been used for decades in exploration seismology to invert for subsurface properties,

to our knowledge, no attempts have yet been made to reconcile it with uncertainty quantification. Here, we fill this gap by connecting it to the principal BFGS algorithm and combine it with RSVD probing to approximate the posterior covariance matrix of the inverse problem. Similar to the SRVM-RSVD workflow (Liu & Peter 2020), we present a BFGS-RSVD approach to access and factorize the inverse Hessian. BFGS runs in the framework of L-BFGS. According to Nocedal & Wright (2006), L-BFGS is equivalent to the BFGS algorithm if L-BFGS keeps the same initial Hessian and all the memories. Fortunately, FWI usually takes only tens of iterations to converge. The resulting storage of full-memory L-BFGS vectors becomes thus affordable even for large-scale applications.

The performance of BFGS in FWI may be closely related to the initial Hessian guess (Brossier et al. 2009; Métivier et al. 2013; Yang et al. 2018; Beller & Chevrot 2020). Thus, we investigate the performance of different diagonal approximations of the initial inverse Hessian and compare them in terms of FWI convergence rate and uncertainty quantification maps.

The purpose of this study is to advance L-BFGS and Hessian-related preconditioners for uncertainty quantification in FWI. We start with a brief review of the FWI optimization problem and recall the theory of the L-BFGS quasi-Newton algorithm. We continue with a presentation of computationally inexpensive diagonal preconditioners used as initial guesses for the L-BFGS approximation of the inverse Hessian. We then discuss retrieval of the inverse Hessian after BFGS-based FWI has converged. Subsequently, we propose a BFGS-RSVD workflow to achieve a faster and cheaper BFGS-based Hessian retrieval. Finally, we verify our method with numerical examples to demonstrate the applicability of preconditioned BFGS-based FWI for uncertainty quantification.

2 THEORY & METHOD

2.1 BFGS in FWI

Seismic FWI aims to iteratively minimize a misfit function $f(\mathbf{m})$ between observed and synthetic data, $\mathbf{d}$ and $\mathbf{s}(\mathbf{m})$, respectively. In principle, FWI consists of three consecutive steps: (i) misfit and gradient computations, e.g., with the adjoint-state method (Tromp et al. 2005; Plessix 2006), (ii) a search direction update based on gradients by optimization (Liu & No-

cedal 1989; Métivier et al. 2013; Liu et al. 2019), and (iii) a linear step search, e.g., using the Wolfe conditions (Wolfe 1969; Nocedal & Wright 2006), along the search direction. In this section, we focus on the BFGS optimization algorithm, discuss its applicability in FWI, and its potential for uncertainty quantification.

At iteration k , the quasi-Newton search direction $\mathbf{p}_k$ is given by

$$\mathbf{p}_k = -\mathbf{B}_k \nabla f(\mathbf{m}_k), \quad (1)$$

where $\mathbf{g}_k = \nabla f(\mathbf{m}_k)$ denotes the gradient, and $\mathbf{B}_k$ the inverse Hessian or its approximation. The gradient $\mathbf{g}_k$ can be efficiently computed using the adjoint-state method (Tromp et al. 2005; Plessix 2006). For practical applications, direct computation and storage of the Hessian or its inverse are prohibitive. Instead, L-BFGS provides an efficient matrix-free and iterative approach, which first approximates the Hessian in a rank-two update (Broyden 1970; Fletcher 1970; Goldfarb 1970; Shanno 1970), namely,

$$\mathbf{H}_{k+1} = \mathbf{H}_k + \frac{\mathbf{y}_k \mathbf{y}_k^T}{\mathbf{y}_k^T \mathbf{s}_k} - \frac{\mathbf{H}_k \mathbf{s}_k \mathbf{s}_k^T \mathbf{H}_k^T}{\mathbf{s}_k^T \mathbf{H}_k \mathbf{s}_k}, \quad (2)$$

with $\mathbf{y}_k = \mathbf{g}_{k+1} - \mathbf{g}_k$ and $\mathbf{s}_k = \mathbf{m}_{k+1} - \mathbf{m}_k$, and then applies the Sherman-Morrison formula (Sherman & Morrison 1950) to get the inverse Hessian,

$$\mathbf{B}_{k+1} = \left(\mathbf{I} - \frac{\mathbf{s}_k \mathbf{y}_k^T}{\mathbf{y}_k^T \mathbf{s}_k} \right) \mathbf{B}_k \left(\mathbf{I} - \frac{\mathbf{y}_k \mathbf{s}_k^T}{\mathbf{y}_k^T \mathbf{s}_k} \right) + \frac{\mathbf{s}_k \mathbf{s}_k^T}{\mathbf{y}_k^T \mathbf{s}_k}, \quad (3)$$

from which we observe that $\mathbf{B}_{k+1}$ remains in a matrix form, impractical for storage. L-BFGS, which approximates BFGS in a limited amount of memory by considering past gradients and model updates, runs in a two-loop recursion algorithm, shown in Alg. (1), with m being the memory value, and γ_k a scaling factor (Morales & Nocedal 2000).

Algorithm 1 L-BFGS search direction computation

```

1:  $\mathbf{q} = \mathbf{g}_k$ 

2: for  $i \leftarrow k - 1$  to  $k - m$  do

3:    $r_i = 1/\mathbf{y}_i^T \mathbf{s}_i$ 

4:    $a_i = r_i \mathbf{s}_i^T \mathbf{q}$ 

5:    $\mathbf{q} = \mathbf{q} - a_i \mathbf{y}_i$ 

6: end for

7:  $\gamma_k = (\mathbf{s}_{k-1}^T \mathbf{y}_{k-1})/(\mathbf{y}_{k-1}^T \mathbf{y}_{k-1})$ 

8:  $\mathbf{z} = \gamma_k \mathbf{B}_k^0 \mathbf{q}$ 

9: for  $i \leftarrow k - m$  to  $k - 1$  do

10:   $b_i = r_i \mathbf{y}_i^T \mathbf{z}$ 

11:   $\mathbf{z} = \mathbf{z} + \mathbf{s}_i(a_i - b_i)$ 

12: end for

13:  $\mathbf{z} = -\mathbf{z}$ 

```

Alg. (1) outputs the scaled search direction $\mathbf{z} = -\gamma_k \mathbf{B}_k \mathbf{g}_k$ towards a local minimizer with linear memory requirements. Besides being efficient and inexpensive, this recursion has the advantage that the initial inverse Hessian $\mathbf{B}_k^0$ at iteration k is included, but isolated from the two-loop recursion. The factor γ_k attempts to estimate an effective scaling for $-\mathbf{B}_k \mathbf{g}_k$, so that a unit step length is accepted for most iterations. $\mathbf{B}_k^0$ opens a window for the preconditioner trials we will conduct for different initial Hessian estimates.

Liu et al. (2019) and Liu and Peter (2019) discuss the feasibility of Square-Root Variable Metric (SRVM) based FWI and uncertainty quantification. Both SRVM and BFGS belong to the family of quasi-Newton methods. They only differ in that SRVM originates from the Davidon-Fletcher-Powell (DFP) algorithm (Davidon 1959; Fletcher & Powell 1963), the dual of BFGS. In an iterative manner, DFP produces a direct approximation to the inverse Hessian while BFGS first approximates the Hessian and then takes its inverse based on the Sherman-Morrison formula. Thus, similar to DFP, we can reconstruct the inverse Hessian from BFGS-based FWI.

L-BFGS, a variant of BFGS, has been the state-of-the-art optimization framework for decades in exploration seismology. Nocedal & Wright (2006) state that L-BFGS is equivalent to BFGS

if all the memories are kept, while the initial guess of the inverse Hessian remains unchanged. Theoretically, BFGS can capture second-order derivative information from start to finish without dropping histories. As a result, $\mathbf{B}_k^0$ becomes independent of iterations, so we can recast $\mathbf{B}_k^0$ as $\mathbf{B}_0$. Thus, there is no need to update the initial Hessian during subsequent iterations. Modrak & Tromp (2016) also illustrated that it is unnecessary to regularly update the preconditioner for FWI within a chosen frequency band.

2.2 Preconditioners in BFGS

Ideally, the inverse Hessian in elastic FWI can be used to reforge the gradient to directly account for parameter trade-offs as well as source-receiver illumination (Pratt 1999; Virieux et al. 2009; Métivier et al. 2013; Yang et al. 2018; Beller & Chevrot 2020). However, explicit computation and storage of the Hessian and its inverse in practical applications are prohibitive. FWI is formulated as a nonlinear minimization of the waveform mismatch between observed and synthetic data via an iterative procedure, indirectly accounting for the inverse Hessian. For large-scale problems (Fichtner et al. 2018; Lei et al. 2020), however, gradients estimated by the adjoint method are also expensive, such that we can only afford a handful of iterations (Tromp 2019). As suggested in Alg. (1), the inverse of an initial Hessian guess can be applied to partially improve performance and reduce the cost of the computationally demanding nonlinear-optimization procedure. Therefore, we first need a direct or indirect estimate of the initial Hessian.

Regarding initial Hessian estimation, in general there are three types of categories: (i) iterative Gauss-Newton, (ii) point-spread function, and (iii) diagonal Hessian estimates. Category (i) involves a Hessian update on the fly (Demanet et al. 2012; Métivier et al. 2013); category (ii) is based on point-spread functions for an initial Hessian estimation (Zhu et al. 2016); category (iii), which is most popular in exploration geophysics, constructs an initial diagonal Hessian using Gauss-Newton methods (Claerbout 1971; Shin et al. 2008; Rickett 2003; Yang et al. 2018). We mainly focus on category (iii) due to its effectiveness and ease of implementation, which will be a good starting point for a more sophisticated algorithm, such as L-BFGS, for inverse Hessian construction.

In this study, we take the category (iii) approach. Given $\mathbf{J} = \delta \mathbf{s} / \delta \mathbf{m}$ as the first-order Fréchet

derivative, to approximate the initial inverse Hessian, $\mathbf{B}_0$ in Alg. (1) can be defined as

$$\mathbf{B}_0 = [\text{diag}(\mathbf{J}^\dagger \mathbf{J}) + \lambda \mathbf{I}]^{-1} = (\mathbf{H}_0 + \lambda \mathbf{I})^{-1}, \quad (4)$$

in which a superscript $\dagger$ denotes the adjoint, $\mathbf{H}_0 = \text{diag}(\mathbf{J}^\dagger \mathbf{J})$ indicates the diagonal initial Hessian, and the regularization term $\lambda \mathbf{I}$ exists for stabilization. For convenience, we call preconditioners the initial diagonal approximations of the Hessian although their inverses enter the algorithms.

We can account for the diagonal Hessian in two different ways: simple data illumination (Claerbout 1971; Rickett 2003; Shin et al. 2008) and the Gauss-Newton method (Yang et al. 2018). Each has two alternatives: consideration of only the source geometry or both the source and receiver geometries. Given elastic FWI for P- and S-wave speeds α and β , let us start from the well-known source-illumination map (Rickett 2003)

$$\mathbf{H}_0(\mathbf{x}) = \int \partial_t \mathbf{v}(\mathbf{x}, t) \cdot \partial_t \mathbf{v}(\mathbf{x}, t) dt, \quad (5)$$

in which $\mathbf{v}(\mathbf{x}, t)$ denotes the source wavefield particle velocity. We impose the same $\mathbf{H}_0(\mathbf{x})$ over the gradients of α and β to obtain our first kind of preconditioner $\mathbf{P}_1 = \{\mathbf{H}_0(\mathbf{x}), \mathbf{H}_0(\mathbf{x})\}$. Luo (2012) derives a similar form called the “ray density” preconditioner. Later, we will show in our examples that even this kind of simple preconditioner can lead to a significant improvement in FWI convergence. However, $\mathbf{H}_0$ fails to account for the acquisition geometry, so we introduce a modified form of eq. (5), namely,

$$\bar{\mathbf{H}}_0(\mathbf{x}) = \left| \int \partial_t \mathbf{v}(\mathbf{x}, t) \cdot \partial_t \bar{\mathbf{v}}(\mathbf{x}, t) dt \right|, \quad (6)$$

in which $\mathbf{v}(\mathbf{x}, t)$ and $\bar{\mathbf{v}}(\mathbf{x}, t)$ are the source and receiver velocity wavefields, respectively. Following Luo (2012), we take the absolute value in eq. (5) to ensure positive definiteness of the approximate initial Hessian. As a result, our second preconditioner is $\mathbf{P}_2 = \{\bar{\mathbf{H}}_0(\mathbf{x}), \bar{\mathbf{H}}_0(\mathbf{x})\}$. The computation of $\mathbf{P}_1$ only involves the source wavefield, whereas that of $\mathbf{P}_2$ involves both the source and receiver wavefields. Therefore, it is expected to see a better performance of $\mathbf{P}_2$ than $\mathbf{P}_1$ in compensation for uneven source-receiver data coverage.

Although $\mathbf{P}_1$ and $\mathbf{P}_2$ can be useful in accelerating FWI, an approximate “elastic” initial Hessian should have better performance due to accommodating inherent nonlinearity in elastic inversions. A mathematical derivation for the multi-parameter initial Hessian can be found in

Appendix A. Eq. (A21) shows that the elastic initial Hessian can be expressed as

$$\begin{aligned} H_{\alpha\alpha}(\mathbf{x}) &= \int \left[\frac{\partial \mathbf{C}}{\partial \alpha} \mathbf{D} \mathbf{v}(\mathbf{x}, t) \right] \cdot \left[\frac{\partial \mathbf{C}}{\partial \alpha} \mathbf{D} \mathbf{v}(\mathbf{x}, t) \right] dt, \\ H_{\beta\beta}(\mathbf{x}) &= \int \left[\frac{\partial \mathbf{C}}{\partial \beta} \mathbf{D} \mathbf{v}(\mathbf{x}, t) \right] \cdot \left[\frac{\partial \mathbf{C}}{\partial \beta} \mathbf{D} \mathbf{v}(\mathbf{x}, t) \right] dt, \end{aligned} \quad (7)$$

where $\mathbf{v}(\mathbf{x}, t)$ denotes the particle velocity of the source wavefield, $\mathbf{C}$ the (α, β, ρ) related stiffness matrix, and $\mathbf{D}$ a combination of differential operators. Details of $\mathbf{C}$ and $\mathbf{D}$ can be found after eq. (A7). The resulting preconditioner of the third kind can be expressed as $\mathbf{P}_3 = \{H_{\alpha\alpha}(\mathbf{x}), H_{\beta\beta}(\mathbf{x})\}$. Again, as discussed in eq. (A22) of Appendix A, we further extend eq. (7) to account for the acquisition geometry:

$$\begin{aligned} \bar{H}_{\alpha\alpha}(\mathbf{x}) &= \left| \int \left[\frac{\partial \mathbf{C}}{\partial \alpha} \mathbf{D} \mathbf{v}(\mathbf{x}, t) \right] \cdot \left[\frac{\partial \mathbf{C}}{\partial \alpha} \mathbf{D} \bar{\mathbf{v}}(\mathbf{x}, t) \right] dt \right|, \\ \bar{H}_{\beta\beta}(\mathbf{x}) &= \left| \int \left[\frac{\partial \mathbf{C}}{\partial \beta} \mathbf{D} \mathbf{v}(\mathbf{x}, t) \right] \cdot \left[\frac{\partial \mathbf{C}}{\partial \beta} \mathbf{D} \bar{\mathbf{v}}(\mathbf{x}, t) \right] dt \right|. \end{aligned} \quad (8)$$

As a result, the preconditioner of the fourth kind becomes $\mathbf{P}_4 = \{\bar{H}_{\alpha\alpha}(\mathbf{x}), \bar{H}_{\beta\beta}(\mathbf{x})\}$. It is expected that $\mathbf{P}_4$ outperforms $\mathbf{P}_3$. As for the computational burden, if we compute the initial Hessian separately, $\mathbf{P}_1$ and $\mathbf{P}_3$ come at the cost of one wavefield simulation, and $\mathbf{P}_2$ and $\mathbf{P}_4$ come at the cost of one gradient computation. Thus, the computations are all cheap compared to the total cost of FWI. When it comes to large-scale applications, we can further minimize the computational cost of the preconditioners by computing them together with the gradient. Also, for each positive-definite preconditioner, a large ratio between its maximum and minimum (similar to a large condition number) can result in numerical instability, which can be alleviated by smoothing and damping (Rickett 2003). Here we use the inverse of the smoothed, damped initial diagonal Hessian $\text{diag}(\mathbf{J}^\top \mathbf{J})$ in the form of $\mathbf{P}_1$, $\mathbf{P}_2$, $\mathbf{P}_3$, and $\mathbf{P}_4$, respectively, to take the role of $\mathbf{B}_0$ in Alg. (1).

2.3 Preconditioned BFGS-based Uncertainty Quantification

With the BFGS algorithm and its preconditioners in place, when elastic FWI converges after n iterations, $\mathbf{B}_{n+1}$ can approximate the inverse Hessian from the past n histories (Tarantola 2005; Nocedal & Wright 2006) as

$$\mathbf{H}^{-1} = \mathbf{B}_{n+1}. \quad (9)$$

From Alg. (1), we see that the reconstruction of $\mathbf{B}_{n+1}$ involves $\mathbf{B}_0$, $\mathbf{s}_i$, and $\mathbf{y}_i$ ($i = 0, 1, \dots, n$). After preconditioning with $\mathbf{B}_0$, the wavefield related pieces of information are mainly embed-

ded in $\mathbf{B}_{n+1}\mathbf{B}_0^{-1}$, which starts from the identity $\mathbf{I}$ for stabilization. Note that $\mathbf{B}_0$ has been estimated and kept, and $\mathbf{B}_{n+1}\mathbf{B}_0^{-1}$ rather than $\mathbf{B}_{n+1}$ is retrieved from $\mathbf{s}_i$ and $\mathbf{y}_i$.

The retrieval of $\mathbf{B}_{n+1}$ from $\mathbf{B}_0$, $\mathbf{s}_i$, and $\mathbf{y}_i$ can be accomplished with a unit pulse probing vector $\hat{\mathbf{e}}_j = (0, 0, \dots, 0, 1, 0, \dots, 0, 0)$ in which the unit pulse “1” is located at the target column (row) index j ($j = 1, \dots, M$, with M being the model dimension). This is a straight forward procedure, however we will see that it is not a very efficient method. Let us extract the j -th row/column elements from $\mathbf{B}_{n+1}$ as depicted in Alg. (2).

Algorithm 2 $\mathbf{B}_{n+1}$ probing

```

1:  $\mathbf{q} = \hat{\mathbf{e}}_j$ 
2: for  $i \leftarrow n$  to 0 do
3:    $r_i = 1/\mathbf{y}_i^T \mathbf{s}_i$ 
4:    $a_i = r_i \mathbf{s}_i^T \mathbf{q}$ 
5:    $\mathbf{q} = \mathbf{q} - a_i \mathbf{y}_i$ 
6: end for
7:  $\mathbf{z} = \mathbf{B}_0 \mathbf{q}$ 
8: for  $i \leftarrow 0$  to  $n$  do
9:    $b_i = r_i \mathbf{y}_i^T \mathbf{z}$ 
10:   $\mathbf{z} = \mathbf{z} + \mathbf{s}_i(a_i - b_i)$ 
11: end for
12:  $\mathbf{z} = \mathbf{z} - \mathbf{B}_0 \hat{\mathbf{e}}_j$ 

```

The two-loop algorithm Alg. (2) outputs the matrix-vector product $\mathbf{z} = (\mathbf{B}_{n+1} - \mathbf{B}_0) \hat{\mathbf{e}}$. However, it is not flexible for arbitrary element extraction. For example, given a model of size M , we need M such operations in extracting the $\mathbf{B}_{n+1}$ diagonals, which is expensive for uncertainty quantification of large-scale applications.

Fortunately, randomized SVD (Liberty et al. 2007; Halko et al. 2011) provides a more efficient eigendecomposition of large matrices, especially for those with low rank. The method proposed in Halko et al. (2011) can probe a matrix only with one set of random vectors in a much simpler implementation. Given an elastic FWI with model size M converges after N_r iterations, the inverse Hessian $\mathbf{B} = \mathbf{H}^{-1} = \mathbf{B}_{n+1}$ can span a $M \times M$ full matrix, which is pro-

hibitive in terms of computation and storage. However, we can probe it with a set of $M \times N_r$ independent random vectors $\mathbf{X}$, similar to the approach of Liu & Peter (2019), to extract its low-rank form. Single-pass randomized SVD allows for an efficient eigendecomposition of the matrix $\mathbf{B}$ of size $M \times M$, as shown in Alg. (3).

Algorithm 3 Single-pass randomised SVD algorithm

1:	$\mathbf{E} = \mathbf{B}\mathbf{X}$	▷ Sampling $\mathbf{B}$ with $\mathbf{X}$
2:	$\mathbf{Q}\mathbf{R} = \mathbf{E}$	▷ QR decomposition on $\mathbf{E}$
3:	$\Omega(\mathbf{Q}^T \mathbf{X}) = \mathbf{Q}^T \mathbf{E}$	▷ Solve for Ω
4:	$\mathbf{U}\Lambda\mathbf{U}^T = \Omega$	▷ SVD on Ω
5:	$\mathbf{V} = \mathbf{Q}\mathbf{U}$	
6:	$\mathbf{Z} = \mathbf{V}\Lambda\mathbf{V}^T$	
-

When setting the target as $\mathbf{B} = \mathbf{H}^{-1}$, we use $\mathbf{X}$ to probe $\mathbf{B}_{n+1} - \mathbf{B}_0$, yielding $\mathbf{E} = \mathbf{B}_{n+1}\mathbf{X} - \mathbf{B}_0\mathbf{X}$. Here $\mathbf{X}$ consists of N_r independent random vectors, with $N_r \ll M$. Note that we never write out the full matrix $\mathbf{B}$. We simply need to replace the $\hat{\mathbf{e}}_j$ in Alg. (2) with x_k , ($x_k \in \mathbf{X}$, and $k = 1, \dots, N_r$), to have $\mathbf{E}$. When performing the QR decomposition, we only need to keep $\mathbf{Q}$ of size $M \times N_r$. Then, except that $\mathbf{V}$ is of size $M \times N_r$, all other matrices are of size $N_r \times N_r$. Finally, Alg. (3) outputs the inverse Hessian in an SVD form as

$$\mathbf{H}^{-1} = \mathbf{V}\Lambda\mathbf{V}^T + \mathbf{B}_0, \quad (10)$$

with $\mathbf{V}$ being an $M \times N_r$ eigenvector matrix, Λ the eigenvalue matrix with N_r diagonal entries, and $\mathbf{B}_0$ the damping term. Eq. (10) provides a convenient way to access arbitrary elements of the inverse Hessian.

For a linearized inverse problem with Gaussian priors, the relation between the model prior and posterior covariance matrices $\mathbf{C}_M$ and $\mathbf{C}_m$ can be expressed as (Tarantola 2005; Rawlinson et al. 2014)

$$\mathbf{C}_M = (\mathbf{J}^\dagger \mathbf{C}_d^{-1} \mathbf{J} + \mathbf{C}_m^{-1})^{-1}, \quad (11)$$

where $\mathbf{J}$ denotes the Jacobian matrix, and $\mathbf{C}_d$ the data covariance matrix. $\mathbf{C}_m$ can function as a regularization term in inverse problems (Tarantola 2005). When solving tomographic inverse problems with the least-squares formula, a “tuning” scalar for balancing the data-misfit and

model-misfit contributions can be estimated based on the relative scaling between $\mathbf{C}_m$ and $\mathbf{C}_d$. Without information gain from the data, we have $\mathbf{C}_M \rightarrow \mathbf{C}_m$. Otherwise, for a well-converged inversion, the information gain from the data dominates eq. (11) as $\mathbf{C}_M \rightarrow (\mathbf{J}^\dagger \mathbf{C}_d^{-1} \mathbf{J})^{-1}$. Therefore, when seismic tomography converges, following Tarantola (2005), Rawlinson et al. (2014), and Liu & Peter (2019), we choose to express eq. (11) in a practical way as

$$\mathbf{C}_M = \mathbf{C}_m^{1/2} (\mathbf{C}_m^{1/2} \mathbf{J}^\dagger \mathbf{C}_d^{-1} \mathbf{J} \mathbf{C}_m^{1/2} + \mathbf{I})^{-1} \mathbf{C}_m^{1/2}, \quad (12)$$

In doing so, $\mathbf{C}_m^{1/2}$ will appear as a transformation matrix along with $\mathbf{J}$ and $\mathbf{J}^\dagger$, yielding the approximate inverse Hessian as

$$\mathbf{H}^{-1} = (\mathbf{C}_m^{1/2} \mathbf{J}^\dagger \mathbf{C}_d^{-1} \mathbf{J} \mathbf{C}_m^{1/2} + \mathbf{I})^{-1}, \quad (13)$$

which can be obtained using Alg. (1), noting that $\mathbf{y}_k$ changes with $\mathbf{g}_{k+1}$ and $\mathbf{g}_k$. Then, based on Alg. (2), we have

$$\mathbf{C}_M = \mathbf{C}_m^{1/2} (\mathbf{V} \mathbf{\Lambda} \mathbf{V}^T + \mathbf{B}_0) \mathbf{C}_m^{1/2}. \quad (14)$$

$\mathbf{C}_m$ and $\mathbf{C}_M$ denote the prior and posterior model distributions (Bui-Thanh et al. 2013; Zhu et al. 2016; Liu & Peter 2019) and can be used to assess the model null-space (Liu & Peter 2020). Note that even without data information gain from the iterative inversion, $\mathbf{B}_0$ can serve as a rough estimate of uncertainty. The prior $\mathbf{C}_m$ may be estimated based on geological information and interpretation, well logs, or seismic imaging (Fomel & Claerbout 2003). Prior information has a significant impact on the inversion. A prior $\mathbf{C}_m$ that is too general provides little useful information, but one that is too strict may lead to biased inversion results (Tarantola 2005). To simplify the remainder of this paper, we equate $\mathbf{C}_m$ to a scaled identity matrix $\epsilon \mathbf{I}$, where ϵ is related to the standard deviation of the model prior. The square-root diagonals of $\mathbf{C}_M$, known as the standard deviations, provide a quantitative measure of the posterior distribution (Tarantola 2005). We observe from eqs (12) and (14) that, given a simple $\mathbf{C}_m$, the main features of $\mathbf{C}_M$ are reflected by $\mathbf{H}^{-1}$. The approximation of $\mathbf{H}^{-1}$ via preconditioned-BFGS is the focus of this paper.

3 NUMERICAL EXAMPLES

Let us consider the isotropic elastic Marmousi benchmark model, with a parametrization of P- and S-wave speeds and density, as a well-studied application in exploration seismology. We only account for the P- and S- wave speeds V_p and V_s , but fix the density in the inversion due to its low-sensitivity in FWI (Virieux & Operto 2009a; Blom et al. 2017) without limiting the applicability of our approach. In the following, we aim to show the feasibility of BFGS and preconditioned-BFGS algorithms in elastic FWI, and compare them in terms of convergence rates and uncertainty quantification maps.

3.1 2D elastic Marmousi benchmark

The 2D elastic Marmousi benchmark (Martin et al. 2006) is popular in exploration geophysics due to its substantial structure complexities and wavespeed variations that pose nonlinear challenges in FWI. We use a modified Marmousi here, whose model dimensions are 9,200 m in the horizontal and 3,000 m in the vertical directions, without a water layer. Fig. 1 shows the true and initial elastic models, respectively. Fig. 1b is smoothed from Fig. 1a with a Gaussian blur wide enough to remove discontinuities and distort traveltimes. We run forward and adjoint simulations with a 2D spectral-element code, i.e., SPECFEM2D (Komatitsch & Vilotte 1998; Komatitsch & Tromp 1999), using absorbing boundaries (Stacey 1988; Komatitsch & Martin 2007) around all model edges to mimic wavefield propagation in an infinite domain. The observation system located 10 m underground consists of evenly-distributed 32 shots between 279 m and 8,921 m, and 500 multi-component geophones between 100 m and 9,100 m, respectively. The source time function is a Ricker wavelet of 4 Hz peak frequency. For the simulations, the time step is 0.9 ms, and the recording duration is 7.2 s. As a FWI workflow tool, we use SeisFlows (Modrak & Tromp 2016).

We start from running BFGS-based elastic FWI with full-memory L-BFGS. To validate the effectiveness of BFGS in FWI, we compare its inversion results in Fig. 2b with those by L-BFGS with 5 memories in Fig. 2a (a value of 5 is commonly used in L-BFGS approaches). We see that they yield almost identical image results in V_P and V_S . Fig. 3 further compares the convergence behavior of BFGS and L-BFGS based FWI in terms of normalized data and model misfits. It shows that for our particular example and acquisition setup, BFGS exhibits a

slightly faster convergence rate than L-BFGS, but at the price of using the full-history memory of vectors $\mathbf{s}_i$ and $\mathbf{y}_i$ mentioned above. Similar L-BFGS and SRVM comparisons can be found in Liu et al. (2019), where the SRVM algorithm resembles BFGS in keeping all the memory states. Note that $\mathbf{s}_i$ and $\mathbf{y}_i$ sizes are the same as the model size, and their storage increases linearly with iterations.

Generally, geophysical inverse problems can encounter high nonlinearities, for example a salt-body contrast. In such cases, L-BFGS approaches benefit from lower memory values (Modrak & Tromp 2016). However, to alleviate the nonlinearity of the inverse problem, one can resort to approaches using multi-scale inversion strategies (Bunks et al. 1995), adaptive waveform inversion (Warner & Guasch 2016), or optimal-transport metrics (Métivier et al. 2016). Also, it is worth mentioning that a main target of our study is to apply this method in global waveform tomography, which often settles into superlinear convergence rates (Tromp 2019).

The aim of preconditioning in FWI is to accelerate convergence rates. In theory, the optimal preconditioner is the true inverse Hessian itself. Thus, preconditioning aims to find good initial Hessian approximations. Regarding the four preconditioners mentioned above, P1 and P3 only need the forward wavefield while P2 and P4 involve both the source and receiver wavefields. All preconditioners are computed before entering the FWI iterations. Smoothing and damping are required to avoid numerical instabilities. We smooth the Hessian the same way as the gradient, and then stabilize the inverse operation with the median of the smoothed Hessian. Note that one could try to tune the stabilizing coefficient for performance improvements, whereas we fix it based on the minimum and maximum values of the smoothed Hessian. Fig. 4 shows the inversion results of preconditioned BFGS-based FWI with P1, P2, P3, and P4, respectively. Comparison of Figs 2b and 4 indicates that the preconditioned BFGS algorithms perform almost identical, or even better than pure BFGS in driving waveform tomography. The convergence comparisons in Fig. 5 further highlight significant convergence speedups, thanks to the preconditioners. Fig. 5a shows that P3 and P4 outperform P1 and P2, respectively, in reaching similar data misfits but with fewer iterations, as expected. These cross-comparisons stress the importance of accounting for the limited acquisition geometry in preconditioner estimations. Besides the stored preconditioners, we keep the set of $\mathbf{s}_i$ and $\mathbf{y}_i$ vectors for the inverse Hessian reconstruction. Even with full-memory histories, $\mathbf{s}_i$ and $\mathbf{y}_i$ inflict manageable storage burdens.

We will see later that when accessing the approximate inverse Hessian, we only need to fetch small segments of the stored files into memory per operation using a memory-map technique, e.g., with a MemMap algorithm in NumPy, to mitigate the peak memory cost. Afterwards, we can factorize the inverse Hessian by running randomized SVD over the stored preconditioner and the s_i and y_i vectors. During this process, we never write out the full matrix of the inverse Hessian thanks to randomized SVD. As discussed previously and in Liu & Peter (2019), the rank of the approximate inverse Hessian equals the number of iterations N_{iter} . The resulting benefit is that we only need N_{iter} such independent random vectors of model size for the random probing $E = ZX$ to start Alg. (3), which subsequently outputs the inverse Hessian in SVD form. The eigen-order number of the factorized matrix is then the same as the iteration number.

The top of Figs. 6 to 10 shows the eigen-spectrum of the inverse Hessians from the BFGS-based elastic FWIs. Each of them starts from values with similar magnitudes, but ends differently depending on the level of convergence. Their corresponding eigenvector groups are displayed following each eigen-spectrum. Each group has five eigenvectors of increasing eigen-orders, including the 1st, 5th, 10th, 20th, and final one. All eigenvectors per group are orthogonal. In each eigenvector group, we see that as the eigen-order increases, the energies begin to move up towards the observation system. We derive this tendency of the energy moving from the inverse of the inverse Hessian B , i.e., the Hessian H , whose eigen-energies fade away from the observation system as the eigen-order increases (Bui-Thanh et al. 2013; Zhu et al. 2016). Another intriguing observation is that there is some random noise in Figs. 7f to 10f. When looking at the corresponding eigen-spectra, we interpret this as due to rank deficiency. It means that the actual rank of this matrix is less than the pre-estimated number, which has the advantage that one can safely truncate the SVD to further save storage.

With the inverse Hessians in SVD form, we can efficiently extract their standard deviations for uncertainty quantification maps, as shown in Fig. 11. Overall, they all look similar and the features of the uncertainty maps resemble those by SRVM-based FWI (Liu & Peter 2019) and Ensemble Kalman Filtering (EnKF) (Thurin et al. 2019). Noticeable is that the methods reported here are based on the most popular quasi-Newton method, L-BFGS, without much additional cost. To understand the uncertainty maps from a physical perspective, again, we can

consider the data coverage, which directly reflects the Hessian. Data coverage decreases due to geometrical spreading and, specifically, a high-wavespeed anomaly may deflect energy. We can thus infer that the appearance of the uncertainty maps should counteract the Hessian characteristics. The uncertainty maps in Fig. 11 reflect multiple features of the inverse Hessian, in particular: (1) uncertainty increases as data coverage decreases; (2) high-wavespeed structures are often related to relatively high uncertainties, especially for the two salts in the corners. In short, the uncertainty map indicates that the model space with better data coverage will have more information gain, i.e., smaller uncertainties.

For a detailed investigation of the uncertainty maps, we first look at Fig. 11 horizontally to compare the V_p row with the V_s row, seeing that the V_p model usually has larger uncertainties than V_s . We explain this from the perspective that in elastic isotropic media the radiation pattern of V_p is isotropic while that of V_s is far-offset dominant. FWI favors far-offset data because it can produce gradients of lower-wavenumber than near-offset data. Such lower-wavenumber components, usually with stronger amplitudes, can further help FWI to overcome local minima, to some extent. The vertical row comparison in Fig. 11 shows uncertainty maps from BFGS-based FWIs with different preconditioners. We note that all preconditioners in BFGS result in almost identical V_p and V_s maps, however at different convergence rates. We conclude that preconditioners estimated from the linear Gauss-Newton method outperform those purely derived from data illumination. Furthermore, the preconditioners which consider the recording geometry outperform those which do not. Finally, we regard the salt bodies as good identifiers to assess the quality of the presented uncertainty maps, because for those areas we expect to see higher uncertainty values. Based on the above considerations, we infer that among the presented preconditioners the P4 preconditioner has the most advantages, providing reasonable uncertainty maps in Fig. 11e, and the fastest convergence rates in Fig. 5.

Uncertainty quantification provides an *a posteriori* assessment of the model. To further validate the uncertainty maps in Fig. 11, we show in Fig. 12 the absolute model difference between the true model in Fig. 1 and one inverted model in Fig. 4d. We observe that Fig. 12 resembles Fig. 11, especially Figs. 11c and 11e with preconditioners P2 and P4.

Overall, we see that preconditioned BFGS algorithms demonstrate considerable speedup in FWI and yield reasonable accompanying uncertainty maps. It is worth mentioning that our

preconditioned BFGS algorithms can readily be incorporated into standard FWI workflows. The extra costs in computation and storage before, during, and after waveform inversion are all manageable even for large-scale inversions. This makes our method suitable for exascale geophysical applications, addressing uncertainty quantification in, e.g., global-scale waveform inversions (Bozdağ et al. 2016; Fichtner et al. 2018; Lei et al. 2020).

4 DISCUSSION

In real-world applications, parts of the model space may have zero data coverage. We do not expect resolution in such places, simply because the acquired data do not sample those parts. Assuming the Hessian is diagonal, a one zero diagonal entry would make it non-invertible. Conceptually, this means an “infinite uncertainty”: there is no way to constrain anything at this location. That is where regularization comes into play in direct inversion: we add a small regularization term to render the matrix invertible. As for uncertainty quantification, we do have prior knowledge about the Earth at our disposal (values and some standard deviation). It is this prior that acts as a regularization term in Eq. (11). Again, the “data-misfit Hessian” — assuming a diagonal approximation for simplicity — is nearly zero in areas with zero data sensitivity. That is to say, we do not have data information gain in those parts, and the uncertainty remains the same, and the prior covariance C_m dominates Eq. (11). The remaining challenge is to identify the “zero-sensitivity” places, which usually feature absolutely zero model updates throughout the inversion. We can assign model prior uncertainties to these places in our final uncertainty map.

5 CONCLUSION

Uncertainties inherently exist in geophysical inverse problems, such as waveform tomography, due to limitations in observations, theories, and algorithms. The L-BFGS algorithm has become the most popular optimization method in applied mathematics, including for seismic waveform tomography. We have demonstrated that the BFGS algorithm has the potential to address uncertainty quantification in FWI. Furthermore, well-specified preconditioners can help gain significant computational savings, yielding the same, or even better, inversion results.

The estimation of uncertainty maps only requires the storage of the related preconditioner and a set of memory vectors in L-BFGS. Based on the variable-metric component of the BFGS algorithm during the inversion process, these stored vectors inherently communicate multi-parameter information from the initial model to the inverted one.

Note that in this article we are addressing a *linearized* inverse problem. This means that we linearize the forward modeling operator around the maximum *a posteriori* (MAP) point. To ensure that this approach is appropriate, a good starting model and careful data processing, involving rigorous data quality control and judicious data selection (e.g., frequency blending, multi-taper measurements, double-difference measurements, etc.), may be required. When encountering nonlinearities, one may have to resort to multi-scale inversion strategies. Because the uncertainty map needs a set of L-BFGS memory vectors, we want the inversion to converge smoothly for a given objective function.

The combination of BFGS and randomized SVD methods facilitates the retrieval of a low-rank representation of the inverse Hessian, rather than the Hessian (Bui-Thanh et al. 2013; Zhu et al. 2016), from which standard deviations provide a straightforward assessment of inversion convergence and uncertainty. Our approach is strictly based on a standard FWI workflow, augmenting it to a Bayesian inversion under the assumption of linearized forward modeling and Gaussian model priors. Finally, the presented BFGS-based FWI and uncertainty quantification workflows are fully scalable and readily amenable to seismic exascale applications.

6 ACKNOWLEDGEMENTS

The authors are grateful to editor Carl Tape and reviewer Christian Boehm and an anonymous reviewer for improving the initial manuscript. The authors are grateful to Frederik J. Simons for inspiring discussions. This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725.

7 DATA AVAILABILITY

Data available on request.

REFERENCES

- Bachmann, E. & Tromp, J., 2020. Source encoding for viscoacoustic ultrasound computed tomography, *The Journal of the Acoustical Society of America*, **147**(5), 3221–3235.
- Baysal, E., Kosloff, D. D., & Sherwood, J. W., 1983. Reverse time migration, *Geophysics*, **48**(11), 1514–1524.
- Beller, S. & Chevrot, S., 2020. Probing depth and lateral variations of upper-mantle seismic anisotropy from full-waveform inversion of teleseismic body-waves, *Geophysical Journal International*, **222**(1), 352–387.
- Betancourt, M., 2017. A conceptual introduction to hamiltonian monte carlo, *arXiv:1701.02434*.
- Biswas, R. & Sen, M., 2017. 2d full-waveform inversion and uncertainty estimation using the reversible jump hamiltonian monte carlo, in *SEG Technical Program Expanded Abstracts 2017*, pp. 1280–1285, Society of Exploration Geophysicists.
- Blom, N., Boehm, C., & Fichtner, A., 2017. Synthetic inversions for density using seismic and gravity data, *Geophysical Journal International*, **209**(2), 1204–1220.
- Bozdağ, E., Peter, D., Lefebvre, M., Komatitsch, D., Tromp, J., Hill, J., Podhorszki, N., & Pugmire, D., 2016. Global adjoint tomography: first-generation model, *Geophysical Journal International*, **207**(3), 1739–1766.
- Brossier, R., Operto, S., & Virieux, J., 2009. Seismic imaging of complex onshore structures by 2d elastic frequency-domain full-waveform inversion, *Geophysics*, **74**(6), WCC105–WCC118.
- Broyden, C. G., 1970. The convergence of a class of double-rank minimization algorithms 1. general considerations, *IMA Journal of Applied Mathematics*, **6**(1), 76–90.
- Bui-Thanh, T., Ghattas, O., Martin, J., & Stadler, G., 2013. A computational framework for infinite-dimensional Bayesian inverse problems Part i: The linearized case, with application to global seismic inversion, *SIAM Journal on Scientific Computing*, **35**(6), A2494–A2523.
- Bunks, C., Saleck, F. M., Zaleski, S., & Chavent, G., 1995. Multiscale seismic waveform inversion, *Geophysics*, **60**(5), 1457–1473.
- Chen, K. & Sacchi, M. D., 2017. Elastic least-squares reverse time migration via linearized elastic full-waveform inversion with pseudo-hessian preconditioning elastic lsrtm, *Geophysics*, **82**(5), S341–S358.
- Claerbout, J. F., 1971. Toward a unified theory of reflector mapping, *Geophysics*, **36**(3), 467–481.
- Davidon, W. C., 1959. Variable metric method for minimization, *A.E.C. Research and Development Report*, **ANL-5990**.
- Demanet, L., Létourneau, P.-D., Boumal, N., Calandra, H., Chiu, J., & Snelson, S., 2012. Matrix probing: a randomized preconditioner for the wave-equation hessian, *Applied and Computational Harmonic Analysis*, **32**(2), 155–168.

- Duane, S., Kennedy, A. D., Pendleton, B. J., & Roweth, D., 1987. Hybrid monte carlo, *Physics letters B*, **195**(2), 216–222.
- Fichtner, A. & Simutè, S., 2018. Hamiltonian monte carlo inversion of seismic sources in complex media, *Journal of Geophysical Research: Solid Earth*, **123**(4), 2984–2999.
- Fichtner, A. & Trampert, J., 2011a. Hessian kernels of seismic data functionals based upon adjoint techniques, *Geophysical Journal International*, **185**(2), 775–798.
- Fichtner, A. & Trampert, J., 2011b. Resolution analysis in full waveform inversion, *Geophysical Journal International*, **187**(3), 1604–1624.
- Fichtner, A. & van Leeuwen, T., 2015. Resolution analysis by random probing, *Journal of Geophysical Research: Solid Earth*, **120**(8), 5549–5573.
- Fichtner, A., van Herwaarden, D.-P., Afanasiev, M., Simutè, S., Krischer, L., Çubuk-Sabuncu, Y., Taymaz, T., Colli, L., Saygin, E., Villaseñor, A., et al., 2018. The collaborative seismic earth model: generation 1, *Geophysical research letters*, **45**(9), 4007–4016.
- Fletcher, R., 1970. A new approach to variable metric algorithms, *The computer journal*, **13**(3), 317–322.
- Fletcher, R. & Powell, M. J., 1963. A rapidly convergent descent method for minimization, *The computer journal*, **6**(2), 163–168.
- Fomel, S. & Claerbout, J. F., 2003. Multidimensional recursive filter preconditioning in geophysical estimation problems, *Geophysics*, **68**(2), 577–588.
- French, S. & Romanowicz, B., 2014. Whole-mantle radially anisotropic shear velocity structure from spectral-element waveform tomography, *Geophysical Journal International*, **199**(3), 1303–1327.
- Gebraad, L., Boehm, C., & Fichtner, A., 2019. Bayesian elastic full-waveform inversion using hamiltonian monte carlo, *EarthArXiv*, p. qftn5.
- Goldfarb, D., 1970. A family of variable-metric methods derived by variational means, *Mathematics of computation*, **24**(109), 23–26.
- Guasch, L., Agudo, O. C., Tang, M.-X., Nachev, P., & Warner, M., 2020. Full-waveform inversion imaging of the human brain, *NPJ digital medicine*, **3**(1), 1–12.
- Halko, N., Martinsson, P.-G., & Tropp, J. A., 2011. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions, *SIAM review*, **53**(2), 217–288.
- Komatitsch, D. & Martin, R., 2007. An unsplit convolutional Perfectly Matched Layer improved at grazing incidence for the seismic wave equation, *Geophysics*, **72**(5), SM155–SM167.
- Komatitsch, D. & Tromp, J., 1999. Introduction to the spectral element method for three-dimensional seismic wave propagation, *Geophysical Journal International*, **139**(3), 806–822.
- Komatitsch, D. & Vilotte, J. P., 1998. The spectral-element method: an efficient tool to simulate the seismic response of 2D and 3D geological structures, *Bull. seism. Soc. Am.*, **88**(2), 368–392.

- Krischer, L., Fichtner, A., Boehm, C., & Igel, H., 2018. Automated large-scale full seismic waveform inversion for north america and the north atlantic, *Journal of Geophysical Research: Solid Earth*, **123**(7), 5902–5928.
- Lailly, P., 1983. The seismic inverse problem as a sequence of before-stack migrations, in *Proceedings of the Conference on Inverse Scattering, Theory and Application Expanded Abstracts*, pp. 206–220, Society of Industrial and Applied Mathematics, Philadelphia, PA, USA.
- Lefebvre, M., Chen, Y., Lei, W., Luet, D., Ruan, Y., Bozdag, E., Hill, J., Komatitsch, D., Krischer, L., Peter, D., et al., 2017. 13 data and workflow management for exascale global adjoint tomography, *Exascale Scientific Applications: Scalability and Performance Portability*, p. 279.
- 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), 1–21.
- Liberty, E., Woolfe, F., Martinsson, P.-G., Rokhlin, V., & Tygert, M., 2007. Randomized algorithms for the low-rank approximation of matrices, *Proceedings of the National Academy of Sciences*, **104**(51), 20167–20172.
- Liu, D. C. & Nocedal, J., 1989. On the limited memory BFGS method for large scale optimization, *Mathematical programming*, **45**(1), 503–528.
- Liu, Q., 2019. Acoustic full-waveform inversion and its uncertainty estimation based on a vector-version square-root variable metric method, *Inverse Problems*, **35**(9), 095006.
- Liu, Q. & Peter, D., 2019. Square-root variable metric based elastic full-waveform inversion—part 2: uncertainty estimation, *Geophysical Journal International*, **218**(2), 1100–1120.
- Liu, Q. & Peter, D., 2020. Square-root variable metric-based nullspace shuttle: A characterization of the nonuniqueness in elastic full-waveform inversion, *Journal of Geophysical Research: Solid Earth*, **125**(2), e2019JB018687.
- Liu, Q., Peter, D., & Tape, C., 2019. Square-root variable metric based elastic full-waveform inversion—part 1: theory and validation, *Geophysical Journal International*, **218**(2), 1121–1135.
- Luo, Y., 2012. Seismic imaging and inversion based on spectral-element and adjoint methods, *Ph.D. thesis*.
- Métivier, L., Brossier, R., Virieux, J., & Operto, S., 2013. Full waveform inversion and the truncated Newton method, *SIAM Journal on Scientific Computing*, **35**(2), B401–B437.
- Métivier, L., Brossier, R., Méridot, Q., Oudet, E., & Virieux, J., 2016. Measuring the misfit between seismograms using an optimal transport distance: Application to full waveform inversion, *Geophysical Supplements to the Monthly Notices of the Royal Astronomical Society*, **205**(1), 345–377.
- Modrak, R. & Tromp, J., 2016. Seismic waveform inversion best practices: regional, global and exploration test cases, *Geophysical Journal International*, **206**(3), 1864–1889.

- Morales, J. L. & Nocedal, J., 2000. Automatic preconditioning by limited memory quasi-newton updating, *SIAM Journal on Optimization*, **10**(4), 1079–1096.
- Nocedal, J. & Wright, S., 2006. *Numerical Optimization*, Springer Series in Operations Research and Financial Engineering, Springer-Verlag, 2nd edn.
- Peter, D., Komatitsch, D., Luo, Y., Martin, R., Le Goff, N., Casarotti, E., Le Loher, P., Magnoni, F., Liu, Q., Blitz, C., et al., 2011. Forward and adjoint simulations of seismic wave propagation on fully unstructured hexahedral meshes, *Geophysical Journal International*, **186**(2), 721–739.
- Plessix, R., 2006. A review of the adjoint-state method for computing the gradient of a functional with geophysical applications, *Geophys. J. Int.*, **167**, 495–503.
- Polychronopoulou, K., Lois, A., Martakis, N., Chevrot, S., Sylvander, M., Diaz, J., Villaseñor, A., Calassou, S., Collin, M., Masini, E., et al., 2018. Broadband, short-period or geophone nodes? quality assessment of passive seismic signals acquired during the maupasacq experiment, *First Break*, **36**(4), 71–76.
- Pratt, R. G., 1999. Seismic waveform inversion in the frequency domain, part 1: Theory and verification in a physical scale model, *Geophysics*, **64**(3), 888–901.
- Rawlinson, N., Fichtner, A., Sambridge, M., & Young, M. K., 2014. Chapter one-seismic tomography and the assessment of uncertainty, *Advances in Geophysics*, **55**, 1–76.
- Rickett, J. E., 2003. Illumination-based normalization for wave-equation depth migration, *Geophysics*, **68**(4), 1371–1379.
- Shanno, D. F., 1970. Conditioning of quasi-newton methods for function minimization, *Mathematics of computation*, **24**(111), 647–656.
- Sherman, J. & Morrison, W. J., 1950. Adjustment of an inverse matrix corresponding to a change in one element of a given matrix, *The Annals of Mathematical Statistics*, **21**(1), 124–127.
- Shin, C., Jang, S., & Min, D.-J., 2008. Improved amplitude preservation for prestack depth migration by inverse scattering theory, *Geophysical prospecting*, **49**(5), 592–606.
- Stacey, R., 1988. Improved transparent boundary formulations for the elastic-wave equation, *Bulletin of the Seismological Society of America*, **78**(6), 2089–2097.
- 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**(1), 433–462.
- Tarantola, A., 1984. Inversion of seismic reflection data in the acoustic approximation, *Geophysics*, **49**, 1259–1266.
- Tarantola, A., 1988. Theoretical background for the inversion of seismic waveforms, including elasticity and attenuation, in *Scattering and attenuations of seismic waves, part i*, pp. 365–399, Springer.
- Tarantola, A., 2005. *Inverse problem theory and methods for model parameter estimation*, Society for Industrial and Applied Mathematics, Philadelphia, Pennsylvania, USA.

- Tarantola, A. & Valette, B., 1982. Generalized nonlinear inverse problems solved using the least squares criterion, *Reviews of Geophysics*, **20**(2), 219–232.
- Thurin, J., Brossier, R., & Métivier, L., 2019. Ensemble-based uncertainty estimation in full waveform inversion, *Geophysical Journal International*, **219**(3), 1613–1635.
- Trampert, J., Fichtner, A., & Ritsema, J., 2012. Resolution tests revisited: the power of random numbers, *Geophysical Journal International*, **192**(2), 676–680.
- Tromp, J., 2019. Seismic wavefield imaging of earth’s interior across scales, *Nature Reviews Earth & Environment*, pp. 1–14.
- Tromp, J. & Bachmann, E., 2019. Source encoding for adjoint tomography, *Geophysical Journal International*, **218**(3), 2019–2044.
- Tromp, J., Tape, C., & Liu, Q., 2005. Seismic tomography, adjoint methods, time reversal and banana-doughnut kernels, *Geophysical Journal International*, **160**(1), 195–216.
- Vigh, D., Jiao, K., Watts, D., & Sun, D., 2014. Elastic full-waveform inversion application using multicomponent measurements of seismic data collection, *Geophysics*, **79**(2), R63–R77.
- Virieux, A. & Operto, S., 2009a. An overview of full-waveform inversion in exploration geophysics, *Geophysics*, **74**(6), WCC1–WCC26.
- Virieux, J. & Operto, S., 2009b. An overview of full-waveform inversion in exploration geophysics, *Geophysics*, **74**(6), WCC1–WCC26.
- Virieux, J., Operto, S., Ben-Hadj-Ali, H., Brossier, R., Etienne, V., Sourbier, F., Giraud, L., & Haidar, A., 2009. Seismic wave modeling for seismic imaging, *The Leading Edge*, **28**(5), 538–544.
- Warner, M. & Guasch, L., 2016. Adaptive waveform inversion: Theory, *Geophysics*, **81**(6), R429–R445.
- Wolfe, P., 1969. Convergence conditions for ascent methods, *SIAM review*, **11**(2), 226–235.
- Yang, P., Brossier, R., Métivier, L., Virieux, J., & Zhou, W., 2018. A time-domain preconditioned truncated newton approach to visco-acoustic multiparameter full waveform inversion, *SIAM Journal on Scientific Computing*, **40**(4), B1101–B1130.
- Zhang, J. & McMechan, G. A., 1995. Estimation of resolution and covariance for large matrix inversions, *Geophysical Journal International*, **121**(2), 409–426.
- Zhu, H., Bozdağ, E., Peter, D., & Tromp, J., 2012. Structure of the european upper mantle revealed by adjoint tomography, *Nature Geoscience*, **5**(7), 493–498.
- Zhu, H., Li, S., Fomel, S., Stadler, G., & Ghattas, O., 2016. A Bayesian approach to estimate uncertainty for full-waveform inversion using a priori information from depth migration, *Geophysics*, **81**(5), R307–R323.

Figure 1. (a) True and (b) smoothed, initial elastic Marmousi models used in elastic FWI. (Left column: V_P ; right column: V_S .)

Figure 2. Inverted models from elastic FWI using (a) L-BFGS and (b) BFGS after 82 and 65 iterations. L-BFGS runs with 5 memories; BFGS runs in the L-BFGS flow but with all the memories. The results look similar even with different iterations. (Left column: V_P ; right column: V_S .)

Figure 3. Convergence rates of L-BFGS- and BFGS-based elastic FWI. Plotted are (a) data-misfit and (b) model-misfit comparisons.

Figure 4. From top to bottom: inverted results of BFGS-based FWI with four different preconditioners after 81, 81, 67, 80 iterations. Among them, (d) is closest to the true model. (Left column: V_P ; right column: V_S .)

Figure 5. Convergence rate comparisons of preconditioned BFGS-based FWI, with the pure BFGS-based inversion in black for reference. Plotted in (a) and (b) are data-misfit and model-misfit convergence curves, respectively.

Figure 6. Eigendecomposition of the inverse Hessian from BFGS-based elastic FWI using randomized SVD. Plotted in (a) are the eigenvalues on a logarithmic scale. Plotted in the remainder are the 1st, 5th, 10th, 20th, and the final eigenvectors, respectively. (Left column: V_P ; right column: V_S .)

Figure 7. Eigendecomposition of the inverse Hessian from BFGS & P1 elastic FWI using randomized SVD. Plotted in (a) are the eigenvalues on a logarithmic scale. Plotted in the remainder are the 1st, 5th, 10th, 20th, and the final eigenvectors, respectively. (Left column: V_P ; right column: V_S .)

Figure 8. Eigendecomposition of the inverse Hessian from BFGS & P2 elastic FWI using randomized SVD. Plotted in (a) are the eigenvalues on a logarithmic scale. Plotted in the remainder are the 1st, 5th, 10th, 20th, and the final eigenvectors, respectively. (Left column: V_P ; right column: V_S .)

Figure 9. Eigendecomposition of the inverse Hessian from BFGS & P3 elastic FWI using randomized SVD. Plotted in (a) are the eigenvalues on a logarithmic scale. Plotted in the remainder are the 1st, 5th, 10th, 20th, and the final eigenvectors, respectively. (Left column: V_P ; right column: V_S .)

Figure 10. Eigendecomposition of the inverse Hessian from BFGS & P4 elastic FWI using randomized SVD. Plotted in (a) are the eigenvalues on a logarithmic scale. Plotted in the remainder are the 1st, 5th, 10th, 20th, and the final eigenvectors, respectively. (Left column: V_P ; right column: V_S .)

Figure 11. Uncertainty quantification maps for elastic FWI using non-preconditioned and preconditioned BFGS algorithms, as shown from top to bottom. Please see more details in Appendix C on how to estimate an absolute scaling range for the UQ maps. (Left column: V_P ; right column: V_S .)

Figure 12. Absolute model difference between the true model in Fig. 1a and the inverted model in Fig. 4d for reference. The model differences shown here resemble the uncertainty quantification maps in Fig. 11. (Left column: V_P ; right column: V_S .)

APPENDIX A: DETERMINATION OF THE DIAGONAL HESSIAN

FWI aims to minimize the misfit function

$$f(\mathbf{m}) = \frac{1}{2} \|\mathbf{s}(\mathbf{m}) - \mathbf{d}\|_2^2, \quad (\text{A1})$$

in which $\mathbf{d}$ are the observed data and $\mathbf{s}(\mathbf{m}) = \mathbf{R}\mathbf{u}(\mathbf{m})$ the simulated data, with $\mathbf{u}(\mathbf{m})$ the simulated wavefield from model $\mathbf{m}$ and $\mathbf{R}$ the acquisition sampling operator. Using the Born approximation for a model perturbation $\mathbf{m} = \mathbf{m}_0 + \delta\mathbf{m}$, we have

$$\mathbf{s}(\mathbf{m}_0 + \delta\mathbf{m}) = \mathbf{s}(\mathbf{m}_0) + \mathbf{J}\delta\mathbf{m} + O(\delta\mathbf{m}^2), \quad (\text{A2})$$

with Jacobian matrix $\mathbf{J} = \delta\mathbf{s}/\delta\mathbf{m}$. Taking the first-order derivative of eq. (A1) yields

$$\frac{\partial f}{\partial \mathbf{m}} = \mathbf{J}^\dagger [\mathbf{J}(\mathbf{m}_0 + \delta\mathbf{m}) - \mathbf{d}] = \mathbf{J}^\dagger (\mathbf{J}\delta\mathbf{m} - \delta\mathbf{d}), \quad (\text{A3})$$

with $\delta\mathbf{d} = \mathbf{d} - \mathbf{J}\mathbf{m}_0$ and $\dagger$ the adjoint of an operator. When $\partial f/\partial \mathbf{m} = 0$, we have a direct solution for model perturbation $\delta\mathbf{m}$, namely,

$$\delta\mathbf{m} = (\mathbf{J}^\dagger \mathbf{J})^{-1} \mathbf{J}^\dagger \delta\mathbf{d}. \quad (\text{A4})$$

However, the Hessian $\mathbf{H} = \mathbf{J}^\dagger \mathbf{J}$ is computationally too expensive to calculate, store, and invert. Hence, in practical applications, we simply calculate the model update as

$$\delta\mathbf{m}^* = \mathbf{J}^\dagger \delta\mathbf{d}. \quad (\text{A5})$$

For convenience, let us use a velocity-stress formulation. The 2D velocity-stress isotropic elastic wave equation can be written as (Vigh et al. 2014; Chen & Sacchi 2017)

$$\begin{pmatrix} \rho \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \partial_t \mathbf{u} - \begin{pmatrix} \mathbf{0} & \mathbf{D}^T \\ \mathbf{CD} & \mathbf{0} \end{pmatrix} \mathbf{u} = \mathbf{f}, \quad (\text{A6})$$

where $\mathbf{f}$ is the source term, $\mathbf{I}$ the identity matrix, and

$$\begin{aligned} \mathbf{u} &= (\mathbf{v}, \sigma), \quad \mathbf{v} = (v_x, v_z)^T, \quad \sigma = (\sigma_{xx}, \sigma_{xz}, \sigma_{zz})^T, \\ \mathbf{C} &= \begin{pmatrix} \lambda + 2\mu & \lambda & 0 \\ \lambda & \lambda + 2\mu & 0 \\ 0 & 0 & \mu \end{pmatrix}, \quad \mathbf{D} = \begin{pmatrix} \partial_x & 0 \\ 0 & \partial_z \\ \partial_z & \partial_x \end{pmatrix} \end{aligned} \quad (\text{A7})$$

with $\mathbf{v}$ containing the particle velocities, σ the stress elements, $\mathbf{C}$ the isotropic elastic tensor with λ and μ being the Lamé parameters, $\mathbf{D}$ an operator of spatial derivatives, and initial

condition of the form $\mathbf{u}(\mathbf{x}, t \mid t < 0) = \mathbf{0}$. To avoid clutter, we drop the spatial and temporal dependence $\mathbf{u} = \mathbf{u}(\mathbf{x}, t)$ and the spatial dependence of $\mathbf{C} = \mathbf{C}(\mathbf{x})$. Eq. (A7) is the state equation when FWI runs as an optimal control problem (Tromp et al. 2005; Plessix 2006). For simplification, the abstract form of eq. (A7) reads

$$\mathbf{A}\mathbf{u} = \mathbf{f}, \quad (\text{A8})$$

and its first-order derivative over $\mathbf{m}$ reads

$$\mathbf{A} \frac{\partial \mathbf{u}}{\partial \mathbf{m}} = - \frac{\partial \mathbf{A}}{\partial \mathbf{m}} \mathbf{u}. \quad (\text{A9})$$

Now the Jacobian $\mathbf{J}$ in eq. (A2) can be recast as

$$\mathbf{J} = \frac{\partial \mathbf{s}}{\partial \mathbf{m}} = \mathbf{R} \frac{\partial \mathbf{u}}{\partial \mathbf{m}} = - \mathbf{R} \mathbf{A}^{-1} \frac{\partial \mathbf{A}}{\partial \mathbf{m}} \mathbf{u}, \quad (\text{A10})$$

in which $\mathbf{A}^{-1}$ is the Green's operator given that $\mathbf{u} = \mathbf{A}^{-1}\mathbf{s}$. The gradient from the data residual can be rewritten in detail (Tarantola 1988) as

$$\delta \mathbf{m}^* = \mathbf{J}^\dagger \delta \mathbf{d} = \left(- \mathbf{R} \mathbf{A}^{-1} \frac{\partial \mathbf{A}}{\partial \mathbf{m}} \mathbf{u} \right)^\dagger \delta \mathbf{d} = - \left(\frac{\partial \mathbf{A}}{\partial \mathbf{m}} \mathbf{u} \right)^\dagger (\mathbf{A}^\dagger)^{-1} \mathbf{R}^\dagger \delta \mathbf{d}, \quad (\text{A11})$$

which represents the mapping from data residual $\delta \mathbf{d}$ to the gradient $\mathbf{J}$ via the adjoint operator $\mathbf{J}^\dagger$. By introducing the adjoint-state variable $\tilde{\mathbf{u}}$ from the following equation

$$\mathbf{A}^\dagger \tilde{\mathbf{u}} = \mathbf{R}^\dagger \delta \mathbf{d}, \quad (\text{A12})$$

its corresponding elastic wave equation can be expressed as

$$\begin{pmatrix} \rho \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} (\partial_t)^\dagger \tilde{\mathbf{u}} - \begin{pmatrix} \mathbf{0} & \mathbf{D}^T \\ \mathbf{CD} & \mathbf{0} \end{pmatrix}^\dagger \tilde{\mathbf{u}} = \mathbf{R}^\dagger \delta \mathbf{d}, \quad (\text{A13})$$

where $\mathbf{R}^\dagger \delta \mathbf{d}$ acts as the adjoint source. After the operators of the adjoint † are applied, Eq. (A13) can be recast as

$$\begin{pmatrix} \rho \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} (-\partial_t \tilde{\mathbf{u}}) + \begin{pmatrix} \mathbf{0} & \mathbf{D}^T \mathbf{C} \\ \mathbf{D} & \mathbf{0} \end{pmatrix} \tilde{\mathbf{u}} = \mathbf{R}^\dagger \delta \mathbf{d}, \quad (\text{A14})$$

The term $\delta \sigma$ in $\delta \mathbf{d} = (\delta \mathbf{v}, \delta \sigma)$ remains zero because we only observe $\delta \mathbf{v}$ in practice. The final form of the adjoint-state equation (Vigh et al. 2014) is

$$\begin{pmatrix} \rho \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} (-\partial_t \tilde{\mathbf{u}}) + \begin{pmatrix} \mathbf{0} & \mathbf{D}^T \\ \mathbf{CD} & \mathbf{0} \end{pmatrix} \tilde{\mathbf{u}} = \mathbf{R}^\dagger \begin{pmatrix} \delta \mathbf{v} \\ \mathbf{0} \end{pmatrix}. \quad (\text{A15})$$

Let us go back to the gradient in eq. (A5), which can be further expressed as

$$\delta \mathbf{m}^* = - \left(\frac{\partial \mathbf{A}}{\partial \mathbf{m}} \mathbf{u} \right)^\dagger \tilde{\mathbf{u}}, \quad (\text{A16})$$

or in an explicit form as

$$\delta \mathbf{m}(\mathbf{x})^* = - \int \left(\frac{\partial \mathbf{A}}{\partial \mathbf{m}} \mathbf{u} \right)^\dagger \tilde{\mathbf{u}} dt = - \int \left(\frac{\partial \mathbf{A}}{\partial \mathbf{m}} \mathbf{u} \right) \tilde{\mathbf{u}}^T dt, \quad (\text{A17})$$

with $\tilde{\mathbf{u}}(\mathbf{x}, t)^T = \tilde{\mathbf{u}}(\mathbf{x}, t)^\dagger$ applied. Eq. (A17) is well known in reverse-time migration (Baysal et al. 1983) with the adjoint method.

When computing the gradients, we may have different parametrizations, such as (α, β, ρ) or (λ, μ, ρ) , with α, β, ρ being the P- and S-wave wavespeeds, and densities, λ, μ being Lamé parameters, respectively. We express the gradients with respect to (α, β) as follows

$$\begin{aligned} \delta \alpha^* &= \int \left(\frac{\partial \mathbf{C}}{\partial \alpha} \mathbf{D} \mathbf{v} \right) \cdot \tilde{\sigma} dt, \\ \delta \beta^* &= \int \left(\frac{\partial \mathbf{C}}{\partial \beta} \mathbf{D} \mathbf{v} \right) \cdot \tilde{\sigma} dt. \end{aligned} \quad (\text{A18})$$

Considering eq. (A4), we aim to accelerate FWI with the Hessian in a Gauss-Newton approximation, which has the form

$$\mathbf{H} = \mathbf{J}^\dagger \mathbf{J} = \left(\mathbf{R} \mathbf{A}^{-1} \frac{\partial \mathbf{A}}{\partial \mathbf{m}} \mathbf{u} \right)^\dagger \left(\mathbf{R} \mathbf{A}^{-1} \frac{\partial \mathbf{A}}{\partial \mathbf{m}} \mathbf{u} \right), \quad (\text{A19})$$

in which $\mathbf{A}^{-1}$ resembles the Green's operator and $\mathbf{u}$ the source wavefield. Regardless of the band-limited source wavelet, $\mathbf{u}$ resembles the source-side Green's functions $\mathbf{G}_S$, and at the same time, $\mathbf{R} \mathbf{A}^{-1}$ the receiver-side Green's functions $\mathbf{G}_R$.

The full computation and storage of $\mathbf{H}$ are prohibitive. Therefore, Shin et al. (2008) propose to forget $\mathbf{G}_R$ to save computational cost, and only compute the Hessian diagonals via zero-lag cross-correlation of the source wavefields as follows:

$$\mathbf{H}(\mathbf{x}, \mathbf{x}) = \int \left[\frac{\partial \mathbf{C}}{\partial m_i} \mathbf{D} \mathbf{v}(\mathbf{x}) \right] \cdot \left[\frac{\partial \mathbf{C}}{\partial m_i} \mathbf{D} \mathbf{v}(\mathbf{x}) \right] dt. \quad (\text{A20})$$

When it comes to the parametrization (α, β) , we have

$$\begin{aligned} \mathbf{H}_{\alpha\alpha} &= \int \left(\frac{\partial \mathbf{C}}{\partial \alpha} \mathbf{D} \mathbf{v} \right) \cdot \left(\frac{\partial \mathbf{C}}{\partial \alpha} \mathbf{D} \mathbf{v} \right) dt = 8\rho^2 \alpha^2 \int (\partial_x v_x + \partial_z v_z)^2 dt, \\ \mathbf{H}_{\beta\beta} &= \int \left(\frac{\partial \mathbf{C}}{\partial \beta} \mathbf{D} \mathbf{v} \right) \cdot \left(\frac{\partial \mathbf{C}}{\partial \beta} \mathbf{D} \mathbf{v} \right) dt = 16\rho^2 \beta^2 \int [(\partial_x v_x)^2 + (\partial_z v_z)^2] dt + 4\rho^2 \beta^2 \int (\partial_x v_z + \partial_z v_x)^2 dt \\ \mathbf{H}_{\rho\rho} &= \int (\partial_t \mathbf{v}) \cdot (\partial_t \mathbf{v}) dt = \int [(\partial_t v_x)^2 + (\partial_t v_z)^2] dt, \end{aligned} \quad (\text{A21})$$

with $\mathbf{v}$ being the source wavefield in particle velocity. The third equation in Eq. (A21) corre-

sponds to a generalized “ray density” Hessian kernel (Luo 2012) regarding ρ in the parametrization (λ, μ, ρ) . Eq. (A21), however, does not involve the receiver-side Green’s functions. Following Luo (2012) and Modrak & Tromp (2016), we modify eq. (A21) to include $\mathbf{G}_R$ as

$$\begin{aligned}\bar{H}_{\alpha\alpha} &= \left| \int \left(\frac{\partial \mathbf{C}}{\partial \alpha} \mathbf{D} \mathbf{v} \right) \cdot \left(\frac{\partial \mathbf{C}}{\partial \alpha} \mathbf{D} \bar{\mathbf{v}} \right) dt \right| = 8\rho^2 \alpha^2 \left| \int (\partial_x v_x + \partial_z v_z) (\partial_x \bar{v}_x + \partial_z \bar{v}_z) dt \right|, \\ \bar{H}_{\beta\beta} &= \left| \int \left(\frac{\partial \mathbf{C}}{\partial \beta} \mathbf{D} \mathbf{v} \right) \cdot \left(\frac{\partial \mathbf{C}}{\partial \beta} \mathbf{D} \bar{\mathbf{v}} \right) dt \right| \\ &= 16\rho^2 \beta^2 \left| \int (\partial_x v_x \partial_x \bar{v}_x + \partial_z v_z \partial_z \bar{v}_z) dt \right| + 4\rho^2 \beta^2 \left| \int (\partial_x v_z + \partial_z v_x) (\partial_x \bar{v}_z + \partial_z \bar{v}_x) dt \right|, \\ \bar{H}_{\rho\rho} &= \left| \int (\partial_t \mathbf{v}) \cdot (\partial_t \bar{\mathbf{v}}) dt \right| = \left| \int (\partial_t v_x \partial_t \bar{v}_x + \partial_t v_z \partial_t \bar{v}_z) dt \right|,\end{aligned}\tag{A22}$$

with $\bar{\mathbf{v}}$ denoting the receiver wavefield in particle velocity. Note that although both wavefields originate at the receivers, $\bar{\mathbf{v}}$ differs from $\tilde{\mathbf{v}}$ in that the former indicates the wavefield due to the simulated data while the latter the wavefield due to the data residual. We take their absolute values to ensure positive definiteness of the initial Hessian.

APPENDIX B: VALIDATION WITH THE ELASTIC OVERTHRUST MODEL

This section is to validate the aforementioned preconditioned-BFGS based elastic FWI and uncertainty quantification workflow using the elastic Overthrust model, which can be found in Liu et al. (2019). The simulations are with sources and receivers 10-m deep and using absorbing boundaries (Stacey 1988). For simplicity, we only use preconditioner P4. Shown in Fig. B1 are the true, initial, and inverted elastic Overthrust models. Also shown are the associated absolute model differences between the true and inverted model together with the corresponding uncertainty maps. We observe that Fig. B1d resembles Fig. B1e.

Figure B1. From top to bottom: (a) true and (b) initial and (c) inverted elastic Overthrust models, (d) the absolute model differences between the true and the inverted, (e) the uncertainty maps estimated from the preconditioned-BFGS FWI workflow with P4. (Left column: V_P ; right column: V_S .)

Figure C1. Data misfits with respect to the UQ-map-based model perturbations. The UQ map here refers to that in Fig. 11e, from which m_{pert} is scaled to have a maximum of 250 m/s . The fact that the data-misfit curve remains flat within the range of $[-m_{pert}, m_{pert}]$ indicates that we can roughly estimate the scaling of the UQ map using the posterior sampling methods stated in Appendix C.

APPENDIX C: THE SCALING OF THE UNCERTAINTY MAPS

The scaling of the UQ maps in Fig. 11, which mainly reflects the square-root diagonals of the $\mathbf{H}^{-1}$ term in Eqs. 12 and 13, ranges between $[0,1]$ because all the gradients have been rescaled with a constant factor to a physically reasonable range. This factor comes from a separate trial once and for all per model, in which we determine the constant by performing a parabolic line search with the 1st unpreconditioned, smoothed gradient. We can, however, choose to scale for the UQ maps like the way in the SRVM-based null-space shuttle (Liu & Peter 2020). The data misfits with respect to the UQ-map-based perturbations are shown in Fig. C1, in which one m_{pert} is scaled from Fig. 11e to be with a maximum of 250 m/s .