INVERSIONSON: FULLY AUTOMATED SEISMIC WAVEFORM INVERSIONS

Solvi Thrastarson[*]
ETH Zürich
Department of Earth Sciences
Institute of Geophysics
soelvi.thrastarson@erdw.ethz.ch

Dirk-Philip van Herwaarden
ETH Zürich
Department of Earth Sciences
Institute of Geophysics
dirkphilip.vanherwaarden@erdw.ethz.ch

Andreas Fichtner
ETH Zürich
Department of Earth Sciences
Institute of Geophysics
andreas.fichtner@erdw.ethz.ch

March 3, 2021

ABSTRACT

We present *Inversionson*, a Python package that fully automates modern full-waveform inversions (FWI). It supports traditional FWI, which uses the same set of events and a single simulation mesh in each iteration, as well as more advanced workflows that exploit the use of dynamic mini-batches and wavefield-adapted meshes. These recently introduced advancements can be time-consuming and challenging to implement, but offer major benefits, such as much lower compute cost in favourable cases and the ability to use much larger datasets, while having built-in optimal experimental design. Inversionson, in essence, acts as a control center that steers the functions of several other packages, making them communicate and work together to perform seismic inversions. This enables scaling of full-waveform inversion problems well on the most powerful High Performance Computing (HPC) clusters, while requiring no user input during run time. Widespread adoption of real data full-waveform inversion workflows to image regional to global 3-D Earth structure has so far been limited to a handful of research groups, with most researchers focusing on 2-D synthetic problems. This probably mainly stems from the challenging implementation, as well as a lack of computational and human resources. Inversionson is built on top of various tools and mitigates each of these challenges, by exploiting the widely used waveform solver package Salvus, supporting efficient modern workflows, and by fully automating the time-consuming steps. With this contribution, we hope that we can make FWI a more interesting option for a much larger group of researchers, thereby accelerating progress in the field.

Keywords Seismology · Software · Inverse Theory

1 Introduction

Over the past half a century, a number of three-dimensional seismic Earth models have been constructed with a process called seismic tomography. While the initial (pioneering) applications were based on relatively little data and limited (ray-based) theories [Aki and Lee, 1976, Aki et al., 1977, Dziewonski et al., 1977], advances in computational power and an increase in seismic data availability, have led to substantial progress in the field.

Through time, new methods of seismic inversion have been developed, including finite-frequency methods like normal-mode inversions [e.g Woodhouse and Dziewonski, 1984, Giardini et al., 1987] or Fresnel zone inversions [e.g Yomogida, 1992, Dahlen et al., 2000].

In the past decade, using the seismic waveforms directly to infer the structure of the Earth has become a more feasible technique. Inverting for the complete waveform, full-waveform inversion (FWI) [Bamberger et al.,

---

[*]Corresponding author: `soelvi.thrastarson@erdw.ethz.ch`

1977, 1982, Lailly, 1983, Tarantola, 1984], requires the numerical solution of the wave equation through complex, heterogeneous media.

Initial applications of FWI were on an exploration scale [e.g. Igel et al., 1996, Pratt, 1999]. The first applications on a regional scale came a decade later [e.g. Chen et al., 2007a,b, Fichtner et al., 2009, Tape et al., 2009], all done with manual workflows using relatively small data volumes. As data volumes used in an inversion increase, the process becomes increasingly fragile and tedious, which is the reason why manual workflows do not scale to larger inversions. In order to allow for increasing data volumes, creating software to make the inversions scalable was essential.

Not all seismic data is usable and finding the windows of usable data is a crucial part of performing an FWI study. Having thousands of waveforms per source in a dataset makes window selection a bottleneck in the scaling of FWI. FLEXWIN [Maggi et al., 2009] is an automatic window selection algorithm which opens up that bottleneck, greatly reducing human workload in FWI workflows. LASIF [Krischer et al., 2015a] provides a good framework for setting up and organizing a local to global 3-D FWI study and taking care of all the time consuming tasks, exploiting established ObsPy [Beyreuther et al., 2010, Krischer et al., 2015b], NumPy [Harris et al., 2020] and SciPy [Virtanen et al., 2020] functionalities. The Adaptable Seismic Data Format (ASDF) [Krischer et al., 2016] was developed to better handle large data volumes by merging seismic data into one file using the HDF5 format, greatly reducing the I/O (input/output) overhead of inversions. Thrastarson et al. [2021b] updated LASIF to natively support ASDF and the Salvus waveform modelling tools [Afanasiev et al., 2019] to make it more scalable. PySIT [Hewett and Demanet, 2013] is a similar tool to LASIF which has less of an emphasis on the large-scale problems, and provides optimization routines internally. SimPEG [Cockett et al., 2015] provides a framework for geophysical inversions, not limited to seismology, including a simulation tool. SeisFlows [Modrak et al., 2018] implements inversion workflows which include communication with HPC systems and is designed with large-scale inversions in mind. Pyatoa [Chow et al., 2020] is developed to perform similar tasks to LASIF and in combination with SeisFlows it has the capability of automating the standard FWI workflow.

Inversionson and LASIF together have a similar purpose to SeisFlows and Pyatoa, except that Inversionson includes more advanced workflows [e.g. Thrastarson et al., 2020, van Herwaarden et al., 2020], and it works in combination with different FWI software packages. The workflows incorporated into Inversionson, utilise stochastic optimization schemes as well as a more efficient discretisation of the wave equation. In this paper we will begin by introducing FWI in section 2, with a special emphasis on quasi-Newton optimization routines (section 2.2). The workflows currently implemented into Inversionson will be introduced in section 3, the details regarding the software implementations will be presented in section 4 and finally a synthetic project demonstrating the capabilities of the software will be displayed in section 5.

2 Full-Waveform Inversion

FWI is a Partial-Differential-Equation (PDE) constrained optimization problem with the PDE being a wave-equation. PDE constrained optimizations are often computationally demanding as they require many solutions of, often expensive to solve, PDEs. FWI requires multiple solutions of the wave equation, which grow more costly with either increasing modelled frequencies or domain size.

In this section, the essence of FWI will be explained along with the required theoretical background to understand the workflows which Inversionson offers.

2.1 Computing a Gradient

With FWI being an iterative method (although probabilistic approaches are emerging [e.g. Gebraad et al., 2020]), the aim of each iteration is to compute how to improve the model, which is inferred from the data, for the next iteration. With potentially millions of model parameters, this becomes a non-trivial task, and is done by computing a gradient of the waveform misfit with respect to the model parameters.

We can write the governing equation of FWI, the discrete wave-equation, in a simplified manner,

$$\mathbf{L}(\mathbf{m})\mathbf{u} = \mathbf{f}, \tag{1}$$

where $\mathbf{L}$ is the wave equation operator, $\mathbf{m}$ is a vector of model parameters (size M), $\mathbf{u}$ is the (forward) wavefield vector, and $\mathbf{f}$ is a vector representing the source term (usually an earthquake). From now on, bold face upper-case characters represent matrices, bold face lower-case characters represent vectors, subscript i represents an index of a vector, and subscript n denotes the iteration. A continuous wave equation can, for example, be discretised using the spectral-element

method [Faccioli et al., 1996]. In an iterative optimization scheme we need to define an objective function, which we iteratively try to minimize. In FWI, the objective function is defined as a waveform misfit which measures the difference between synthetically computed waveforms (i.e. the results of solving Eq. (1)) and seismograms recorded by seismometers or geophones. The choice of waveform misfit is an important one, and there are multiple options available [e.g. Tromp et al., 2005, Fichtner et al., 2008, Bozdağ et al., 2011]. We can define an objective function $\chi[\mathbf{u}(\mathbf{m})]$, where χ can represent any waveform misfit function. As previously mentioned, each iteration requires computing the gradient of χ with respect to the model parameters $\mathbf{m}$, which, using the chain rule, can be written as

$$\frac{\partial \chi}{\partial m_i} = \frac{\partial \chi}{\partial \mathbf{u}}^T \frac{\partial \mathbf{u}}{\partial m_i}. \tag{2}$$

The problem which becomes apparent in Eq. (2) is that the computation of $\frac{\partial \mathbf{u}}{\partial m_i}$ would require solving Eq. (1) once for each model parameter m_i, with $i = 1, 2, \ldots, M$, in order to compute a finite-difference type gradient. This becomes practically impossible quickly as the number of model parameters increases, making any realistic FWI problem unsolvable in practice.

Fortunately there is a clever solution to this problem, called the adjoint-state method [Tarantola, 1984, 1988, Tromp et al., 2005, Fichtner et al., 2006], which enables the solution of Eq. (2) at the cost of one additional waveform simulation. We will derive the discrete-adjoint formulation introduced to seismology in the frequency domain by Pratt et al. [1998] and Pratt [1999] and applied in the time domain by van Driel et al. [2020] and Thrastarson et al. [2020]. The discrete-adjoint formulation is a necessary component of FWI to enable the multi-mesh workflow, which will be introduced in section 3.3.

A misfit function χ can be minimized by simply making the synthetic waveforms equal to the recorded waveforms. The problem is however not that simple, as the construction of the synthetic waveforms has to adhere to the wave equation (Eq. (1)). We can construct an augmented objective function ψ using Lagrange multipliers to create a constrained optimization problem. The function we wish to mimimize is

$$\psi(\mathbf{m}, \mathbf{u}, \boldsymbol{\lambda}) = \chi(\mathbf{m}, \mathbf{u}) - \boldsymbol{\lambda}^T (\mathbf{L}\mathbf{u} - \mathbf{f}), \tag{3}$$

where $\boldsymbol{\lambda}$ is the Lagrange multiplier. The augmented objective function, ψ, is minimized by forcing all partial derivatives to zero. We will go through each of the partial derivatives one by one. Firstly, the partial derivative of Eq. (3) with respect to $\mathbf{m}$ is,

$$\frac{\partial \psi}{\partial m_i} = \frac{\partial \chi}{\partial m_i} - \boldsymbol{\lambda}^T \frac{\partial \mathbf{L}}{\partial m_i} \mathbf{u} - \boldsymbol{\lambda}^T \mathbf{L} \frac{\partial \mathbf{u}}{\partial m_i} \stackrel{!}{=} 0, \tag{4}$$

which can be rewritten using the chain rule,

$$\frac{\partial \psi}{\partial m_i} = \frac{\partial \chi}{\partial \mathbf{u}}^T \frac{\partial \mathbf{u}}{\partial m_i} - \boldsymbol{\lambda}^T \mathbf{L} \frac{\partial \mathbf{u}}{\partial m_i} - \boldsymbol{\lambda}^T \frac{\partial \mathbf{L}}{\partial m_i} \mathbf{u}$$
$$= \left(\frac{\partial \chi}{\partial \mathbf{u}} - (\boldsymbol{\lambda}^T \mathbf{L})^T \right)^T \frac{\partial \mathbf{u}}{\partial m_i} - \boldsymbol{\lambda}^T \frac{\partial \mathbf{L}}{\partial m_i} \mathbf{u}. \tag{5}$$

The second condition of the Lagrangian optimization is the partial derivative of ψ with respect to the synthetic wavefield $\mathbf{u}$,

$$\frac{\partial \psi}{\partial \mathbf{u}} = \frac{\partial \chi}{\partial \mathbf{u}} - (\boldsymbol{\lambda}^T \mathbf{L})^T \stackrel{!}{=} \mathbf{0}, \tag{6}$$

which forces the first term of Eq. (5) to zero, removing the, previously mentioned, problematic term from Eq. (2), $\frac{\partial \mathbf{u}}{\partial m_i}$. Eq. (6) is called the discrete adjoint equation, and it can be written as,

$$\mathbf{L}^T \boldsymbol{\lambda} = \frac{\partial \chi}{\partial \mathbf{u}}, \tag{7}$$

where $\boldsymbol{\lambda}$ is called the adjoint field, and it can be computed at the cost of a single waveform simulation. The final constraint of the Lagrangian optimization is,

$$\frac{\partial \psi}{\partial \boldsymbol{\lambda}} = \mathbf{L}\mathbf{u} - \mathbf{f} \stackrel{!}{=} \mathbf{0} \tag{8}$$
$$\implies \mathbf{L}\mathbf{u} = \mathbf{f}, \tag{9}$$

which ensures that the wave equation (Eq. (1)) holds.
By substituting Eq. (9) into Eq. (3), we can see that, at the stationary points of Eq. (3), where Eq. (9) holds, $\psi = \chi$. The

three partial derivatives of the augmented objective function, ψ, thus lead to three equations, which need to be fulfilled in order to compute a gradient.

$$\frac{\partial \chi}{\partial m_i} = -\boldsymbol{\lambda}^T \frac{\partial \mathbf{L}}{\partial m_i} \mathbf{u}, \qquad \text{(Gradient)}$$

$$\mathbf{L}^T \boldsymbol{\lambda} = \frac{\partial \chi}{\partial \mathbf{u}}^T, \qquad \text{(Adjoint equation)} \qquad (10)$$

$$\mathbf{L}\mathbf{u} = \mathbf{f}, \qquad \text{(Wave equation)}.$$

By solving Eqs. (10) together, one can obtain the gradient of χ with respect to $\mathbf{m}$ at the cost of two wave propagation simulations and a convolution of the adjoint field and the synthetic wavefield.

2.2 Optimization

As previously mentioned, the gradient is used to compute a model update which results in a misfit reduction between iterations. The process of iteratively improving the model, or to be exact, reducing the waveform misfit is called optimization.

There are multiple optimization algorithms available, and it has been an active research topic for decades. Different optimization algorithms work well for different applications, and there is no way to cover them all. For a good overview, we recommend Nocedal and Wright [2006]. In this section we will give an introduction to quasi-Newton methods of optimization as that is the family of optimization schemes, Inversionson was designed for.

2.2.1 Gradient-Based Optimization

Optimization problems like FWI are designed with the goal of minimizing functions. The function to be minimized is called an objective function (χ), and what makes the minimization a tedious task is that there is no explicit expression for this, generally highly non-quadratic, function. In gradient-based optimization one starts with an initial guess of model parameters $\mathbf{m}_0$, evaluates the objective function $\chi(\mathbf{m}_0)$, and iteratively improves the guess with the goal of lowering the value resulting from the evaluation of $\chi(\mathbf{m})$ until data are fit within the noise. The objective function can also be maximized in a similar way [e.g. Bösch et al., 2020].

As the objective function is evaluated, a decision needs to be made of how to modify the current model to reduce the value of the objective function. In order to do so, an approximation of the objective function in the local vicinity of the current model is computed using a Taylor expansion. The simplest of which being a first-order (linear) Taylor expansion

$$\chi(\mathbf{m}) \approx \tilde{\chi}(\mathbf{m}) = \chi(\mathbf{m}_n) + \nabla\chi(\mathbf{m})|_{\mathbf{m}=\mathbf{m}_n}^T (\mathbf{m} - \mathbf{m}_n), \qquad (11)$$

where $\mathbf{m}_n$ is the current best guess in model space. Eq. (11) is a linear approximation of χ, and as χ is seldomly linear in FWI problems, the first-order Taylor expansion quickly becomes inaccurate as we move away from $\mathbf{m}_n$. However, in the local vicinity of $\mathbf{m}_n$, the information contained in Eq. (11) can still be used to improve the model. The gradient of $\tilde{\chi}(\mathbf{m})$ with respect to model parameters is $\nabla\chi(\mathbf{m})|_{\mathbf{m}=\mathbf{m}_n}$ and as the gradient of a function points in the direction of its steepest ascent, a good approach to improve the current best guess would thus be to move in the opposite direction of the gradient,

$$\mathbf{m}_{n+1} = \mathbf{m}_n - \alpha_n \nabla\chi(\mathbf{m})|_{\mathbf{m}=\mathbf{m}_n}. \qquad (12)$$

As Eq. (11) is just a local approximation of χ, it depends on how good the local approximation is, how far it is wise to step in the local negative gradient direction. The update is thus scaled with a step-length parameter, α_n, which depends on how well Eq. (11) approximates χ. A good value for α_n can be estimated with a line search, where different values of α_n are tested in order to search for the optimal one, or with a trust-region approach which will be describe in section 2.2.4.

Eq. (12) is the update formula for one of the most commonly used optimization algorithms, steepest descent, where the model is updated in the opposite direction of local steepest ascent every iteration.

An option to potentially improve upon steepest descent, is to improve the local approximation of χ by higher order Taylor expansions. Let us expand $\tilde{\chi}$ to the second order,

$$\chi(\mathbf{m}) \approx \tilde{\chi}(\mathbf{m}) = \chi(\mathbf{m}_n) + \nabla\chi(\mathbf{m})|_{\mathbf{m}=\mathbf{m}_n}^T (\mathbf{m} - \mathbf{m}_n) + \frac{1}{2}(\mathbf{m} - \mathbf{m}_n)^T \mathbf{B}_n (\mathbf{m} - \mathbf{m}_n), \qquad (13)$$

where $\mathbf{B}_n$ is the local Hessian (or 2nd derivative) of $\chi(\mathbf{m})$ evaluated at $\mathbf{m}_n$. The function $\tilde{\chi}(\mathbf{m})$ can be minimized by

$$\mathbf{m} - \mathbf{m}_n = -\mathbf{B}_n^{-1}\nabla\chi(\mathbf{m})|_{\mathbf{m}=\mathbf{m}_n}, \tag{14}$$

which shows that the modification of $\mathbf{m}_n$ needed to minimize the local approximate objective function is the negative inverse local Hessian times the local gradient. It is thus a good search direction, giving rise to the 2nd-order update function,

$$\mathbf{m}_{n+1} = \mathbf{m}_n - \mathbf{H}_n\nabla\chi(\mathbf{m})|_{\mathbf{m}=\mathbf{m}_n}, \tag{15}$$

where $\mathbf{H}_n = \mathbf{B}_n^{-1}$ is the inverse Hessian matrix. If χ is a quadratic function, Eq. (15) minimizes χ with one update. Most realistic objective functions are not perfectly quadratic though, often making it necessary to scale the update of $\mathbf{m}$ as the one step minimization of Eq. (15) only holds locally for $\tilde{\chi}$ and not globally for χ. The scaled version of Eq. (15) can be written as

$$\mathbf{m}_{n+1} = \mathbf{m}_n - \beta_n\mathbf{H}_n\nabla\chi(\mathbf{m})|_{\mathbf{m}=\mathbf{m}_n}, \tag{16}$$

where β_n represents a step length scalar which depends on the quality of the local approximation of χ. Estimating a good value for β_n is discussed in section 2.2.4. Optimization methods using Eq. (16) are called Newton's methods. Unfortunately, it is often the case that 2nd-order derivatives are expensive to compute and/or have prohibitively large memory requirements. This happens when the model space is large, as typically in FWI. One cure for that problem are quasi-Newton methods which approximate the Hessian or the inverse Hessian. In the following sections we will introduce a quasi-Newton method which works well for FWI problems.

2.2.2 BFGS

The Broyden-Fletcher-Goldfarb-Shanno (BFGS)[Broyden, 1970, Fletcher, 1970, Goldfarb, 1970, Shanno, 1970] method is a quasi-Newton method of optimization, meaning that it approximates the second derivative of the misfit, the Hessian, in order to make better updates than first-order methods, at a lower cost than Newton methods.

The Hessian gives information on the local curvature of the objective function which in turn can give valuable information regarding optimal (component-wise) scaling of the gradient. If curvature is high, taking a long step is risky, as the gradient changes rapidly, while with a low curvature, taking a long step is safer. The optimal gradient scaling is thus inversely proportional to the Hessian. Unfortunately, as previously stated, the Hessian is expensive to compute and has intensive memory requirements for large model spaces. The BFGS method computes an approximation of the inverse Hessian through monitoring how the gradient changes as the model space is explored.

If we define vectors $\mathbf{s}_n = \mathbf{m}_{n+1} - \mathbf{m}_n$ (change in model) and $\mathbf{y}_n = \nabla\chi_{n+1} - \nabla\chi_n$ (change in gradient) we can write the secant equation,

$$\mathbf{B}_{n+1}\mathbf{s}_n = \mathbf{y}_n, \tag{17}$$

stating that the curvature of χ maps changes in model space to changes in the gradient of χ. For the inverse Hessian we have

$$\mathbf{H}_{n+1}\mathbf{y}_n = \mathbf{s}_n. \tag{18}$$

The approximate inverse Hessian and the gradient can be used to compute an update in model space for iteration $n+1$,

$$\mathbf{m}_{n+1} = \mathbf{m}_n - \beta_n\tilde{\mathbf{H}}_n\nabla\chi_n, \tag{19}$$

where β_n is a step length which should be computed to satisfy the Wolfe condition [Wolfe, 1969] and $\tilde{\mathbf{H}}_n$ is a positive definite matrix satisfying Eq. (18). The approximation of the inverse Hessian is updated every iteration. When updating $\tilde{\mathbf{H}}_n$, there are still infinitely many positive definite solutions which satisfy Eq. (18). It has shown to be a good choice to define the updated inverse Hessian, $\tilde{\mathbf{H}}_{n+1}$, as the solution to the minimization problem

$$\mathbf{H}_{n+1} \approx \tilde{\mathbf{H}}_{n+1} = \underset{\mathbf{H}}{\mathrm{argmin}}\,\|\mathbf{H} - \tilde{\mathbf{H}}_n\|$$
$$\text{subject to} \quad \mathbf{H} = \mathbf{H}^T \quad \text{and} \quad \mathbf{H}\mathbf{y}_n = \mathbf{s}_n \tag{20}$$

where the norm is a weighted Frobenius norm. The solution to Eq. (20) is the BFGS formula for updating the inverse Hessian,

$$\tilde{\mathbf{H}}_{n+1} = \mathbf{V}_n^T\tilde{\mathbf{H}}_n\mathbf{V}_n + \rho_n\mathbf{s}_n\mathbf{s}_n^T, \tag{21}$$

where

$$\mathbf{V}_n = \mathbf{I} - \rho_n\mathbf{y}_n\mathbf{s}_n^T, \tag{22}$$

and

$$\rho_n = \frac{1}{\mathbf{y}_n^T\mathbf{s}_n}. \tag{23}$$

5

Doing the same sort of derivation for the approximate Hessian results in a formula called the Davidon-Fletcher-Powell updating formula (DFP),

$$\tilde{\mathbf{B}}_{n+1} = \mathbf{V}_n \tilde{\mathbf{B}}_n \mathbf{V}_n^T + \rho_n \mathbf{y}_n \mathbf{y}_n^T. \tag{24}$$

For a more extensive derivation of the BFGS approach we refer to Nocedal and Wright [2006].

2.2.3 L-BFGS

The BFGS method becomes hard to use in practice once the model space gets large. If the model space contains M parameters, the Hessian requires $(M^2 - M)/2$ indices, which is too memory intensive for large-scale problems. Nocedal [1980] introduced a variant of the BFGS method, called the limited memory BFGS or L-BFGS method. It does not require storing the inverse Hessian in memory but rather computes its approximation from series of previous gradients and model updates on the fly. The BFGS update of the inverse Hessian is described in Eq. (21), but the previous update of the inverse Hessian is in the first term and thus needs to be stored in memory. This is what L-BFGS is designed to avoid.

In L-BFGS, a history of x previous gradients and model updates is used to approximate the inverse Hessian at every iteration.

$$
\begin{aligned}
\tilde{\mathbf{H}}_n = \quad & (\mathbf{V}_{n-1}^T \cdots \mathbf{V}_{n-x}^T) \tilde{\mathbf{H}}_n^0 (\mathbf{V}_{n-x} \cdots \mathbf{V}_{n-1}) \\
& + \rho_{n-x} (\mathbf{V}_{n-1}^T \cdots \mathbf{V}_{n-x+1}^T) \mathbf{s}_{n-x} \mathbf{s}_{x-n}^T (\mathbf{V}_{n-x+1} \cdots \mathbf{V}_{n-1}) \\
& + \rho_{n-x+1} (\mathbf{V}_{n-1}^T \cdots \mathbf{V}_{n-x+2}^T) \mathbf{s}_{n-x+1} \mathbf{s}_{x-n+1}^T (\mathbf{V}_{n-x+2} \cdots \mathbf{V}_{n-1}) \\
& + \cdots \\
& + \rho_{n-1} \mathbf{s}_{n-1} \mathbf{s}_{n-1}^T
\end{aligned}
\tag{25}
$$

One has to choose an initial inverse Hessian approximation $\tilde{\mathbf{H}}_n^0$ which can vary as the optimization goes on. Nocedal and Wright [2006] show an efficient two loop recursive algorithm to compute the descent direction from Eq. (14) which is then put into Eq. (19), and β_n can be found with line search satisfying the Wolfe condition. Performing a line-search is expensive when each realization of the objective function is costly, which makes it attractive to look at trust-region methods to find a reasonable step length.

2.2.4 Trust-Region Methods

Gradient based optimization methods can be separated into two distinct groups, line-search methods and trust-region methods. The fundamental difference between the groups is the philosophy of how to find the optimal step length for model updates. Line-search methods optimize the step length based on a pre-computed descent direction while the trust-region methods define a region where the objective function can confidently be approximated, and optimize the update within that region.

Line-search is done by evaluating χ at a few neighbouring points in model space and picking a step length which satisfies the Wolfe conditions [Wolfe, 1969]. In FWI, evaluating χ is a costly exercise which would ideally be kept to a minimum and trust-region methods provide a cheaper alternative.

Often a simple Taylor expansion around the current best guess is used to approximate χ, either first order, or second order where Hessian information is available. In each step, the optimization is done for the approximate objective function $\tilde{\chi}$

$$\tilde{\chi}(\mathbf{m}) = \chi(\mathbf{m}_n) + \nabla \chi(\mathbf{m})|_{\mathbf{m}=\mathbf{m}_n}^T (\mathbf{m} - \mathbf{m}_n) + \frac{1}{2} (\mathbf{m} - \mathbf{m}_n)^T \mathbf{B}_n (\mathbf{m} - \mathbf{m}_n). \tag{26}$$

The local approximation, $\tilde{\chi}(\mathbf{m})$, is minimized subject to a condition regarding the Euclidean distance between $\mathbf{m}$ and $\mathbf{m}_n$

$$||\mathbf{m}_n - \mathbf{m}||_2 \leq \delta, \tag{27}$$

where δ is called the trust-region radius.

In each iteration, the predicted value of the new $\tilde{\chi}(\mathbf{m}_{n+1})$ is compared to the actual value of $\chi(\mathbf{m}_{n+1})$ and if the ratio is close to one, that gives an indication that $\tilde{\chi}(\mathbf{m})$ was a good local approximation of $\chi(\mathbf{m})$, giving a reason to try and increase the trust-region radius in the next iteration. If however the ratio is far from 1, the trust-region should be shrunk. Additionally if the misfit does not decrease between iterations, δ should be shrunk and the previous model update tried again.

2.2.5 Trust-Region L-BFGS

As an observant reader may have noticed, the Hessian matrix is used in the Taylor expansion in Eq. (26), which is not tractable to compute for large-scale problems. One way to mitigate this issue is to use a first-order Taylor expansion, limiting the quality of the approximation which in turn limits the feasible trust-region radii. Another way to mitigate it is to use an approximate Hessian like the one computed in Eq. (24), or to go even further and use the trick of the L-BFGS approach where the approximate inverse Hessian is computed based on a series of updates of the model and its corresponding gradients. The trick of the L-BFGS inverse Hessian approximation can also be applied to the Hessian approximation, and as this computation is rather cheap, it is well worth the additional accuracy it provides in approximating the objective function within the trust-region. Eq. (24) can be computed with the limited memory approach just like Eq. (21) is computed using Eq. (25). The limited memory approximate Hessian update is

$$
\begin{aligned}
\tilde{\mathbf{B}}_n = \ & (\mathbf{V}_{n-1}\cdots\mathbf{V}_{n-x})\tilde{\mathbf{B}}_n^0(\mathbf{V}_{n-x}^T\cdots\mathbf{V}_{n-1}^T) \\
& + \rho_{n-x}(\mathbf{V}_{n-1}\cdots\mathbf{V}_{n-x+1})\mathbf{y}_{n-x}\mathbf{y}_{x-n}^T(\mathbf{V}_{n-x+1}^T\cdots\mathbf{V}_{n-1}^T) \\
& + \rho_{n-x+1}(\mathbf{V}_{n-1}\cdots\mathbf{V}_{n-x+2})\mathbf{y}_{n-x+1}\mathbf{y}_{x-n+1}^T(\mathbf{V}_{n-x+2}^T\cdots\mathbf{V}_{n-1}^T) \\
& + \cdots \\
& + \rho_{n-1}\mathbf{y}_{n-1}\mathbf{y}_{n-1}^T,
\end{aligned}
\tag{28}
$$

and it can also be computed in an efficient two loop recursive algorithm like the descent direction in L-BFGS.

By combining L-BFGS with a trust-region approach, one has a good approximation of the local objective function while a line-search for a good step-length is not necessary anymore, significantly reducing the cost of an FWI update.

2.3 Stochastic FWI

Inversionson has the option of using the standard non-stochastic FWI workflow. On top of that, it offers the usage of the dynamic mini-batch workflow introduced by van Herwaarden et al. [2020]. Here we will give a brief overview of the difference between a standard optimization scheme and a stochastic one. For a detailed description of the stochastic FWI we refer to van Herwaarden et al. [2020].

In a standard FWI, the objective function to be minimized is the sum of the waveform misfits,

$$
\theta(\mathbf{m}) = \sum_{i=1}^{N}\theta_i(\mathbf{m}),
\tag{29}
$$

where N is the number of sources in the entire dataset. The value of this objective function needs to be computed in each iteration and monitored in a way that the misfit keeps decreasing. Another way to define an objective function is to use an approximate misfit function of the full dataset. Rather than using the whole dataset to compute the value of the objective function, the dataset is approximated with a subset of it (a batch), and the sample average of the waveform misfit becomes the new objective function,

$$
\chi(\mathbf{m}) \approx \chi_b(\mathbf{m}) = \frac{1}{B}\sum_{i=1}^{B}\chi_i(\mathbf{m}),
\tag{30}
$$

where the subscript b stands for the (mini-)batch which is currently used to approximate the dataset, and B stands for the number of sources in batch b. In order to explore the full dataset, the batch needs to change between iterations. At the same time, the optimization procedure needs to have some measure of misfit evolution from iteration n to iteration $n + 1$. The approach used in Inversionson, is to split the batch into two parts, control group sources and non-control group sources. The control group consists of the sources which stay in the batch from iteration n to $n + 1$, and can thus be used to monitor misfit changes, while the other sources are replaced. The control group selection process is explained in Algorithm 1. The relationship between the dataset, batch and control group is described in Fig. 1 where C_n and B_n are the control group and batch for iteration n respectively, and A is the full dataset.

Creating an optimal dataset for a seismic inversion is a subjective process and the trade-off between computational cost and quality of inversion is non-trivial. Using a stochastic approach to the inversion enables much larger datasets without having to worry about the additional computational cost which comes with it, given that the datasets contain a certain level of redundancy, as is common in seismic datasets. The quasi-stochastic source selection process
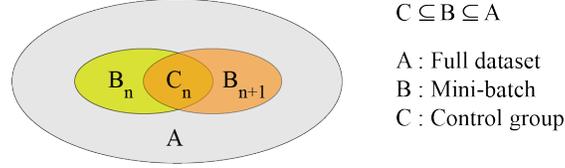
Figure 1: Relationship between control group, batch and dataset. Iterations are denoted by n subscript.

finds the most dominant sources in the dataset and gives them relatively more attention than the sources which are less dominant. This gives the dynamic mini-batch approach an element of optimal experimental design [Martiartu et al., 2017, Maurer et al., 2017, Vinard et al., 2018]. Each stochastic model update is cheaper than a full dataset model update, and with the dynamic mini-batch approach [van Herwaarden et al., 2020], which is implemented in Inversionson, the inversion itself figures out how many sources are needed in a batch in order to make a meaningful model update.

In FWI, one has to deal with imperfect data and has to be careful to fit the synthetic seismograms to the observed ones in a reasonable way. It is important to make measurements for which we expect that the gradient of it will guide us in the right descent direction. Therefore, we window parts of the seismogram for which meaningful measurements can be made. In a standard FWI workflow, like the one described in section 3.1, this is done for all sources at the start of the inversion, and it is not done again without restarting the whole inversion process. As an inversion progresses and the model improves, the number of possible windows to do safe measurements increases, but with a standard workflow, this cannot be exploited without restarting the optimization. That is one of the reasons why, traditionally, FWI is highly dependent on its starting model. In practice though, inversions are split into multiple frequency bands, and the data are re-windowed every time the inversion moves into a new band, which partially remedies the issue of windows being dependent on the starting model. It may, however, slow down convergence, as a restart of the inversion, either to pick new windows, add more data, or move to a new frequency band, breaks the quasi-Newton part of the optimization routine, which benefits from multiple prior iterations to approximate the inverse Hessian of the objective function.

One of the main benefits of stochastic seismic inversions is how naturally data can be re-windowed regularly and how new data can be added into an ongoing inversion [van Herwaarden et al., 2021]. Sources leave the batch and re-enter it regularly. With such a system, the windowing process can be repeated whenever a source (re-)enters the batch, making stochastic FWI much less dependent on its starting model.

3   Workflows

As described in the previous section, FWI is an iterative optimization process, and guiding large data volumes through complex optimization algorithms requires great care. Inversionson is able to handle various levels of workflow complexities which will be described in this section. Each workflow contains two phases, the preparation phase, which is only done once, and the iterative phase which is repeated until arbitrary criteria are reached. The workflows will be described with figures, showing how the various levels of complexities stack up.

3.1   Standard FWI Workflow

Inversionson always assumes the trust-region L-BFGS optimization scheme as the base level optimization. This can be changed relatively easily but for now, let us assume that the standard FWI workflow revolves around the trust-region L-BFGS optimization scheme. The workflow is described in Fig. 2 a).

Preparation phase: Synthetic waveforms, windows, misfits and gradients are computed for each source in the dataset, the preferred regularization is applied before moving on to the iterative phase.

Iterative phase: After updating the model based on the gradient of the previous iteration, synthetic waveforms are computed using the new model. The synthetic waveforms are compared to the data, and if the misfits between the two have dropped compared to the misfits of the previous iteration, a gradient can be computed and regularized before starting a new iteration. However, if the misfits did not drop, the trust-region needs to be reduced and the model updated again based on the new trust-region. The iteration is then restarted with the new model. It is best to avoid this happening as it is computationally expensive to re-compute all the synthetic waveforms. Usually, when everything is adequately set up, this does not happen often, if at all.
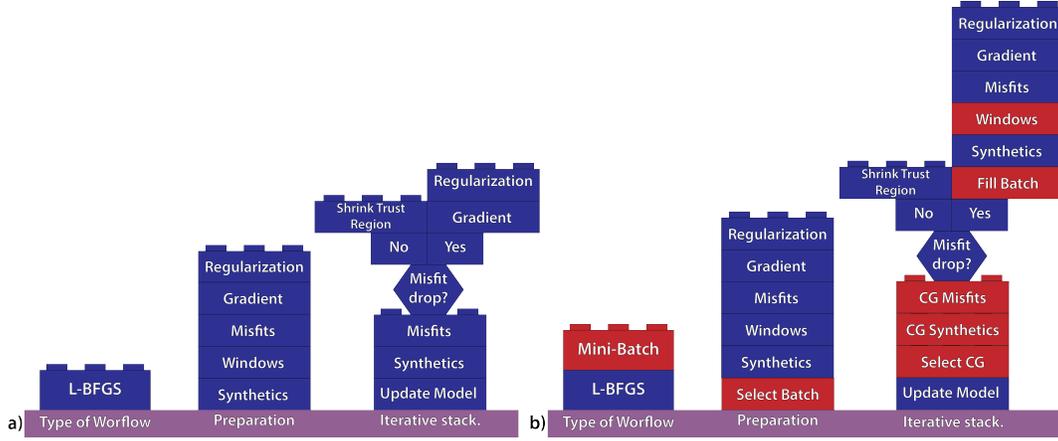
Figure 2: **a)** The standard trust-region L-BFGS FWI workflow illustrated with Lego©-type block structures. In all the workflow figures, it starts by moving from bottom to top in the preparation stack and then climbing the iterative stack repeatedly. In this workflow, synthetics are computed for all events, using the initial model, those synthetics are windowed, misfits computed, and a gradient is computed. With the preparation phase completed, the model is updated, and new synthetics and misfits computed. If the misfit reduced, a gradient is computed, otherwise the trust region is reduced and model updated again. **b)** The additional complexities of the Mini-Batch workflow (red cubes) have been added to the standard workflow. CG stands for control group. In this case, the synthetics are initially only computed on the first batch of events. These events are windowed, and a gradient computed based on them. In the iterative stack, only the control group events are used to quantify misfit reduction. If the misfit was reduced, the batch is filled up with new events, which are windowed and a gradient is computed based on the full batch.

3.2 Mini-Batches

The addition of the mini-batch approach to the standard workflow of section 3.1 adds a layer of complexity as displayed in Fig.2 b), where the red cubes have been added to the workflow of Fig. 2 a). Along with the extra layer of complexity in the workflow, comes a significant reduction in computational cost required to converge to a good model [van Herwaarden et al., 2020].

The main concept behind the mini-batch method is to not use every earthquake in the dataset in every iteration, but to rather use subsets at a time. This comes with subjective choices like how to pick sources to use in each iteration, and Inversionson has its own way of automatically dealing with that (Algorithm 1), thus reducing the subjectivity between individual inversions.

The mini-batch workflow can be split into two phases like the standard one. It will not be described from the ground up but only compared to the previously described one.

Preparation phase: Rather than computing initial synthetics for all events, they are only computed for a subset of events, the number of which is defined by the user. The events are selected to maximize spatial coverage.

Iterative phase: Upon initialization of each new iteration, a control group is selected as described in Algorithm 1. Initially, synthetics and misfits are computed for the control group. If the misfit increases, the trust region is reduced, and the iteration is restarted. If the misfit drops, the batch of the iteration is filled up, as described in Algorithm 1, and synthetics are computed for the new events. Misfits for the new events are computed using new windows and a gradient is computed for all events of the iteration

Benefits: The mini-batch approach includes multiple benefits, the most obvious one being greatly reduced computational cost to converge to a reasonable model. A less obvious benefit is the reduction of initial model dependence by (re-)windowing data from events that (re-)enter the inversion. The mini-batch approach allows us to expand the dataset during an inversion as shown in van Herwaarden et al. [2021], meaning that the choice of the initial dataset is not as crucial and binding as in the normal approach. Local minima are less of a problem in the mini-batch approach, as the objective function changes between iterations.
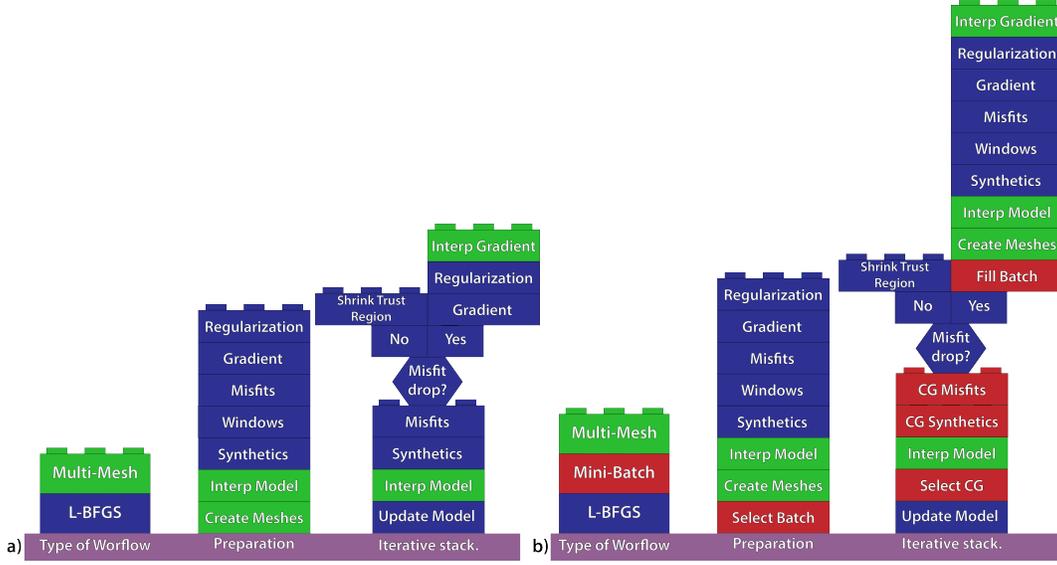
9

Figure 3: **a)** The relevant blocks corresponding to the multi-mesh workflow (green cubes) have been added to the standard workflow from Fig. 2 a). The difference in this case, is that new meshes need to be created for each event, and the model and gradient of each iteration need to be interpolated between discretizations. **b)** By combining the multi-mesh and mini-batch additions into one workflow, the workflow gets quite complex. New meshes get created every time a new event enters the workflow, and the previously mentioned interpolations have been added to the workflow displayed on Fig. 2 b).

Drawbacks: The only real drawback of the mini-batch approach is that it complicates the workflow, making it harder to implement. However, it is readily available for usage in Inversionson.

Algorithm 1: Event selection process

X = *max_angular_difference*
if *preparation phase* **then**
  | Pick N events to maximize coverage
else
  | Gradually remove events from previous iteration to select control group
  | **while** *angular_difference < X* **do**
  | | Remove event with the gradient with the smallest angular effect
  | **end**
  | **if** *unused events in dataset* **then**
  | | Pick unused events to maximize coverage
  | **else**
  | | Pick events quasi randomly, based on previous removal orders and spatial coverage
  | **end**
end

3.3 Multi-Mesh Approach

The multi-mesh approach was initially implemented by Thrastarson et al. [2020], where each seismic source is modelled on a specific mesh which is optimally designed for the respective source location [van Driel et al., 2020]. The meshes are adapted to the expected complexity of the wavefield prior to simulating, exploiting how the wavefield complexity is lower in the azimuthal dimension with respect to the source and thus needing fewer grid-points to resolve the wavefield in that dimension. As demonstrated by van Driel et al. [2020], these meshes can greatly reduce the compute cost of both wavefield and adjoint simulations, and as demonstrated by Thrastarson et al. [2020], the meshes

can be used in FWI to provide equally good results as standard meshes at a much lower cost given that enough elements are used to cover the azimuthal dimension.

Using wavefield-adapted meshes creates a complication, which results in the need to separate simulation meshes and the inversion parameterization (which can be a mesh as well). Apart from the need for multiple interpolations between different parameterizations, the multi-mesh approach does not complicate that standard workflow to a great extent, as can be seen in Fig. 3 a).

Preparation phase: Initially all the meshes need to be created and model parameters interpolated onto the new meshes. The workflow continues as in the standard workflow.

Iterative phase: After the model is updated, it has to be interpolated onto the meshes, and as the gradient has been processed, the individual gradients need to be interpolated onto the inversion parameterization.

Benefits: Simulation cost can be reduced by an order of magnitude. It can be reduced by up to a factor 15-20 in favourable cases [Thrastarson et al., 2020]. With simulation cost being the main expense factor of FWI, this is an important benefit.

Drawbacks: The meshes used in this method assume laterally smooth media making it perfect for finding smooth velocity variations, but not applicable to reflection-based studies. The workflow requires many mesh generations and interpolations. These actions need to be efficient in order to maximize the potential of the multi-mesh approach. Inversionson has efficient algorithms for meshing and interpolating, the latter of which exploits the Gauss-Lobatto-Legendre parameterization of each element of the mesh, to perform accurate interpolations.

3.4 Mini-batch Combined with Multi-mesh

Inversionson was initially written to automate this combined workflow and is thus optimized to facilitate it efficiently. The combined workflow is illustrated in Fig. 3 b). The mini-batch and the multi-mesh approaches complement each other. The benefits of the workflows go well together as the former needs fewer simulations to converge and the latter makes each of the simulations an order of magnitude cheaper than usually. An application of this workflow is shown in section 5.

4 Software

Inversionson is all written in Python [Rossum, 1995], with a large emphasis on NumPy [Harris et al., 2020]. It is designed to automate the workflows described in section 3, but also to facilitate the automation of similar FWI workflows which the user wants to implement.
Much work has already been spent on designing software specializing in specific tasks related to FWI. Inversionson takes pre-existing software and does everything which is required to make the other software packages work together in an efficient way to automate FWI.

The software which Inversionson ties together are:

- Salvus [Afanasiev et al., 2019] is a central part of Inversionson. It runs forward and adjoint simulations, it creates meshes, is the optimization manager, communicates with the computational cluster. Inversionson is designed to work with Salvus, exploiting its functionalities to streamline inversions, and cannot run without it. For automatic inversions without using Salvus, we recommend Chow et al. [2020].
- LASIF [Krischer et al., 2015a, Thrastarson et al., 2021b] is the inversion framework, meaning that it takes care of data downloading and managing, window selections, misfit and adjoint source computations.
- MultiMesh [Thrastarson et al., 2021a] is only needed if one wants to use the multi-mesh part of the implemented workflows. MultiMesh handles all interpolations for Inversionson.

The automatic inversion procedure works in a way that LASIF is used to create a LASIF project, define a domain, and download data. A user has to define certain configurations for both LASIF and Inversionson, and as an inversion is initialized, Inversionson communicates with the optimization routines in Salvus to get a task to perform. The Inversionson AutoInverter class (see Fig. 4) then delegates tasks to the relevant objects, shown in Fig. 4. How these tasks are performed and delegated depends on which workflow the user picks. As the task is completed, Inversionson reports back to the optimization module of Salvus to get a new task, and by repeating this workflow the whole inversion can be performed fully automatically.
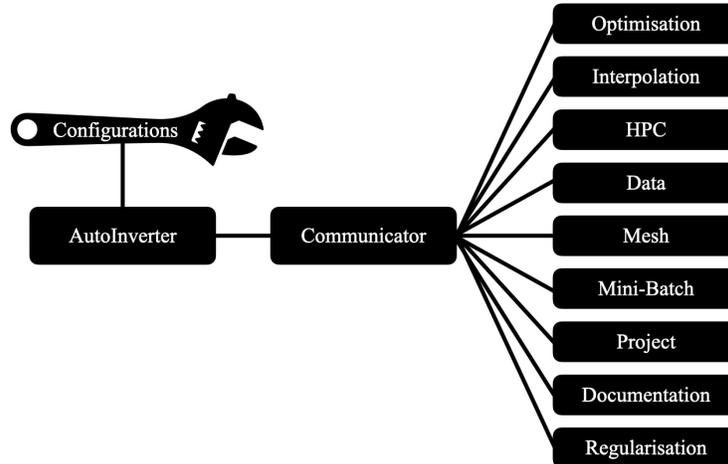
Figure 4: Software buildup. The wrench represents a configuration file filled in by the user, while each of the boxes represent individual classes.

5 Project

In order to demonstrate that Inversionson works for the kinds of large-scale seismic FWI problems that it was designed for, we show results of a synthetic test case. As large-scale FWI is a computationally expensive exercise, we constructed a case of relatively low compute cost, resolving low frequency wave propagation for a small amount of sources. The project was a long-period isotropic elastic global-scale FWI using the workflow described in section 3.4. Inversionson has been tested for all the workflows described in section 3, however, only the most complex workflow will be shown here. The number of sources used in the inversion was 29. The synthetic dataset was created using cubed-sphere meshes [Ronchi et al., 1996] with a period range from 150 s to 250 s. The selected waveform misfit function was the time-frequency phase misfit [Fichtner et al., 2008]. The test case is far from an optimal case for the workflow, as its benefits increase with higher frequencies and more sources. It does, however, serve as a good initial test case, and this section of the paper will be updated as the work of more large-scale cases will be published.

Inversionson ran uninterrupted for 21 mini-batch iterations starting from a 1-D model with the goal of recovering a Gaussian chequerboard model of up to 7% deviations (Fig. 5 d)) which is the order of deviations one can expect in the mantle.

The evolution from initial to final model with a comparison of the true model can be seen in Fig. 5. The recovery is already quite good in most regions of sufficient coverage (equator, N-America) while it is a bit worse where the coverage is inadequate (N-Atlantic, southern hemisphere). Getting a better recovery by running more iterations or expanding the dataset would have been an option, but this project was simply done to illustrate the workings of Inversionson, so we did not see a reason to do so. Inversionson generated an automatic documentation file during the inversion which gives information on which sources are used in each iteration and their respective misfit values.

5.1 Cost Analysis

Using the relatively small dataset in this test case, the cost of a full iteration is relatively low. Using the standard workflow from section 3.1, the total number of simulations required for an iteration is 58.

The cost benefit of the mini-batch approach from section 3.2 scales with the size of the dataset so for this study the benefit is unusually small. Making a robust cost comparison between the standard workflow and the mini-batch one is difficult as a mini-batch iteration is not equivalent in benefit compared to a full iteration so it is important to keep that in mind. The average number of simulations per mini-batch iteration was 36.8 resulting in a factor of 1.6 reduction in compute cost per iteration from using the mini-batch approach. This factor increases as the dataset gets larger since the cost of a full iteration scales linearly with number of used earthquakes, but the mini-batch approach automatically adjusts the size of the batch in each iteration, making it somewhat independent of the number of earthquakes. van Herwaarden et al. [2020] show an example for a reasonably sized dataset of 125 sources used to invert for the subsurface structure of the African continent which has a factor 5.3 reduction in the number of required
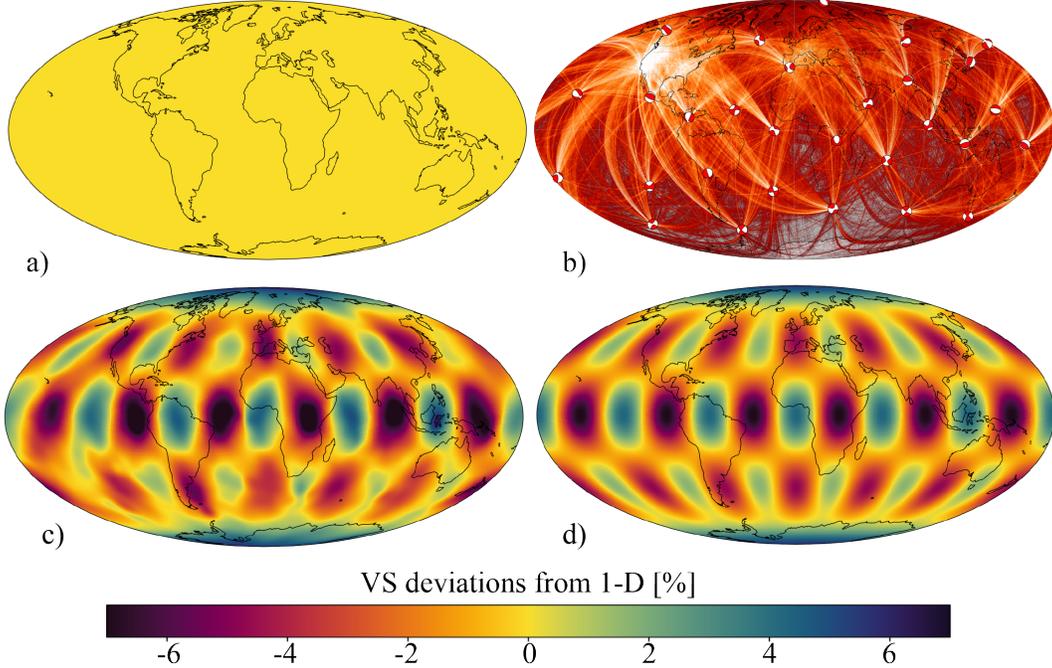
Figure 5: S-wave velocity deviations from 1-D model (PREM [Dziewonski and Anderson, 1981] with no crustal layer):
a) Starting model. **b)** Distribution of sources, with rays connecting sources and receivers, the lighter the ray color, the more ray crossings. **c)** Final model after 21 iterations. **d)** True model.

simulations to reach convergence.

Using the multi-mesh approach from section 3.3, results in each simulation becoming cheaper. As demonstrated by Thrastarson et al. [2020], the factor by which the simulation cost is reduced is proportional to the resolved frequency. This was a long period study, with a minimum period of 150 s to be resolved, meaning that the benefit of the multi-mesh approach is on the lower end of its potential. A standard, cubed sphere mesh, for the inversion settings includes 170,496 elements, while a wavefield adapted mesh for the same settings contains 30,696 elements. That results in a factor of 5.6 reduction in simulation cost which would increase with higher frequencies. Thrastarson et al. [2020] show a similar case, but resolving higher frequency wave propagation (50 s period), where the cost reduction factor is 14.1, demonstrating how the cost benefits increase with frequency.

Putting the benefits together by using the combined workflow from section 3.4 we can say that the 21 mini-batch iterations were computed at an equivalent cost of 2.4 full iterations using the standard FWI workflow.

6 Conclusions

We have presented, Inversionson, a Python based software package which, in combination with other scientific libraries, can perform FWIs fully automatically. It includes various advanced workflows, such as the dynamic mini-batch (section 3.2) and the multi-mesh (section 3.2) workflows, which greatly reduce the required computational cost of an inversion.

We showed an overview of the Trust-Region L-BFGS optimization method in the context of FWI, as well as four different FWI workflow options. One of which was the most commonly used standard approach, while the other ones have recently been developed. Finally, we showcased an example of a project done using Inversionson's capabilities.

Inversionson is currently being used in several, ongoing, real data applications, which will be material for future publications. The current applications include a global FWI, as well as continental scale inversions of both Africa and Eastern Asia.

All in all, Inversionson eases the process of performing an FWI study of a region, making it much

more approachable to any interested researcher. The code is fully open-source and available at `https://github.com/solvithrastar/Inversionson`.

7 Acknowledgements

References

M. Afanasiev, C. Boehm, M. van Driel, L. Krischer, M. Rietmann, D. A. May, M. G. Knepley, and A. Fichtner. Modular and flexible spectral-element waveform modelling in two and three dimensions. *Geophysical Journal International*, 216(3):1675–1692, 2019.

K. Aki and W. Lee. Determination of three-dimensional velocity anomalies under a seismic array using first p arrival times from local earthquakes: 1. a homogeneous initial model. *Journal of Geophysical research*, 81(23):4381–4399, 1976.

K. Aki, A. Christoffersson, and E. S. Husebye. Determination of the three-dimensional seismic structure of the lithosphere. *Journal of Geophysical Research*, 82(2):277–296, 1977.

A. Bamberger, G. Chavent, and P. Lailly. Une application de la théorie du contrôle à un problème inverse sismique. *Ann. Geophys.*, 33:183–200, 1977.

A. Bamberger, G. Chavent, C. Hemon, and P. Lailly. Inversion of normal incidence seismograms. *Geophysics*, 47(5): 757–770, 1982.

M. Beyreuther, R. Barsch, L. Krischer, T. Megies, Y. Behr, and J. Wassermann. Obspy: A python toolbox for seismology. *Seismological Research Letters*, 81(3):530–533, 2010.

C. Bösch, T. Dubček, F. Schindler, A. Fichtner, and M. Serra-Garcia. Discovery of topological metamaterials by symmetry relaxation and smooth topological indicators. *Physical Review B*, 102(24):241404, 2020.

E. Bozdağ, J. Trampert, and J. Tromp. Misfit functions for full waveform inversion based on instantaneous phase and envelope measurements. *Geophysical Journal International*, 185(2):845–870, 2011.

C. G. Broyden. The convergence of a class of double-rank minimization algorithms 1. general considerations. *IMA Journal of Applied Mathematics*, 6(1):76–90, 1970.

P. Chen, T. H. Jordan, and L. Zhao. Full three-dimensional tomography: a comparison between the scattering-integral and adjoint-wavefield methods. *Geophysical Journal International*, 170(1):175–181, 2007a.

P. Chen, L. Zhao, and T. H. Jordan. Full 3d tomography for the crustal structure of the los angeles region. *Bulletin of the Seismological Society of America*, 97(4):1094–1120, 2007b.

B. Chow, Y. Kaneko, C. Tape, R. Modrak, and J. Townend. An automated workflow for adjoint tomography—waveform misfits and synthetic inversions for the north island, new zealand. *Geophysical Journal International*, 223(3): 1461–1480, 2020.

R. Cockett, S. Kang, L. J. Heagy, A. Pidlisecky, and D. W. Oldenburg. Simpeg: An open source framework for simulation and gradient based parameter estimation in geophysical applications. *Computers & Geosciences*, 85: 142–154, 2015.

F. Dahlen, S.-H. Hung, and G. Nolet. Fréchet kernels for finite-frequency traveltimes—i. theory. *Geophysical Journal International*, 141(1):157–174, 2000.

A. M. Dziewonski and D. L. Anderson. Preliminary reference earth model. *Physics of the earth and planetary interiors*, 25(4):297–356, 1981.

A. M. Dziewonski, B. H. Hager, and R. J. O'Connell. Large-scale heterogeneities in the lower mantle. *Journal of Geophysical Research*, 82(2):239–255, 1977.

E. Faccioli, F. Maggio, A. Quarteroni, and A. Tagliani. Spectral-domain decomposition methods for the solution of acoustic and elastic wave equations. *Geophysics*, 61(4):1160–1174, 1996.

A. Fichtner, H.-P. Bunge, and H. Igel. The adjoint method in seismology: I. theory. *Physics of the Earth and Planetary Interiors*, 157(1-2):86–104, 2006.

A. Fichtner, B. L. Kennett, H. Igel, and H.-P. Bunge. Theoretical background for continental-and global-scale full-waveform inversion in the time–frequency domain. *Geophysical Journal International*, 175(2):665–685, 2008.

A. Fichtner, B. L. Kennett, H. Igel, and H.-P. Bunge. Full seismic waveform tomography for upper-mantle structure in the australasian region using adjoint methods. *Geophysical Journal International*, 179(3):1703–1725, 2009.

R. Fletcher. A new approach to variable metric algorithms. *The computer journal*, 13(3):317–322, 1970.

L. Gebraad, C. Boehm, and A. Fichtner. Bayesian elastic full-waveform inversion using hamiltonian monte carlo. *Journal of Geophysical Research: Solid Earth*, 125(3):e2019JB018428, 2020.

D. Giardini, X.-D. Li, and J. H. Woodhouse. Three-dimensional structure of the earth from splitting in free-oscillation spectra. *Nature*, 325(6103):405–411, 1987.

D. Goldfarb. A family of variable-metric methods derived by variational means. *Mathematics of computation*, 24(109): 23–26, 1970.

C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, et al. Array programming with numpy. *Nature*, 585(7825):357–362, 2020.

R. Hewett and L. Demanet. the pysit team, 2013, pysit: Python seismic imaging toolbox v0. 5. *Release 0.6*, 2013.

H. Igel, H. Djikpéssé, and A. Tarantola. Waveform inversion of marine reflection seismograms for p impedance and poisson's ratio. *Geophysical Journal International*, 124(2):363–371, 1996.

L. Krischer, A. Fichtner, S. Zukauskaite, and H. Igel. Large-scale seismic inversion framework. *Seismological Research Letters*, 86(4):1198–1207, 2015a.

L. Krischer, T. Megies, R. Barsch, M. Beyreuther, T. Lecocq, C. Caudron, and J. Wassermann. Obspy: A bridge for seismology into the scientific python ecosystem. *Computational Science & Discovery*, 8(1):014003, 2015b.

L. Krischer, J. Smith, W. Lei, M. Lefebvre, Y. Ruan, E. S. de Andrade, N. Podhorszki, E. Bozdağ, and J. Tromp. An adaptable seismic data format. *Geophysical Supplements to the Monthly Notices of the Royal Astronomical Society*, 207(2):1003–1011, 2016.

P. Lailly. as a sequence of before stack migrations. In *Conference on Inverse Scattering–Theory and Application*, volume 11, page 206. Siam, 1983.

A. Maggi, C. Tape, M. Chen, D. Chao, and J. Tromp. An automated time-window selection algorithm for seismic tomography. *Geophysical Journal International*, 178(1):257–281, 2009.

N. K. Martiartu, C. Boehm, N. Vinard, I. J. Balic, and A. Fichtner. Optimal experimental design to position transducers in ultrasound breast imaging. In *Medical Imaging 2017: Ultrasonic Imaging and Tomography*, volume 10139, page 101390M. International Society for Optics and Photonics, 2017.

H. Maurer, A. Nuber, N. K. Martiartu, F. Reiser, C. Boehm, E. Manukyan, C. Schmelzbach, and A. Fichtner. Optimized experimental design in the context of seismic full waveform inversion and seismic waveform imaging. In *Advances in geophysics*, volume 58, pages 1–45. Elsevier, 2017.

R. T. Modrak, D. Borisov, M. Lefebvre, and J. Tromp. Seisflows—flexible waveform inversion software. *Computers & geosciences*, 115:88–95, 2018.

J. Nocedal. Updating quasi-newton matrices with limited storage. *Mathematics of computation*, 35(151):773–782, 1980.

J. Nocedal and S. Wright. *Numerical Optimization*. Springer Science & Business Media, 2006. doi: 10.1007/ 978-0-387-40065-5.

R. G. Pratt. Seismic waveform inversion in the frequency domain; part 1, theory and verification in a physical scale model. *Geophysics*, 64(3):888–901, 1999.

R. G. Pratt, C. Shin, and G. Hick. Gauss–newton and full newton methods in frequency–space seismic waveform inversion. *Geophysical Journal International*, 133(2):341–362, 1998.

C. Ronchi, R. Iacono, and P. S. Paolucci. The "cubed sphere": a new method for the solution of partial differential equations in spherical geometry. *Journal of computational physics*, 124(1):93–114, 1996.

G. Rossum. Python reference manual. 1995.

D. F. Shanno. Conditioning of quasi-newton methods for function minimization. *Mathematics of computation*, 24(111): 647–656, 1970.

C. Tape, Q. Liu, A. Maggi, and J. Tromp. Adjoint tomography of the southern california crust. *Science*, 325(5943): 988–992, 2009.

A. Tarantola. Inversion of seismic reflection data in the acoustic approximation. *Geophysics*, 49(8):1259–1266, 1984.

A. Tarantola. Theoretical background for the inversion of seismic waveforms, including elasticity and attenuation. In *Scattering and attenuations of seismic waves, part i*, pages 365–399. Springer, 1988.

S. Thrastarson, M. van Driel, L. Krischer, C. Boehm, M. Afanasiev, D.-P. van Herwaarden, and A. Fichtner. Accelerating numerical wave propagation by wavefield adapted meshes. part II: full-waveform inversion. *Geophysical Journal International*, 221(3):1591–1604, feb 2020. doi: 10.1093/gji/ggaa065. URL `https://doi.org/10.1093%2Fgji%2Fggaa065`.

S. Thrastarson, D.-P. van Herwaarden, and A. Fichtner. solvithrastar/MultiMesh: MultiMesh - Python-based interpolations between discretizations, Feb. 2021a. URL `https://doi.org/10.5281/zenodo.4564523`.

S. Thrastarson, D.-P. van Herwaarden, L. Krischer, and A. Fichtner. LASIF: LArge-scale seismic inversion framework, an updated version. Jan. 2021b. doi: 10.31223/x5nc84. URL `https://doi.org/10.31223%2Fx5nc84`.

J. Tromp, C. Tape, and Q. Liu. Seismic tomography, adjoint methods, time reversal and banana-doughnut kernels. *Geophysical Journal International*, 160(1):195–216, 2005.

M. van Driel, C. Boehm, L. Krischer, and M. Afanasiev. Accelerating numerical wave propagation using wavefield adapted meshes. part i: forward and adjoint modelling. *Geophysical Journal International*, 221(3):1580–1590, 2020.

D. P. van Herwaarden, C. Boehm, M. Afanasiev, S. Thrastarson, L. Krischer, J. Trampert, and A. Fichtner. Accelerated full-waveform inversion using dynamic mini-batches. *Geophysical Journal International*, 221(2):1427–1438, feb 2020. doi: 10.1093/gji/ggaa079. URL `https://doi.org/10.1093%2Fgji%2Fggaa079`.

D. P. van Herwaarden, M. Afanasiev, S. Thrastarson, and A. Fichtner. Evolutionary full-waveform inversion. *Geophysical Journal International*, 224(1):306–311, 2021.

N. Vinard, N. K. Martiartu, C. Boehm, I. J. Balic, and A. Fichtner. Optimized transducer configuration for ultrasound waveform tomography in breast cancer detection. In *Medical Imaging 2018: Ultrasonic Imaging and Tomography*, volume 10580, page 105800I. International Society for Optics and Photonics, 2018.

P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, et al. Scipy 1.0: fundamental algorithms for scientific computing in python. *Nature methods*, 17(3):261–272, 2020.

P. Wolfe. Convergence conditions for ascent methods. *SIAM review*, 11(2):226–235, 1969.

J. H. Woodhouse and A. M. Dziewonski. Mapping the upper mantle: Three-dimensional modeling of earth structure by inversion of seismic waveforms. *Journal of Geophysical Research: Solid Earth*, 89(B7):5953–5986, 1984.

K. Yomogida. Fresnel zone inversion for lateral heterogeneities in the earth. *Pure and Applied Geophysics*, 138(3):391–406, 1992.