# LASIF: LArge-scale Seismic Inversion Framework
## An updated version

**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

**Lion Krischer**
ETH Zürich
Department of Earth Sciences
Institute of Geophysics
lion.krischer@erdw.ethz.ch

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

March 11, 2022

## Abstract

Recent methodological advances and increases in computational power have made it feasible to perform full-waveform inversions (FWI) of large domains while using more sources. This trend, along with the increasing availability of seismic data has led to an explosion of the data volumes that can, and should, be used within an inversion. Similar to machine learning problems, the incorporation of more data can result in more robust and higher quality models. In this contribution, we present the new version of LASIF, an open-source LArge-scale Seismic Inversion Framework, which helps to automate many of the historically labor-intensive tasks that were bottlenecks in earlier FWI workflows and prevented the use of the larger datasets. Among other things, the framework automates data selection, data acquisition from public web services, and data processing. It also defines an inversion project structure that organizes the data and documents the progress of the inversion. The code is open-source and available on Github. Features are available through a graphical user interface (GUI), a command-line interface (CLI), and an application programming interface (API). While we will show examples for use of LASIF with the Salvus wave equation solver, the API makes it possible to use the features of LASIF for any type of wave equation solver as long as the LASIF file formats are adhered to.

***Keywords*** Seismology · Inverse Problems · Software · Open-Source

## 1 Introduction

Full-Waveform inversion (FWI) was first introduced in the late 1970s and early 1980s, and has since then progressed into a state-of-the-art methodology for constructing seismic velocity models [Bamberger et al., 1977, 1982, Lailly, 1983, Tarantola, 1984, Virieux and Operto, 2009]. Model updating typically requires the calculation of the gradient through the adjoint state method [e.g. Tromp et al., 2005, Fichtner et al., 2006] or the scattering-integral method [Chen et al., 2007a], after which the model is updated through iterative gradient based optimization techniques [e.g. Nocedal and Wright, 2006]. For each iteration of the waveform inversion, the gradient calculation requires the solution of the 3D wave equation for each source. This makes the method computationally expensive. With an increase in computational resources the method gained traction in regional- and continental-scale studies [Chen et al., 2007a,b, Fichtner et al.,

---

*Equal contribution of first two authors ()

2009, Tape et al., 2009].

Along with the increase in the availability of computational resources, the available amount of global seismic waveform data has also been doubling every 3 to 6 years [IRIS, 2020]. The manual workflows of the aforementioned studies would prove difficult to scale for usage with very large data volumes.

Krischer et al. [2015a] initially developed the LArge-scale Seismic Inversion Framework (LASIF). It was built on top of Obspy [Beyreuther et al., 2010, Megies et al., 2011, Krischer et al., 2015b] to provide a stable framework to carry out local- to global-scale FWI studies and it has been used in various applications [e.g. Simutė et al., 2016, Çubuk-Sabuncu et al., 2017, Krischer et al., 2018]. As the datasets grew in size, several of the original algorithms started to create bottlenecks in the workflow. The previous framework also did not offer support for the new wave equation solver Salvus [Afanasiev et al., 2019]. This motivated an update of the original LASIF package.

In this contribution, we describe the modifications and improvements that have taken place over the last years. We also aim to provide a basic overview of the most commonly performed tasks and give examples of work that was done with the current version of LASIF.

## 1.1 What is new?

In this section, we briefly list the most important changes to the initial package. Readers who are unfamiliar with the old code can safely skip this section.

- All seismic data is now stored in the ASDF seismic data format [Krischer et al., 2016]. ASDF is a container format, that enables multiple files to be merged into a single file. By merging files, the file count is reduced, which in turn is important for the file systems that are used on many High Performance Computing (HPC) systems

- Measurement windows, that define where synthetic waveforms can be compared to data, are now stored in the form of SQLite databases [Owens, 2006], to allow for quick queries even when millions of windows are used.

- Station weightings can now be automatically generated to down-weigh clusters of stations. Intuitively, this helps speed up convergence by balancing the kernels and model updates. For a 2D case, this has already been demonstrated to have significant benefits [Ruan et al., 2019].

- Windows or station weights no longer have to be copied between iterations but are now defined as sets, that can be specified upon the calculation of misfits and adjoint sources.

- All functions have become available through an API, which enables convenient scripting, for example within a Jupyter Notebook [Kluyver et al., 2016] (see the example in section 4).

- Window selection has become an order of magnitude faster. Previously, correlation coefficients and other measurements were computed for a sliding window that stepped through each time step of the pair of waveform recordings. Instead, the windows now slide by a fraction of the minimum period of interest. This greatly reduces the total number of computations required.

- All compute time-intensive tasks, such as selecting windows, processing data, or computing adjoint sources, are now performed in parallel and come with progress bars.

- Irregularly shaped domains can be recognized from unstructured meshes. Data which fits into the domain can be automatically downloaded.

- In the previous version, caches were used to establish relationships between StationXMLs and events. In the new version, data is structured per event. This makes it trivial to add and remove events from a project.

- The internal earthquake database, provided by the GCMT catalog [Dziewonski et al., 1981, Ekström et al., 2012] can now be automatically updated through the API to include the most recent events.

- Plots are now generated using Cartopy [Met Office, 2010 - 2015] rather than, the now deprecated, Basemap.

- The GUI has been updated to enable the browsing through iterations, the visualization of raw and processed data, and the visualization of varying window sets. It is now possible to compare waveforms between different iterations.

- Extensive support for the Salvus wave equation solver has been added.

- We no longer use XML internally and now use TOML. This makes it much easier to read and manually edit files.

- An automatic clean-up function has been added to remove data that for example lies outside of the current domain of interest. This enables the copying of data from an old project to a new project with a different domain.

Figure 1: The directory structure of LASIF. Red boxes expand with each iteration and green boxes with each earthquake. The Sets directory contains window sets and station weights. The Iteration directory contains information about each iteration, such as which events were used, simulation settings and the computed misfit values.

## 2 Tasks performed by LASIF

LASIF is designed with the motivation to streamline the process of seismic inversions. An inversion, or project, is defined by its folder structure where all the key parameters are defined in a configuration file. As a project is created, a directory tree (see Figure 1) is built which has a natural way of growing new branches as the inversion accumulates iterations or earthquakes are added to the project. Within a seismic inversion, several tasks need to be carried out. Many of these tasks are quite similar between inversions and there is not much research required at each step. Figure 2 shows a simplistic view of an FWI workflow.

Often the most (human) time-consuming tasks tend to be some of the green boxes on the figure and those are the ones that LASIF performs on its own. Regarding the green boxes that can vary between inversion workflows (i.e. way of picking windows or computing adjoint sources), LASIF provides function templates that can easily be modified if the user wants to experiment. The blue boxes in Figure 2 are ones which LASIF helps with but does not do on its own. It has options for defining simplistic domains but works best when a mesh is used for the domain, it gathers and prepares all the files needed for creating synthetic waveforms or computing gradients, and it helps out with dispatching the simulations to a high-performance-computing center of choice. During an inversion it can be helpful to inspect data and synthetics by eye as human intuition for data is hard to replace. LASIF thus provides a graphical user interface where waveforms can be compared between iterations.

LASIF is not designed as a stand-alone library to do FWI. It is designed as a framework to aid with the most tedious and labor-intensive parts of FWI and as such is meant to be a research enabling product. For that reason, LASIF does not ship with regularization or optimization routines involved in updating models between iterations.

In section 4 the capabilities of LASIF are further demonstrated in an example project/tutorial form.

## 3 Structure

The architecture of LASIF is displayed in a diagram in Figure 3 where it can be seen that the core is built on 11 key classes which all have their distinct roles. To facilitate communication between the key objects (components), there is a Communicator class that has access to the methods and attributes of the components. The Communicator class allows
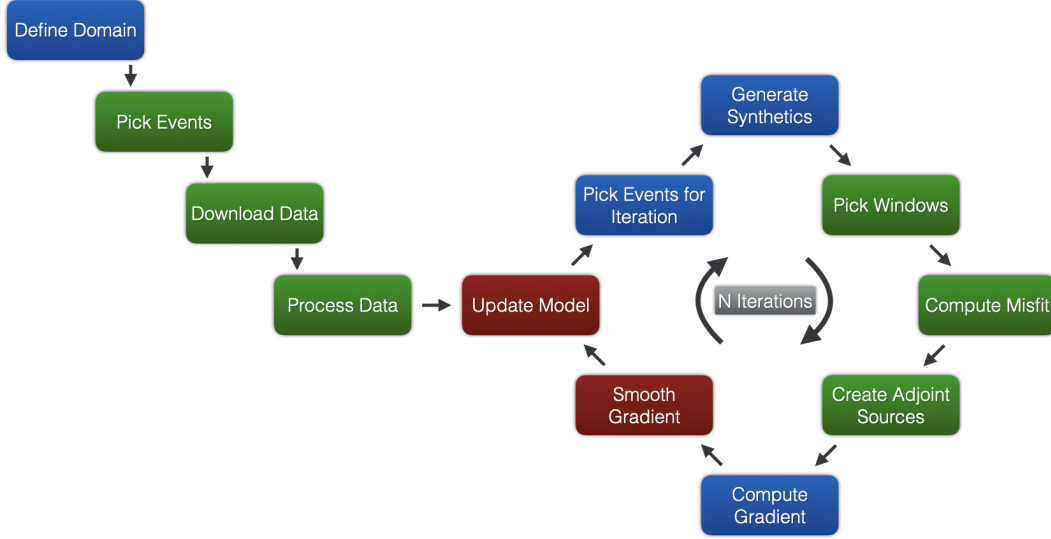
Figure 2: A typical FWI workflow diagram. Green boxes represent tasks that LASIF performs. Blue boxes are tasks that LASIF helps with, generating input files for an external waveform solver for example. Red boxes are tasks where LASIF does not currently get involved. Each LASIF project contains a folder structure for storing the files for all of these things in an organized manner.

for accessing the same instance of each class from within the instances of all the other classes from Fig. 3. That means that if at some point an object variable is changed, every other object will have access to that changed variable, making sure that the objects always have consistent information between them. Any information, such as the events or stations that were used within a particular iteration can be queried from this one Communicator object. The Communicator object is also the entry point for each user interface. A user can interact with LASIF in four ways:

1. **Graphical User Interface (GUI)**: An interface for the user to look at waveforms, windows and to do manual window selection.

2. **Application Programming Interface (API)**: Directly accessible from a Python script, includes all the high-level functionalities of LASIF and makes it easy to incorporate LASIF functionalities into a Python script.

3. **Command Line Interface (CLI)**: The CLI is for those who prefer to call the functions directly from the terminal. Each CLI function simply calls the API functions so it does not add anything on top of the API aside from being another way to interact with LASIF.

4. **Communicator**: From the API it is possible to return an instance of the Communicator class. This is needed for more advanced scripting beyond the capabilities of the previously mentioned interfaces.

## 4    Example project

In this section, we will show several code snippets (Listings). These snippets together create a full script and need to be used together to function, unless they are broken into cells using a Jupyter notebook [Kluyver et al., 2016]. The code snippets are not guaranteed to work as the code evolves through time but they are displayed here as a way to illustrate what LASIF does. To demonstrate a real-world use case for LASIF, we will describe an FWI project where the library was used. It is an inversion of the African continent by van Herwaarden et al. [2020]. The project was done by combining LASIF and Salvus [Afanasiev et al., 2019], introducing a new stochastic L-BFGS approach to FWI. Such work is precisely what LASIF is designed for, it is a basic framework for FWI which is created to be built on top of.
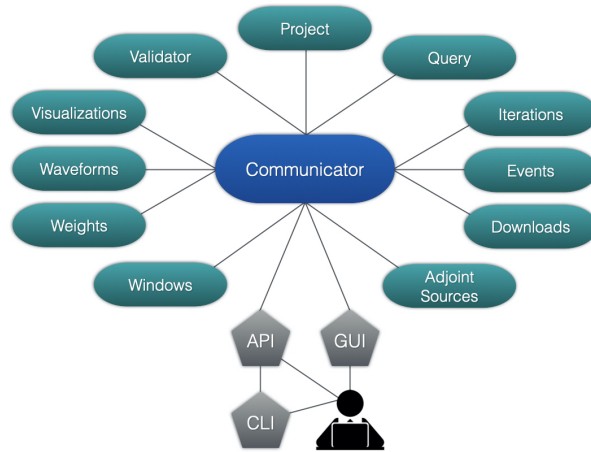
Figure 3: A schematic representation of the LASIF code structure. The blue and greenblue shapes represent the core of LASIF while the hexagons represent the user interface.

## 4.1 Preparation

Preparing for an inversion is a process that can take a long time. Adjusting the mesh, finding the right earthquakes to model, and gathering all data for the respective earthquakes into a consistent format. With only three commands, LASIF does most of that work for the user, saving valuable time. After providing LASIF with a computational mesh that defines the domain, preparing the whole inversion can be done swiftly in the following way:

```python
import lasif

lasif_project_path = "." # Set project path
num_events = 250 # Number of events to use in inversion
min_mag = 5.5 # Minimum magnitude to filter events
max_mag = 6.7 # Max magnitude to filter events
min_distance = 50.0 # Minimum distance between events in km

# Creating figure 4 (left):
lasif.api.plot_domain(lasif_project_path)

# Add events from GCMT catalog
lasif.api.add_gcmt_events(
    lasif_project_path,
    count=num_events,
    min_mag=min_mag,
    max_mag=max_mag,
    min_dist=min_distance,
)
# Creating figure 4 (center):
lasif.api.plot_events(lasif_project_path)

# Download data
lasif.api.download_data(lasif_project_path)

# Creating figure 4 (right):
lasif.api.plot_raydensity(lasif_project_path)

# Process data (remove instrument response, bandpass filter, etc.)
lasif.api.process_data(lasif_project_path)
```

Listing 1: Preparing the inversion project. Define parameters of which earthquakes to use and download the available data, and then process it.

All data were downloaded and organized into individual files per earthquake using ASDF [Krischer et al., 2016]. The data was processed and organized in a way such that LASIF knows how to find it and use it for future tasks.
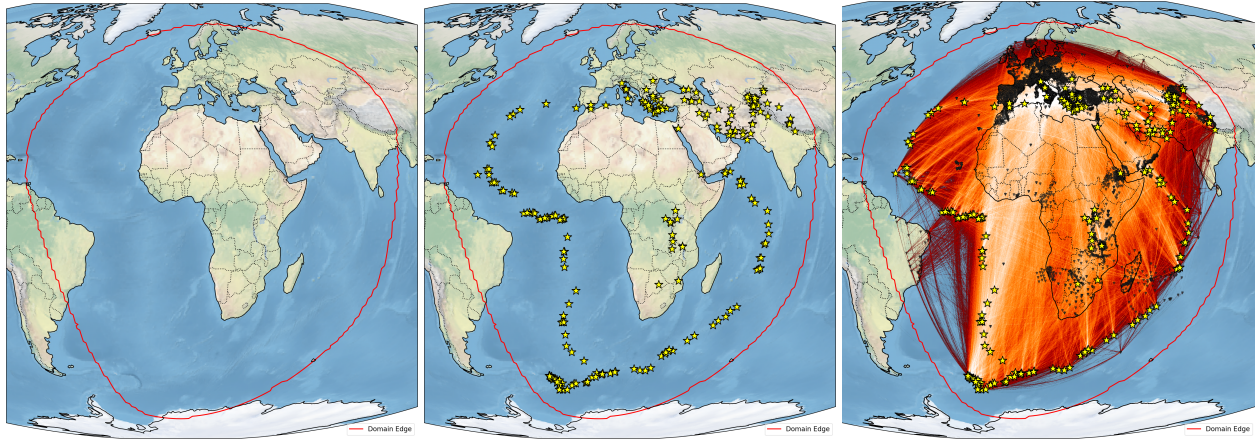


Figure 4: This figure shows the step by step preparation of a LASIF project. From left to right. First, the domain of interest is defined. Afterwards, sources are added, indicated by the yellow stars, and finally data are downloaded for the respective sources. Station locations are indicated by black triangles.

## 4.2 First iteration

As preparation is done, it is time to prepare and run the first iteration. In the stochastic workflow of van Herwaarden et al. [2020], each iteration uses a distinct set of events and we thus need to get a subsample of the dataset to use in this iteration.

```python
# Define number of earthquakes to use in iteration
n_events_in_iteration = 5

# Define iteration name (will be used in folder structure)
iteration_name = "it0000"

# Find all the earthquakes available in the project
events_in_project = lasif.api.list_events(lasif_project_path)

# Get a subset of the full dataset of earthquakes while
# keeping good spatial distribution
event_subset = lasif.api.get_subset(
    lasif_project_path,
    events=events_in_project,
    count=n_events_in_iteration,
)
# Expand the folder structure from Figure 1 to incorporate the new iteration
lasif.api.set_up_iteration(
    lasif_project_path,
    iteration=iteration_name,
    events=event_subset
)
```

Listing 2: Set up the first iteration. Define which earthquakes to use in the iteration and then expand the folder structure from Figure 1 to incorporate the new iteration.

The folder structure of the project has now been expanded to allow for iteration `it0000`, and the project is ready to compute its first synthetic seismograms by submitting waveform simulations to an HPC cluster.

### 4.2.1 Wavefield simulation

LASIF has built-in utility functions to prepare simulation objects in a way that Salvus needs them, but similar utility functions for other solvers can be created. Salvus simulation objects are created and dispatched to an HPC cluster in Listing. 3.

```python
from lasif import salvus_utils  # Import utility functions to work with Salvus

# We extract the communicator object of the project
lasif_communicator = lasif.api.find_project_comm(lasif_project_path)

# Find the earthquakes which are used in the current iteration
events = lasif.api.list_events(lasif_project_path, iteration=iteration_name)

# Create Salvus simulation objects. One for each earthquake
simulations = []
for event in events:
    simulations.append(
        salvus_utils.create_salvus_forward_simulation(
            comm=lasif_communicator,
            event=event,
            iteration=iteration_name,
            side_set="r1"
        )
    )

# Dispatch the simulations to a supercomputing cluster defined in the config file
salvus_utils.submit_salvus_simulation(
    comm=lasif_communicator,
    simulations=simulations,
    events=events,
    iteration=iteration_name,
    sim_type="forward"
)
```

Listing 3: Create simulation objects compatible with Salvus and submit the simulation objects to be executed on a supercomputing cluster.

### 4.2.2 Gradient computation

When using data from real seismometers, they are often heterogeneously distributed, noisy, or dominated by local effects. LASIF combats this with a window selection algorithm described in Krischer et al. [2015a], and a spatial weighting of the receivers. To prepare for the gradient computation, LASIF comes equipped with internal routines to do each of these tasks and to compute the actual adjoint sources:

```python
# Download the computed synthetic waveforms from the supercomputing cluster
salvus_utils.retrieve_salvus_simulations(
    comm=lasif_communicator,
    events=events,
    iteration=iteration_name,
    sim_type="forward",
)

# Compute spatial station weighting to speed up convergence
station_weights = f"station_weights_{iteration_name}"
lasif.api.compute_station_weights(
    lasif_project_path,
    weight_set=station_weights,
    iteration=iteration_name,
)

# Select measurement windows into data where synthetics and data are comparable
window_set_name = f"windows_{iteration_name}"
lasif.api.select_windows(
    lasif_project_path,
```

```
    iteration=iteration_name,
    window_set=window_set_name,
)

# Compute adjoint sources and misfits between data and synthetics within the
#  measurement windows and weighted by the station weights
lasif.api.calculate_adjoint_sources(
    lasif_project_path,
    iteration=iteration_name,
    window_set=window_set_name,
    weight_set=station_weights,
)
```

Listing 4: Retrieve the computed synthetics from the supercomputing cluster, compare to data by finding measurement windows where the data is comparable, spatially weight the receivers, and compute misfits and adjoint sources.

The type of the adjoint source is specified in the configuration file and it is possible to add custom functions at will. As the adjoint sources are computed one can go ahead and initiate the gradient computations:

```
# Create the adjoint simulation objects, one for each earthquake
simulations = []
for event in events:
    simulations.append(
        salvus_utils.create_salvus_adjoint_simulation(
            comm=lasif_communicator,
            event=event,
            iteration=iteration_name,
        )
    )

# Dispatch the simulations to a super-computing cluster
salvus_utils.submit_salvus_simulation(
    comm=lasif_communicator,
    simulations=simulations,
    events=events,
    iteration=iteration_name,
    sim_type="adjoint",
)
```

Listing 5: Create adjoint simulation objects which are compatible with Salvus and submit the objects to be executed on a super-computing cluster. The adjoint simulations are used to compute gradients.

As the adjoint simulations have been completed, LASIF can download the resulting gradients:

```
# Download the gradients resulting from the adjoint simulations
salvus_utils.retrieve_salvus_simulations(
    comm=lasif_communicator,
    events=events,
    iteration=iteration_name,
    sim_type="adjoint",
)
```

Listing 6: Download the gradients, resulting from the adjoint simulations, from the super-computing cluster. The gradients will be placed into dedicated folders where LASIF knows how to find them

The individual gradient should now be in its relevant folders.

### 4.2.3 Updating model

As previously stated, LASIF does not assist with updating the model. From here on, it is the individual researcher's decision on how to update the model between iterations. For an overview of non-linear optimization techniques, we refer to Nocedal and Wright [2006].
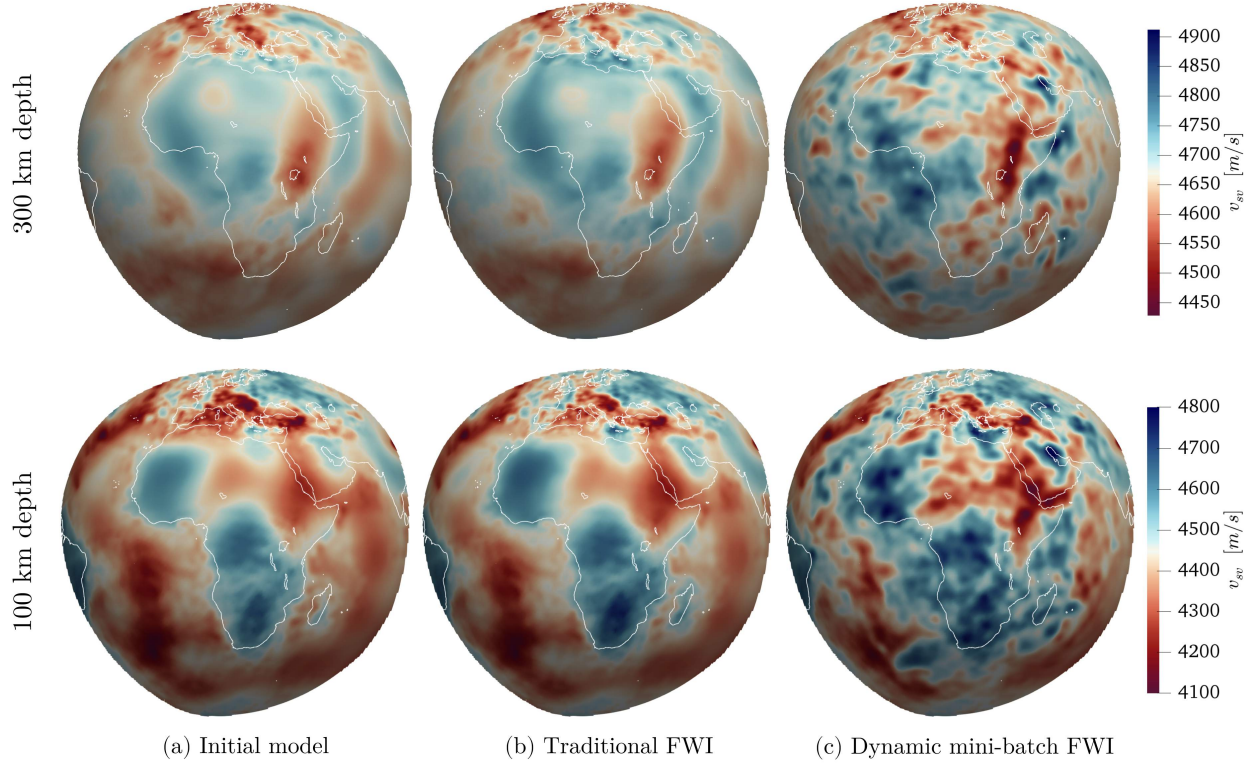
Figure 5: The figure shows the $v_{sv}$ velocity model depth slices at 300 km (upper) and 100 km (lower) depths. Column **(a)** shows the model which was used to compute the initial synthetic seismograms [Fichtner et al., 2018], and column **(b)** shows the retrieved model after 3 iterations of standard FWI. Column **(c)** displays the model which was retrieved using the same number of simulations as in (b), but using the stochastic workflow of van Herwaarden et al. [2020], running 40 dynamic mini-batch iterations. Figure taken from van Herwaarden et al. [2020].

### 4.3 Future iterations

After the model has been updated, one can repeat the steps listed in section 4.2 as many times as needed to make the results satisfactory. The results of the example project can be seen in Figure 5. From this example, it is quite evident how LASIF can be used to write a fully automatic FWI script in combination with an optimization routine to update the model between iterations. This has already been done for example in Thrastarson et al. [2020] and van Herwaarden et al. [2021].

## 5 Conclusions

We have presented an updated version of the LASIF framework. Through its new structure, it has become easier to scale the inversion to include more data. Many of the most time-consuming tasks can be performed completely automatically. Because of the open-source nature and the updated structure with the API, it becomes simple to use the same framework for different types of projects than the one we presented. As this is a software package that continues to be in development, some of the features here may change in the future. We invite those that work on similar projects to collaborate, such that LASIF may eventually evolve into a community-developed code. For any questions, feel free to contact the authors directly by e-mail or through Github. The LASIF documentation is online at lasif.net and its repository can be found at github.com/dirkphilip/LASIF_2.0.

## 6 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. *Geophys. J. Int.*, 216(3):1675–1692, 2019. doi: 10.1093/gji/ggy469.

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, and J. Wassermann. ObsPy: A Python toolbox for seismology. *Seis. Res. Lett.*, 81:47–58, 2010.

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.

Y. Çubuk-Sabuncu, T. Taymaz, and A. Fichtner. 3-d crustal velocity structure of western turkey: Constraints from full-waveform tomography. *Physics of the Earth and Planetary Interiors*, 270:90–112, 2017.

A. Dziewonski, T.-A. Chou, and J. H. Woodhouse. Determination of earthquake source parameters from waveform data for studies of global and regional seismicity. *Journal of Geophysical Research: Solid Earth*, 86(B4):2825–2852, 1981.

G. Ekström, M. Nettles, and A. Dziewoński. The global cmt project 2004–2010: Centroid-moment tensors for 13,017 earthquakes. *Physics of the Earth and Planetary Interiors*, 200:1–9, 2012.

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. Full seismic waveform tomography for upper-mantle structure in the australasian region using adjoint methods. *Geophysical Journal International*, 179(3):1703–1725, 2009.

A. Fichtner, D.-P. van Herwaarden, M. Afanasiev, S. Simutė, L. Krischer, Y. Çubuk-Sabuncu, T. Taymaz, L. Colli, E. Saygin, A. Villaseñor, et al. The collaborative seismic earth model: generation 1. *Geophysical research letters*, 45 (9):4007–4016, 2018.

IRIS. 2020 iris at-a-glance. `https://www.iris.edu/hq/publications/annual_reports`, December 2020.

T. Kluyver, B. Ragan-Kelley, F. Pérez, B. E. Granger, M. Bussonnier, J. Frederic, K. Kelley, J. B. Hamrick, J. Grout, S. Corlay, et al. Jupyter notebooks-a publishing format for reproducible computational workflows. In *ELPUB*, pages 87–90, 2016.

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. doi: 10.1088/1749-4699/8/1/014003.

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.

L. Krischer, A. Fichtner, C. Boehm, and H. Igel. Automated large-scale full seismic waveform inversion for north america and the north atlantic. *Journal of Geophysical Research: Solid Earth*, 123(7):5902–5928, 2018.

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

T. Megies, M. Beyreuther, R. Barsch, L. Krischer, and J. Wassermann. Obspy - what can it do for data centers and observatories? *Annals Geophys.*, 54:47–58, 2011.

Met Office. *Cartopy: a cartographic python library with a Matplotlib interface*. Exeter, Devon, 2010 - 2015. URL `https://scitools.org.uk/cartopy`.

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

M. Owens. *The definitive guide to SQLite*. Apress, 2006.

Y. Ruan, W. Lei, R. Modrak, R. Örsvuran, E. Bozdağ, and J. Tromp. Balancing unevenly distributed data in seismic tomography: a global adjoint tomography example. *Geophysical Journal International*, 219(2):1225–1236, 08 2019. ISSN 0956-540X. doi: 10.1093/gji/ggz356. URL `https://doi.org/10.1093/gji/ggz356`.

S. Simutė, H. Steptoe, L. Cobden, A. Gokhberg, and A. Fichtner. Full-waveform inversion of the japanese islands region. *Journal of Geophysical Research: Solid Earth*, 121(5):3722–3741, 2016.

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.

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, 2020.

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.

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, 2020.

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

J. Virieux and S. Operto. An overview of full waveform inversion in exploration geophysics. *Geophysics*, 74:WCC127–WCC152, 2009.