



LAWRENCE  
LIVERMORE  
NATIONAL  
LABORATORY

LLNL-TR-839591

# Report on the LLNL Global Seismic Waveform Tomography Modeling Project

N. A. Simmons, C. Morency, A. Chiang, S. C.  
Myers

September 6, 2022

## **Disclaimer**

---

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

# Report on the LLNL Global Seismic Waveform Tomography Modeling Project

Nathan Simmons, Christina Morency, Andrea Chiang,  
Steve Myers

September 9, 2022

Lawrence Livermore National Laboratory is operated by Lawrence Livermore National Security, LLC, for the U.S. Department of Energy, National Nuclear Security Administration under Contract DE-AC52-07NA27344.



## Introduction

Earth models have important applications for seismic event monitoring including the location and characterization of potential underground nuclear tests. Seismic waves generated by earthquakes and man-made events are altered by variations in Earth's subsurface properties; and these distortions effect the arrival times of energy packets (phases) and obscure the characteristics of the original energy source (event mechanism). Three-dimensional models of Earth's seismic properties can help explain and predict the distortions to the seismic wavefield and reveal the properties of a source that generated the waves, including the event location using the model-predicted timing of the phases and source mechanism using the model-predicted waveform characteristics.

The US National Laboratories have a long history in the research and development of Earth models designed for monitoring purposes, including 3-D and simpler models. One of the most successful examples is the Regional Seismic Travel Time (RSTT) model developed as a multi-laboratory effort to account for regional travel time bias due to strong shallow upper mantle heterogeneities that lead to seismic event location inaccuracy (Phillips et al., 2007; Steck et al., 2009; Myers et al., 2010; Begnaud et al., 2020). RSTT models includes velocity variations in the crust and lateral variations of the velocity gradients in the shallow upper mantle and are now becoming routinely used by multiple monitoring agencies for accurate event location. Research at the National Laboratories has since evolved into the development of fully 3-D models including several examples on the global-scale including the early LLNL series of models (e.g. Simmons et al., 2012; 2015), the SALSA3D model (e.g. Ballard et al., 2016), and the SPiRaL model (Simmons et al. 2021). Like the earlier models, the SPiRaL model was developed primarily for the purpose of predicting seismic travel times for accurate event location prediction using travel times at regional ( $\sim 200$ -2000km) and teleseismic ( $>\sim 2000$ km) distances. However, SPiRaL differs from the previous models since it was constructed with surface wave dispersion estimates as well as body wave travel times (see the next section for more details).

We now seek to, not only utilize travel times and dispersion to create 3-D models but construct a more comprehensive global model via waveform inversion which exploits more information in the seismograms than arrival times of specific seismic phases. Waveform inversion research is well-established and has led to numerous studies on regional and global scales (e.g., Lekić and B. Romanowicz, 2011; Yuan et al., 2014; French and Romanowicz, 2015; Afansiev et al., 2016; Bozdag et al., 2016; Fichtner et al., 2018; Krischer et al., 2018; Tao et al., 2018; Lei et al., 2020; Örsvuran et al., 2020; van Herwaarden et al., 2021). However, higher-fidelity models are rare, particularly those that can predict complete waveforms with periods of  $<< 30$  seconds over large regions, which limits their applicability to monitoring applications including moment tensor inversion of smaller events.

Exploiting more information contained in complete waveforms, and possibly fitting them “wiggle-for-wiggle”, has a twofold advantage. Firstly, the shape and character of the wave trains contain important information about the underlying Earth structure, and accounting for these characteristics lead to more accurate Earth models. This is especially significant in regions

with sparse coverage as opposed to highly sampled regions where travel time data could appreciably resolve gross Earth structure. Secondly, fitting seismograms wiggle-for-wiggle lead to increased model-based predictability of complete waveforms which has applications in source characterization via 3-D moment tensor inversion. This report describes the progress and plans to improve LLNL's SPiRaL global model via application of full-waveform inversion techniques. An area of emphasis is on the development of a streamlined process to execute the complex and computationally-intensive workflow and apply it to the global problem which presents unique challenges including: the physical size of the Earth model in memory, the amount of data that must be considered, and the ability to continuously improve the model as new events of interest emerge (including smaller events) without increasing the computational burden which can easily snowball into an unmanageable problem.

### The Starting Model (SPiRaL)

The SPiRaL model (Simmons et al., 2021) is a global-scale tomography model of seismic velocities (P and S) and anisotropy variations (vertical transverse isotropy) from the crust to the core. **SPiRaL (S-, P-, Rayleigh, and Love waves)** is a follow-up to the LLNL-G3D series of travel-time-based models, including the LLNL-G3Dv3 P-wave model (Simmons et al., 2012) and the LLNL-G3D-JPS joint P and S wave velocity model (Simmons et al., 2015). One of the key differences from our previous models is the inclusion of surface wave constraints based on global surface wave phase and group velocity maps for Rayleigh and Love waves from Ma & Masters (2014) and Ma et al. (2014). Collectively, these surface wave maps provide 25–200 second dispersion profiles covering the Earth and therefore important additional constraints on upper mantle wave speeds that body wave travel times cannot provide. SPiRaL includes modeled terms that account for 3-D variations of radial anisotropy in the crust and mantle, which is necessitated by simultaneous inversion of Rayleigh and Love waves. Finally, high-resolution regions in the crust and upper mantle were included in SPiRaL beneath North America and Eurasia (minimum node spacing of  $\sim 0.25^\circ$ ) to capture details in these well-sampled regions (see Figure 1).

Although SPiRaL is a *travel-time-based* model, we demonstrated that the model is on par with independent 3-D *waveform-based* models when predicting waveforms for a few key events in the period band of 10–50 seconds in multiple regions (Simmons et al., 2021) including the western United States, northeast Asia, and the Middle East (see Figure 2). An additional study performed under an LLNL Laboratory Directed Research and Development (LDRD) project showed that the *travel-time-based* SPiRaL model is at least on par with regional *waveform-based* models for predicting waveforms in the western United States (Rodgers et al., 2022). These findings demonstrate that the SPiRaL global model is a valuable starting model for full-waveform inversion which exploits seismograms wiggle-for-wiggle rather than only considering the onset times of arriving seismic phases which is done in travel-time-based tomography methods.

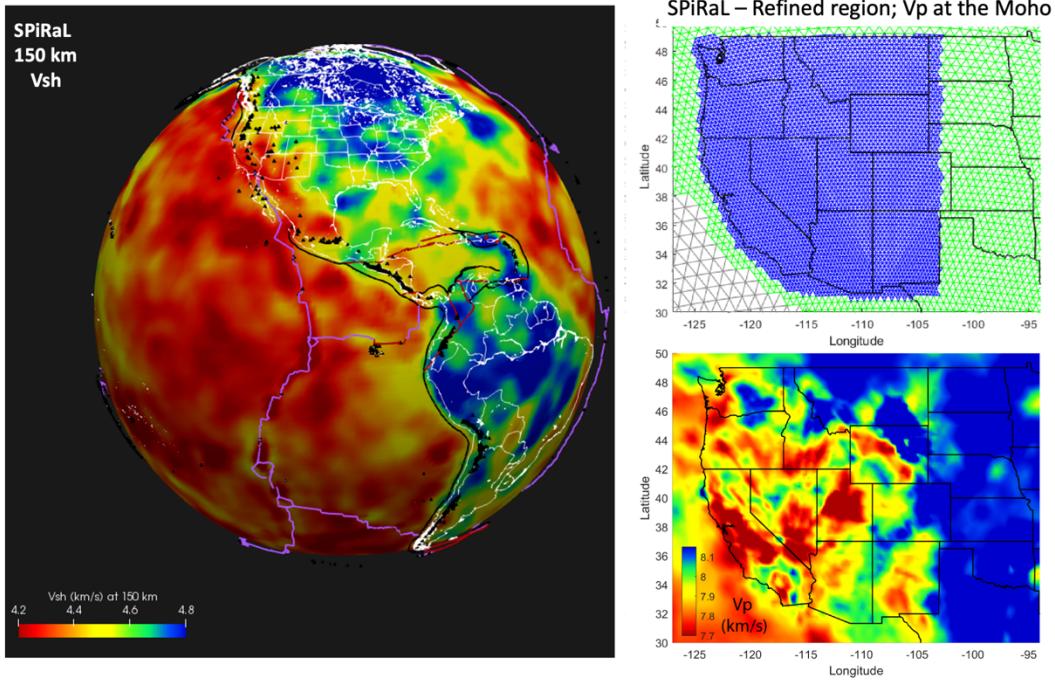


Figure 1. The SPIRAL global multi-resolution seismic model. The left panel shows a depth slice of the horizontally polarized shear wave speeds ( $V_{sh}$ ). The panels on the right show the multi-resolution grid in the western United States (top) and the compressional wavespeed ( $V_p$ ) in the shallow upper mantle (bottom).

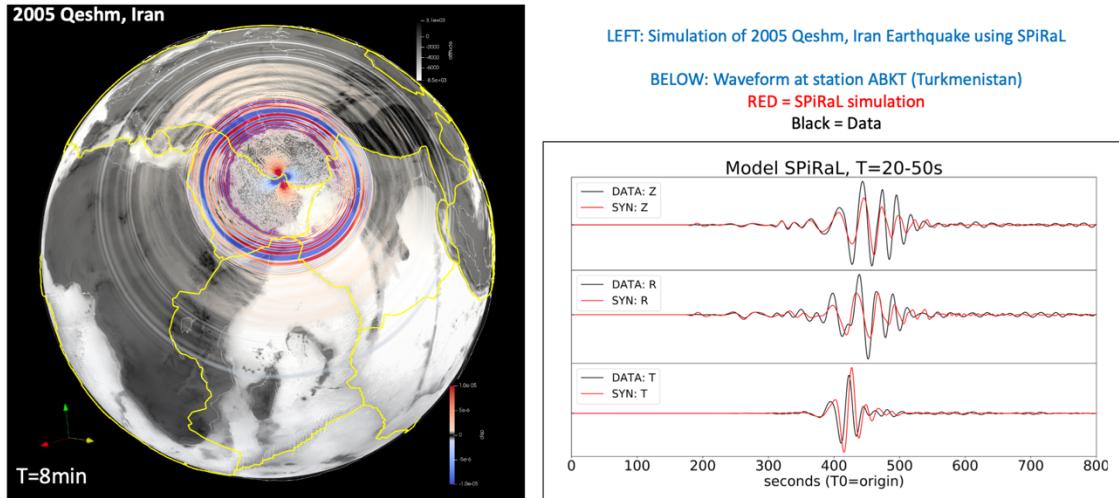


Figure 2. Seismic wavefield simulations using the SPIRAL travel-time-based model for an earthquake occurring beneath Qeshm Island, Iran in 2005. The figure on the left shows a snapshot of the displacement wavefield  $\sim 8$  minutes after the event and the figure on the right shows the simulations (red lines) compared to data (black lines) at station ABKT in Turkmenistan. These waveform comparisons were used in Simmons et al. (2021) to benchmark SPIRAL with independent waveform-based models.

## FWI Process and Workflow Development

Full-waveform inversion (FWI) is via the adjoint method is a well-established method used for seismic tomography of Earth's subsurface. See Tarantola 1984, Tromp et al. 2005, and Tape et al. 2007 for some key papers that developed and described the adjoint process and application methods. The adjoint method involves 3-D seismic wavefield simulations in a forward sense (forward through time from seismic sources to a set of receivers) and reversed-time simulations of “adjoint sources” propagating from the receivers back to the sources. The “adjoint sources” are derived by some measurement of the misfit between the predicted wavefield and the simulations. The interaction of the forward and adjoint wavefields integrated over time effectively projects the misfit into the model volume and model sensitivity and needed model changes are derived from this interaction (Figure 3). See the textbook by Fichtner (2011) for a comprehensive examination of full waveform modeling methods including the adjoint inversion method.

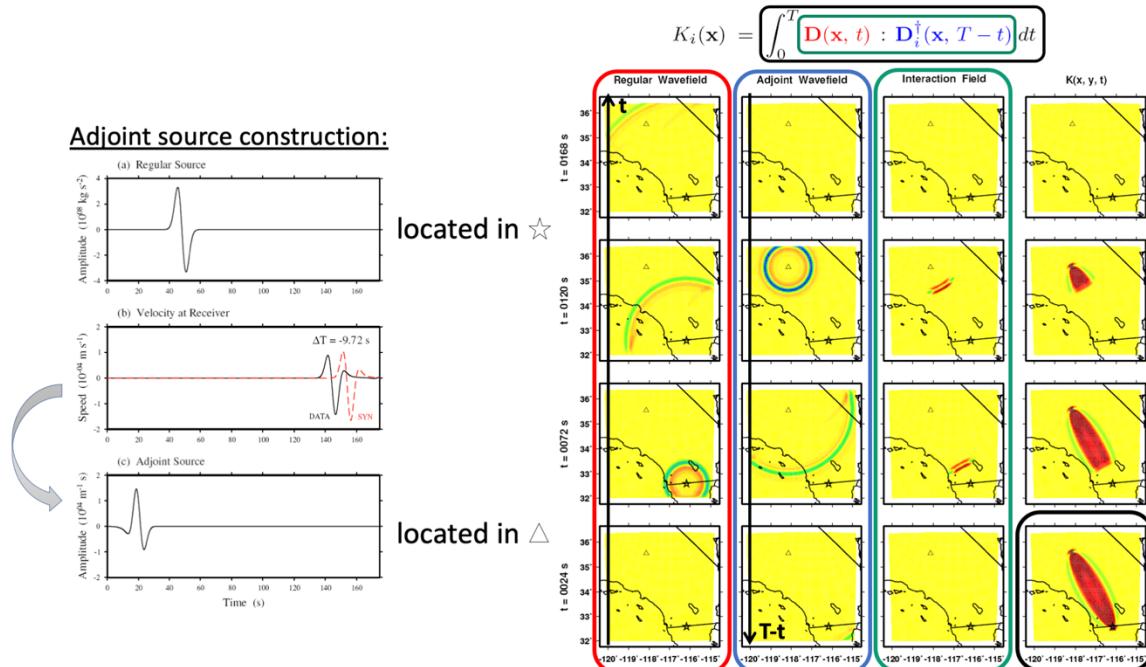


Figure 3. The construction of adjoint sources and model sensitivity kernels (adapted from Tape et al. 2007). (LEFT) A simple seismic source function (top panel) is propagated to a seismic receiver and compared to the data (middle panel). The difference between the observed (DATA) and predicted (SYN) waveforms is taken and time-reversed to construct the adjoint source (bottom panel). (RIGHT) The regular wavefield (1<sup>st</sup> column) and adjoint wavefield (2<sup>nd</sup> column) are pointwise multiplied at discrete time steps to create the interaction wavefield (3<sup>rd</sup> column). The sensitivity kernels are developed through time integrating the interaction wavefield (4<sup>th</sup> column).

The adjoint inversion method is computationally intensive due to the necessary full wavefield simulations necessary and requires supercomputing in many 3-D studies, especially a global-scale study undertaken here. The inversion process is highly non-linear and therefore requires multiple iterations (and simulations) further amplifying the computational requirements.

Ultimately, millions of CPU-hours are often needed for larger problems including global-scale inversions.

In addition to the computational burden, a complex iterative workflow is required to properly develop misfits and stabilize the inversion for convergence (see Figure 4). Specifically, sophisticated data windowing techniques must be applied to select time windows where the data and synthetics are similar enough to develop a reliable misfit function that will ultimately lead to reliable model kernels. These methods include FLEXWIN (Maggi et al. 2009) and python-based version called Pyflex (Krischer and Casarotti, 2015). After the raw kernel construction, pre-conditioning must be applied to the kernels so that the model changes are not too complex and sensitivities are not focused only at the source and receivers. Kernel (gradient) pre-conditioning includes masking the source region and applying the diagonal of the Hessian matrix for each event kernel which prevents the projection of the misfit to very complex source effects which are likely not reasonable or resolved (see Modrak and Tromp, 2016). Summation of event kernels and smoothing further simplifies the resulting model gradient leading to smooth model changes at each step (see Figure 5 for a synthetic example).

Our current custom workflow begins with the base SeisFlows seismic waveform inversion workflow framework (Modrak et al. 2018). While SeisFlows provides a basic framework, several significant modifications have been made including the elements discussed above and diagrammed in Figure 4. Modifications also include changes in SLURM batch job messaging across compute nodes and the development of a custom SPECFEM3D\_Globe (Komatitsch and Tromp, 2002a, 2002b; Komatitsch et al., 2015) interface to properly run forward/adjoint simulations, construct kernels, and execute a kernel smoothing algorithm. We initially used a Gaussian kernel smoother that is applied across all of the elements; however, the Gaussian smoother required to much computational resources. For example, the global mesh can take >24 hours using >300 processors to apply the Gaussian-based smoother. One way to reduce the compute time is to use a GPU-accelerated implementation (rather than using CPU-based systems), however we decided to integrate and use a Laplacian-based smoother since our tests concluded that it could be applied at more than an order magnitude faster than the Gaussian-based version on a more standard CPU-based compute platform.

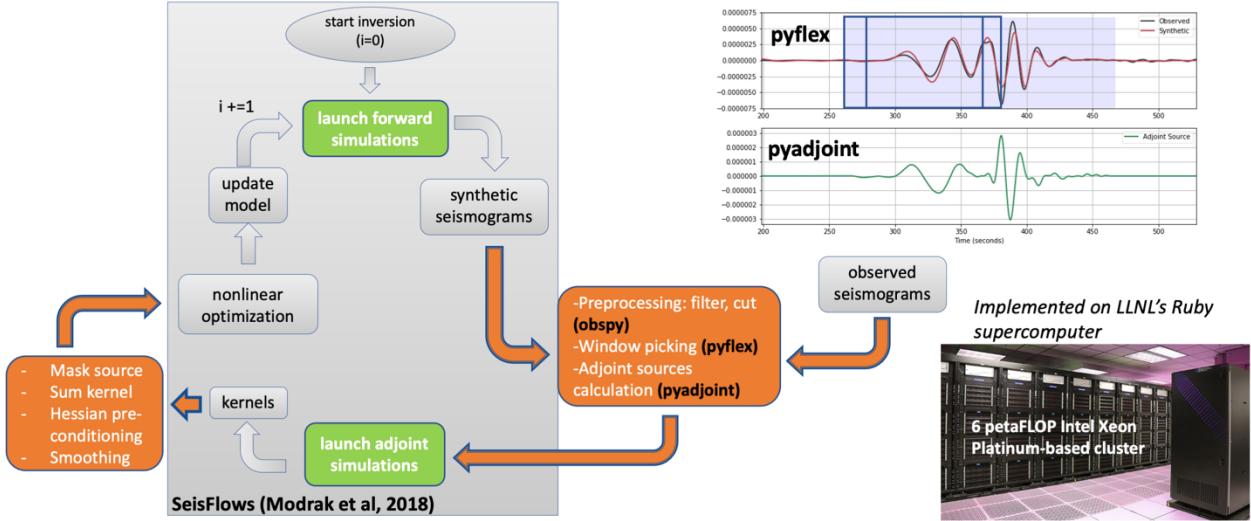


Figure 4. Our custom adjoint FWI workflow is a significant modification to SeisFlows (Modrak et al. 2018). Some of the integrated features include window picking and adjoint source calculations with Pyflex (ref) and Pyadjoint (ref), a custom SPECFEM3D\_Globe interface and modifications required for execution on the LLNL Ruby supercomputer.

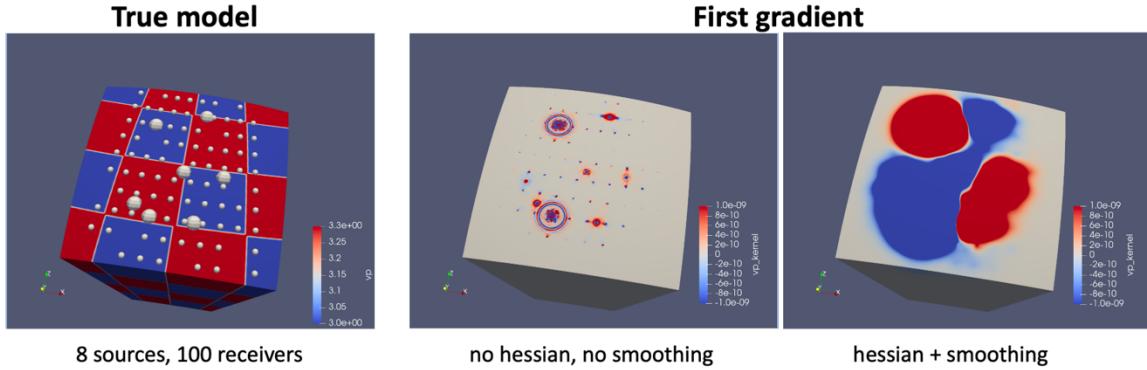


Figure 5. Workflow testing using a single chunk in the SPECFEM3D\_Globe. (LEFT) A synthetic checkerboard model with 7 sources (large spheres) and 82 receivers (small spheres). (CENTER) The first model gradient calculation without pre-conditioning. (RIGHT) The first model gradient with Hessian pre-conditioning and smoothing. Note that the model will change in the opposite direction of the gradient (negative sign applied).

Our modified waveform inversion workflow has been tested with synthetic setups using SPECFEM2D, SPECFEM3D (Cartesian) and SPECFEM3D\_Globe (cubed-sphere) with a single chunk (1 part of the 6 sided cubed-sphere). Each of the tests have increased in size in both the scale of the models and the number of stations and synthetically derived data. Working with real data and the full-scale global model presents further challenges. Compared to synthetically derived data, real seismic data are notoriously more difficult to handle for numerous reasons. For example, real seismic data are often plagued with issues including metadata errors, missing segments, various noise conditions at different stations, and instabilities when deconvolving the instrument response. Data issues can be handled with thorough pre-processing steps which we have developed and refined, but other issues arise in the real problem due to the increased amount of data, simulation time, and the size of the data files that must be created and

modified. Here is a list of some key factors involved with the global-scale problem that present some unique challenges:

- 1) All seismic events are potentially relevant (enormous amount of available data, even with careful event selection)
- 2) All seismic stations are potentially relevant (>1,000 available stations globally in the most recent years)
- 3) Thorough pre-processing of large datasets is required to identify and handle corrupt or incomplete data streams
- 4) Signals are easily seen at large distances requiring long time series simulations (~1 hour long records)
- 5) The physical memory size of the spectral element model is large (A model mesh with the shortest resolved period of 30 seconds requires ~1.2 M elements and ~80 M model points)

Moreover, in the case of the SPiRaL model, there are 5 parameters at each point (vertical compressional and shear wave speeds with anisotropy) giving ~400 M total model property values for the whole Earth. The physical memory for this model and mesh properties is ~14 Gbytes. For a spectral element mesh applicable for a minimum resolved period of 30 seconds, a single forward simulation requires ~200 CPU-hours to generate 1-hour long seismic records. Therefore, to make the problem tractable, hundreds of processors are required for each event simulation (our current setup uses 216 CPUs and 4-6 compute nodes for each simulation on our current computing platforms). The run time for the full inversion workflow can easily exceed 24 hours for a single iteration, depending on how many events are considered in a single run. This run time can easily approach maximum allowable run times ("walltime") without priority or special access to the computer systems which can lead to shutdown of the inversion workflow prior to completion. One way to deal with this issue is to select the right number of events that can be considered in a single run while providing the coverage needed to construct gradients that create reasonable model updates. A staged approach to consider only batches of events in sequence is attractive as discussed in a section below. Finally, the amount of data that must be read and written (I/O) requires stable file systems with sufficient communication across compute nodes to limit run failures and data loss. We have encountered this latter issue on multiple occasions and are seeking ways to mitigate the problems.

## Large-Magnitude Event Dataset

For regional or local tomographic studies, we are often seeking all available data to provide enough coverage of the study area (see Rodgers et al. 2022 for a recent example) and source-receiver paths outside of the regions are justifiably omitted from the studies. For global-scale tomographic studies however, *all* seismic data recorded at instruments around the globe are potentially useful for constraining structure. Therefore, any seismic event that has ever occurred is potentially relevant to the study. Obviously, it is not feasible to consider all data and an event selection process must narrow down to only the most valuable and reliable data. Larger events tend to provide the most valuable seismic recordings for global studies and this section development of this primary dataset. Large events are most valuable for three main reasons:

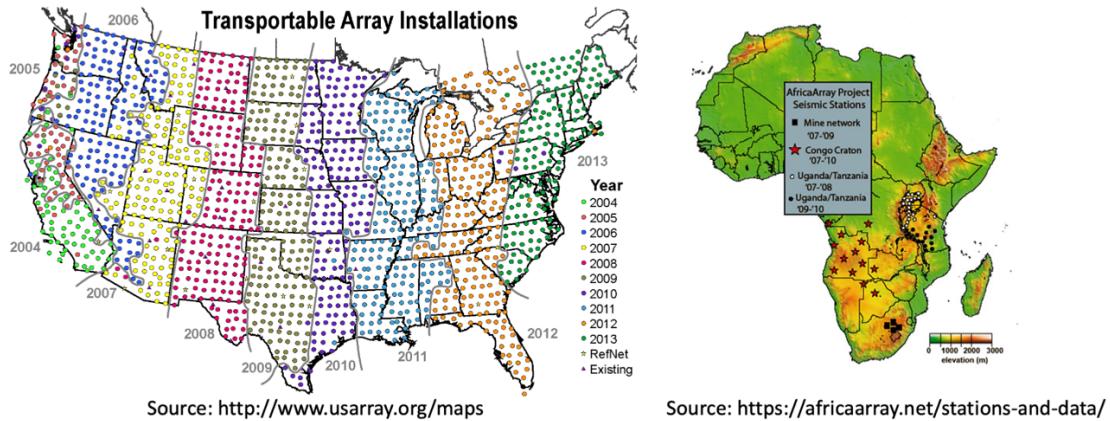
1. they generate strong long-period signals (>>20 seconds) that help resolve deep structure,
2. signals can be easily observed above the noise level at very long distances (>10,000 km),
3. source parameters derived from moment tensor inversion are reliable

Therefore, the magnitude of the events provides the first criteria for down-selecting to develop a large-event global waveform dataset.

The Global Centroid Moment Tensor project (GCMT, <https://www.globalcmt.org>) was established in 1982 by Adam Dziewonski and others at Harvard University to determine source parameters (moment tensors) for earthquakes occurring from 1976 to the present, including rapid determinations for events  $M > 5.5$ . See Ekström et al. (2012) for technical details about the project. The GCMT catalog consists of >25,000 events and moment tensor solutions and forms the basis for our large-event global event selection. We further reduced the GCMT catalog to only include events in the magnitude range of  $M = 5.8$ -6.8. The lower threshold ( $M = 5.8$ ) was selected based on our experience in evaluating waveforms for events with  $M < 5.8$  for waveform cross-correlation measurements for travel-time-based tomography (e.g. Grand 2002) where we found that several events in this range did not always yield the highest-quality long-period seismograms at very far distances (>10,000 km). Moreover, limiting to these larger events for the large-event dataset reduces the number of events to a more manageable set. The higher threshold ( $M = 6.8$ ) was selected for a very different reason. Specifically, events approaching  $M = 7$  and larger often have very complex sources due to the large earthquake rupture zones (long rupture planes) and long rupture durations. The complexity of these very large sources introduces directivity and source-time-function issues that require special care for wavefield simulations and comparisons to data. Limiting the magnitude to  $M \sim 6.8$  avoids many of these complexities.

Seismic instrumentation and the availability of public broadband data has dramatically increased over the past few decades. There are often data available for >1,500 seismic stations around the world provided to the public by the Incorporated Research Institutions for Seismology (IRIS) Data Management Center (DMC), which is our source for the large-event

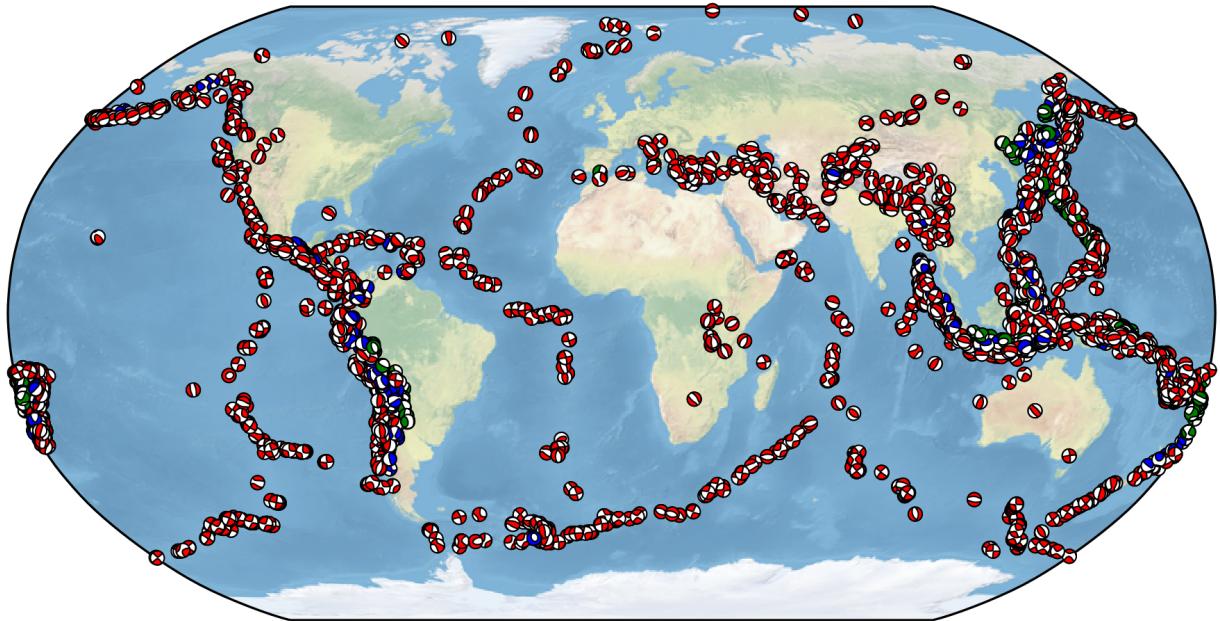
dataset compiled in this study. However, going back to the mid 1990's, the number of seismic stations with available data from IRIS drops to <200 and this number drops even further prior to the 1990's. For this study and other global waveform tomography studies, there is a lot of value added for events with the most available recordings (i.e. stations) since every modeled events requires multiple 3-D wavefield simulations which are computationally expensive.



*Figure 6. Examples of temporary broadband seismic station deployments and schedules. (left) The USAArray Transportable Array deployment in the United States during the 2004-2013 timeframe. (right) The temporary deployments of the AfricaArray deployed in portions of Africa during the 2007-2010 timeframe. Note that is only a portion of the important deployments within the AfricaArray.*

Based on the point made above, it may seem reasonable to only consider the most recent events given the increase of available of seismic data through the years. However, there are two important reasons to consider events spanning at least several decades. Firstly, many seismic station deployments are temporary and therefore the global station configurations change each year providing unique coverage of the planet. One of the best examples is the USAArray Transportable Array (TA) funded by the US National Science Foundation that marched hundreds of stations across the United States from 2004-2013 (Figure 6). However, there are several other great examples including the temporary deployments in Africa under the AfricaArray project that cover this under sampled part of the Earth below the African continent (AfricaArray was operated by the University of Witwatersrand in South Africa, the South African Geological Survey, and Pennsylvania State University). The second important reason to consider events spanning a few decades is that earlier events can fill spatial gaps in event coverage not provided by the more recent events. These include rarer events that may occur on mostly aseismic portions of mid-ocean ridges or rare intracontinental events occurring in regions without major tectonic plate margins.

To balance the available number of stations in more recent years and the coverage provided by temporary deployments in prior years we chose to consider events in the years spanning 1997-2021. This down-selection of the GCMT catalog (from 1976-present to 1997-2021) reduces the event pool for the large-event dataset down to 4,526 events (Figure 7).



*Figure 7. Locations and moment tensor solutions (plotted as “beach balls”) for all GCMT catalog events with magnitude range of  $M=5.8-6.8$  in the years spanning 1997-2021. The catalog is reduced from  $>25,000$  events down to 4,526 events in this plot. The color of the beach balls corresponds to the depth of the events: red=0-100 km, blue=100-400 km, green=400+ km.*

The locations of the events plotted in Figure 7 are often very similar for clusters of events in the high-seismicity regions, most notably along the Pacific Ocean margins. To reduce the redundancy and the number of events that must be simulated-modeled, we further down-sampled the events based on interspatial distances. To do this, we first selected the last event occurring in 2021 from the pool of events shown in Figure 7 and went backwards in time selected the next event that was at least  $1^\circ$  from any previous event ( $\sim 111$  km separation required). Shallow events (in the crust and shallow upper mantle) often spatially overlap events in the deeper upper mantle and transition zone ( $\sim 400-700$  km). The deep upper mantle and transition zone events are much rarer than crustal and shallow upper mantle events and provide important coverage of subducted slabs (sunken tectonic plates). To guarantee selection of deeper events, we performed the  $1^\circ$  sampling procedure described above in multiple event depth bins including:

- 1) 0-50 km (~crust)
- 2) 50-150 km (shallow upper mantle)
- 3) 150-400 km (deep upper mantle – rare)
- 4) 400-500 km (shallow transition zone)
- 5) 500-600 km (middle transition zone)
- 6) 600+ km (deep transition zone)

This down-selection process reduces the number events from 4,526 (in Figure 7) to 1,198 events with a very similar spatial coverage which can be seen in Figure 8. Note that the number

of selected events (1,198) is comparable to the number of events considered in the development of the GLAD\_M25 global adjoint waveform model (1,480) (Lei et al. 2020).

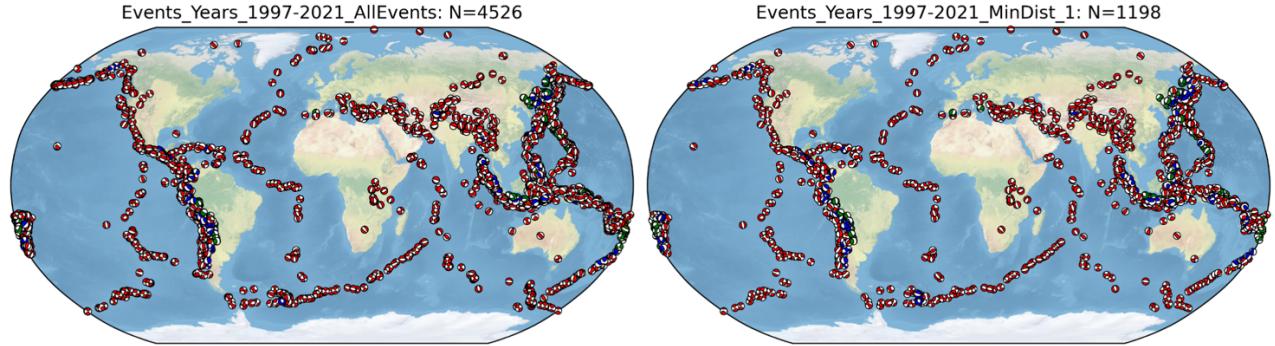


Figure 8. Comparison of the initial selection of 4,526 GCMT events with  $M=5.8-6.8$  between the years 1997-2021 (left) with the spatial down-sampling to 1,198 events (right). The down-selected event coverage is very similar to the initial selection since redundant events are effectively removed from the event set. Events in regions with lower large-event seismicity remain in the down-selection while events in high-seismicity regions are thinned. The color of the beach balls corresponds to the depth of the events: red=0-100 km, blue=100-400 km, green=400+ km.

Our down-selection process steps backwards in time from 2021 to 1997. This selection process results in more events in the later years and fewer events per year as we go backwards in time to 1997. The later events are the most well-recorded (more available seismic stations) while the earlier events take advantage of prior station configurations while also filling in spatial gaps in regions with less seismicity and perhaps very rare large events. The distribution of the 1,198 down-selected events by year is shown in Figure 9 and some example station configurations by year are shown in Figure 10.

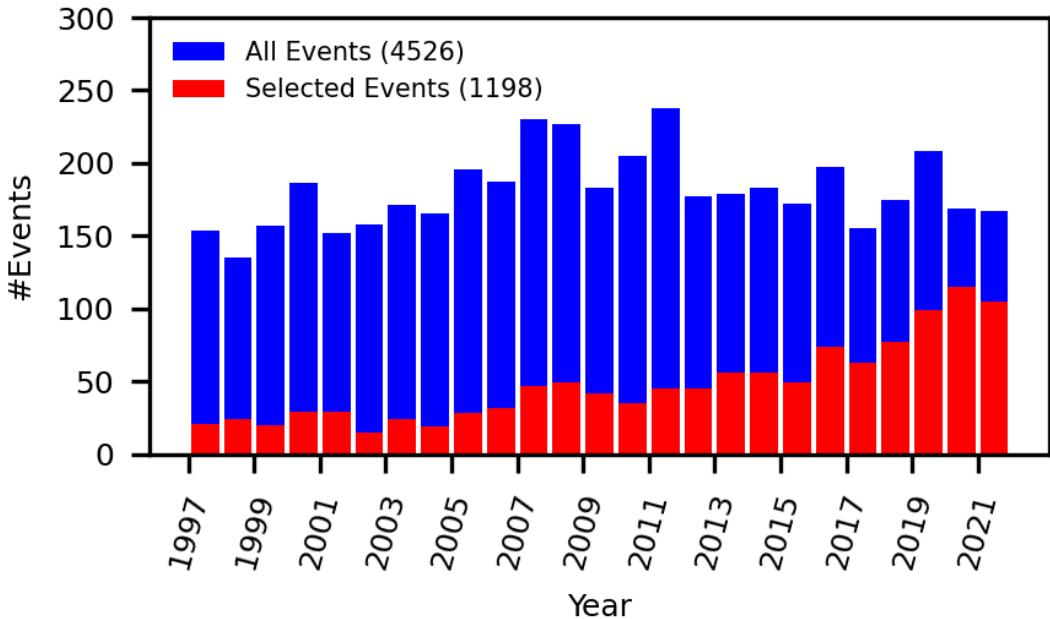


Figure 9. The number events selected by year from GCMT catalog. The blue bars show the number of events per year in the GCMT catalog with the magnitude range of  $M=5.8-6.8$  (4,526 events;  $\sim 181$  events per year). The red bars show the number of events per year resulting from the down-selection stepping backwards from 2021 to 1997 (1,198 events;  $\sim 48$  events per year on average).

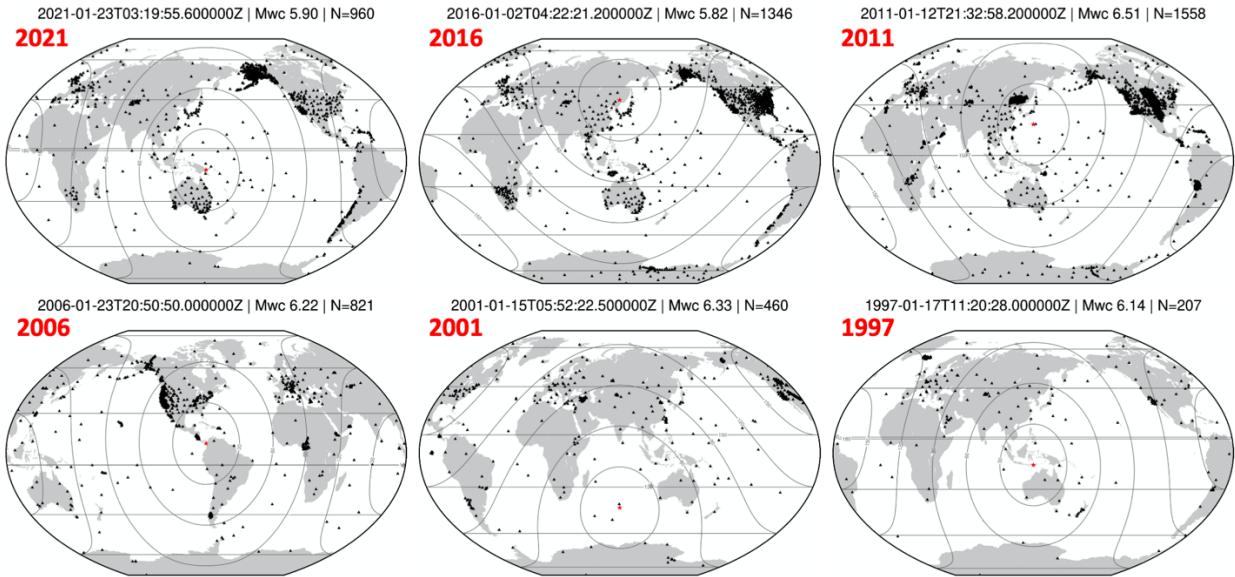


Figure 10. Example station configurations for randomly selected events from the 1,198 large-event ( $M=5.8\text{-}6.8$ ) dataset. Stations with available public data from the IRIS DMC are plotted for the selected events spanning from 1997 to 2021. The red labels identify the year the event occurred, the red dots are the event locations, the black lines are equidistance rings from the event, the black dots are station locations, and the value of “N” in the title of each frame indicates the number of stations for that event. Note the progression of the USArray, the temporary nature of several arrays, and the overall decline of the number of available stations going backwards in time. Prior to 1997, the number of stations drops below 200.

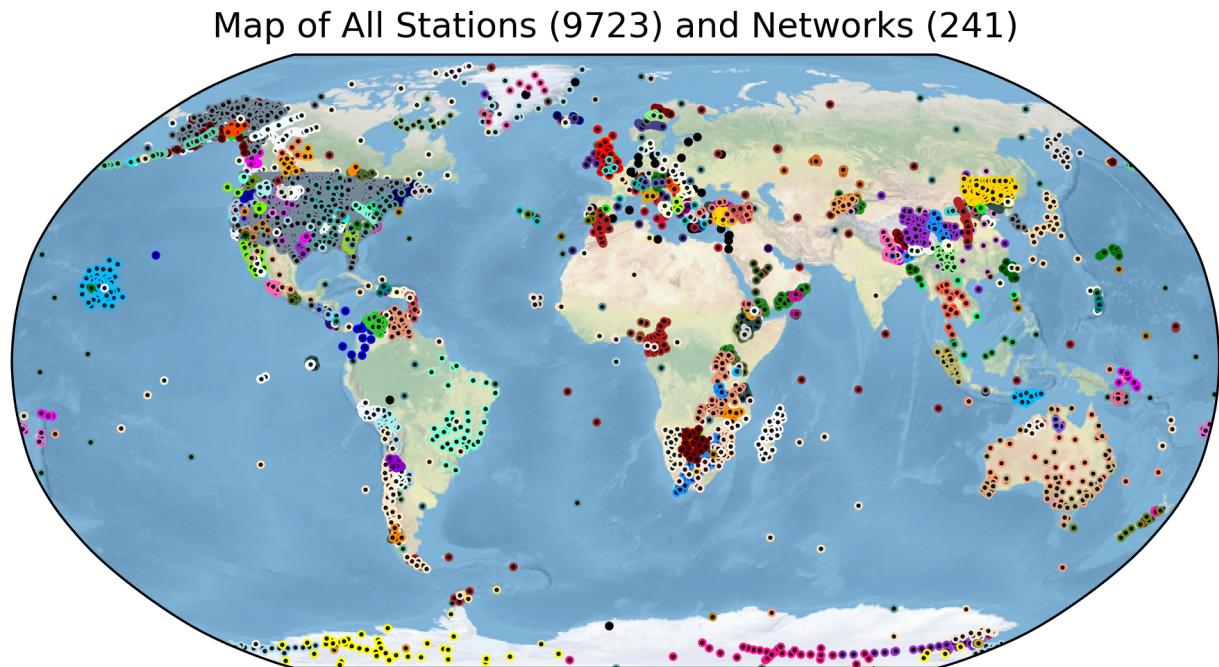


Figure 11. Map of stations (9,723) and networks (241) with available data from the IRIS DMC for the 1,198 selected events spanning the years 1997–2021. Although, the oceans are largely under-instrumented, the collection of networks/stations deployed in that time frame provide reasonable coverage of the planet.

The network configurations vary by year, largely due to the transient nature of temporary seismic deployments. If we were to consider only events within a narrow timespan, constraints on Earth structure would be limited to available stations in that timeframe (see Figure 10). However, by considering events spanning multiple decades, we can take advantage of more stations and provide a much more comprehensive coverage of the planet. Figure 11 illustrates the stations (9,723) and seismic networks (241) with available public data from the IRIS DMC for the 1,198 selected large-event global data set.

We downloaded all available public data for the 1,198 events via the IRIS DMC/FDSN services using customized Python/Obspy utilities. The data includes all broadband channels (BHE, BHN, BHZ, BH1, BH2) for all stations for the entire planet (0-180° distance). The data consists of >4.1 million seismic traces and we included time windows sufficient to record the event for the given event-station distance (time series range from 700 to 10,000 seconds). The data were collected as native miniseed files and converted to Adaptable Seismic Data Format (ASDF). Data for each event were subsequently archived into event-based files using Hierarchical Data Format (HDF5) containers for convenience. Finally, the HDF5 archive data were pre-processed using custom parallel processing scripts includes time synching, clipping, down-sampling to 10 samples/second, and removing the instrument response to create displacement seismograms (see Figure 12 for an example). The final storage size of the global database is ~1.1 terabytes and summary is provided in Table 1.

Note that we recognize that smaller events are significant as well, particularly in regions without large events (M>5.8), and we intend on developing an additional “Special Event” dataset which include large events of interest that were not randomly selected as well as small events in some low-seismicity regions including the eastern United States.

**Table 1. Large-event dataset summary statistics**

Property	Value
<b>Number of events</b>	<b>1198</b>
<b>Number of seismic traces</b>	<b>~4.1 million</b>
<b>Event-Station distances</b>	<b>0 to 180 degrees</b>
<b>Time series lengths</b>	<b>700 to 10,000 seconds</b>
<b>Size of dataset</b>	<b>~1.1 terabytes</b>
<b>Dataset format</b>	<b>ASDF format stored in HDF5 containers</b>
<b>Number of stations per event</b>	<b>169 to 1,804</b> <b>Average = 1,148</b>
<b>Total number of unique stations</b>	<b>9,723</b>
<b>Total number of networks</b>	<b>241</b>

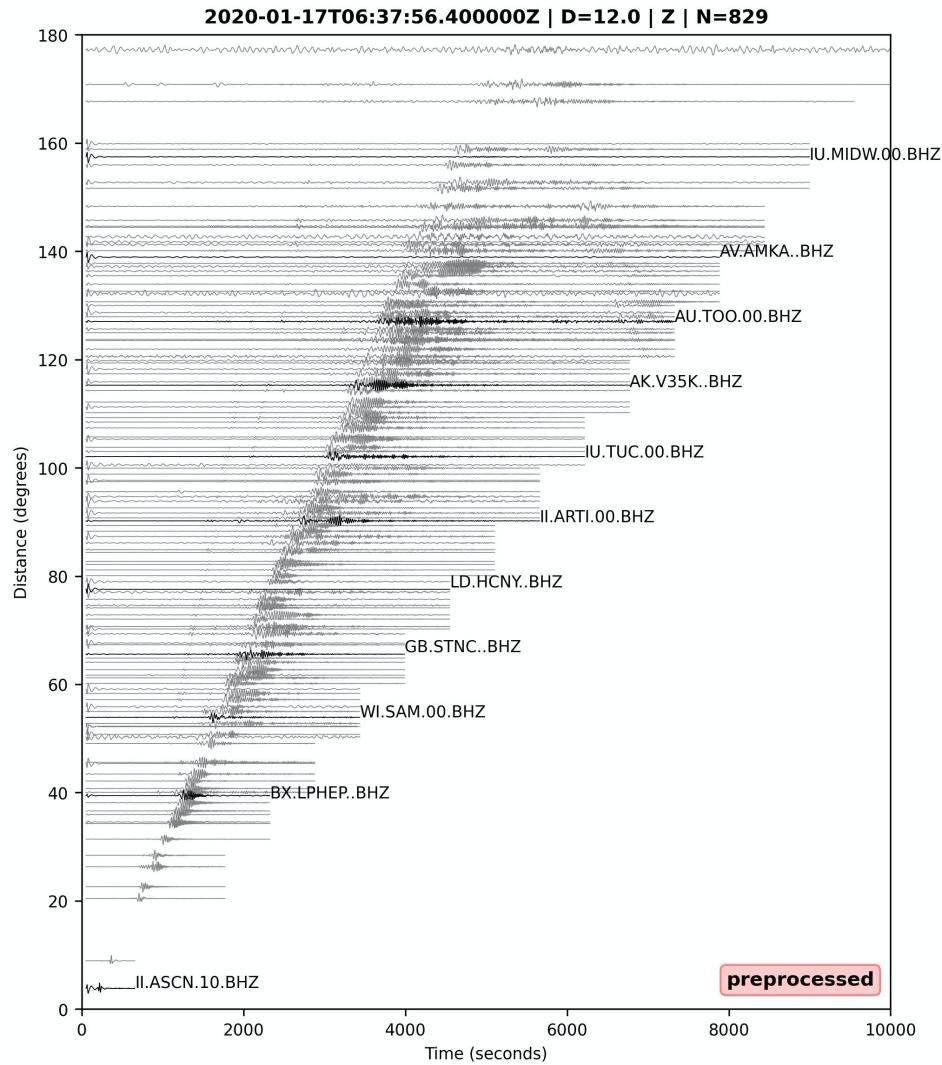


Figure 12. Example of seismograms for an event occurring in January of 2020 beneath the Ascension Islands in the South Atlantic Ocean. Network and stations names are labeled next to several of the traces and the length of the stored time series depends on the event-station distance. This example shows vertical component broadband seismograms (BHZ channel) and the data are pre-processed with standard seismic signal processing steps to synchronize, down-sample, instrument response removal, etc. The pre-processed displacement seismograms are stored in ASDF format in HDF5 archives (1 archive per event).

#### Staged Approach for Model Updating

The vast amount of data that are required for the global waveform tomography problem presents several significant challenges that differ from a travel-time-based approach, including with respect to the relatively complex approach used to construct the SPiRaL model. As described in a previous section, the workflow to update a model using full waveforms involves the orchestration of several processes including numerous forward and adjoint 3-D wavefield simulations. These simulations are very computationally expensive, especially in the global case where simulations must be run for up to 1-hour simulations (1-hr time series). Therefore, we

reduce the global set of events through a down-selection process to make global waveform inversion a more tractable problem. However, the number of required events and the amount of data needed to ingest and analyze (compute data misfits and adjoint sources) remains substantial (1,198 large events and 1.1 terabytes of data in the current study).

Given the computational burden, an attractive approach is to construct global waveform models by increasing the number of events during the progression of the model construction or using subsets of an pre-determined event pool. A recent example of the former method (increased number of events during the progression) is in the development of the GLAD\_M25 model (Lei et al. 2020). The development of GLAD\_M25 extends to development of a prior model GLAD\_M15 (Bozdağ et al. 2016) which served as starting model. GLAD\_M15 was developed using 253 earthquakes and the first iterations of GLAD\_M25 were derived using an expanded set of 520 global events. The event set was subsequently expanded through the iterations from 520 to 1,040 events, eventually finishing with 1,480 events as the model began to become more refined. Note that this final set of events used in GLAD\_M25 is comparable to our large-event data set selection (1,198 events).

Another recent approach to deal with large waveform datasets and computational issues is a concept of “dynamic mini-batches” (van Herwaarden et al., 2021). In this approach, subsets of sources are selected in a quasi-random fashion from a pool of events. The selection of events at each stage is designed to reduce redundancy of the information based on the properties of the model update gradients that each event would provide. Overlap between sets of events from one stage to the next is achieved by using a control group of events from a previous optimization iteration, providing a smooth transition to the next mini-batch set of events used to construct the new model gradient and update. Importantly, a major benefit of this approach is that new events can be introduced into the inversion without inverting the entire dataset.

Our proposed approach is to consider event subsets and is most similar to the mini-batch approach described above (Figure 13). Subsets of 100-200 events for the pool of 1,198 large events will be considered at each optimization stage. After 3-5 iterations have been performed for a given subset, the final model gradient will be stored. Initiation of the next stage (with a new subset of events) will take an average of the prior and new model gradients (computed with the new subset) to initiate the next stage of the optimization. This process will provide some overlap from one event subset to the next, similar to the mini-batch approach. Similar to the mini-batch approach, our proposed approach will also allow insertion of new events of interest at any stage of the inversion. These may include new large events of interest or smaller events that are of regional importance.

The periods of the seismograms will decrease from a minimum period of 40 seconds until convergence (Phase 1), then the minimum period will be reduced to 30 seconds in Phase 2. The tentative plan for Phase 3 is to return to a SPiRaL-like travel time inversion to ensure onset times (travel times) are still satisfactorily matched. However, it may be determined that alternative events of interest or regional refinements with additional waveform data is more attractive research pursuit for Phase 3.

Recall that the 1,198 large events were selected across a 25-year timespan (1997-2021). Given the natural recurrence of earthquakes, we would expect that if we selected all events occurring in a given month (across all years), that we would get roughly equal number events occurring in each month ( $1,198/12 \approx 100$ ). We indeed find  $\approx 100$  events per month selected for the large-event set (Figure 14). The breakdown of events by month provides a convenient way to create event subsets since each month provides the target number of events in our proposed FWI Data Scheme (Figure 13) and each subset is a quasi-random set of events spanning all years (takes advantage of all network configurations).

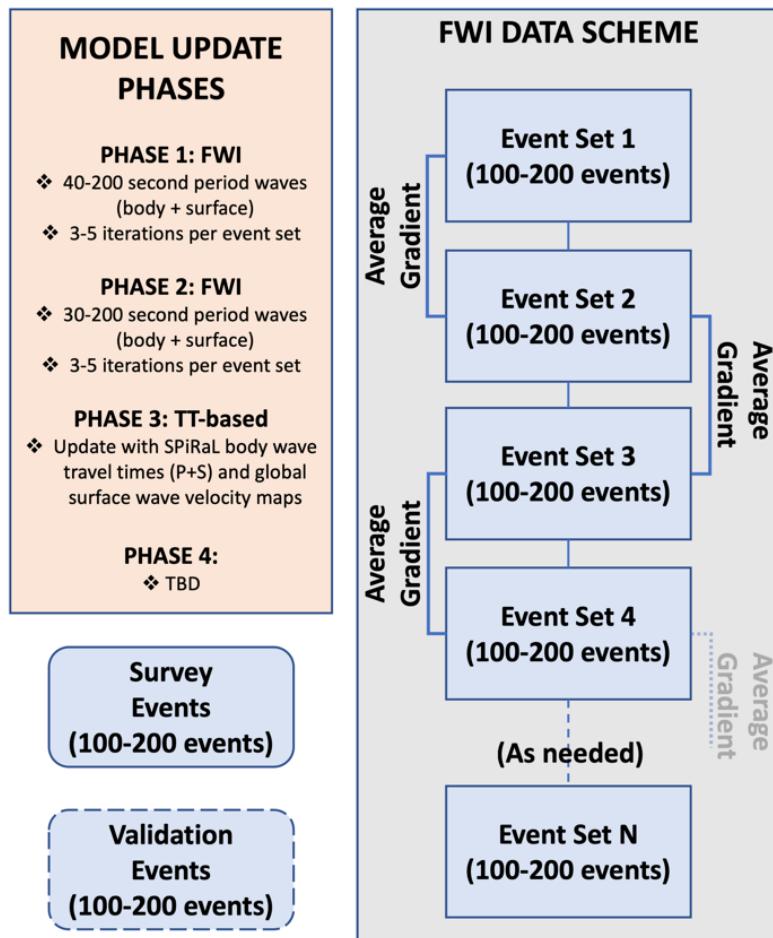


Figure 13. Our approach is to update the SPiRaL model using a progression of subset events/data for updating. Subsets of events (100-200 events) will be considered at each stage, and then proceeding to the next by averaging model gradients determined from the prior and new subsets. A standard approach of reducing the minimum period will follow in Phases. A “Survey” and “Validation” set will serve to evaluate the model. See the text for more details.

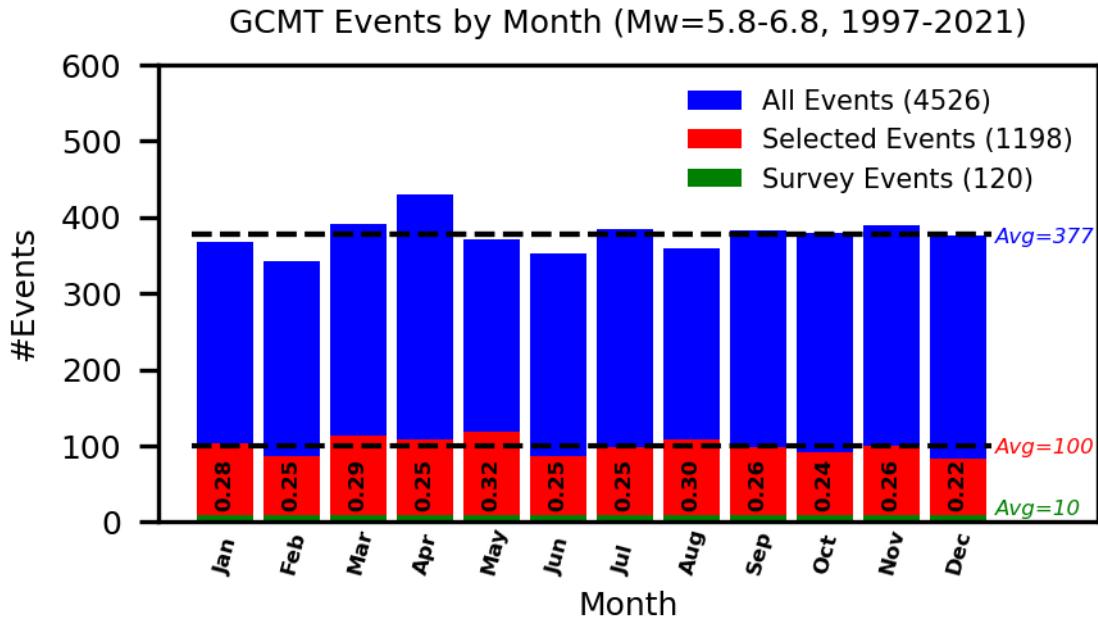
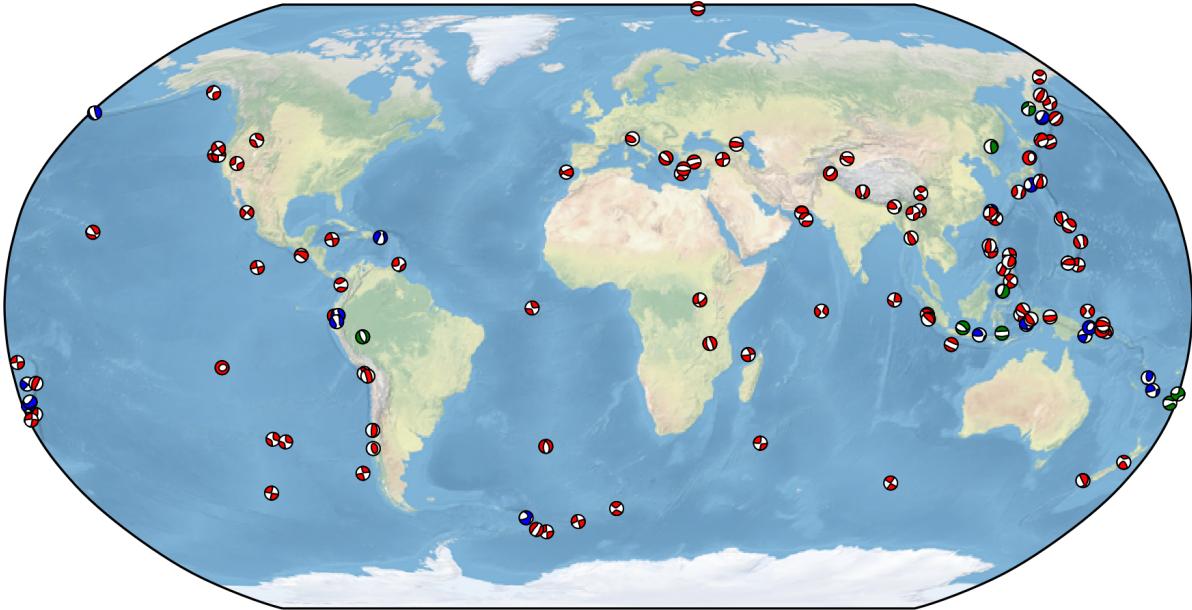


Figure 14. Subsets of the large-event data set broken down by month. The entire pool of events with  $M=5.8-6.8$  from 1997-2021 available in the GCMT catalog are shown in the blue bars. The red bars show the number of events in the large-event set of selected events (1,198 in total) broken down by month (spanning all years between 1997-2021). The numbers on the red bars show the fraction of the total (All/Selected). The green bars illustrate the 10 randomly selected events from each monthly subset to serve as the “survey” set.

In addition to the monthly subsets, we also define a set of “survey” events for evaluation of the model progression at various stages via forward simulations of the current model and comparison to data. The selection of survey events was done by randomly sampling 10 events from each monthly subset (see Figure 14). This sampling produces a global set of random events with a suitable global distribution (Figure 15). Unlike validation events, some survey events will be included in the inversion at each stage. Evaluating the improvement of fit to the waveforms for these events will reveal, not only how the updated model fits the data that went into the inversion, but how well the model will be predicting events that have yet to be included in the process. For example, if our first subset of  $\sim 100$  events used in the inversion consists only of events occurring in January, the fit to the 10 events that were included in the inversion will be evaluated as well as the remaining 110 events from February-December that were not. As the inversion progression moves to the February subset of events, we will also reveal how the fit changes to the January events that are now excluded. Therefore, the survey event set evaluation will illustrate how the fit evolves as we cycle through the data and give us an indication of overall convergence.



*Figure 15. The locations and moment tensors (beach balls) of the random selection of large ( $M=5.8-6.8$ ) events to serve as the “survey” event set. Ten events were randomly selected from each of the monthly subsets (spanning the years 1997-2021) to evaluate the fit progression of the updated model at each stage of the inversion. The color of the beach balls corresponds to the depth of the events: red=0-100 km, blue=100-400 km, green=400+ km.*

### Conclusions/Next Steps

Global scale seismic tomography modeling at LLNL has evolved from the travel-time-based SPiRaL model (Simmons et al. 2021) to full-waveform inversion to inform the SPiRaL model by exploiting more information contained in waveforms beyond the arrival times of specific seismic phases. The expectation is that a singular model can predict accurate travel times as well as waveforms for monitoring applications including event location and 3-D model-based moment tensor inversion. Regional-scale models and data calibration methods are potentially quite valuable for these purposes in limited areas with substantial empirical information. However, advanced global scale models are desirable for providing model-based predictions at long distances, in understudied regions, remote regions around the globe, and with self-consistency using a singular model.

The next phase of this research will involve execution of our custom workflow with a small subset of events to mitigate bottlenecks and optimize the process for scale-up to at least several hundred seismic events ( $>>100,000$  seismograms considered at a single stage). In parallel, we plan to statistically evaluate the waveform prediction accuracy of current global scale models including SPiRaL. Previous results suggest SPiRaL is on par with some models for predicting waveforms in multiple regions, but this is not yet confirmed on the global scale. This will help determine the appropriate minimum waveform periods to target in the early stages of the inversion workflow with the full dataset. The following step will be to launch the workflow with the large datasets and execute the staged model updating approach described in previous

sections. In parallel to this scale-up, we plan to develop additional datasets of “special” events which will include smaller events than those included in our large-event dataset (M<5.8) and target regions where large events are rare or non-existent.

### Acknowledgements

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. This is LLNL contribution LLNL-\*\*\*\*-\*\*\*\*\*.

### References

Afanasiev M., Peter D., Sager K., Simutė S., Ermert L., Krischer L., Fichtner A., 2016. Foundations for a multiscale collaborative global Earth model, *Geophys. J. Int.*, 204, 39–58, doi: <https://doi.org/10.1093/gji/ggv439>.

Ballard, S., Hipp, J.R., Begnaud, M.L., Young, C.J., Encarnacao, A.V., Chael, E.P. & Phillips, W.S., 2016. SALSA3D: a tomographic model of compressional wave slowness in the Earth’s mantle for improved travel-time prediction and travel-time prediction uncertainty, *Bull. seism. Soc. Am.*, 106(6), 2900–2916.

Begnaud, M. L., S. C. Myers, B. Young, J. R. Hipp, D. Dodge, and W. S. Phillips (2020). Updates to the Regional Seismic Travel Time (RSTT) Model: 1. Tomography, *Pageoph*, 24 pp, doi:10.1007/s00024-020-02619-5.

Bozdağ, E. D. Peter, M. Lefebvre, D. Komatitsch, J. Tromp, J. Hill, N. Podhorszki, and D. Pugmire (2016), Global adjoint tomography: first-generation model, *Geophysical Journal International*, 207(3), 1739–1766, <https://doi.org/10.1093/gji/ggw356>.

Ekstrom, G., M. Nettles, A. M. Dziewonski (2012), The global CMT project 2004–2010: Centroid-moment tensors for 13,017 earthquakes, *Phys. Earth Planet. Int.*, 200–201, 1–9, <https://doi.org/10.1016/j.pepi.2012.04.002>

Fichtner, A. (2010). Full seismic waveform modelling and inversion. Berlin, Heidelberg: Springer.

Fichtner, A., van Herwaarden, D. P., Afanasiev, M., Simutė, S., Krischer, L., Çubuk-Sabuncu, Y., Taymaz, T., Colli, L., Saygin, E., Villaseñor, A., Trampert, J., Cupillard, P., Bunge, H. P., & Igel, H. (2018), The Collaborative Seismic Earth Model: Generation 1. *Geophysical research letters*, 45(9), 4007–4016. <https://doi.org/10.1029/2018GL077338>

French, S.W., and B. Romanowicz (2015), Broad plumes rooted at the base of the earth’s mantle beneath major hotspots, *Nature*, vol. 525, no. 7567. doi:10.1109/IPDPS.2015.58

Grand, S.P. 2002. “Mantle Shear-Wave Tomography and the Fate of Subducted Slabs.” *Phil. Trans. R. Soc. Lond.* 360:2475–2491. <https://doi.org/10.1098/rsta.2002.1077>.

Komatitsch, D.; Tromp, J. (2002a), Spectral-element simulations of global seismic wave propagation-I. Validation, *Geophys. J. Int.*, 149 (2) , 390-412, doi: 10.1046/j.1365-246X.2002.01653.x.

Komatitsch, D.; Tromp, J. (2002b), Spectral-element simulations of global seismic wave propagation–II. Three-dimensional models, oceans, rotation and self-gravitation, *Geophys. J. Int.*, 150 (1) , 303-318.

Komatitsch, D.; Villette, J.-P.; Tromp, J.; Afanasiev, M.; Bozdag, E.; Charles, J.; Chen, M.; Goddeke, D.; Hjorleifsdottir, V.; Labarta, J.; Le Goff, N.; Le Loher, P.; Liu, Q.; Maggi, A.; Martin, R.; McRitchie, D.; Messmer, P.; Michea, D.; Nissen-Meyer, T.; Peter, D.; Rietmann, M.; de Andrade, S.; Savage, B.; Schuberth, B.; Sieminski, A.; Strand, L.; Tape, C.; Xie, Z.; Zhu, H. (2015), SPECFEM3D GLOBE v7.0.0 [software], *Computational Infrastructure for Geodynamics*.

Krischer, L., and E. Casarotti, (2015, September 30). pyflex: 0.1.4 (Version 0.1.4). Zenodo. Doi: 10.5281/zenodo.31607.

Krischer, L., Igel, H., Fichtner, A., 2018. Automated large-scale full seismic waveform inversion for North America and the North Atlantic. *Journal of Geophysical Research* 123. <https://doi.org/10.1029/2017JB015289>

Lei, W., Ruan, Y., Bozdag, E., Peter, D, Lefebvre, M., Komatitsch, D., Tromp, J., Hill, J. Podhorszki, N. and D. Pugmire (2020), Global adjoint tomography-model GLAD-M25, *Geophys. J. Int.*, 233,1-21. doi: 10.1093/gji/ggaa253

Lekić, V., and B. Romanowicz, 2011, Inferring upper-mantle structure by full waveform tomography with the spectral element method. *Geophys. J. Int.* 185, 799-831. doi:10.1111/j.1365-246X.2011.04969.x

Ma, Z., and G. Masters (2014). A New Global Rayleigh- and Love-Wave Group Velocity Dataset for Constraining Lithosphere Properties, *Bull. Seism. Soc. Amer.*, 104, 2007-2026. DOI: 10.1785/0120130320

Ma, Z., G. Masters, G. Laske and M.E. Pasyanos (2014), A comprehensive dispersion model of surface wave phase and group velocity, *Geophys. J. Int.*, 199, 113-135. doi: [10.1093/gji/ggu246](https://doi.org/10.1093/gji/ggu246)

Maggi, A., C. Tape, M. Chen, D. Chao, and J. Tromp (2009). An automated time-window selection algorithm for seismic tomography. *Geophys. J. Int.*, 178(1), 257-281, doi: 10.1111/j.1365-246X.2009.04099.x.

Modrak, R., J. Tromp (2016), Seismic waveform inversion best practices: regional, global and exploration test cases, *Geophys. J. Int.*, 206, 1864-1889, doi: [10.1093/gji/ggw202](https://doi.org/10.1093/gji/ggw202)

Modrak, R. T., D. Borisov, M. Lefebvre, and J. Tromp (2018). SeisFlows-Flexible waveform inversion software. *Comp. Geosci.*, 115, 88-95, doi: 10.1016/j.cageo.2018.02.004.

Myers, S. C., M. L. Begnaud, S. Ballard, M. E. Pasyanos, W. S. Phillips, A. L. Ramirez, M. S. Antolik, K. D. Hutchenson, G. S. Wagner, J. J. Dwyer, C. A. Rowe, and D. R. Russell (2010). A crust and upper mantle model of Eurasia and North Africa for Pn travel time calculation, *Bull. Seism. Soc. Am.* **100**, 2, 640-656. doi:10.1785/0120090198

Örsvuran, R., E. Bozdağ, R. Modrak, W. Lei, and Y. Ruan, 2020. Double-difference measurements in global full-waveform inversions, *Geophysical Journal International*, 220 (1), 661-680, doi:10.1093/gji/ggz444.

Phillips, W. S., M. L. Begnaud, C. A. Rowe, L. K. Steck, S. C. Myers, M. Pasyanos, and S. Ballard (2007). Accounting for lateral variations of the upper mantle gradient in Pn tomography studies, *Geophys. Res. Lett.* **34**, L14312. doi:10.1029/2007GL029338

Rodgers, A., L. Krischer, M. Afanasiev, C. Boehm, C. Doody, A. Chiang and N. Simmons (2022). WUS256: An Adjoint Waveform Tomography Model of the Crust and Upper Mantle of the Western United States for Improved Waveform Simulations, *J. Geophys. Res.*, 2022JB024549R, <https://doi.org/10.1029/2022JB024549>

Simmons, N. A., S. C. Myers, G. Johannesson, and E. Matzel (2012), LLNL-G3Dv3: Global P wave tomography model for improved regional and teleseismic travel time prediction, *J. Geophys. Res.*, 117, B10302. doi:10.1029/2012JB009525

Simmons, N.A, B. Schuberth, S.C. Myers, and D.K. Knapp, Resolution and Covariance of the LLNL-G3D-JPS Global Seismic Tomography Model: Applications to Travel Time Uncertainty and Tomographic Filtering of Geodynamic Models, *Geophys. J. Int.*, 217, 1543-1557. doi.org/10.1093/gji/ggz102, 2019

Simmons N. A., S. C. Myers, C. Morency, A. Chiang, and D. R. Knapp (2021). SPiRaL: A multi-resolution global tomography model of seismic wave speeds and radial anisotropy variations in the crust and mantle, *Geophys. J. Int.*, 227(2), 1366-1391. doi:<https://doi.org/10.1093/gji/ggab277>.

Steck, L.K., W.S. Phillips, K. Mackey, M.L. Begnaud, R.J. Stead, and C.A. Rowe, 2009, Seismic tomography of crustal P and S across Eurasia, *Geophys. J. Int.*, 177, 81-92, 2009.

Tao, K., Grand, S. P., & Niu, F. (2018). Seismic structure of the upper mantle beneath eastern Asia from full waveform seismic tomography. *Geochemistry, Geophysics, Geosystems*, 19, 2732– 2763. <https://doi.org/10.1029/2018GC007460>

Tape, C., Q. Liu, J. Tromp (2007), Finite-frequency tomography using adjoint methods- Methodology and examples using membrane surface waves. *Geophys. J. Int.*, 168, 1105- 1129, doi: [10.1111/j.1365-246X.2006.03191.x](https://doi.org/10.1111/j.1365-246X.2006.03191.x)

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

Tromp, J., C. Tape, and Q. Liu (2005). Seismic tomography, adjoint methods, time reversal, and banana-doughnut kernels. *Geophys. J. Int.*, 160, 195-216, doi: 10.1111/j.1365-246X.2004.02453.x.

van Herwaarden, D. P., C. Boehm, M. Afansieve, S. Thrastarson, L. Krischer, J. Trampert, A. Fichtner (2020), Accelerated full-waveform inversion using dynamic mini-batches. *Geophys. J. Int.* 221, 1427-1438. Doi: [10.1093/gji/ggaa079](https://doi.org/10.1093/gji/ggaa079)

Yuan, H., S. French, P. Cupillard, B. Romanowicz (2014), Lithospheric expression of geological units in central and eastern North America from full waveform tomography. *Earth Planet Sc Lett.* 402, 176–186. <https://doi.org/10.1016/j.epsl.2013.11.057>