

An automated workflow for adjoint tomography—waveform misfits and synthetic inversions for the North Island, New Zealand

Bryant Chow¹, Yoshihiro Kaneko^{2,3}, Carl Tape⁴, Ryan Modrak⁵ and John Townend¹

¹*School of Geography, Environment and Earth Sciences, Victoria University of Wellington, PO Box 600, Wellington 6140, New Zealand. E-mail: bryant.chow@vuw.ac.nz*

²*GNS Science, PO Box 30368, Lower Hutt 5011, New Zealand*

³*Department of Geophysics, Kyoto University, Sakyo-ku, Kyoto 606-8501, Japan*

⁴*Geophysical Institute, University of Alaska Fairbanks, 2156 Koyukuk Drive, Fairbanks, AK 99775, USA*

⁵*Los Alamos National Laboratory, Bikini Atoll Rd., SM 30 Los Alamos, NM 87545, USA*

Accepted 2020 August 12. Received 2020 August 4; in original form 2020 April 7

SUMMARY

We develop and verify an automated workflow for full-waveform tomography based on spectral element and adjoint methods. We choose the North Island, New Zealand as a study area because of its high seismicity, extensive seismic network, and the availability of a candidate ray tomography starting model. To assess the accuracy of this model, we simulated 250 regional earthquakes using a spectral element solver, and compared the resulting synthetics with recorded waveforms. In a 10–30 s passband, reasonable cross-correlation phase and amplitude misfits exist between data and synthetics, whereas at 2–30 s, waveform misalignment is severe enough that meaningful cross-correlation measurements are no longer possible. To improve the velocity model at these short periods, we created an automated inversion framework based on existing tools for signal processing, phase measurement, nonlinear optimization, and workflow management. To verify the inversion framework, we performed a realistic synthetic inversion for 3-D checkerboard structure and analyzed model recovery, misfit reduction, and waveform improvement. The results of this analysis show that the source–receiver distribution within the chosen domain is capable of resolving velocity anomalies in regions of sufficient data coverage, and of magnitudes comparable to those expected in a real seismic inversion. Along with this finding, the relative ease of use and reliability of the workflow motivates future efforts targeting a high-resolution (2–30 s), large-scale (>50 000 measurements) seismic inversion for the North Island. Updated models from such an inversion are expected to improve ground motion predictions, constrain complex velocity structures, and advance understanding of New Zealand tectonics.

Key words: Seismic tomography; Computational seismology; Crustal imaging.

1 INTRODUCTION

Seismic tomography offers a powerful tool for understanding the dynamics and evolution of our planet. Computational advances in the 2000s opened the door for advanced tomographic methods, including full-waveform modeling and inversion based on spectral element (Komatitsch & Vilotte 1998; Komatitsch *et al.* 2002) and adjoint methods (Tarantola 1984; Talagrand & Courtier 1987; Luo & Schuster 1991). Nevertheless, such inversions remain challenging not only in terms of computational cost, but also due to the complex software required to perform them.

Full-waveform adjoint tomography has led to detailed tomographic images from regional (e.g. Chen *et al.* 2007b; Tape *et al.*

2010; Lee *et al.* 2014b; Miyoshi *et al.* 2017), to continental (e.g. Fichtner *et al.* 2009, 2013; Zhu *et al.* 2015; Chen *et al.* 2015), to global scales (e.g. Bozdağ *et al.* 2016). In this study, we define regional scale as domains spanning less than 10°. Inversions at continental scales and larger typically fit long-period waveforms (e.g. >8 s as in Fichtner *et al.* 2013), require spherical meshes to honor Earth structure, and examine upper mantle and long-wavelength crustal structure. In comparison, regional scale inversions typically use Cartesian meshes and examine shallow crustal structure related to short-period waveforms (>2 s as in Tape *et al.* 2010). One constraint that limits the applicability of adjoint tomography is the assumption that the starting velocity model be close enough to true

Earth structure to avoid spurious local minima in the inversion procedure. Assessments of candidate models, based on waveform fits, can be carried out to judge their suitability in this respect (Lin *et al.* 2011; Lee *et al.* 2014a; Gao & Shen 2015; Bao & Shen 2016; Taborda *et al.* 2016). Additionally, sufficient data coverage within the domain is required, as tomographic models from earthquake-based inversions can be spatially biased due to uneven distributions of sources and receivers. Ambient noise tomography (e.g. Shapiro *et al.* 2005; Lin *et al.* 2007, 2008; Yao *et al.* 2010; Rawlinson *et al.* 2016) can improve upon these difficulties, but typically limits shortest periods to >5 s to capture surface wave dispersion. Therefore, to attain high-resolution images of shallow crustal structure and complex tectonic features in seismically active regions, it is crucial to develop modern tools to address the unique challenges presented by earthquake-based adjoint tomography.

The domain encompassed by the central and southern North Island and the northeastern South Island, New Zealand (Fig. 1) is a prime candidate for performing regional-scale adjoint tomography. The domain exhibits high seismicity rates due to active subduction of the Pacific Plate below the Australian Plate (Reyners 1998; Wallace *et al.* 2004, 2012), in a region known as the Hikurangi subduction zone. The Hikurangi trench is directly offshore (~ 100 km) and the plate interface below land is shallow, at roughly 15 km below the east coast of the North Island (Wallace *et al.* 2009; Barnes *et al.* 2010; Williams *et al.* 2013). Intraplate seismicity within the Pacific slab, interplate seismicity along the principal décollement and upper plate seismicity associated with volcanism in the central North Island and strike-slip faulting along the length of the Hikurangi margin, collectively provide a geographically and kinematically diverse set of earthquake sources. More than a decade of earthquake moment tensors (Ristau 2008) and broadband seismic data are available through the permanent seismic network of New Zealand (GeoNet). Additionally, many years of research by Eberhart-Phillips and colleagues have produced a detailed 3-D velocity model (Eberhart-Phillips *et al.* 2010, 2014, 2015, 2017, 2020; Eberhart-Phillips & Fry 2017, 2018) that can be used to generate synthetic seismograms for the region. These qualities make this area a good choice for developing methods for adjoint tomography, and applying them to a region of complex structure and substantial seismic hazard (Stirling *et al.* 2012).

In this paper, we present a methodological framework for full-waveform tomography using adjoint methods. The purpose of this paper is threefold:

- (i) To use 3-D wavefield simulations to evaluate a candidate starting velocity model of New Zealand.
- (ii) To expand the capabilities of an open-source automated workflow package for adjoint tomography (SeisFlows), by providing a modular interface for misfit assessment (Pyatoo).
- (iii) To demonstrate the potential of this software suite for large-scale adjoint tomography through complete synthetic inversions using New Zealand station and earthquake distributions.

In developing an automated workflow, we aim to promote transparency and reproducibility in seismic tomography, and encourage community development through expansion of existing open-source software packages. We start with analysis of synthetics computed using the latest ray tomography model of New Zealand (Section 2). We then introduce and describe an automated inversion workflow to be used in current and future inversions (Section 3). This is followed by an application of the workflow to realistic synthetic inversions (Section 4), designed to illustrate the capability for resolution testing and assessment of methodological choices in large-scale inversions.

The paper concludes with a discussion looking toward full-scale adjoint tomography of the North Island, New Zealand, and applications elsewhere (Section 5).

2 MISFIT ASSESSMENT OF AN INITIAL 3-D VELOCITY MODEL

We evaluate a New Zealand velocity model to assess its applicability as a candidate starting model for adjoint tomography. Using ray-based body-wave traveltime tomography, Eberhart-Phillips *et al.* (2010) jointly computed earthquake hypocentres and seismic velocities (V_p , V_p/V_s) in a New Zealand-wide 3-D velocity model ('NZ-wide'). These velocity models incorporate traveltime measurements from local earthquake data and shots from active seismic experiments, recorded on permanent and temporary seismic stations throughout New Zealand. Portions of the model have since been improved through inversions for Rayleigh-wave group velocity maps (Eberhart-Phillips & Fry 2017), and joint inversions of local earthquake and teleseismic data (Eberhart-Phillips & Fry 2018). Attenuation models (Q_p , Q_s) have also been constructed by fitting spectral decay for body waves from local earthquakes (Eberhart-Phillips *et al.* 2014, 2015, 2017, 2020).

In this study, we use a portion of the NZ-wide velocity model, which we refer to as 'NZ-North'. Fig. 2 shows map views and cross sections of the NZ-North V_p and V_s models. Tectonic features of the North Island and the Hikurangi subduction interface (Wallace *et al.* 2004, 2009; Clark *et al.* 2019) are identifiable by the stark velocity contrasts. The lowest velocities within the model are found in the accretionary prism of the subduction forearc, which is akin to a sedimentary basin, shown to exhibit extremely long durations of ground motion following seismic events (Kaneko *et al.* 2019). In the central North Island, the modern (2 Myr) portion of the volcanic arc, the Taupō Volcanic Zone, exhibits high heat flow and geothermal activity, extensional faulting and associated seismicity, underlain by low-velocity, highly attenuating crust and mantle (Wilson *et al.* 1995, 2009; Behr *et al.* 2011; Rowland *et al.* 2010; Eberhart-Phillips *et al.* 2020). Cross sections show the subducting Pacific Plate and overlying Australian Plate, as well as low velocity sediments in the forearc region. The northeastern extent of the South Island is included in this domain, as seismicity there is high in the southward transition of the subduction margin from oblique convergence to strike-slip behaviour. The southwest corner of NZ-North contains the 2016 M_w 7.8 Kaikōura earthquake (Hamling *et al.* 2017; Holden *et al.* 2017) and its extensive aftershock sequence.

2.1 Simulation framework

We use the NZ-North velocity model to carry out spectral element wave propagation simulations to assess its applicability for adjoint tomography based on waveform fits. To generate synthetic seismograms, we use the spectral element solver SPECFEM3D Cartesian (Komatitsch & Tromp 2002a, b) to simulate ground displacement from local earthquakes. Using the SPECFEM3D internal mesher, a rectangular mesh was created with roughly 600 000 elements and a minimum element spacing of 1 km. Topography and bathymetry, necessary for accurate propagation of surface waves, are interpolated from SRTM-30P (Becker *et al.* 2009). A water layer is not used in this mesh. Mesh boundaries are chosen to maximize the number of local earthquakes within the domain, while limiting unused domain space. A mesh depth of 400 km is chosen to avoid numerical reflections from the bottom boundary.

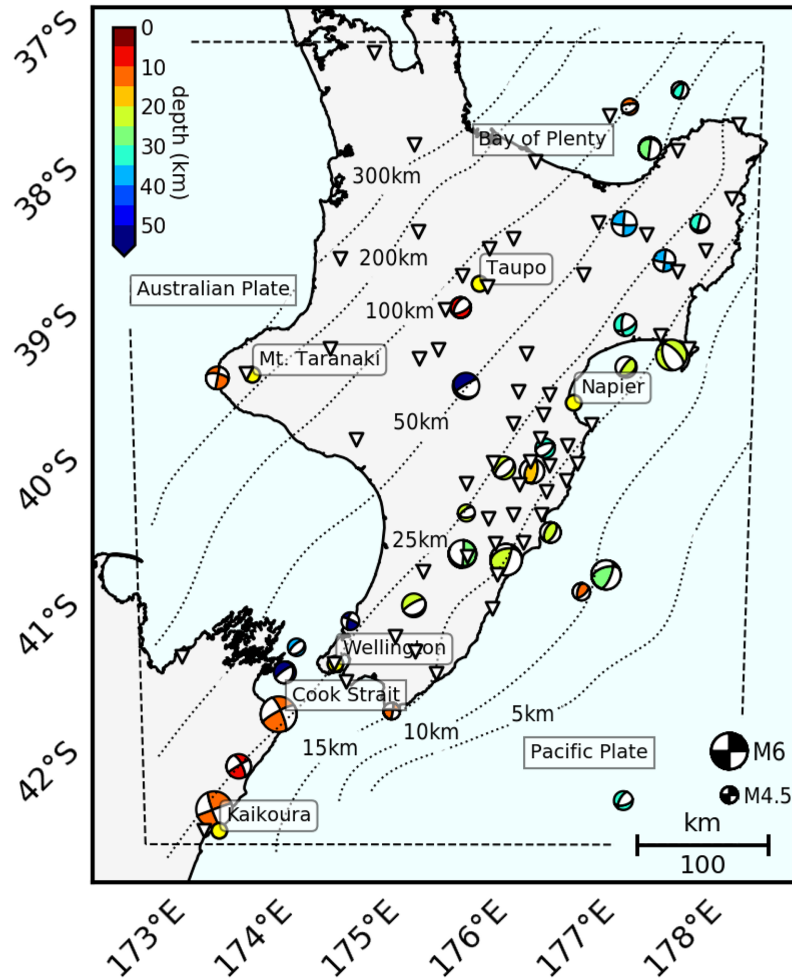


Figure 1. Map view of the study region, containing the central and southern North Island and the northeastern tip of the South Island, New Zealand. 30 earthquakes and 58 receivers used in synthetic inversions are shown as focal mechanisms and inverted triangles, respectively. Earthquakes are chosen to provide a varied representation of available focal mechanisms within the domain. Beachball size corresponds to magnitudes $4.8 \leq M_w < 6$; depth information coded by colour. The simulation domain is outlined by the dashed black line. Select cities or landmarks are shown as yellow circles. Named regions and tectonic plates are shown with square name plates. Dotted lines represent the approximate depth contours of the subduction interface, with depths denoted by inline annotations (Williams *et al.* 2013).

Our earthquake catalog consists of 250 moment tensors from local earthquakes of moment magnitudes $4.4 \leq M_w < 6.0$ that occurred between January 2005 and November 2019 at depths shallower than 60 km. The depth and lower magnitude bound are chosen to discard events that will not excite strong surface wave signals, while the upper magnitude bound is set to avoid the need for finite-fault source representations. Moment tensors are routinely calculated by GeoNet for the New Zealand region (Ristau 2008, 2013) using a time domain moment tensor inversion code (Dreger 2003) and New Zealand-specific 1-D velocity models. These regional moment tensors have been shown to compare well with Global CMT solutions (Ekström *et al.* 2012) and first motion focal mechanisms (Robinson & Webb 1996). We calculate synthetic seismograms at locations corresponding to 45 GeoNet broadband stations.

2.2 Misfit assessment

We assess the degree to which the NZ-North velocity model generates synthetic seismograms which accurately reproduce seismic observations. We do this by performing forward simulations for

all events in our catalog and quantifying waveform accuracy using a time windowing algorithm, which compares similar sections of waveforms and excludes noise or poor-fitting data. Here, time windows are chosen using the FLEXWIN algorithm (Maggi *et al.* 2009), which identifies wave-group like signals using a short-term average/long-term average (STA/LTA) ratio. The selection algorithm identifies acceptable time windows based on waveform fit, characterized by peak cross-correlation, time-shift, and amplitude ratio. Amplitude ratio is defined by Dahlen & Baig (2002) as

$$\Delta \ln(A) = \ln \left(\frac{A_{\text{obs}}}{A_{\text{syn}}} \right) = 0.5 \ln \left[\frac{\int d^2(t) dt}{\int s^2(t) dt} \right], \quad (1)$$

where d and s are the data and synthetic waveforms, respectively. User-defined parameters allow control over the quality of waveforms contained within the chosen windows. For this region, we find adequate window selection when absolute time-shift does not exceed 8 s, absolute amplitude ratio is less than 2, and peak cross-correlation is greater than 0.7. For reference, a window around two identical waveforms has time-shift and amplitude ratio equal to 0, and peak cross-correlation equal to 1. A summary of windowing parameters used in this study is provided in Table A1.

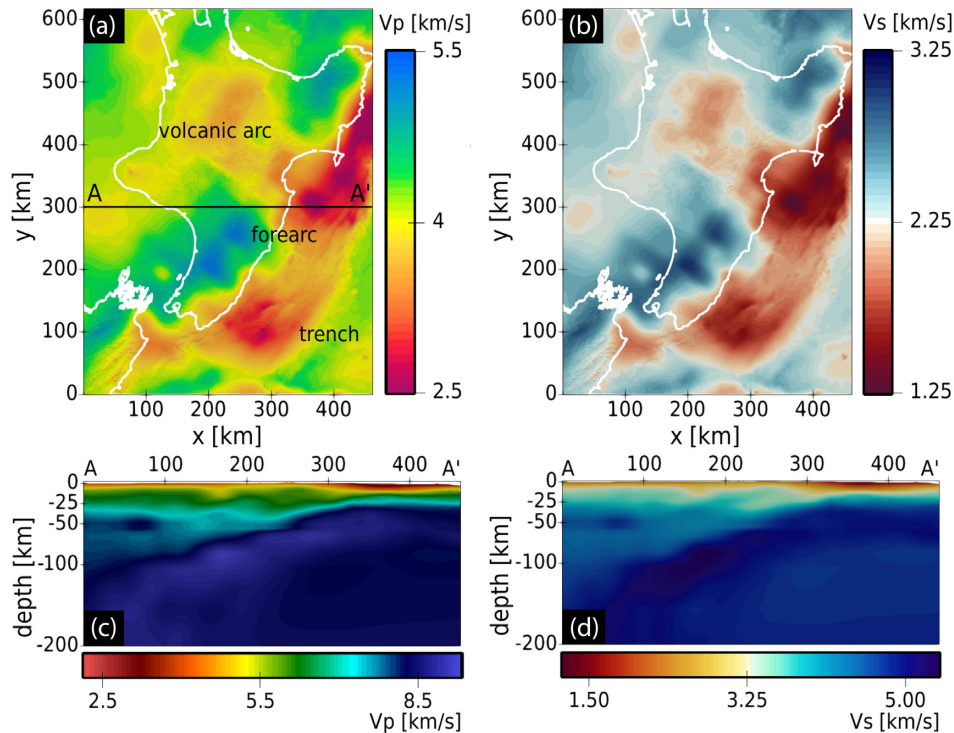


Figure 2. (a) Map view of the NZ-North V_p model. Coastline shown as white line. Line A–A' corresponds to the surface trace of cross sections in (c–d). Text labels provide general locations of Hikurangi subduction zone features. (b) Map view of the NZ-North V_s model. (c–d) Cross sections of the V_p and V_s models, corresponding to the line A–A' in (a).

2.2.1 Selected waveforms

Here we show a representative set of average and below-average vertical-component waveforms from various source–receiver paths to provide a sense of the varying degrees of misfit for the candidate velocity model. Waveforms are shown for a period band of 10–30 s (Figs 3a–c) and 2–30 s (Figs 3a'–c'), corresponding to the labelled source receiver paths in Fig. 3(d). Window annotations provide a reference for misfit through values of time-shift (dT), amplitude ratio (dA), and peak cross-correlation (cc). By convention, a positive time-shift corresponds to synthetic phases arriving early with respect to observed phases ($dT = T_{\text{obs}} - T_{\text{syn}}$), and a positive amplitude ratio corresponds to synthetic amplitudes that are smaller than observed amplitudes.

In Fig. 3, waveform A at 10–30 s represents average misfit. Cross-correlation is high for both misfit windows, however time-shift is large for the later window. At 2–30 s periods, a stark increase in misfit is illustrated by low cross-correlation values and no selected windows around the later arriving phases. This average misfit may be explained by the ray path crossing an area of minimal data coverage. Waveform B of Fig. 3 traverses the Taupō Volcanic Zone, characterized by average time-shift but large amplitude difference at 10–30 s periods. At 2–30 s periods, waveform peak amplitudes are more than a factor 10 different, contributing to a short selected time window. This large amplitude difference may be due to inaccuracies in the attenuation model for this complex volcanic region. Waveform C of Fig. 3 traverses parallel to the Hikurangi forearc region. Observations of surface waves in this area show increased durations of ground motions due to the basin-like resonance effects of the low-velocity accretionary prism. Inaccuracy in the initial velocity model is observable in both period bands, where synthetic seismograms do not fully capture the resonant waveforms shown in the observations.

The same analysis for radial and transverse components (Fig. A3) show similar waveform behaviour and window selections. Synthetics were also generated using a North-Island-specific 1-D velocity model (Ristau 2008). Comparisons between the resultant synthetic seismograms (Fig. A1) show that NZ-North is more capable of capturing accurate waveform fits and travel times, especially for later arrivals. Although waveform amplitudes are better fit with the 1-D velocity model of Fig. A1, this is deemed less important due to our intended use of phase-based objective functions (Section 4.3).

2.2.2 Bulk misfit analysis

We analyze misfit in bulk for all collected measurements at a period band of 10–30 s. Misfit histograms detail time-shift, amplitude ratios, peak cross-correlation, and relative start time for 250 events recorded on three components of 45 broadband stations, with a total of 26 676 measurements.

Time-shifts (Fig. 4a) show that, on average, synthetic phases arrive 2.0 ± 3.65 s earlier than observed phases, meaning the model velocities are slightly fast with respect to observations. Amplitude ratios (Fig. 4b) show a mean and standard deviation of -0.07 ± 0.68 , meaning that synthetic amplitudes are approximately 7 per cent larger than observed amplitudes. Peak cross correlation (Fig. 4c) shows a roughly uniform distribution of waveform fits up to a value of 0.85, with progressively lower counts above that. Figs 4(d)–(f) show time-shifts for vertical, radial and transverse components, respectively. Radial and vertical components show a comparable number of measurements and average time-shift, while a lower number and larger average time-shift for the transverse component suggest less accurate synthetics for transverse-sensitive phases like SH and




Figure 3. A representative selection of vertical-component data-synthetic comparisons. (a–c) Misfit comparisons for the 10–30 s period band. Observed waveforms are shown in black, synthetic waveforms in red, grey boxes represent misfit windows chosen by the time windowing algorithm. Measured time-shift (dT), amplitude ratio (dA), and peak cross-correlation (cc) are indicated within each window. Waveforms are shown in units of displacement. (a'–c') The same waveforms as (a–c), but filtered at a period band of 2–30 s, and windows chosen with slightly altered windowing parameters (Table A1). (d) Map of source–receiver locations with ray paths, focal mechanisms, and labels corresponding to general subduction zone features.

Love waves. Figs 4(g)–(i) relate to the relative start time of measurement windows with respect to the event origin time. The median start time is 86 s, which suggests that a majority of measurements are made for direct arrivals and small source–receiver distances. Cross correlation values for measurements made before 100 s start time show a relatively uniform distribution, while measurements made after 100 s show progressively lower counts for cross-correlation values greater than 0.8. Phases associated with later arrivals, that is surface waves and large source–receiver distances, show increased misfit, which can be interpreted either as accumulation of misfit for large propagation distances, or high misfit for surface waves stemming from a candidate velocity model primarily derived using body waves.

The overall low time-shifts and amplitude ratios, high peak cross correlation values, and even distribution of measurements among components, suggest that NZ-North is a reasonable starting model for adjoint tomography. This is further evidenced through comparisons with bulk misfit analysis of synthetics generated using a 1D North Island velocity model (Fig. A2). Despite less restrictive windowing parameters for the 1D synthetics (Table A1), only 9834 measurements could be made, with considerably larger average time-shifts (6.11 ± 5.03 s), and amplitude ratios (0.55 ± 0.73), illustrating significant improvements gained using the region-specific 3D velocity model. Additionally, analysis of the ray-based 3D starting model of southern California by Tape *et al.* (2010) show comparable overall time-shifts (1.06 ± 1.43 s) and amplitude ratios (0.34 ± 0.39) for 36 553 measurements at 6–30 s periods. We argue that the larger average misfit for NZ-North is attributed more to the complicated nature of New Zealand tectonics with respect to southern California, rather than any substantial differences in accuracy for their respective starting models. The improvements in waveform fit seen in Tape *et al.*, as well as the accuracy gained using NZ-North, in contrast to a 1D velocity model, provides optimism that this is a reasonable candidate starting model for adjoint tomography.

3 AUTOMATED INVERSION WORKFLOW

Adjoint tomography consists of many near-identical processes manipulating large amounts of data in an iterative procedure. It therefore benefits strongly from automation, parallelization and structured organization. Here we present our workflow, which is comprised of new and existing open-source tools that address and automate each step of an earthquake-based adjoint tomography workflow. In this section, we briefly outline the components of a typical workflow (Fig. 5), followed by explanations of the individual software packages comprising our automated workflow.

Adjoint tomography typically begins with data gathering, characterized by selection of a region and collection of seismic data consisting of moment tensors, waveforms, and source and receiver metadata. An initial velocity model is generated for the region using external information, and the forward problem is solved by simulating seismic wave propagation for each source. The resultant synthetic waveforms are compared to corresponding seismic data through some objective function, in a misfit quantification procedure, or function evaluation. Here, time windowing can be used to restrict the data evaluated in the objective function, preferable as data quality can not always be well controlled in an automated framework. To solve the inverse problem, we seek to minimize the objective function through perturbations of either velocity model or seismic source parameters. This is accomplished by determining the gradient of the objective function, with respect to these parameters, using the adjoint-state method (Tarantola 1984). In the adjoint simulation, the forward wavefield interacts with a secondary (adjoint) wavefield to produce the gradient of the objective function for a given source–receiver pair. This adjoint wavefield is produced by an adjoint source, which is a source time function, input at receiver locations, weighted by misfit for a given source–receiver pair. Using gradient information from many source–receiver pairs, a nonlinear optimization algorithm searches for a local minima in the objective

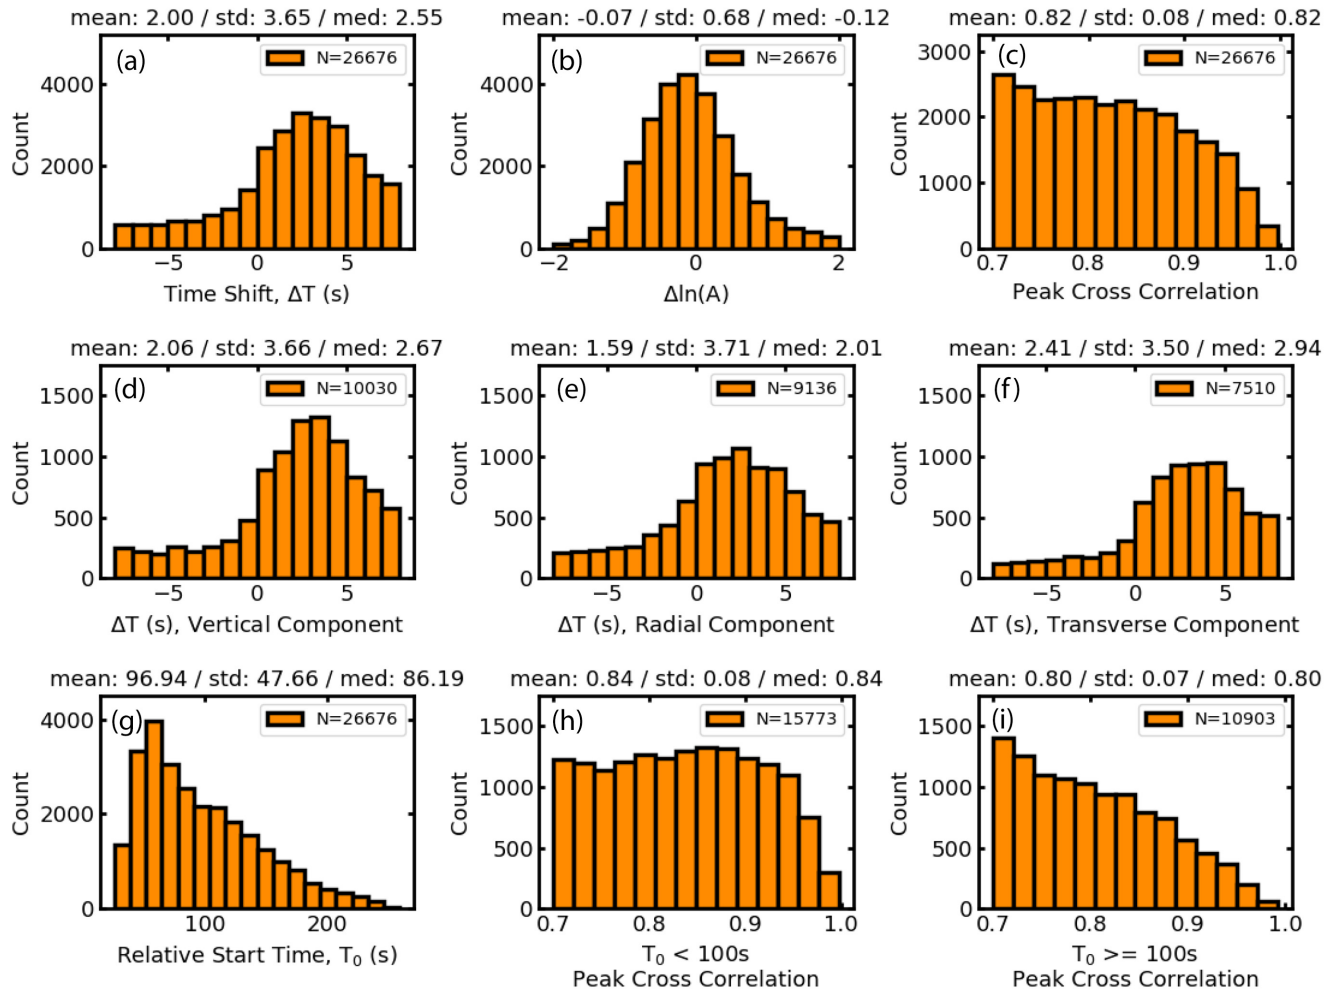


Figure 4. Misfit histograms for forward simulations using the NZ-North velocity model. Measurements from 250 events recorded on three components of 45 broadband stations, with 26 676 total measurements. Mean, standard deviation, and median are given in the titles of each histogram. The number of measurements for each histogram is provided in the respective legend. (a) Time-shift corresponding to peak cross correlation between data and synthetics, in units of seconds. (b) Data-synthetic amplitude difference $\Delta \ln(A)$ (Eq. 1). (c) Peak cross correlation value. (d–f) Time-shift for vertical, radial, and transverse component measurements. (g) Measurement start time relative to event origin time. (h) Peak cross correlations for measurements made before 100 s relative start time. (i) Peak cross correlations for measurements made at or after 100 s relative start time.

function through perturbations of model or source parameters. The forward problem is solved again with the perturbed model or source and evaluated for misfit reduction. This process repeats iteratively until some pre-defined convergence criteria are met. Updated models can be appraised through resolution tests, and comparisons with external information such as known geology, tectonics, or tomographic models.

Our aim in automating the adjoint tomography workflow is to increase the speed and efficiency with which an inversion can be initialized and performed, while thereby reducing the possibility of user error. We emphasize open-source code development to promote scientific reproducibility and increased transparency in seismic inversions. Desirable aspects of such a workflow include having as few programming languages as possible, the ability to implement varying degrees of manual control, and a flexible design that can be applied to a wide range of scales and systems. In the following sections, we describe our chosen workflow tool, SeisFlows, as well as its complement package for earthquake tomography, Pyatoa.

3.1 SeisFlows

For automation of the adjoint tomography workflow, we use SeisFlows (Modrak *et al.* 2018), a flexible waveform inversion package written in the Python programming language. SeisFlows provides custom routines for seismic inversions with modules that address, for example, workflow management, compute system interactions, and nonlinear optimization. Additionally, it integrates the functionalities of external numerical solvers, like SPECFEM3D Cartesian, using generalized wrapper functions. SeisFlows automates all components of a seismic inversion, from job submissions on high performance computing (HPC) systems to model updates, and has been used for active source full-waveform inversions (e.g. Smith *et al.* 2019) and synthetic inversion studies (e.g. Rusmanugroho *et al.* 2017; Borisov *et al.* 2018; Matharu & Sacchi 2018; Tromp & Bachman 2019a, b). SeisFlows has yet to be used for earthquake-based inversions, and therefore important aspects of earthquake tomography, such as seismic data handling and detailed misfit quantification, are underdeveloped with respect to the package’s more robust components. This has led to the development of Pyatoa,

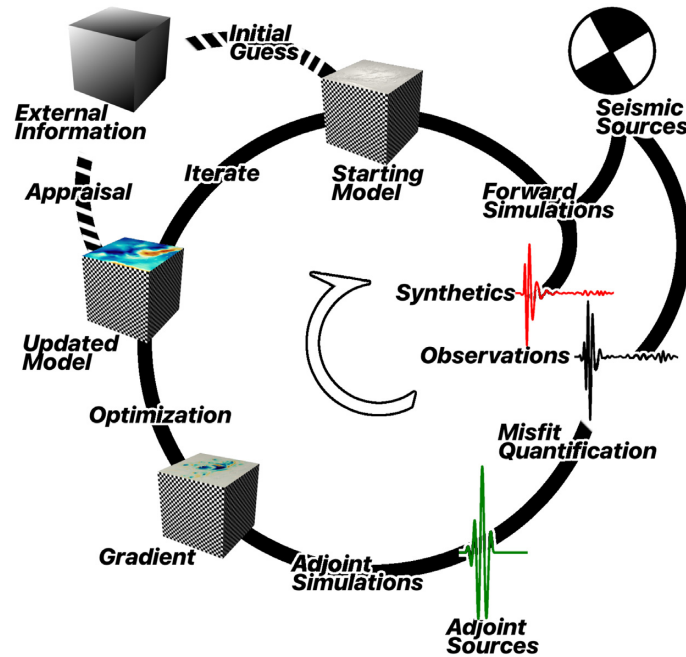


Figure 5. A diagrammatic illustration of the adjoint tomography workflow, adapted from Fig. 2 of Liu & Gu (2012). The workflow tools introduced in this paper are capable of automating all components of the adjoint tomography workflow, except for the initial guess, creation of the starting model and the appraisal of the final model using external information. Inversions can be run fully-automated, or with various levels of manual intervention introduced throughout the inversion.

a complementary package to enhance SeisFlows' capabilities for earthquake-based adjoint tomography.

3.2 Pyatoo

Python's Adjoint Tomography Operations Assistant (Pyatoo) is an open-source tool designed to handle seismic data within a seismic inversion workflow. It primarily addresses the misfit quantification procedure, through gathering and processing of seismic data, and evaluation of objective functions. Pyatoo is built around the widely-used seismological processing package, ObsPy (Beyreuther *et al.* 2010; Krischer *et al.* 2015b), and benefits from the seismic data handling and gathering capabilities contained therein. Using custom methods and data structures, Pyatoo abstracts, or simplifies, routines defined by ObsPy and other existing Python packages to facilitate automation of misfit evaluation in a seismic inversion. Pyatoo additionally provides a highly structured data organization scheme, and easy-to-use measurement aggregation tool, to allow for efficient storage, retrieval, and assessment of inversion data. The aim in developing Pyatoo is to codify the typically custom-made, hard-coded subroutines that are accessory, but nonetheless essential, when performing earthquake-based adjoint tomography. We provide a general explanation of the structure and functionalities of Pyatoo here, and leave more technical explanations for the documentation (github.com/bch0w/pyatoo). A flowchart detailing the role of Pyatoo within a SeisFlows workflow is given in Fig. A4.

3.2.1 Structured data storage

Data gathered at each function evaluation are stored in custom file formats defined by the Adaptable Seismic Data Format (ASDF, Krischer *et al.* 2016). ASDF data sets are structured file formats built to store large amounts of seismic data. The use of data sets

removes the need for cumbersome and error-prone directory structures commonly used in inversion workflows. This facilitates storage, transport, and analysis of inversion data, and increases the potential for reproducible seismic inversions. ASDF data sets are self-describing, as they contain both data and metadata. For Pyatoo, this means data sets provide all necessary information to understand and repeat previous function evaluations. ASDF data sets are built to store seismic waveforms and metadata as ObsPy objects, removing the need for data conversion in the transition from storage to processing. Pyatoo provides additional functionality to ASDF, allowing for storage of misfit windows, adjoint sources, and configuration parameters. Some pertinent applications of this added functionality: reusing previously derived misfit windows in a new function evaluation, or quickly re-evaluating misfit using new windowing parameters or objective functions.

3.2.2 Misfit quantification

At each function evaluation and for each source–receiver pair, observations and synthetics must be standardized, preprocessed, cut into time windows and evaluated using an objective function. Pyatoo streamlines and automates these operations through custom workflow and data structures. Fig. 6 shows an example end product of a Pyatoo function evaluation. Here we refer to standardization as the process of equating two time series, for example by start time and sampling rate, while pre-processing refers primarily to instrument response removal and frequency domain filtering. Through a custom data object which relates matching waveforms and metadata, Pyatoo wraps the processing functionalities of ObsPy to automatically standardize and pre-process related waveform data. A series of internal checks ensure that inappropriate waveforms are not included in the function evaluation, for example, in the presence of data gaps.

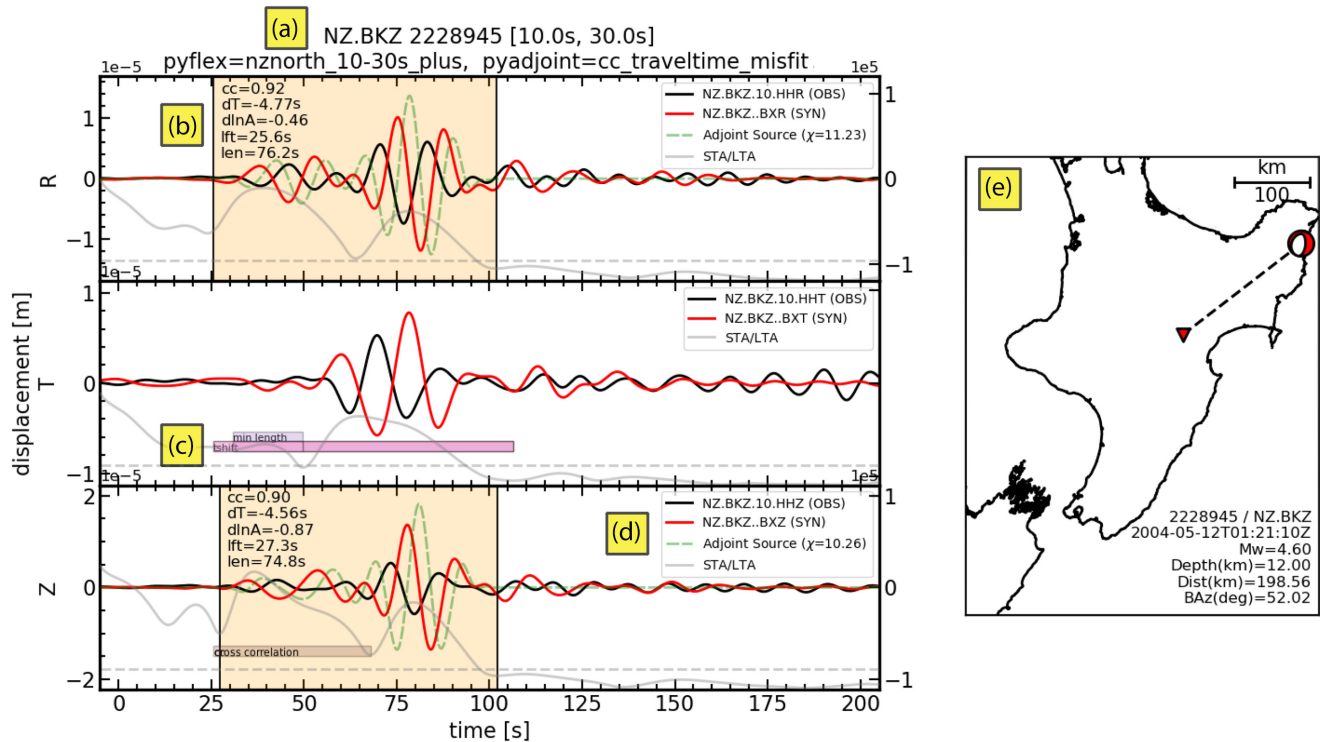


Figure 6. Misfit assessment for one source–receiver pair, generated using Pyatoa. Yellow labels are added for descriptive purposes. (a) Waveform title, which displays relevant information like processing parameters. (b) Time windows are shown with measurement information for quick assessment of waveforms and misfit. By default, cross-correlation (cc), time-shift in seconds (dT), amplitude anomaly (dlnA), window start time in seconds (lft), and window length in seconds (len) are provided. (c) Rejected time windows are shown as colour-coded bars. Here preliminary time windows are rejected for unacceptable time-shift (shift) and minimum length (min length). (d) The legend provides component identification and total calculated misfit for a single component (χ). The grey short-term-average over long-term-average waveform (STA/LTA) is used to determine preliminary windows, and is shown alongside a waterlevel (dashed grey) used for the internal rejection criteria. (e) A corresponding map with useful information pertaining to the given source–receiver pair.

Time window selection is accomplished using FLEXWIN, an automatic selection algorithm which identifies distinct energy arrivals using an STA/LTA ratio waveform, and isolates waveform similarities using a comprehensive set of rejection criteria. Pyatoa uses the FLEXWIN algorithm through wrapping functions of the Python port, Pyflex (Krischer 2015b). Additional windowing criteria, such as weights based on geographic location (e.g Ruan *et al.* 2019), can readily be inserted into the windowing procedure via Pyatoa. In Fig. 6, the grey STA/LTA waveforms and associated time windows in orange show the outputs of Pyflex. Here the peaks of the STA/LTA waveform highlight distinct arrivals in the synthetic waveform. The annotations displayed in the time windows provide useful information for misfit assessment. Time windows that were discarded due to the FLEXWIN internal rejection criteria are also shown (Fig. 6c), giving the user a more informed understanding of waveform misfit in the absence of selected time windows.

To evaluate objective functions and generate adjoint sources, Pyatoa abstracts the functionalities of Pyadjoin (Krischer 2015a). Within chosen misfit windows, Pyadjoin is used to evaluate the objective function and derive adjoint sources; the green dashed waveforms in Fig. 6 are examples of adjoint sources. Pyadjoin provides a template structure that allows for the use of any objective function relating data and synthetics to be used to define misfit. Pyatoa provides an additional level of organization for Pyadjoin results to facilitate assessment of misfit in bulk. Using the functionality of Pyadjoin, adjoint sources can be stored in ASDF data sets, and/or exported into SPECSEM3D-ready file formats for use in adjoint simulations. Pyatoa provides additional utility functions to produce

auxiliary files required by SPECSEM3D to launch a subsequent adjoint simulation.

3.2.3 Assessment of inversion results

Seismic inversions produce a large amount of data which are prohibitively time-consuming to manually parse, organize, and understand. To help address this, Pyatoa makes use of the Python-based data analysis library Pandas (Wes McKinney 2010; Pandas Development Team 2020). Time windows and associated misfit values are collected and organized, at each iteration, into a useful spreadsheet-like data structure, providing an efficient tool for aggregating measurements and assessing misfit at all scales, from the single time window, to event- or station-based comparisons, to bulk analysis of each iteration. Additional automated routines allow Pyatoa to visualize waveforms, time windows and adjoint sources (e.g. Fig. 6), while pre-defined statistical plots summarize bulk assessment of waveform quality, window choice and misfit characteristics. These tools are designed to inform the user on inversion behaviour and to facilitate the assessment of inversion quality.

4 SYNTHETIC INVERSIONS

We perform realistic synthetic inversions to demonstrate the workflow, provide resolution information and evaluate methodological choices through repeated testing. A large-scale real-data inversion has not yet been undertaken and is the subject of on-going research;

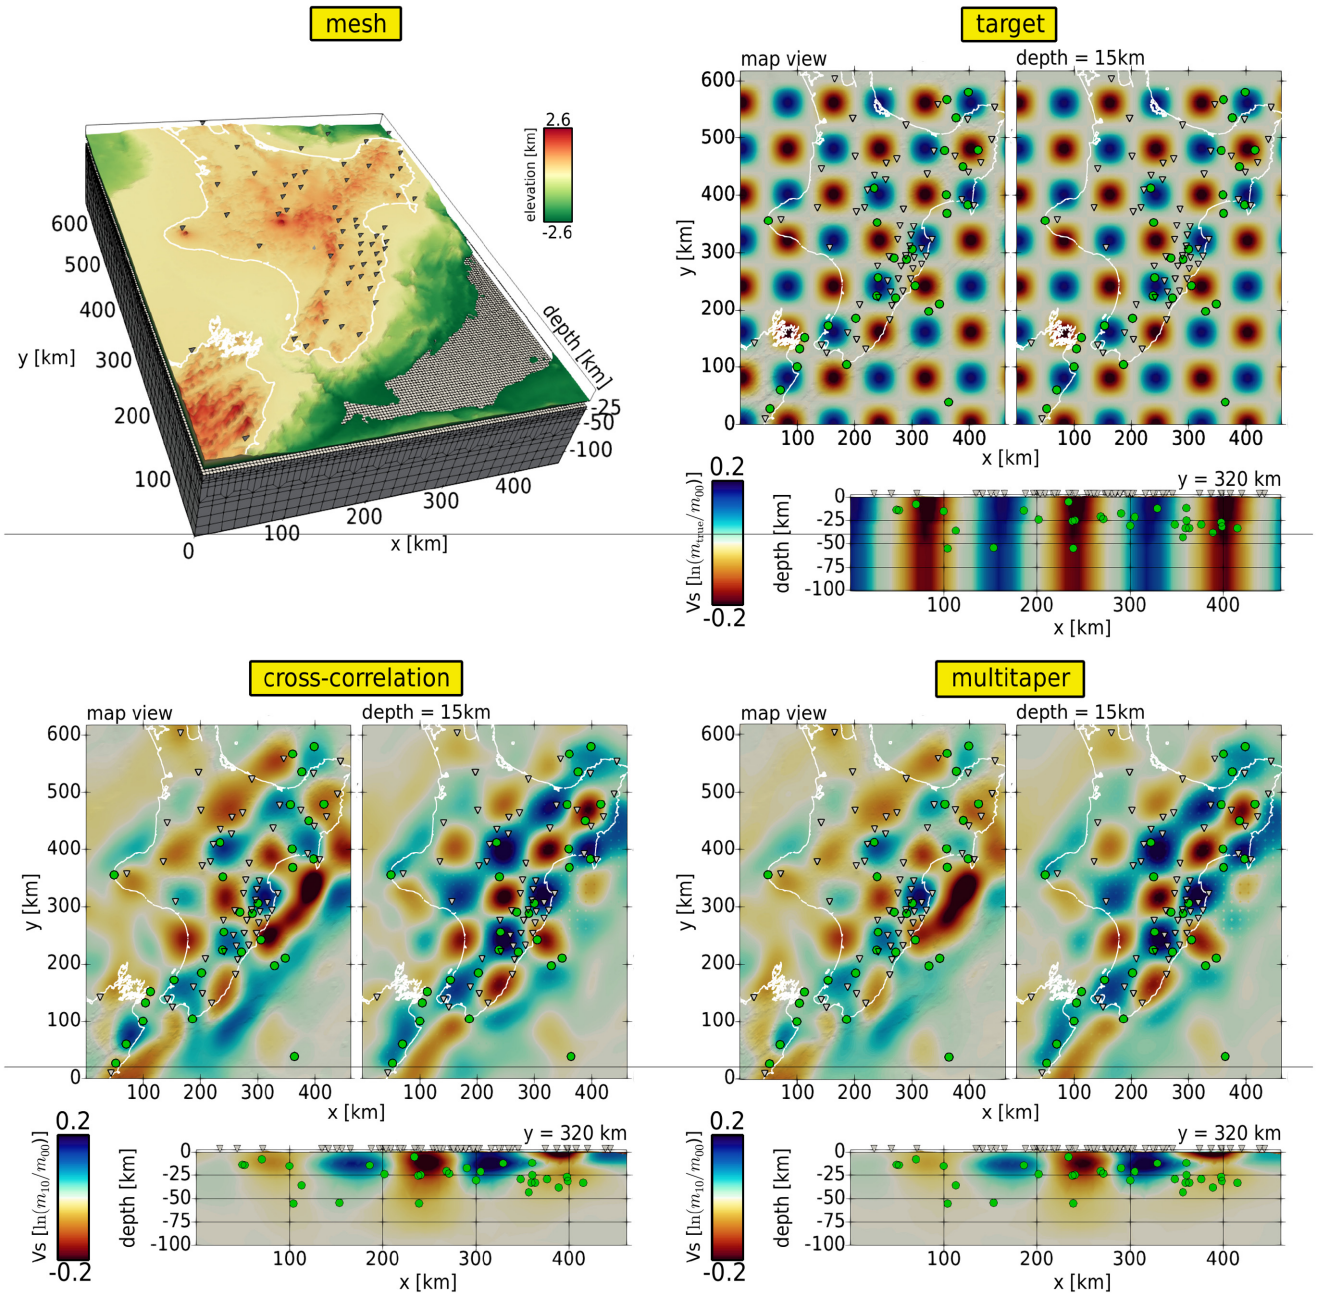


Figure 7. Summary figure showing setup and results of the synthetic inversions. The Cartesian numerical mesh is shown in the upper-left. In the top portion, elevation is coded by colour; below, the numerical mesh is shown from depths of 20–100 km. Bathymetry is purposely cut off below -2.6 km to reveal the smallest elements with a size of 2 km. Two coarsening layers can be seen at approximately 25 and 50 km depths, with a largest element size of 40 km. Though not shown, the mesh continues to 400 km depth. The target velocity perturbations (upper-right) and recovered velocity perturbations (bottom row) are shown in map view, with a horizontal slice at 15 km depth, and in cross section. Cross-sections are shown at a fixed Y -axis value of 320 km, which traverses through the centre of one set of 80-km-wide checkers, where perturbations of the NZ-North velocity model are at their maximum. Colours represent net model update of V_s , which saturates at the largest perturbation values of ± 0.2 . In map view and horizontal slice, green circles show event epicentral locations. In cross section, green circles show event depths. Receiver locations are shown as inverted triangles.

however, the aim of the synthetic experiments is to lay the foundation for a real-data inversion and provide confidence that methodological choices work well for our study area.

4.1 Inversion setup

We perform a synthetic 3-D inversion using a realistic numerical mesh and target perturbation checkerboard model (Fig. 7).

The checkerboard consists of 80 km wide, Gaussian-tapered vertical columns, with alternating peak central amplitudes of ± 20 per cent perturbation. A width of 80 km was chosen as a reasonable size for resolvable features for the target period band of 10–30 s. The perturbation amplitude is chosen to be 20 per cent following the results of Tape *et al.* (2010), who observed maximum absolute velocity changes of 30 per cent with respect to their initial ray-based tomography model of southern California. The perturbations are

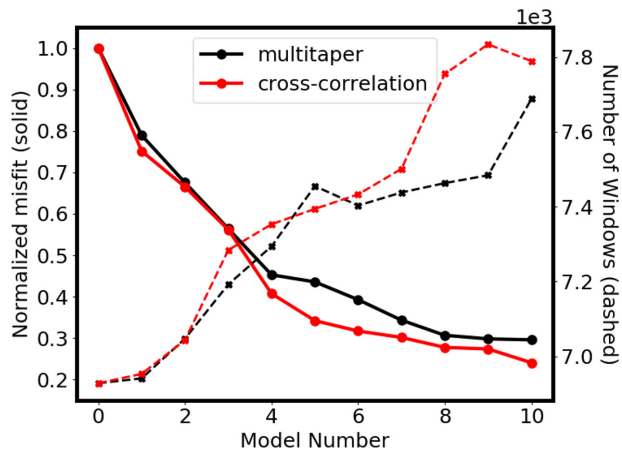


Figure 8. Synthetic inversion convergence behaviour comparing the traveltimes cross-correlation misfit and multitaper misfit objective functions. Misfit normalized to start at 1. Comparative behaviour is almost identical for this specific synthetic test case, unsurprising for a comparison of two phase-based objective functions. New misfit windows are evaluated at each function evaluation, showing an on-average increase in measurement number throughout the inversion. This is interpreted as a sign of increasing waveform improvement for each model update.

chosen to be columnar in depth, since the test data (10–30 s surface waves) are mostly sensitive to lateral variations. The perturbation checkerboard is overlain on the NZ-North velocity model and used to create realistic synthetic data for 30 events gathered at 58 receiver locations (Fig. 1).

In the iterative inversion procedure, model updates are computed using the L-BFGS algorithm (Liu & Nocedal 1989), which involves a local quadratic approximation of the objective function. We use a safeguarded backtracking line search (Modrak & Tromp 2016) to determine the step length of model perturbation. At each iteration, V_p and V_s are updated, but density is not estimated due to the limited sensitivity of surface waves to density structure (Nazarian & Stokoe 1984). Although adjoint source–parameter inversion (Kim *et al.* 2011) is feasible within the context of the workflow, it is not addressed, and misfit between data and synthetics is attributed to inaccuracies in the velocity models only.

Smoothing the gradient by convolution with a Gaussian is used to suppress non-uniqueness in the updated models. Based on the perturbation scales and target bandpass, we choose horizontal and vertical half-widths of the Gaussian to be 17.5 and 4 km, respectively. Additionally, misfit windows are re-evaluated at each function evaluation, allowing the objective function to change with each iteration. Because the introduction of misfit windows is gradual, changes in the misfit function are accounted for over multiple L-BFGS updates. Finding that L-BFGS always returned a descent direction and noting that L-BFGS restarts can be quite costly (Modrak & Tromp 2016), we refrained from restarting the optimization algorithm. To be shown, the model converges in a well-behaved manner, suggesting these are acceptable choices for this specific inversion scenario.

Updated models are analysed based on recovery of perturbations, waveform improvement and overall reduction in the objective function. For a given model i , perturbation recovery is shown relative to the starting model as the net model update $\ln(\mathbf{m}_i/\mathbf{m}_{00})$, which, to first order, approximates fractional differences $(\mathbf{m}_i/\mathbf{m}_{00} - 1)$. For large differences, the logarithmic misfit provides a more suitable

representation of compounding changes accrued during an iterative inversion.

4.2 Computational expense

To provide a sense of computational scale for these synthetic inversions, we outline the demands entailed by our preliminary analysis. To limit the overall computational expense, inversions are stopped after 10 iterations with each simulation lasting 300 s seismogram length. The numerical mesh described in Section 2.1 is accurate to 2 s periods, which is oversampled for our target period band of 10–30 s. We consequently remake the mesh with an accuracy of 10 s, which corresponds to 124 000 elements and a minimum element spacing of 2 km (Fig. 7). Each simulation requires roughly 10 min on 80 physical cores of the New Zealand eScience Infrastructure’s (NeSI) Cray XC50 supercomputer, named Maui. In the first few iterations, additional forward simulations are required by the L-BFGS algorithm to properly scale the step length, which typically results in three to five simulations per event per iteration. Once the search direction is well scaled, iterations typically require only two simulations (one forward, one adjoint) per event. An additional smoothing operation is required per iteration. Overall, a 10 iteration inversion required roughly 8000 CPU-hours, with a walltime of approximately 20 hr. After the initial manual tasks of mesh generation and parameter setup, all tasks ran fully automated until the end of the inversion.

4.3 Resolution and parameter testing

The synthetic inversions performed are used not only as a form of resolution testing, but also to illustrate the capability of our workflow to quantitatively guide methodological decisions for full-scale inversions. Choices of, for example, starting model, nonlinear optimization algorithm, bandpass corners, windowing criteria, smoothing lengths and objective function, will have significant effects on the behaviour, efficiency and robustness of an inversion. Our automated workflow facilitates the use of synthetic inversions to test the effects of various input configurations on inversion behaviour. Here, as an example, we show the results of inversions run using two different objective functions. One objective function is the cross-correlation traveltimes misfit,

$$\chi(\mathbf{m}) = \frac{1}{2} \left[\frac{T^{\text{obs}} - T(\mathbf{m})}{\sigma} \right]^2, \quad (2)$$

where T^{obs} is the observed traveltimes, $T(\mathbf{m})$ the corresponding synthetic traveltimes for a model \mathbf{m} and σ an optional weighting of the measurement uncertainty (Tromp *et al.* 2005). The other objective function is the multitaper misfit,

$$\chi(\mathbf{m}) = \frac{1}{2} \int_0^W W_p(\omega) \left| \frac{\tau^d(\omega) - \tau^s(\omega, \mathbf{m})}{\sigma_p(\omega)} \right|^2 d\omega, \quad (3)$$

where $\tau^d(\omega)$ is an observed frequency-dependent phase measurement, $\tau^s(\omega, \mathbf{m})$ is the synthetic phase measurement for model \mathbf{m} , σ_p is a measurement uncertainty and W_p is a tapered frequency-domain window (Tape 2009). Both objective functions are defined for a single time window. Although we only show results from two very-similar synthetic inversions, we emphasize that this capability can be extended to compare the effects of any range of input parameters.




Figure 9. Misfit histograms for a synthetic inversion using the cross-correlation traveltimes misfit function. Comparisons are made between the initial model m00 (orange) and final model m10 (blue). (a) Cross correlation time-shift in seconds. (b) Amplitude ratio $\Delta \ln(A)$ (Eq. 1). (c) Peak cross correlation value. (d) Window start time relative to event origin time. Mean, standard deviation, and median for each quantity and model are provided in the title. The number of measurements for each model is given in the respective legends. Histograms illustrate an overall improvement in waveforms. Results are similar for the multitaper inversion.

4.3.1 Recovered perturbations

We qualitatively analyse perturbation recovery for the two inversions. Fig. 7 shows the target V_s perturbations compared with final recovered perturbations using the cross-correlation traveltimes misfit and multitaper misfit in map view, a horizontal slice at 15 km, and in cross section. Both synthetic inversions are able to recover amplitudes and spatial distribution of perturbations for areas with sufficient data coverage. Visually, the two objective functions exhibit near-identical model recovery. As expected, areas with no source–receiver coverage show little to no change from the starting model. Features are resolved at the surface with pronounced smearing due to the uneven source–receiver distribution and the use of primarily surface wave derived measurements. Features are more sharply resolved at 15 km depth, which we attribute to enhanced coverage from events, which have an average depth of roughly 25 km. Cross sections across a single row of checkers show that V_s perturbations are recovered adequately to 25 km depth, and best below land. Recovery in the oceanic forearc region is constrained mostly to the upper 10 km. As expected, resolution is worse on the western half of the North Island where source and receiver coverage is more sparse. Recovered perturbations suggest that the source–receiver distribution in the chosen domain is capable of recovering features on land and immediately offshore, with depth coverage mostly constrained to the upper 25 km.

4.3.2 Misfit reduction

Convergence behaviour for each model is shown in Fig. 8. For both objective functions, after 10 iterations, the total misfit reduces by

roughly a factor of 3 with respect to the initial model. The cross-correlation inversion shows negligibly better misfit reduction compared to multitaper. Fig. 8 also shows the total number of measurements made for each model. Approximately 5000 measurements would be expected if one time window was chosen for each of the three components of every source–receiver pair. Roughly 7000 initial measurements are made for the initial model, which increases by approximately 700 measurements at model 10. Both objective functions see a similar increase in total number of measurements over the course of the inversion. For this specific synthetic inversion, the convergence behaviour of multitaper and cross-correlation misfits is almost identical. This is unsurprising, however, for two phase-based objective functions, evaluated for a long-period bandpass over a limited number of measurements.

Two additional, identical inversions were performed, where a fixed set of windows, chosen for the initial model, were reused throughout the entire inversion. With respect to the case shown above, the resultant convergence plot and map-views of recovered perturbations show similar comparative behaviour between the two objective functions (Fig. A5). Although misfit reduction is more constrained and surface features less sharply resolved, this can primarily be explained by the restricted number of measurements. We emphasize, however, that the results shown here are unique to our specific synthetic-inversion scenario, and that the effect of objective function, windowing parameters, etc., on inversion behaviour remains contingent on initial conditions.

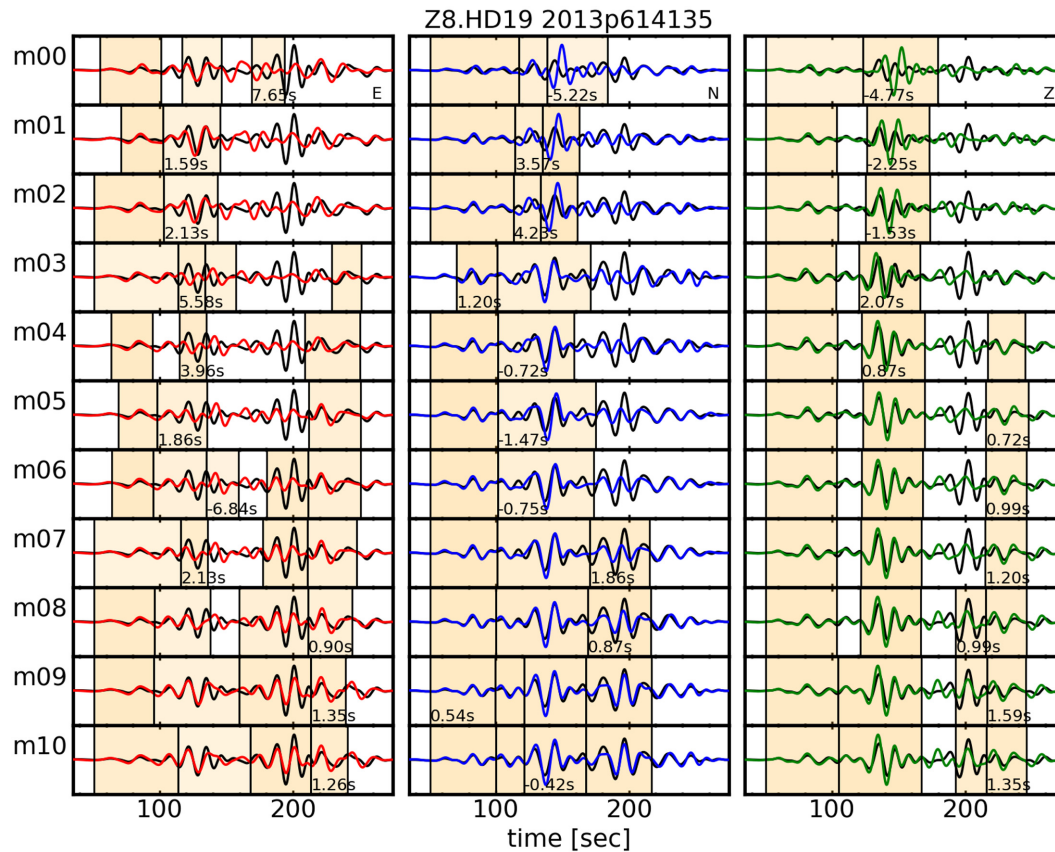


Figure 10. An example of synthetic inversion waveform improvement for the cross-correlation traveltime misfit function. Waveforms are shown from the initial model (m00) to the final model (m10) for three components: east–west (E), north–south (N) and vertical (Z). Observation waveforms shown in black, synthetic waveforms shown in colour. Waveforms shown in units of displacement. Time windows chosen by Pyflex are shown as orange windows. For simplicity, only the maximum time-shift for each pair of waveforms is shown. Time windows are re-evaluated for each new model.

4.3.3 Waveform improvement

Misfit histograms are used to document overall measurement improvement after 10 iterations, from the initial (m00) to the final model (m10). Fig. 9 shows results for the cross-correlation traveltime misfit inversion only, as the multitaper inversion results are very similar. Average time-shift (Fig. 9a) is reduced from -0.7 ± 2.1 to 0.0 ± 1.0 s. Amplitude differences are also reduced in mean and standard deviation; amplitude information is not considered in the phase-based objective function, so we would not expect amplitude differences to reduce to zero. Peak cross correlations are seen to shift towards a peak value of 1.0, and relative start times of measurements increase mostly in measurements made after 100 s, suggesting that an increased number of later arrivals and/or large propagation distances, are chosen using the final model.

A representative example of waveform improvement over the course of the inversion is shown for three components of a source receiver pair in Fig. 10. Observed and synthetic waveforms, along with misfit windows and maximum time-shift values are shown for each progressive model update. Large initial waveform differences, characterized by large time-shifts and amplitude ratios, are improved significantly throughout the inversion. By model 10, almost all available sections of the waveforms are windowed and aligned, and maximum time-shifts are small, or found in the latest arrivals (>200 s).

5 DISCUSSION

5.1 Choice of workflow tools

Our choice of workflow tools is based on the requirements of our earthquake tomography problem. We use existing tools like SPECIFEM3D, because they are accurate, efficient and available. SeisFlows has been chosen as the optimization tool in preference to other contenders such as SEISCOPE (Métivier & Brossier 2016), because it is written in Python and provides built-in integration with SPECIFEM2D/3D/3D.GLOBE. For misfit quantification and workflow management, we considered LASIF (Krischer *et al.* 2015a), which provides a framework for large-scale inversions, however attractive capabilities such as automated launching of simulations on HPC systems, or an accompanying optimization library for model updates, are not featured. The Salvus software suite (Afanasiev *et al.* 2019) provides an all-encompassing approach to full-waveform inversion but contains components that are currently closed-source. Similarly, inversion frameworks used in various published studies (e.g. Chen *et al.* 2007a; Lee *et al.* 2014c) are not currently available. To meet the requirements of the earthquake tomography problem, we have opted to design our own toolbox. We envision the flexible nature of SeisFlows, enhanced with the detailed measurement capabilities of Pyatoo, forming a tool chain applicable to future earthquake tomography problems.




Figure 11. Potential setup for a full-scale inversion within the NZ-North domain (dashed line). 250 earthquakes with magnitudes $4.4 \leq M_w < 6.0$ are shown as moment tensors, scaled by size with colour coded by depth. 100 broadband seismic station locations, consisting of New Zealand's permanent seismic network and temporary networks throughout the North Island, are shown as inverted triangles.

5.2 Potential applications of an automated inversion workflow

The methodological framework developed here is general and should be applicable, with minor modifications, to other regions with suitable data coverage and candidate velocity models. The comparison of synthetic inversions in this study is a simple example to illustrate the powerful capacity of an automated inversion workflow to inform and assess any choice of input parameters for a seismic inversion. The synthetic checkerboard inversion is a useful method for quantifying the resolving power of the source–receiver distribution for a given region, although it should not be seen as a replacement for rigorous quantitative analyses of model resolution, such as the point spread function of Fichtner & Trampert (2011). Though not implemented, the capability for adjoint moment tensor inversions (Kim *et al.* 2011) is feasible within the context of this workflow, and it remains a potential avenue for future efforts towards joint velocity model and source parameter inversions. As with the 2-D synthetic comparisons of optimization schema by Modrak & Tromp (2016), the ability to easily repeat inversions with slight modifications in input parameters opens up a wide range of possibilities for assessing inversion behaviour. For example, with the availability of various objective functions for seismic inversions, such as instantaneous phase differences and envelope ratios (Bozdağ *et al.* 2011), or amplitudes, attenuation and differential traveltimes (Tromp *et al.* 2005), more detailed analyses can be performed to understand the impact of objective functions on tomographic inversions.

5.3 Towards a full-scale North Island inversion

One overarching goal of this work is to set the foundation for a full-scale seismic inversion. Given our analysis of the NZ-North velocity model, as well as the results of our synthetic inversions, we expect a full-scale inversion to result in a substantial overall reduction in full-waveform misfit between data and synthetics, and improved accuracy in crustal images for the North Island. Our efforts will focus towards better resolution of shallow crustal structure, and deeper understanding of the Hikurangi subduction zone, Taupō Volcanic zone, and North Island tectonics. If 2 s waveforms can be fit, then we should expect spatial resolutions in the shallow region of approximately 1–2 km, corresponding to the lowest wave speeds ($\sim 1 \text{ km s}^{-1}$) in our starting model. Interpretations will be directly compared with results of active seismic experiments, geological profiles and tectonic studies of the area.

Our full-scale inversion will begin by fitting large scale features corresponding to long-period waveforms, subsequently moving to shorter periods with a target period band of 2–30 s. Data are available for 250 earthquakes recorded on up to 100 broadband seismic stations, including temporary seismic networks throughout the North Island (Fig. 11). Smoothing at large spatial scales may be necessary in early iterations to promote resolution of large scale structures before attempting to image finer details. If necessary, empirical pre-conditioners can be used to down weight the contribution of near-source and near-receiver structure to model updates, as in Fichtner *et al.* (2009). Additionally, data-category and geographical weighting of measurements can assist with the uneven source–receiver distribution in our domain (Ruan *et al.* 2019). A subset of our event catalog will not be included in the inversion, and retained for *post hoc* analysis of the inversion's ability to resolve true Earth structure. Full-scale resolution testing, using a point spread function, Hessian kernels, and/or full scale checkerboard inversion including vertical variations, can be performed to quantify the resolving capabilities of the full source–receiver distribution. Forward simulations on the numerical mesh described in Section 2.1 require 25 min on 160 physical cores of Maui. Basing our estimates on the synthetic inversions of Section 4.2, with a conservative expectation of 20 iterations to reach convergence, we estimate a full-scale inversion using 200 events to require on the order of 1 000 000 CPU-hours.

6 CONCLUSIONS

We present an automated inversion workflow for earthquake-based tomography using adjoint methods. The workflow is constructed from both newly developed and existing Python tools built around the widely used numerical solver SPECFEM3D Cartesian, providing end-to-end capability for adjoint tomography. We verify the workflow through automatic recovery of 3-D checkerboard perturbations via iterative model updates on an HPC system. We also use these synthetic inversion results to carry out resolution analysis and illustrate parameter testing capabilities. Additionally, synthetic seismograms generated using a candidate 3-D velocity model of New Zealand show reasonable misfit with respect to observed waveforms, illustrating the potential for improvement of the velocity model using adjoint methods. Comparisons with a 1-D North Island velocity model and similarly sized study regions provide optimism that NZ-North provides a strong candidate starting velocity model for adjoint tomography.

The success of the synthetic inversions and the reasonable misfit of the candidate velocity model, provide confidence moving

toward a large-scale seismic inversion of the North Island, New Zealand. Synthetic resolution tests highlight the potential for substantial model improvement both onshore and offshore, with well-resolved velocity changes extending from the surface to roughly 25 km depth. The automated and modular nature of the workflow means that a wide range of input parameters may be tested and chosen, to efficiently guide convergence of the inversion. A full-scale high-resolution seismic inversion of the North Island, New Zealand, using spectral element and adjoint methods, is the subject of on-going research.

ACKNOWLEDGEMENTS

This work was funded by a Rutherford Discovery Fellowship and Marsden Fund, awarded by the Royal Society of New Zealand. We thank editor Andrea Morelli, reviewer Cliff Thurber, and an anonymous reviewer for their comments which helped us improve the manuscript. We thank Donna Eberhart-Phillips for providing the latest New Zealand-wide velocity models. We acknowledge John Ristau for providing the moment tensor catalog. Broadband seismic data used in this study was provided by GeoNet. All numerical simulations were run on the New Zealand eScience Infrastructure (NeSI) HPC system Maui. We thank NeSI for the computational resources, and their assistance in optimizing SPECSEM3D on Maui.

REFERENCES

- Afanasyev, M., Boehm, C., van Driel, M., Krischer, L., Rietmann, M., May, D.A., Knepley, M.G. & Fichtner, A., 2019. Modular and flexible spectral-element waveform modelling in two and three dimensions, *Geophys. J. Int.*, **216**(3), 1675–1692.
- Bao, X. & Shen, Y., 2016. Assessing waveform predictions of recent three-dimensional velocity models of the Tibetan Plateau, *J. geophys. Res.*, **121**, 2521–2538.
- Barnes, P.M., et al., 2010. Tectonic and geological framework for gas hydrates and cold seeps on the Hikurangi subduction margin, New Zealand, *Mar. Geol.*, **272**(1–4), 26–48.
- Becker, J., et al., 2009. Global bathymetry and elevation data at 30 arc seconds resolution: SRTM30_PLUS, *Mar. Geod.*, **32**(4), 355–371.
- Behr, Y., Townend, J., Bannister, S. & Savage, M.K., 2011. Crustal shear wave tomography of the Taupō Volcanic Zone, New Zealand, via ambient noise correlation between multiple three-component networks, *Geochem. Geophys. Geosyst.*, **12**(3).
- Beyreuther, M., Barsch, R., Krischer, L., Megies, T., Behr, Y. & Wassermann, J., 2010. ObsPy: a python toolbox for seismology, *Seismol. Res. Lett.*, **81**(3), 530–533.
- Borisov, D., Modrak, R., Gao, F. & Tromp, J., 2018. 3D elastic full-waveform inversion of surface waves in the presence of irregular topography using an envelope-based misfit function, *Geophysics*, **83**(1), R1–R11.
- Bozdağ, E., Trampert, J. & Tromp, J., 2011. Misfit functions for full waveform inversion based on instantaneous phase and envelope measurements, *Geophys. J. Int.*, **185**(2), 845–870.
- Bozdağ, E., Peter, D., Lefebvre, M., Komatitsch, D., Tromp, J., Hill, J., Podhorszki, N. & Pugmire, D., 2016. Global adjoint tomography: first-generation model, *Geophys. J. Int.*, **207**(3), 1739–1766.
- Chen, M., Niu, F., Liu, Q., Tromp, J. & Zheng, X., 2015. Multiparameter adjoint tomography of the crust and upper mantle beneath East Asia: 1. Model construction and comparisons, *J. geophys. Res.*, **120**(3), 1762–1786.
- Chen, P., Zhao, L. & Jordan, T.H., 2007a. Full 3D tomography for the crustal structure of the Los Angeles region, *Bull. seism. Soc. Am.*, **97**(4), 1094–1120.
- Chen, P., Zhao, L. & Jordan, T.H., 2007b. Full 3D tomography for the crustal structure of the Los Angeles region, *Bull. seism. Soc. Am.*, **97**(4), 1094–1120.
- Clark, K., et al., 2019. Geological evidence for past large earthquakes and tsunamis along the Hikurangi subduction margin, New Zealand, *Mar. Geol.*, **412**, 139–172.
- Dahlen, F. & Baig, A.M., 2002. Fréchet kernels for body-wave amplitudes, *Geophys. J. Int.*, **150**(2), 440–466.
- Dreger, D.S., 2003. Tdmt.inv: time domain seismic moment tensor inversion, in *International Geophysics*, Vol. **81**, p. 1627, Elsevier.
- Eberhart-Phillips, D. & Fry, B., 2017. A new scheme for joint surface wave and earthquake travel-time inversion and resulting 3-D velocity model for the western North Island, New Zealand, *Phys. Earth planet. Inter.*, **269**, 98–111.
- Eberhart-Phillips, D. & Fry, B., 2018. Joint local earthquake and teleseismic inversion for 3-D velocity and Q in New Zealand, *Phys. Earth planet. Inter.*, **283**, 48–66.
- Eberhart-Phillips, D., Reyners, M., Bannister, S., Chadwick, M. & Ellis, S., 2010. Establishing a versatile 3-D seismic velocity model for New Zealand, *Seismol. Res. Lett.*, **81**(6), 992.
- Eberhart-Phillips, D., Bannister, S. & Ellis, S., 2014. Imaging P and S attenuation in the termination region of the Hikurangi subduction zone, New Zealand, *Geophys. J. Int.*, **198**(1), 516–536.
- Eberhart-Phillips, D., Reyners, M. & Bannister, S., 2015. A 3D QP attenuation model for all of New Zealand, *Seismol. Res. Lett.*, **86**(6), 1655–1663.
- Eberhart-Phillips, D., Bannister, S. & Reyners, M., 2017. Deciphering the 3-D distribution of fluid along the shallow Hikurangi subduction zone using P- and S-wave attenuation, *Geophys. J. Int.*, **211**(2), 1032–1045.
- Eberhart-Phillips, D., Bannister, S. & Reyners, M., 2020. Attenuation in the mantle wedge beneath super-volcanoes of the Taupō Volcanic Zone, New Zealand, *Geophys. J. Int.*, **220**(1), 703–723.
- Ekström, G., Nettles, M. & Dziewoński, A., 2012. The global CMT project 2004–2010: centroid-moment tensors for 13,017 earthquakes, *Phys. Earth planet. Inter.*, **200**, 1–9.
- Fichtner, A. & Trampert, J., 2011. Resolution analysis in full waveform inversion, *Geophys. J. Int.*, **187**(3), 1604–1624.
- Fichtner, A., Kennett, B.L., Igel, H. & Bunge, H.-P., 2009. Full seismic waveform tomography for upper-mantle structure in the Australasian region using adjoint methods, *Geophys. J. Int.*, **179**(3), 1703–1725.
- Fichtner, A., Trampert, J., Cupillard, P., Saygin, E., Taymaz, T., Capdeville, Y. & Villasenor, A., 2013. Multiscale full waveform inversion, *Geophys. J. Int.*, **194**(1), 534–556.
- Gao, H. & Shen, Y., 2015. Validation of recent shear wave velocity models in the United States with full-wave simulation, *J. geophys. Res.*, **120**, 344–358.
- Hamling, I.J., et al., 2017. Complex multifault rupture during the 2016 Mw 7.8 Kaikōura earthquake, New Zealand, *Science*, **356**(6334), doi:10.1126/science.aam7194.
- Holden, C., Kaneko, Y., D’Anastasio, E., Benites, R., Fry, B. & Hamling, I., 2017. The 2016 Kaikōura earthquake revealed by kinematic source inversion and seismic wavefield simulations: Slow rupture propagation on a geometrically complex crustal fault network, *Geophys. Res. Lett.*, **44**(22), 11–320.
- Kaneko, Y., et al., 2019. Ultra-long duration of seismic ground motion arising from a thick, low-velocity sedimentary wedge, *J. geophys. Res.*, **124**(10), 10 347–10 359.
- Kim, Y., Liu, Q. & Tromp, J., 2011. Adjoint centroid-moment tensor inversions, *Geophys. J. Int.*, **186**(1), 264–278.
- Komatitsch, D. & Tromp, J., 2002a. Spectral-element simulations of global seismic wave propagation—I. Validation, *Geophys. J. Int.*, **149**(2), 390–412.
- 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. & Vilotte, J.-P., 1998. The spectral element method: an efficient tool to simulate the seismic response of 2D and 3D geological structures, *Bull. seism. Soc. Am.*, **88**(2), 368–392.
- Komatitsch, D., Ritsema, J. & Tromp, J., 2002. The spectral-element method, Beowulf computing, and global seismology, *Science*, **298**, 1737–1742.
- Krischer, L., 2015a. Pyadjoint, <http://krischer.github.io/pyadjoint/>.

- Krischer, L. & Casarotti, E., 2015b. Pyflex: 0.1.4 (Version 0.1.4), <https://github.com/krischer/pyflex>.
- Krischer, L., Fichtner, A., Zukauskaitė, S. & Igel, H., 2015a. Large-scale seismic inversion framework, *Seismol. Res. Lett.*, **86**(4), 1198–1207.
- Krischer, L., Megies, T., Barsch, R., Beyreuther, M., Lecocq, T., Caudron, C. & Wassermann, J., 2015b. ObsPy: a bridge for seismology into the scientific Python ecosystem, *Comput. Sci. Discov.*, **8**(1), 014003.
- Krischer, L., et al., 2016. An adaptable seismic data format, *Geophys. Suppl. Mon. Not. R. Astron. Soc.*, **207**(2), 1003–1011.
- Lee, E.-J., Chen, P. & Jordan, T.H., 2014a. Testing waveform predictions of 3D velocity models against two recent Los Angeles earthquakes, *Seismol. Res. Lett.*, **85**(6), 1275–1284.
- Lee, E.-J., Chen, P., Jordan, T.H., Maechling, P.B., Denolle, M.A. & Beroza, G.C., 2014b. Full-3-D tomography for crustal structure in southern California based on the scattering-integral and the adjoint-wavefield methods, *J. geophys. Res.*, **119**(8), 6421–6451.
- Lee, E.-J., Chen, P., Jordan, T.H., Maechling, P.B., Denolle, M.A.M. & Beroza, G.C., 2014c. Full-3-D tomography for crustal structure in Southern California based on the scattering-integral and the adjoint-wavefield methods, *J. geophys. Res.*, **119**, 6421–6451.
- Lin, F.-C., Ritzwoller, M., Townend, J., Bannister, S. & Savage, M., 2007. Ambient noise Rayleigh wave tomography of New Zealand, *Geophys. J. Int.*, **170**(2), 649–666.
- Lin, F.-C., Moschetti, M. & Ritzwoller, M., 2008. Surface wave tomography of the western United States from ambient seismic noise: Rayleigh and Love wave phase velocity maps, *Geophys. J. Int.*, **173**(1), 281–298.
- Lin, Y.-P., Zhao, L. & Hung, S.-H., 2011. Assessment of tomography models of Taiwan using first-arrival times from the TAIGER active-source experiment, *Bull. seism. Soc. Am.*, **101**(2), 866–880.
- Liu, D.C. & Nocedal, J., 1989. On the limited memory BFGS method for large scale optimization, *Math. Program.*, **45**(1–3), 503–528.
- Liu, Q. & Gu, Y., 2012. Seismic imaging: from classical to adjoint tomography, *Tectonophysics*, **566**, 31–66.
- Luo, Y. & Schuster, G.T., 1991. Wave-equation traveltime inversion, *Geophysics*, **56**, 645–653.
- Maggi, A., Tape, C., Chen, M., Chao, D. & Tromp, J., 2009. An automated time-window selection algorithm for seismic tomography, *Geophys. J. Int.*, **178**(1), 257–281.
- Matharu, G. & Sacchi, M.D., 2018. Source encoding in multiparameter full waveform inversion, *Geophys. J. Int.*, **214**, 792–810.
- Métivier, L. & Brossier, R., 2016. The SEISCOPE optimization toolbox: a large-scale nonlinear optimization library based on reverse communication, *Geophysics*, **81**(2), F1–F15.
- Miyoshi, T., Obayashi, M., Peter, D., Tono, Y. & Tsuboi, S., 2017. Adjoint tomography of the crust and upper mantle structure beneath the Kanto region using broadband seismograms, *Prog. Earth planet. Sci.*, **4**(1), 1–20.
- Modrak, R. & Tromp, J., 2016. Seismic waveform inversion best practices: regional, global and exploration test cases, *Geophys. J. Int.*, **206**(3), 1864–1889.
- Modrak, R.T., Borisov, D., Lefebvre, M. & Tromp, J., 2018. SeisFlows–Flexible waveform inversion software, *Comput. Geosci.*, **115**, 88–95.
- Nazarian, S. & Stokoe, K.H., 1984. Nondestructive testing of pavements using surface waves, *Transport. Res. Record*, **993**, 67–79.
- Pandas Development Team, 2020. pandas-dev/pandas: Pandas.
- Rawlinson, N., Pilia, S., Young, M., Salmon, M. & Yang, Y., 2016. Crust and upper mantle structure beneath southeast Australia from ambient noise and teleseismic tomography, *Tectonophysics*, **689**, 143–156.
- Reyners, M., 1998. Plate coupling and the hazard of large subduction thrust earthquakes at the Hikurangi subduction zone, New Zealand, *N. Z. J. Geol. Geophys.*, **41**(4), 343–354.
- Ristau, J., 2008. Implementation of routine regional moment tensor analysis in New Zealand, *Seismol. Res. Lett.*, **79**(3), 400–415.
- Ristau, J., 2013. Update of regional moment tensor analysis for earthquakes in New Zealand and adjacent offshore regions, *Bull. seism. Soc. Am.*, **103**(4), 2520–2533.
- Robinson, R. & Webb, T.H., 1996. *AMPRAT and MECHTOOL: Programs for Determining Focal Mechanisms of Local Earthquakes*, Vol. **96**, Institute of Geological & Nuclear Sciences Ltd.
- Rowland, J.V., Wilson, C.J. & Gravley, D.M., 2010. Spatial and temporal variations in magma-assisted rifting, Taupō Volcanic Zone, New Zealand, *J. Volc. Geotherm. Res.*, **190**(1), 89–108.
- Ruan, Y., Lei, W., Modrak, R., Örsvuran, R., Bozdağ, E. & Tromp, J., 2019. Balancing unevenly distributed data in seismic tomography: a global adjoint tomography example, *Geophys. J. Int.*, **219**(2), 1225–1236.
- Rusmanugroho, H., Modrak, R. & Tromp, J., 2017. Anisotropic full-waveform inversion with tilt-angle recovery, *Geophysics*, **82**(3), 1–17.
- Shapiro, N., Campillo, M., Stehly, L. & Ritzwoller, M., 2005. High-resolution surface-wave tomography from ambient seismic noise, *Science*, **307**(5715), 1615–1618.
- Smith, J.A., et al., 2019. Tunnel detection at Yuma Proving Ground, Arizona, USA—Part 2: 3D full-waveform inversion experiments, *Geophysics*, **84**(1), B95–B108.
- Stirling, M., et al., 2012. National seismic hazard model for New Zealand: 2010 update, *Bull. seism. Soc. Am.*, **102**(4), 1514–1542.
- Taborda, R., Azizzadeh-Roodpish, S., Khoshnevis, N. & Cheng, K., 2016. Evaluation of the southern California seismic velocity models through simulation of recorded events, *Geophys. J. Int.*, **205**, 1342–1364.
- Talagrand, O. & Courtier, P., 1987. Variational assimilation of meteorological observations with the adjoint vorticity equation. I: theory, *Quart. J. R. Meteorol. Soc.*, **113**(478), 1311–1328.
- Tape, C., Liu, Q., Maggi, A. & Tromp, J., 2010. Seismic tomography of the southern California crust based on spectral-element and adjoint methods, *Geophys. J. Int.*, **180**(1), 433–462.
- Tape, C.H., 2009. Seismic tomography of southern California using adjoint methods, *PhD thesis*, California Institute of Technology.
- Tarantola, A., 1984. Inversion of seismic reflection data in the acoustic approximation, *Geophysics*, **49**(8), 1259–1266.
- Tromp, J. & Bachmann, E., 2019a. Source encoding for adjoint tomography, *Geophys. J. Int.*, **218**, 2019–2044.
- Tromp, J. & Bachmann, E., 2019b. Source encoding for adjoint tomography, *Geophys. J. Int.*, **218**(3), 2019–2044.
- Tromp, J., Tape, C. & Liu, Q., 2005. Seismic tomography, adjoint methods, time reversal and banana-doughnut kernels, *Geophys. J. Int.*, **160**(1), 195–216.
- Wallace, L., et al., 2009. Characterizing the seismogenic zone of a major plate boundary subduction thrust: Hikurangi Margin, New Zealand, *Geochem. Geophys. Geosyst.*, **10**(10), doi:10.1029/2009GC002610.
- Wallace, L.M., Beavan, J., McCaffrey, R. & Darby, D., 2004. Subduction zone coupling and tectonic block rotations in the North Island, New Zealand, *J. geophys. Res.*, **109**(B12), doi.org/10.1029/2004JB003241.
- Wallace, L.M., Beavan, J., Bannister, S. & Williams, C., 2012. Simultaneous long-term and short-term slow slip events at the Hikurangi subduction margin, New Zealand: implications for processes that control slow slip event occurrence, duration, and migration, *J. geophys. Res.*, **117**(B11), doi:10.1029/2012JB009489.
- Wes, McKinney, 2010. Data Structures for Statistical Computing in Python, in *Proceedings of the 9th Python in Science Conference*, pp. 56–61.
- Williams, C.A., Eberhart-Phillips, D., Bannister, S., Barker, D.H., Henrys, S., Reyners, M. & Sutherland, R., 2013. Revised interface geometry for the Hikurangi subduction zone, New Zealand, *Seismol. Res. Lett.*, **84**(6), 1066–1073.
- Wilson, C., Houghton, B., McWilliams, M., Lanphere, M., Weaver, S. & Briggs, R., 1995. Volcanic and structural evolution of Taupo Volcanic Zone, New Zealand: a review, *J. Volc. Geotherm. Res.*, **68**(1–3), 1–28.
- Wilson, C., Gravley, D., Leonard, G. & Rowland, J., 2009. Volcanism in the central Taupo Volcanic Zone, New Zealand: tempo, styles and controls, in *Studies in Volcanology: The Legacy of George Walker*, Special Publications of IAVCEI, Vol. 2, pp. 225–247, eds T. Thordarson, S. Self, G. Larsen, S. K. Rowland, Á. Höskuldsson, Geological Society of London.
- Yao, H., Van Der Hilst, R. & Montagner, J.-P., 2010. Heterogeneity and anisotropy of the lithosphere of SE Tibet from surface wave array tomography, *J. geophys. Res.*, **115**(12), doi:10.1029/2009JB007142.
- Zhu, H., Bozdağ, E. & Tromp, J., 2015. Seismic structure of the European upper mantle based on adjoint tomography, *Geophys. J. Int.*, **201**(1), 18–52.

APPENDIX A: WINDOWING PARAMETERS

FLEXWIN/Pyflex provides a set of user-defined parameters which control the behaviour of the windowing algorithm. Among the available parameters, the most important control the maximum allowable time-shift, amplitude ratio and cross-correlation values. Fine-tuning parameters are used to reject windows from an initial set of chosen windows based on STA/LTA waveform features. Table A1 provides the parameters used for windowing measurements in this study. For detailed descriptions of each parameter, see Maggi *et al.* (2009).

APPENDIX B: COMPARISONS WITH A 1D NORTH ISLAND VELOCITY MODEL

We perform the same analysis of Section 2.2 for a 1-D North Island velocity model defined by Ristau (2008). 250 sources and 45 station

locations are used. Using the same windowing parameters defined for the NZ-North results, a very small number of misfit windows was retrieved. A less conservative set of windowing parameters was defined to collect a more reasonable number of measurements, 9834 total. As expected, misfit histograms (Fig. A2) show that overall time-shift, amplitude differences and peak cross-correlation values are objectively worse for the 1-D velocity model synthetics. Comparisons of representative waveforms (Fig. A1) suggest that although the 1-D velocity model synthetics capture amplitude information better for these specific source–receiver paths, the NZ-North synthetics are more capable of capturing accurate arrival times, and complex phases such as surface waves. These improvements ultimately lead to a more diverse set of waveform measurements, and consequently a more informed tomographic inversion.

Table A1. Pyflex/FLEXWIN windowing parameters used in this study. NZ3D 10–30 parameters were used for the bulk misfit analyses of Section 2, and the synthetic inversions of Section 4.1. NZ3D 2–30 parameters were used to gather misfit windows for the 2–30 s waveforms in Section 2.2.1. NZ1D parameters were used to select windows for the synthetics generated using the 1-D North Island velocity model shown in Section 2.2.2 and Appendix B. The ‘<’ symbol denotes that the given parameter is the same as the value in the column to the left.

Parameter	Description	NZ3D 10–30	NZ3D 2–30	NZ1D
$T_{0,1}$	min, max period	10, 30	2, 30	10, 30
$r_{p,A}$	waveform signal-to-noise	3.5, 3.0	<	<
r_0	window signal-to-noise	3.0	<	<
w_E	STA/LTA water level	0.1	<	0.07
CC_0	CC acceptance	0.7	<	<
ΔT_0	time-shift acceptance	8.0	<	10.0
$\Delta \ln A_0$	amplitude ratio acceptance	2.0	<	<
ΔT_{ref}	time-shift reference	0.0	<	4.0
$\Delta \ln A_{ref}$	amplitude ratio reference	0.0	<	<
c_0	water level rejection	0.7	<	<
c_1	minimum window length	2.0	5.0	2.0
c_2	prominence rejection	0.0	<	<
$c_{3,a,b}$	phase separation	3.0, 2.0	<	<
$c_{4,a,b}$	signal emergence	2.5, 12.0	<	<

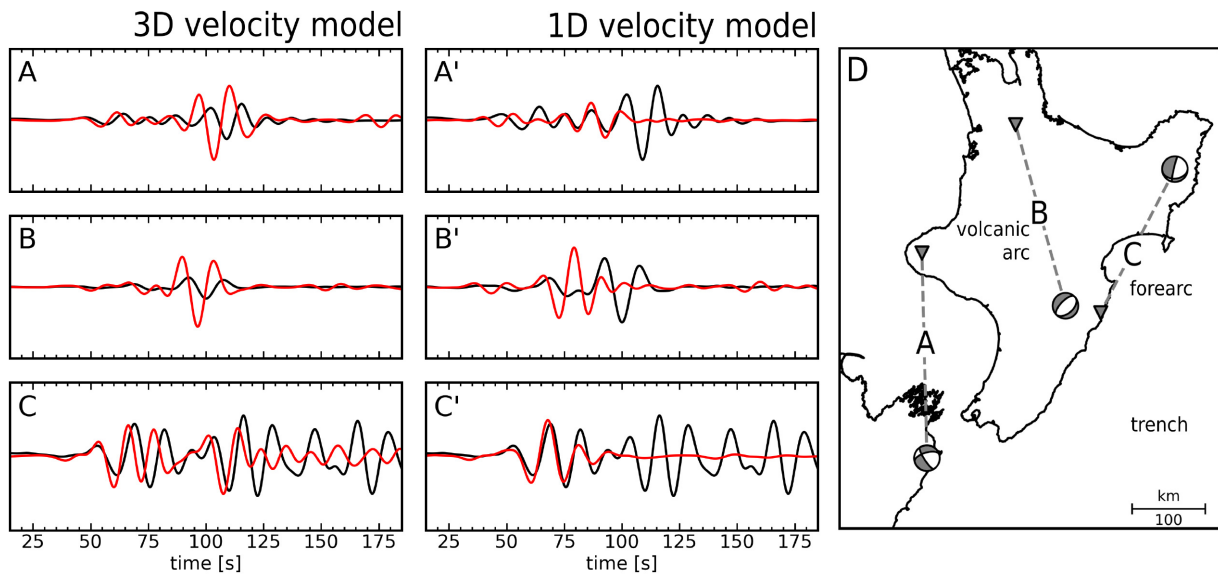


Figure A1. A comparison of vertical-component synthetic waveforms generated using two North Island specific velocity models. Waveforms are shown in units of displacement, filtered at 10–30 s periods. Observations in black, synthetics in red. (a–c) Synthetics from the 3-D NZ-North velocity model of Eberhart-Phillips *et al.* (a’–c’) Synthetics generated using the North Island 1-D velocity model of Ristau (2008). (d) Source–receiver map corresponding to waveform labels, with ray paths, focal mechanisms and labels denoting general subduction zone features.

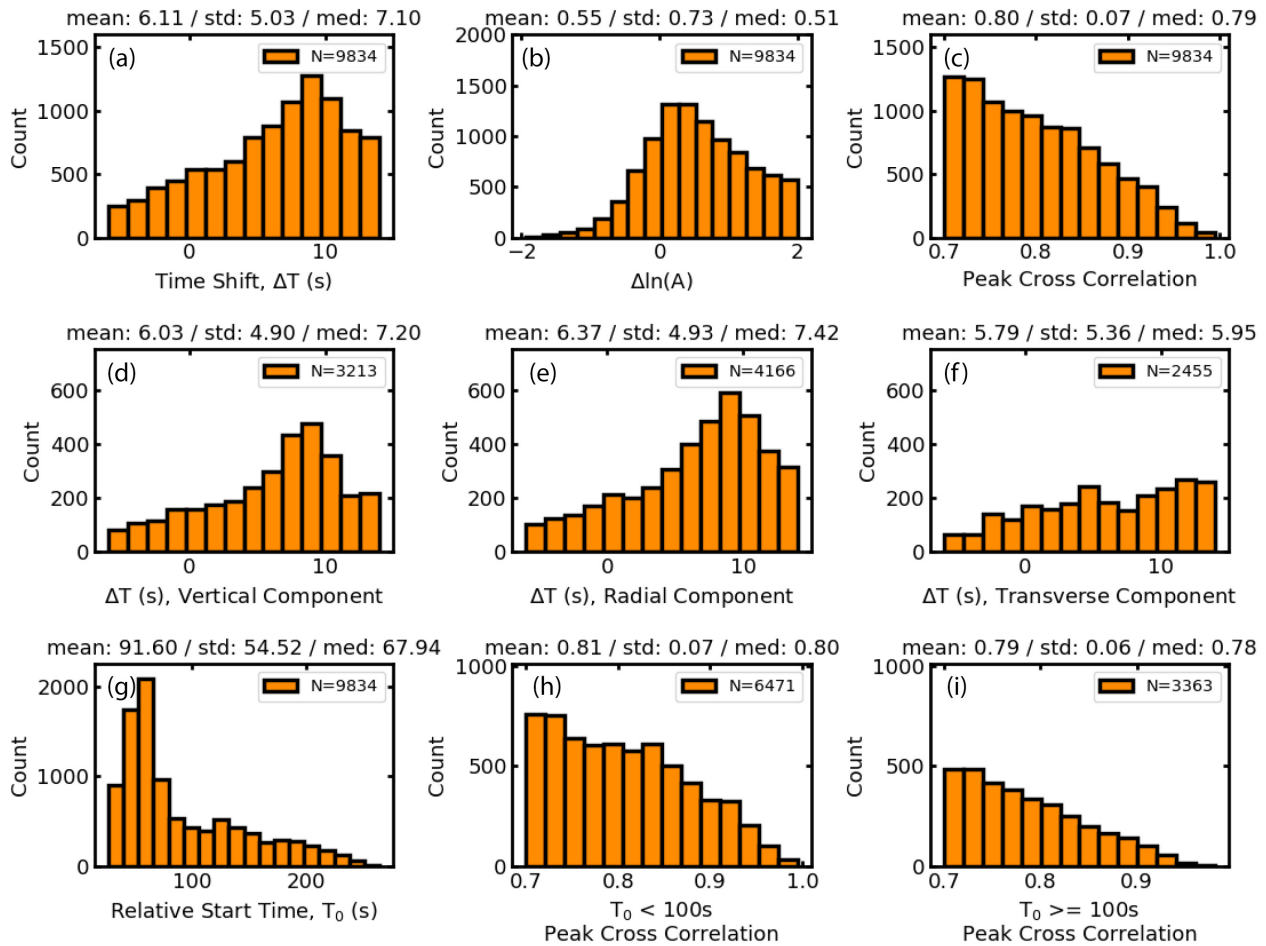


Figure A2. Misfit histograms for forward simulations using the 1-D North Island tomographic velocity model of Ristau (2008), in the same format as Fig. 4. Measurements from 250 events recorded on three components of 45 broadband stations, with 9834 total measurements. Mean, standard deviation, and median given in the titles of each histogram, number of measurements for each provided in the legend. Windowing parameters are less conservative than those used to generate the histograms in Fig. 4. (a) Time-shift corresponding to peak cross correlation between data and synthetics, in units of seconds. (b) Data-synthetic amplitude difference $\Delta \ln(A)$ (Eq. 1). (c) Peak cross correlation value. (d–f) Time-shift for vertical, radial and transverse component measurements. (g) Measurement start time relative to event origin time. (h) Peak cross correlations for measurements made before 100 s relative start time. (i) Peak cross correlations for measurements made at or after 100 s relative start time.

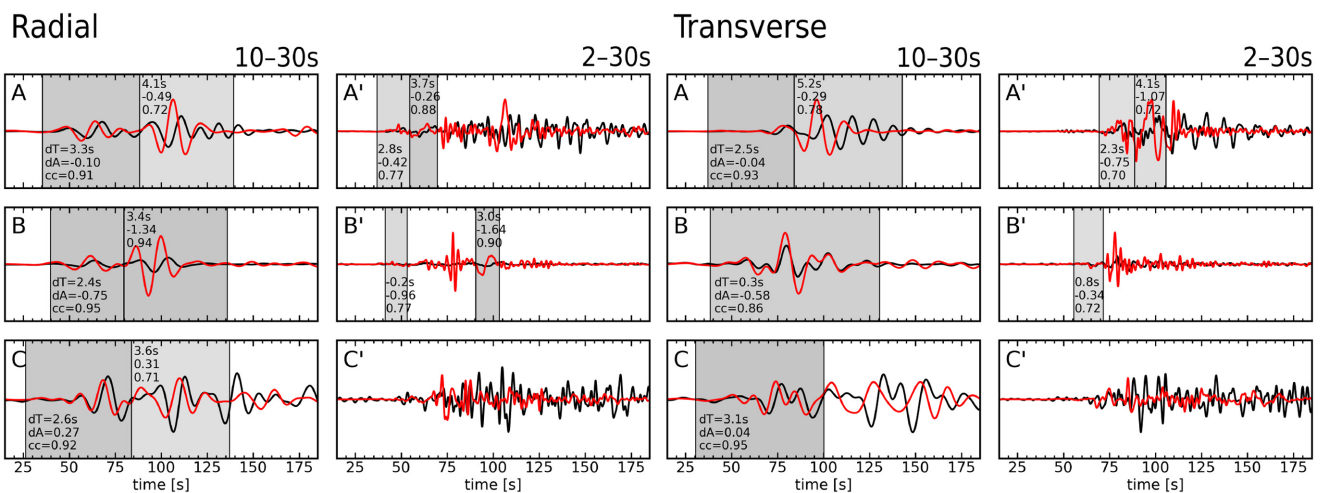


Figure A3. A representative selection of radial and transverse component data-synthetic comparisons. The same format is used as in Fig. 3. No discernible differences in radial or transverse component waveform quality, with respect to the vertical component, for this set of source–receiver pairs. Waveform labels correspond to the map in Fig. A1.

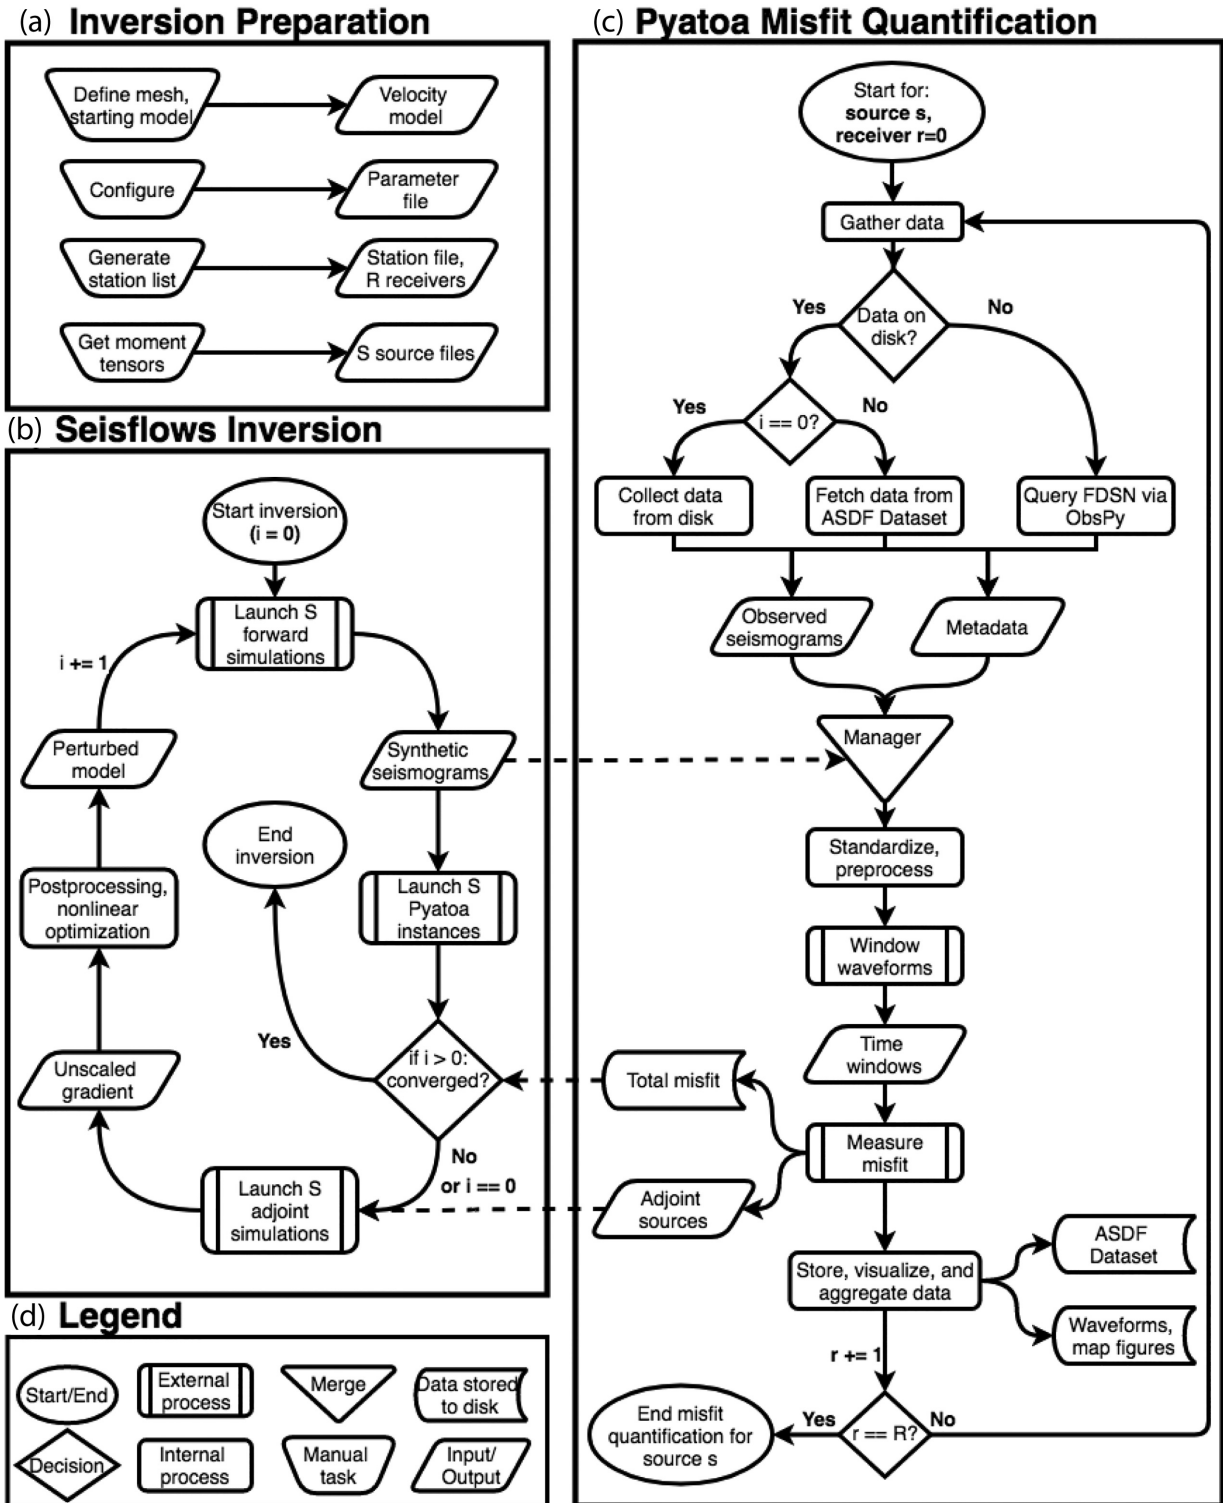


Figure A4. A flowchart representing the inversion workflow. Dashed lines show interactions between SeisFlows and Pyatoa via data stored to disk. (a) A one-time manual preparation is required to generate the velocity model, define a single parameter file shared by SeisFlows and Pyatoa, and create source and station files. (b) A simplified SeisFlows inversion, ignoring the complexities involved in, for example, the optimization schema, workflow management, HPC job handling. SeisFlows calls Pyatoa as a subroutine to calculate misfit and generate adjoint sources. (c) A Pyatoa misfit quantification instance. Data is preferentially gathered from disk, but online web service queries are available via the International Federation of Digital Seismograph Networks (FDSN). Synthetic seismograms from SPECfEM are made accessible by SeisFlows. A custom data structure (Manager) is responsible for collecting, processing and measuring waveforms. Data are saved to disk using ASDF Datasets. Pyatoa provides additional auxiliary input files required for subsequent adjoint simulations launched by SeisFlows. (d) Legend describing the flowchart shapes.

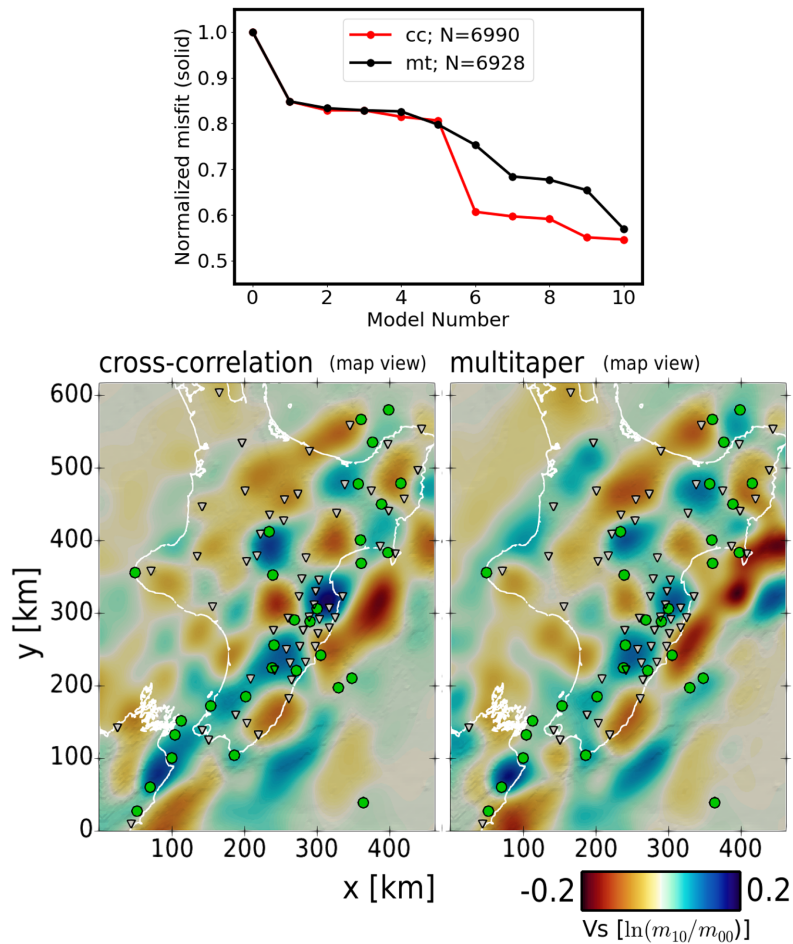


Figure A5. Convergence behaviour and map view of model 10 for synthetic inversions using the traveltime cross-correlation misfit (cc) and multitaper misfit (mt) functions, with a fixed set of time windows. Same as Fig. 8 but with the same set of misfit windows, derived using the initial model, reused throughout the inversion. Misfit normalized to start at 1. As expected, misfit convergence behaviour is shown to be more constrained than the open-window case. As with Fig. 8, the final misfit is negligibly lower using the cross-correlation misfit, for this specific synthetic inversion scenario.