An automated time-window selection algorithm for seismic tomography

Alessia Maggi,[1] Carl Tape,[2] Min Chen,[2,*] Daniel Chao[2] and Jeroen Tromp[2,†]

[1]*Institut de Physique du Globe de Strasbourg (*UMR 7516 CNRS, *Université de Strasbourg/EOST), 5 rue René Descartes, 67084 Strasbourg Cedex, France.*
E-mail: alessia@sismo.u-strasbg.fr
[2]*Seismological Laboratory, California Institute of Technology, Pasadena,* CA 91125, *USA*

SUMMARY
We present FLEXWIN, an open source algorithm for the automated selection of time windows on pairs of observed and synthetic seismograms. The algorithm was designed specifically to accommodate synthetic seismograms produced from 3-D wavefield simulations, which capture complex phases that do not necessarily exist in 1-D simulations or traditional traveltime curves. Relying on signal processing tools and several user-tuned parameters, the algorithm is able to include these new phases and to maximize the number of measurements made on each seismic record, while avoiding seismic noise. Our motivation is to use the algorithm for iterative tomographic inversions, in which the synthetic seismograms change from one iteration to the next. Hence, automation is needed to handle the volume of measurements and to allow for an increasing number of windows at each model iteration. The algorithm is sufficiently flexible to be adapted to many tomographic applications and seismological scenarios, including those based on synthetics generated from 1-D models. We illustrate the algorithm using data sets from three distinct regions: the entire globe, the Japan subduction zone, and southern California.

Key words: Time series analysis; Tomography; Body waves; Surface waves and free oscillations; Seismic tomography.

GJI Seismology

1 INTRODUCTION

Seismic tomography—the process of imaging the 3-D structure of the Earth using seismic recordings—has been transformed by recent advances in methodology. Finite-frequency approaches are being used instead of ray-based techniques, and 3-D reference models instead of 1-D reference models. These transitions are motivated by a greater understanding of the volumetric sensitivity of seismic measurements (Marquering *et al.* 1999; Zhao *et al.* 2000; Dahlen *et al.* 2000) and by computational advances in the forward modelling of seismic wave propagation in fully 3-D media (Komatitsch & Vilotte 1998; Komatitsch *et al.* 2002; Capdeville *et al.* 2003). In the past decade we have learned to calculate analytic sensitivity kernels in 1-D media (e.g. Li & Tanimoto 1993; Dahlen & Baig 2002; Dahlen & Zhou 2006) and numeric sensitivity kernels in 3-D media (e.g. Capdeville 2005; Tromp *et al.* 2005; Zhao *et al.* 2005; Liu & Tromp 2006; Liu & Tromp 2008). The analytic kernels have been taken up rapidly by tomographers, and used to produce new 3-D earth models (e.g. Montelli *et al.* 2004; Zhou *et al.* 2006).

The numeric kernels have opened up the possibility of '3D–3D' tomography, that is, seismic tomography based upon a 3-D reference model, 3-D numerical simulations of the seismic wavefield, and finite-frequency sensitivity kernels (Tromp *et al.* 2005; Chen *et al.* 2007b).

It is common practice in tomography to work only with certain subsets of the available seismic data. The choices made in selecting these subsets are inextricably linked to the assumptions made in the tomographic method. For example, ray-based traveltime tomography deals only with high-frequency body-wave arrivals, whereas great-circle surface-wave tomography must satisfy the path-integral approximation, and only considers surface-waves that present no evidence of multipathing. In both these examples, a large proportion of the information contained within the seismograms is unused. The emerging 3D–3D tomographic methods take advantage of full wavefield simulations and numeric finite-frequency kernels, thereby reducing the data restrictions required when using approximate forward modelling and simplified descriptions of sensitivity. These methods seem to be the best candidates for studying regions with complex 3-D structure, as they permit the use of a larger proportion of the information contained within each seismogram, including complex arrivals not predicted by 1-D approximations of earth structure. To exploit the full power of 3D–3D tomographic methods, we require a new data selection strategy that does not exclude such complex arrivals.

*Now at: Department of Earth, Atmospheric and Planetary Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139, USA.
†Now at: Department of Geosciences, Princeton University, Princeton, NJ 08540, USA.

As data selection strategies for tomography depend so closely on the tomographic technique, there are nearly as many such strategies as there are tomographic methods. Furthermore, many of these strategies have been automated in some way, as larger and larger volumes of data have become available. Body-wave studies that have moved away from using manual traveltime picks or catalogue arrival times generally pick windows around specific seismic phases defined by predicted traveltimes, and include automated tests on arrival time separation and/or the fit of observed to synthetic waveforms to reject inadequate data (e.g. Ritsema & van Heijst 2002; Lawrence & Shearer 2008). Partial automation of the vanDecar & Crosson (1990) multichannel cross-correlation method has led to efficient methods for obtaining highly accurate traveltime (Sigloch & Nolet 2006; Houser *et al.* 2008) and even attenuation (Lawrence *et al.* 2006) measurements. In the surface-wave community, there has been much work done to automate methods for extracting dispersion characteristics of fundamental mode (Trampert & Woodhouse 1995; Laske & Masters 1996; Ekström *et al.* 1997; Levshin & Ritzwoller 2001) and higher mode (Van Heijst & Woodhouse 1997; Debayle 1999; Yoshizawa & Kennett 2002; Beucler *et al.* 2003; Lebedev *et al.* 2005; Visser *et al.* 2007) surface-waves. Recently, Panning & Romanowicz (2006) have described an algorithm to semi-automatically pick body and surface-wavepackets based on the predicted traveltimes of several phases.

Our algorithm is designed for tomographic applications with 3-D Earth reference models. Unlike the techniques discussed above, ours is not tied to arrival time predictions of known phases, and, therefore, is able to accommodate complex phases due to 3-D structure. One promising approach to 3D–3D tomography is based upon adjoint methods (Tarantola 1984; Tromp *et al.* 2005; Liu & Tromp 2006; Tape *et al.* 2007). In 'adjoint tomography' the sensitivity kernels that tie variations in earth model parameters to variations in the misfit are obtained by interaction between the wavefield used to generate the synthetic seismograms (the direct wavefield) and an adjoint wavefield that obeys the same wave equation as the direct wavefield, but with a source term, that is derived from the misfit measurements. The computational cost of such kernel computations for use in seismic tomography depends only on the number of events, and not on the number of receivers nor on the number of measurements. It is therefore to our advantage to make the greatest number of measurements on each seismogram. The adjoint kernel calculation

procedure allows us to measure and use for tomographic inversion almost any part of the seismic signal. We do not need to identify specific seismic phases, as the kernel will take care of defining the relevant sensitivities. However, there is nothing in the adjoint method itself that prevents us from constructing an adjoint kernel from noise-dominated data, thereby polluting our inversion. An appropriate data selection strategy for adjoint tomography should therefore define measurement time windows that cover as much of a given seismogram as possible, whilst avoiding portions of the waveform that are dominated by noise.

From a signal processing point of view, the simplest way to avoid serious contamination by noise is to select and measure strong signals, which in seismology correspond to seismic arrivals. Our strategy is therefore to select time windows on the synthetic seismogram within which the waveform contains a distinct arrival, and then require an adequate correspondence between observed and synthetic waveforms within these windows. This selection paradigm is general, and can be applied to synthetic seismograms regardless of how they have been obtained. It is clear, however, that a synthetic seismogram obtained by 3-D propagation through a good 3-D Earth model will provide a better fit to the observed seismogram over a greater proportion of its length than will be the case for a more approximate synthetic seismogram.

To isolate changes in amplitude or frequency content potentially associated with distinct arrivals, we need to analyse the character of the synthetic waveform itself. This analysis is similar to that used on observed waveforms in automated phase detection algorithms for the routine location of earthquakes. In designing our time-window selection algorithm, we have taken a tool used in this detection process—the short-term average/long-term average ratio—and applied it to the definition of time windows around distinct seismic phases.

The choices made in time-window selection for tomography are interconnected with all aspects of the tomographic inversion process, from the waveform simulation method (direct problem), through the choice of measurement method, to the method used to obtain sensitivity kernels, and the inversion method itself. One of the major difficulties in defining a general data selection strategy is the great range of possible choices open to the tomographer. We have designed a configurable data selection process that can be adapted to different tomographic scenarios by tuning a handful of parameters (see Table 1). Although we have designed our algorithm

Table 1. Overview of standard tuning parameters, and of fine tuning parameters. Values are defined in a parameter file, and the time dependence of those that depend on time is described by user-defined functions.

Standard tuning parameters:	
$T_{0,1}$	bandpass filter corner periods
$r_{P,A}$	signal to noise ratios for whole waveform
$r_0(t)$	signal to noise ratios single windows
$w_E(t)$	water level on short-term:long-term ratio
$CC_0(t)$	acceptance level for normalized cross-correlation
$\Delta\tau_0(t)$	acceptance level for time lag
$\Delta\ln A_0(t)$	acceptance level for amplitude ratio
$\Delta\tau_{ref}$	reference time lag
$\Delta\ln A_{ref}$	reference amplitude ratio
Fine tuning parameters:	
c_0	for rejection of internal minima
c_1	for rejection of short windows
c_2	for rejection of un-prominent windows
$c_{3a,b}$	for rejection of multiple distinct arrivals
$c_{4a,b}$	for curtailing of windows with emergent starts and/or codas
$w_{CC} \quad w_{len} \quad w_{nwin}$	for selection of best non-overlapping window combination

for use in adjoint tomography, its inherent flexibility should make it useful in many data-selection applications.

We have successfully applied our windowing algorithm, the details of which are described in Section 2, to diverse seismological scenarios: local and near regional tomography in Southern California, regional subduction-zone tomography in Japan, and global tomography. We present examples from each of these scenarios in Section 3, and we discuss the use of the algorithm in the context of tomography in Section 4.

2 THE SELECTION ALGORITHM

Our open-source algorithm, called FLEXWIN to reflect its FLEXibility in picking time WINdows for measurement, operates on pairs of observed and synthetic single component seismograms. There is no restriction on the type of simulation used to generate the synthetics, though realistic Earth models and more complete propagation theories yield waveforms that are more similar to the observed seismograms, and thereby allow the definition of measurement windows covering more of the available data. The input seismograms can be measures of displacement, velocity, or acceleration, indifferently. There is no requirement for horizontal signals to be rotated into radial and transverse directions. All the synthetic seismograms presented in this paper have been generated using the SPECFEM3D package (Komatitsch *et al.* 2002, 2004).

The window selection process has five stages, each of which is discussed individually below: *Stage A:* pre-processing; *Stage B:* definition of preliminary measurement windows; *Stage C:* rejection of preliminary windows based on the content of the synthetic seismogram alone; *Stage D:* rejection of preliminary windows based on the differences between observed and synthetic seismograms; *Stage E:* resolution of preliminary window overlaps. The parameters that permit tuning of the window selection towards a specific tomographic scenario are all contained in a simple parameter file (see Table 1). More complexity and finer tuning can be obtained by making some of these parameters time-dependent via user defined functions that can depend on the source parameters (e.g. event location or depth).

2.1 Stage A

The purpose of this stage is to pre-process input seismograms, to reject noisy records, and to set up a secondary waveform (the short-term average/long-term average ratio) derived from the envelope of the synthetic seismogram. This STA:LTA waveform will be used later to define preliminary measurement windows.

We apply minimal and identical pre-processing to both observed and synthetic seismograms: removal of any linear trend, tapering, and bandpass filtering with a non-causal Butterworth filter, whose short and long period corners we denote by T_0 and T_1, respectively. Values of these corner periods should reflect the information content of the data, the quality of the Earth model, and the accuracy of the simulation used to generate the synthetic seismograms. All further references to 'seismograms' in this paper will refer to these filtered waveforms.

Our next step is to reject seismograms that are dominated by noise. This rejection is based on two signal-to-noise criteria that compare the power and amplitude of the signal to those of the background noise (given by the observed waveform before the first *P*-wave arrival). The power signal-to-noise ratio is defined as

$\text{SNR}_P = P_{\text{signal}}/P_{\text{noise}}$, where the time-normalized power in the signal and noise portions of the data are defined, respectively, by

$$P_{\text{signal}} = \frac{1}{t_E - t_A} \int_{t_A}^{t_E} d^2(t) \, dt, \tag{1}$$

$$P_{\text{noise}} = \frac{1}{t_A - t_0} \int_{t_0}^{t_A} d^2(t) \, dt, \tag{2}$$

where $d(t)$ denotes the observed seismogram, t_0 is its start time, t_A is set to be slightly before the time of the first arrival, and t_E is the end of the main signal (a good choice for t_E is the end of the dispersed surface wave). The amplitude signal-to-noise ratio is defined analogously as $\text{SNR}_A = A_{\text{signal}}/A_{\text{noise}}$, where A_{signal} and A_{noise} are the maximum values of $|d(t)|$ in the signal and noise time-spans, respectively. The limits for these two signal-to-noise ratios are given by the parameters r_P and r_A in Table 1. We reject any record for which $\text{SNR}_P < r_P$ or $\text{SNR}_A < r_A$.

Detection and identification of seismic phase arrivals is routinely performed by automated earthquake location algorithms (e.g. Allen 1982; Earle & Shearer 1994; Aster & Rowe 2000; Bai & Kennett 2000; Sleeman & van Eck 2003). We have taken a tool used in most implementations of the automated detection process—the short-term average long-term average ratio (e.g. Withers *et al.* 1998; Bai & Kennett 2001)—and adapted it to the task of defining time windows around seismic phases. Given a synthetic seismogram $s(t)$, we derive an STA:LTA time-series using an iterative algorithm applied to the envelope of the synthetic. If we denote the Hilbert transform of the synthetic seismogram by $\mathcal{H}[s(t)]$, its envelope $e(t)$ is given by

$$e(t) = |s(t) + i\mathcal{H}[s(t)]|. \tag{3}$$

To create the STA:LTA waveform $E(t)$, we discretize the envelope time series with time step Δt, calculate its short term average $S(t_i)$ and its long term average $L(t_i)$ recursively

$$S(t_i) = C_S S(t_{i-1}) + e(t_i), \tag{4}$$

$$L(t_i) = C_L L(t_{i-1}) + e(t_i), \tag{5}$$

and obtain their ratio: $E(t_i) = S(t_i)/L(t_i)$. The constants C_S and C_L determine the decay of the relative weighting of earlier parts of the signal in the calculation of the current average. This decay is necessarily longer for the long-term average than for the short-term average, implying that $C_S < C_L < 1$. The choice of these constants determines the sensitivity of the STA:LTA time-series. Bai & Kennett (2001) used a similar time-series to analyse the character of broad-band waveforms, and allowed the constants C_S and C_L to depend on the dominant period of the waveform under analysis. We have followed their lead in setting

$$C_S = 10^{-\Delta t/T_0} \qquad \text{and} \qquad C_L = 10^{-\Delta t/12T_0}, \tag{6}$$

where the use of T_0, the low-pass corner period of our bandpass filter, substitutes that of the dominant period.

The STA:LTA of a constant signal converges to a constant value when the length of the time-series is greater than the effective averaging length of the long-term average. An energy arrival in $e(t)$ causes $E(t)$ to rise sharply, then drop to below the convergence value after the arrival, before stabilizing again. The maximum height reached by $E(t)$ for a given T_0 depends on the amplitude of the arrival in $e(t)$ and on its duration: higher amplitudes and shorter durations cause higher $E(t)$ values. The depth to which $E(t)$ drops after the end of the arrival depends on the same parameters: higher
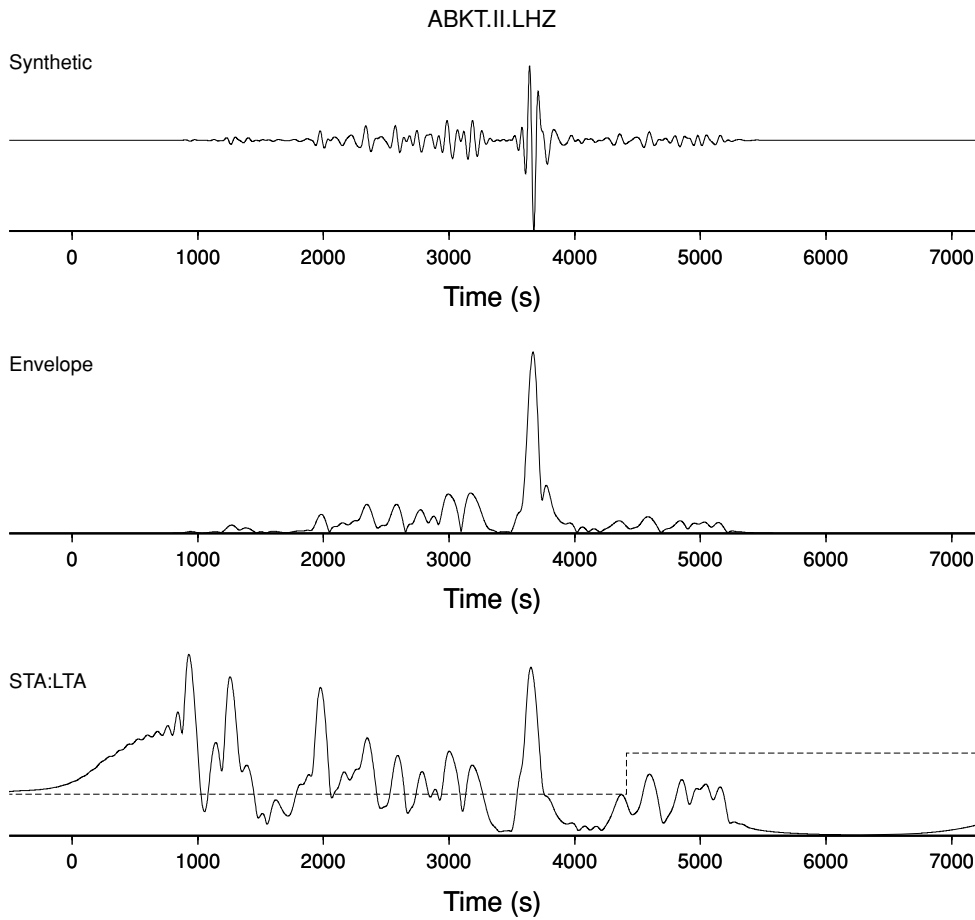
ABKT.II.LHZ

Figure 1. Synthetic seismogram and its corresponding envelope and STA:LTA timeseries. The seismogram was calculated using SPECFEM3D and the Earth model S20RTS (Ritsema *et al.* 2004) for the CMT catalogue event 050295B, whose details can be found in Table 2. The station, ABKT, is at an epicentral distance of 14 100 km and at an azimuth of 44° from the event. The top panel shows the vertical component synthetic seismogram, filtered between periods of 50 and 150 seconds. The center panel shows its envelope, and the bottom panel shows the corresponding STA:LTA waveform. The dashed line overlaid on the STA:LTA waveform is the water level $w_E(t)$.

amplitudes and longer durations cause deeper drops in $E(t)$ that take longer to return to the convergence value for a constant signal.

An example of a synthetic seismogram and its corresponding envelope and STA:LTA time-series $E(t)$ is shown in Fig. 1. The $E(t)$ time-series starts at its value for a constant signal, then rises gradually due to the tapered low level numerical noise on the synthetic. At each seismic arrival, $E(t)$ rises to a local maximum. We can see from Fig. 1 that these local maxima correspond both in position and in width to the seismic phases in the synthetic, and that the local minima in $E(t)$ correspond to the transitions between one phase and the next. In the following sections we shall explain how we use these correspondences to define time windows.

2.2 Stage B

The correspondence between local maxima in the STA:LTA waveform $E(t)$ and the position of the seismic phases in the synthetic seismogram suggests that we should centre time windows around these local maxima. The correspondence between the local minima in $E(t)$ and the transition between successive phases suggests the time windows should start and end at these local minima. In the case of complex phases, there may be several local maxima and minima within a short time-span. To correctly window these com-

plex phases, we must determine rules for deciding when adjacent local maxima should be part of a single window. From an algorithmic point of view, it is simpler to create all possible combinations of adjacent windows and subsequently reject the unacceptable ones, than to consider combining small, single-maximum windows into larger ones.

We start by defining a water level on $E(t)$ via the time-dependent parameter $w_E(t)$ in Table 1. All local maxima that lie above $w_E(t)$ are considered acceptable, and are used for the creation of candidate time windows. The water level shown in Fig. 1 corresponds to $w_E = 0.08$ for the duration of the main seismic signal. Once set for typical seismograms for a given seismological scenario, it is not necessary to change w_E for each seismogram. This is also true of all the other parameters in Table 1: once the system has been tuned, these parameters remain unchanged and are used for all seismic events in the same scenario. A summary of the main considerations the user should take into account to tune these parameters can be found in Appendix A. Functional forms of the time-dependent parameters are defined by the user, can depend on information about the earthquake source and the receiver, and also remain unchanged once the system has been tuned. For the example in Fig. 1, we have required the water level $w_E(t)$ to double after the end of the surface-wave arrivals (as defined by the epicentral distance and a group velocity of 3.2 km s^{-1}) to avoid creating time windows after $R1$.
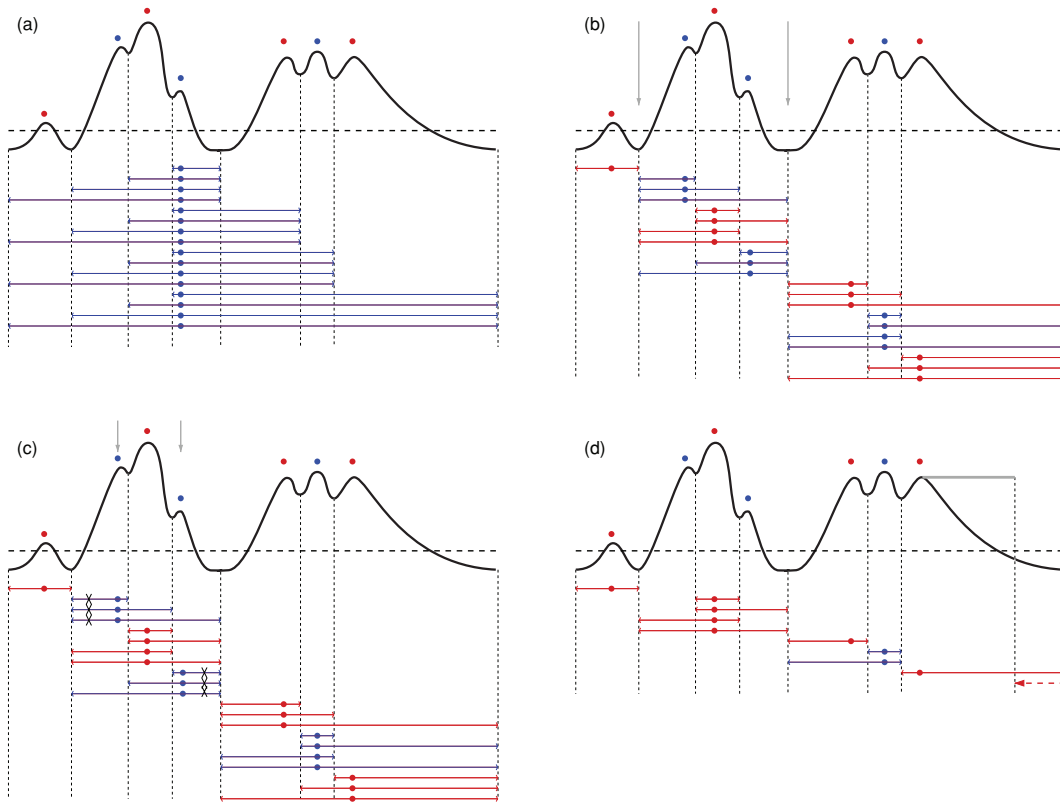
Figure 2. (a) Window creation process. The thick black line represents the STA:LTA waveform $E(t)$, and the thick horizontal dashed line its water level $w_E(t)$. Local maxima are indicated by alternating red and blue dots, windows are indicated by two-headed horizontal arrows. The time of the local maximum used as the window seed t_M is denoted by the position of the dot. Only windows for the fourth local maximum are shown. (b) Rejection of candidate windows based on the amplitude of the local minima. The two deep local minima indicated by the grey arrows form virtual barriers. All candidate windows that cross these barriers are rejected. (c) Rejection of candidate windows based on the prominence of the seed maximum. The local maxima indicated by the grey arrows are too low compared to the local minima adjacent to them. All windows that have these local maxima as their seed are rejected (black crosses over the window segments below the timeseries). (d) Shortening of long coda windows. The grey bar indicates the maximum coda duration $c_{4b}T_0$. Note that after the rejection based on prominence represented in (c) and before shortening of long coda windows represented in (d), the algorithm rejects candidate windows based on the separation of distinct phases, a process that is illustrated in Fig. 3.

We take each acceptable local maximum in turn as a seed maximum, and create all possible candidate windows that contain it, as illustrated by Fig. 2(a). Each candidate window is defined by three times: its start time t_S, its end time t_E and the time of its seed maximum t_M. The start and end times correspond to local minima in $E(t)$. It is important to note that in many of the window rejection algorithms, t_M will be significant. For N local maxima that lie above $w_E(t)$, the number of preliminary candidate windows defined in this manner is

$$N_{\text{win}} = \sum_{n=1}^{N} [nN - (n-1)^2] \sim O(N^3). \quad (7)$$

2.3 Stage C

After having created a complete set of candidate time windows in the manner described above, we start the rejection process. We reject windows based on two sets of criteria concerning, respectively, the shape of the STA:LTA waveform $E(t)$, and the similarity of the observed and synthetic waveforms $d(t)$ and $s(t)$ within each window. Here we describe the first set of criteria; the second set is described in the following section.

The aim of shape-based window rejection is to retain the set of candidate time windows within which the synthetic waveform

$s(t)$ contains well-developed seismic phases or groups of phases. The four rejection criteria described here are parametrized by the constants c_{0-3} in Table 1, and are scaled in time by T_0 and in amplitude by $w_E(t)$. We apply these criteria sequentially.

First, we reject all windows that contain internal local minima of $E(t)$ whose amplitude is less than $c_0 w_E(t)$. We have seen above that local minima of $E(t)$ tend to lie on the transitions between seismic phases. By rejecting windows that span deep local minima, we are in fact forcing partitioning of unequivocally distinct seismic phases into separate time windows (see Fig. 2b). Second, we reject windows whose length is less than $c_1 T_0$. By rejecting short windows, we are requiring that time windows be long enough to contain useful information. Third, we reject windows whose seed maximum $E(t_M)$ rises by less than $c_2 w_E(t_M)$ above either of its adjacent minima. Subdued local maxima of this kind represent minor changes in waveform character, and should not be used to anchor time windows. They may, however, be included within a time window with a more prominent seed maximum (see Fig. 2c). Lastly, we reject windows that contain at least one strong phase arrival that is well separated in time from t_M. The rejection is performed using the following criterion:

$$h/h_M > f(\Delta T/T_0; c_{3a}, c_{3b}), \quad (8)$$

where h_M is the height of the seed maximum $E(t_M)$ above the deepest minimum between itself and another maximum, h is the height of
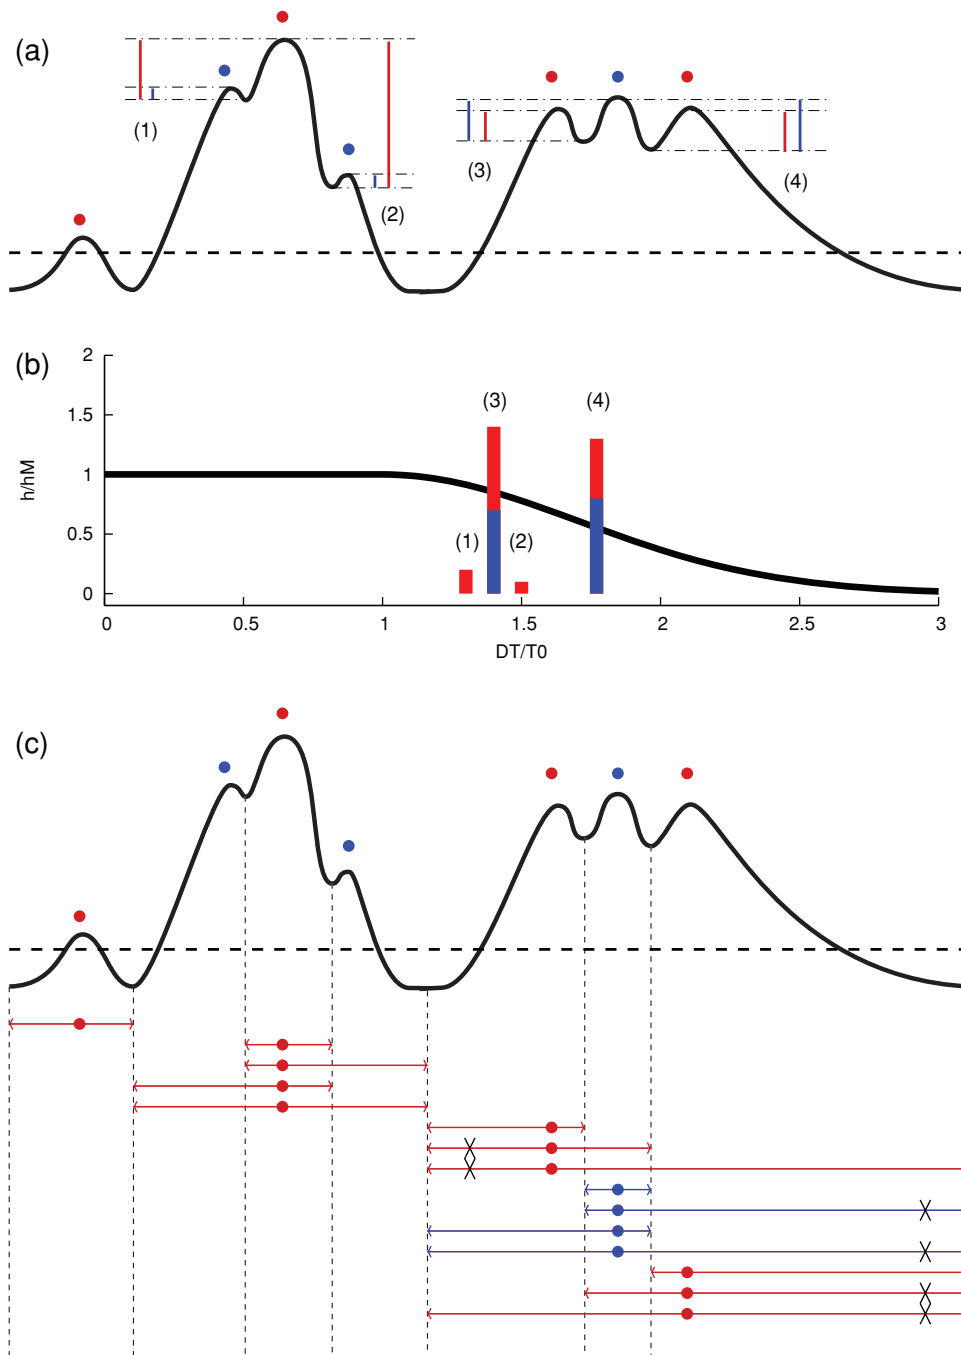
Figure 3. Rejection of candidate windows based on the separation of distinct phases. (a) Heights of pairs of local maxima above their intervening minimum. (b) The black line represents $f(\Delta T/T_0)$ from eq. (9) with $c_{3a} = c_{3b} = 1$. Vertical bars represent h/h_M for each pair of maxima. Their position along the horizontal axis is given by the time separation ΔT between the maxima of each pair. The colour of the bar is given by the colour of the seed maximum corresponding to h_M. Bars whose height exceeds the $f(\Delta T/T_0)$ line represent windows to be rejected. (c) The windows that have been rejected by this criterion are indicated by black crosses.

this other maximum above the same minimum, and f is a function of the time separation ΔT between the two maxima (see Fig. 3). The function $f(\Delta T)$ has the following form:

$$f(\Delta T) = \begin{cases} c_{3a} & \Delta T/T_0 \leq c_{3b}, \\ c_{3a} \exp\left[-(\Delta T/T_0 - c_{3b})^2/c_{3b}^2\right] & \Delta T/T_0 > c_{3b}. \end{cases} \quad (9)$$

If we take as an example $c_{3a} = 1$, this criterion leads to the automatic rejection of windows containing a local maximum that is higher

than their seed maximum; it also leads to the rejection of windows containing a local maximum that is lower than their seed maximum if this local maximum is sufficiently distant in time from t_M. This criterion allows us to distinguish unseparable phase groups from distinct seismic phases.

The candidate windows that remain after application of these four rejection criteria are almost ready to be passed on to the next stage, in which we shall evaluate the similarity between observed
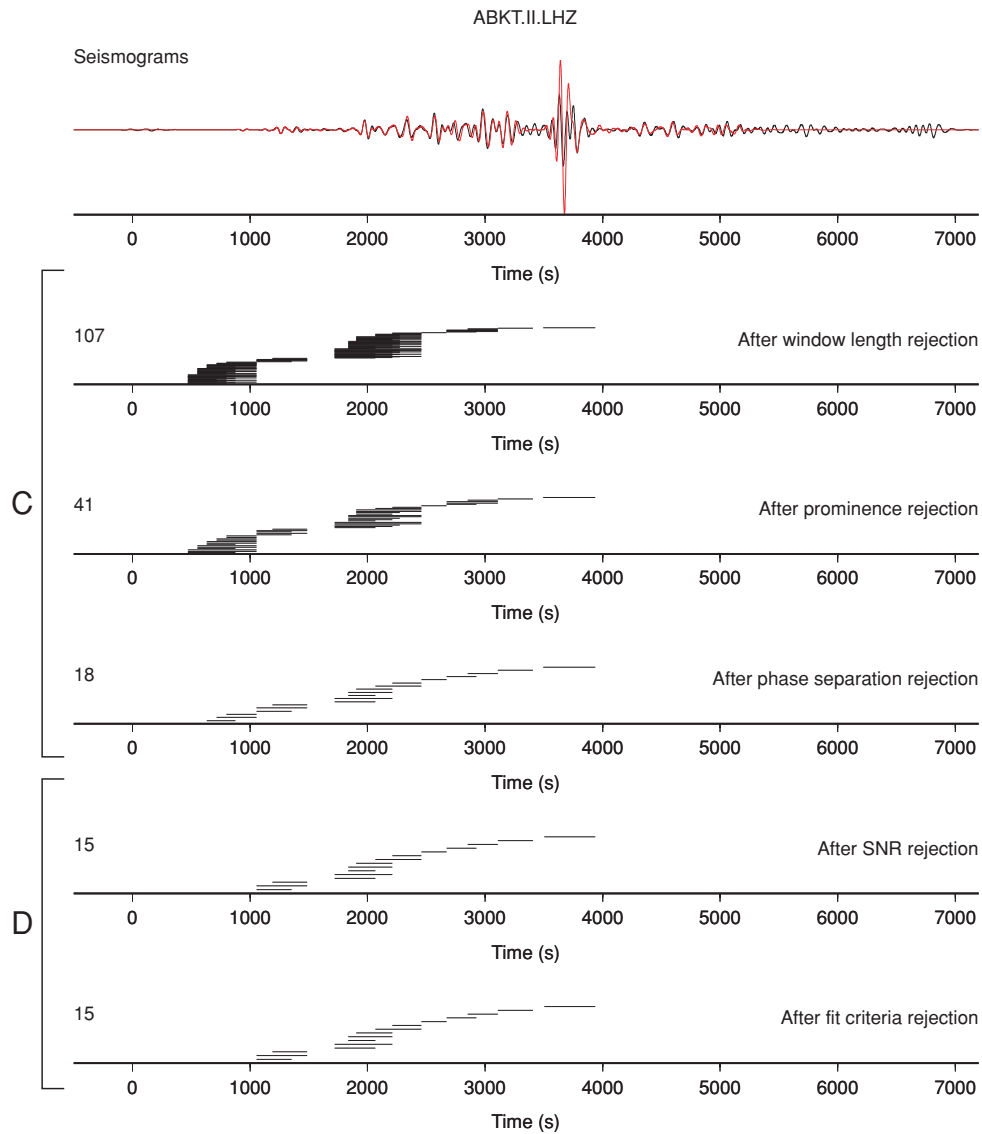
Figure 4. Window rejection applied to real data. Top panel: observed (black) and synthetic (red) seismograms for the 050295B event recorded at ABKT (see Fig. 1). Subsequent panels: candidate windows at different stages, separated into Stage C (shape based rejection) and Stage D (fit based rejection). Each candidate window is indicated by a black segment. The number of windows at each stage is shown to the left of the panel.

and synthetic waveforms within the window limits. Special precautions may have to be taken, however, in the case of windows that contain long coda waves: the details of codas are often poorly matched by synthetic seismogram calculations, as they are essentially caused by multiple scattering processes. To avoid rejecting a nicely fitting phase because of a poorly fitting coda or a poorly fitting emergent start, we introduce the c_4 tuning parameters, which permit shortening of windows starting with monotonically increasing $E(t)$ or ending with monotonically decreasing $E(t)$. These windows are shortened on the left if they start earlier than $c_{4a} T_0$ before their first local maximum, and on the right if they end later than $c_{4b} T_0$ after their last local maximum (see Fig. 2d).

Figs 2 and 3 illustrate the shape-based rejection procedure (Stage C) on a schematic $E(t)$ time-series. Each successive criterion reduces the number of acceptable candidate windows. A similar reduction occurs when this procedure is applied to real $E(t)$ time-series, as shown by the upper portion of Fig. 4.

2.4 Stage D

After having greatly reduced the number of candidate windows by rejection based on the shape of the STA:LTA timeseries $E(t)$, we are now left with a set of windows that contain well-developed seismic phases or groups of phases on the synthetic seismogram. The next stage is to evaluate the degree of similarity between the observed and synthetic seismograms within these windows, and to reject those that fail basic fit-based criteria. For each window, we consider the windowed waveforms $\widetilde{d}(t)$ and $\widetilde{s}(t)$ to be the product of $d(t)$ and $s(t)$ with a boxcar function that is unity between start and end times of the window and zero elsewhere.

The quantities we use to define well-behavedness of data within a window are signal-to-noise ratio SNR_W, normalised cross-correlation value between observed and synthetic seismograms CC, cross-correlation time lag $\Delta \tau$, and amplitude ratio $\Delta \ln A$. The signal-to-noise ratio for single windows is defined as an amplitude
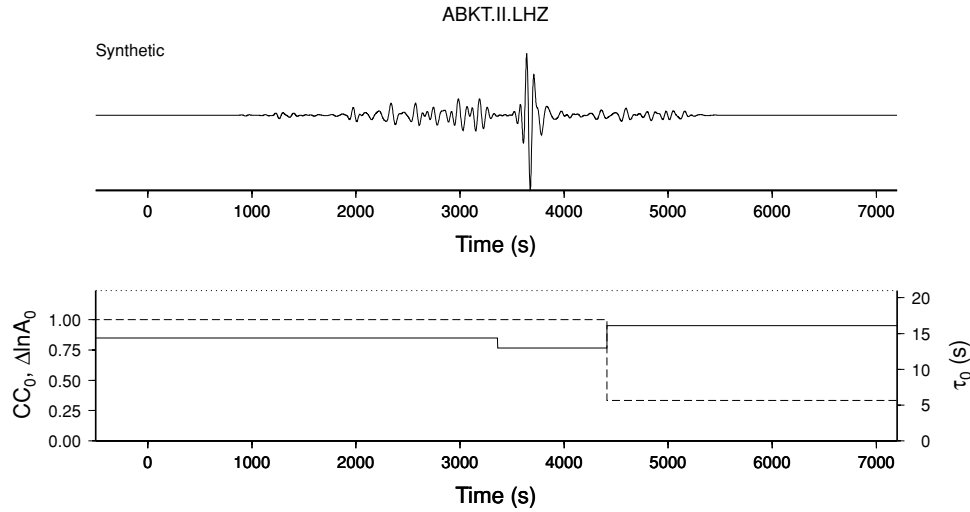
ABKT.II.LHZ

Figure 5. Time-dependent fit based criteria for the 050295B event recorded at ABKT. The time-dependence of these criteria is given by the formulae in Appendix A1.1. The lower limit on acceptable cross-correlation value, CC_0 (solid line), is 0.85 for most of the duration of the seismogram; it is lowered to 0.75 during the approximate surface-wave window defined by the group velocities 4.2 km s^{-1} and 3.2 km s^{-1}, and is raised to 0.95 thereafter. The upper limit on time lag, τ_0 (dotted line), is 21 s for the whole seismogram. The upper limit on amplitude ratio, $\Delta \ln A_0$ (dashed line), is 1.0 for most of the seismogram; it is reduced to 1/3 of this value after the end of the surface-waves.

ratio, $\text{SNR}_W = A_{\text{window}}/A_{\text{noise}}$, where A_{window} is the maximum of $|\widetilde{d}(t)|$, and A_{noise} is the maximum value of $|d(t)|$ in the noise time-span (the noise time-span is the same as that for eq. 2). The cross-correlation value CC is defined as the maximum value of the normalised cross-correlation function, $CC = \max[\Gamma(t)]$, where

$$\Gamma(t) = \frac{\int \widetilde{s}(t')\widetilde{d}(t' - t)\, dt'}{\left[\int \widetilde{s}^2(t')\, dt' \int \widetilde{d}^2(t' - t)\, dt'\right]^{1/2}} \quad (10)$$

quantifies the similarity in shape between the $\widetilde{s}(t)$ and $\widetilde{d}(t)$ waveforms, and the integration limits are the start and end times of the window. The time lag $\Delta \tau$ is defined as the value of t at which Γ is maximal, and quantifies the delay in time between a synthetic and observed phase arrival. The amplitude ratio $\Delta \ln A$ is defined as the amplitude ratio between observed and synthetic seismograms (Dahlen & Baig 2002)

$$\Delta \ln A = \ln(A_{\text{obs}}/A_{\text{syn}}) = 0.5 \ln\left[\frac{\int \widetilde{d}^2(t)\, dt}{\int \widetilde{s}^2(t)\, dt}\right] \quad (11)$$

(note that Dahlen & Baig 2002, eq. 3 is the first-order approximation of eq. 11). The limits that trigger rejection of windows based on the values of these four quantities are the parameters $r_0(t)$, $CC_0(t)$, $\Delta \tau_{\text{ref}}$, $\Delta \tau_0(t)$, $\Delta \ln A_{\text{ref}}$ and $\Delta \ln A_0(t)$ in Table 1. As for the STA:LTA water level $w_E(t)$ used above, the functional forms of the time-dependent parameters are defined by the user, and can depend on source and receiver parameters such as epicentral distance and earthquake depth. Examples of functional forms for these parameters can be found in Appendix A1. Fig. 5 shows the time dependence of CC_0, $\Delta \tau_0$ and $\Delta \ln A_0$ for the example seismogram of Fig. 4.

We only accept candidate windows that satisfy all of the following:

$$\text{SNR}_W \geq r_0(t_M), \quad (12)$$

$$CC \geq CC_0(t_M), \quad (13)$$

$$\Delta \tau_{\text{min}} \leq \Delta \tau \leq \Delta \tau_{\text{max}}, \quad (14)$$

$$\Delta \ln A_{\text{min}} \leq \Delta \ln A \leq \Delta \ln A_{\text{max}}, \quad (15)$$

where

$$\Delta \tau_{\text{min}} \equiv \Delta \tau_{\text{ref}} - \Delta \tau_0(t_M), \quad (16)$$

$$\Delta \tau_{\text{max}} \equiv \Delta \tau_{\text{ref}} + \Delta \tau_0(t_M), \quad (17)$$

$$\Delta \ln A_{\text{min}} \equiv \Delta \ln A_{\text{ref}} - \Delta \ln A_0(t_M), \quad (18)$$

$$\Delta \ln A_{\text{max}} \equiv \Delta \ln A_{\text{ref}} + \Delta \ln A_0(t_M), \quad (19)$$

and t_M is the time of the window's seed maximum. In words, we only accept windows in which the observed signal is sufficiently above the noise level, the observed and synthetic signals are reasonably similar in shape, their arrival time differences are small, and their amplitudes are broadly compatible. The parameters $\Delta \tau_{\text{ref}}$ and $\Delta \ln A_{\text{ref}}$ allow the algorithm to work efficiently when there are systematic differences between data and synthetics. For example, if the Earth model is on average too fast, then $\Delta \tau_{\text{ref}}$ should be set to a positive value. Or if the magnitudes of the synthetic sources lead to systematically non-zero $\Delta \ln A$ values, then $\Delta \ln A_{\text{ref}}$ should be chosen accordingly. In practice, these reference values should designate the approximate center-value of the distribution of measurements (see, for example, Fig. 17c).

When the synthetic and observed seismograms are similar, the fit-based criteria of eqs (13)–(15) reject only a few of the candidate data windows (see lower portion of Fig. 4). They are essential, however, in eliminating problems due secondary events (natural or man-made), diffuse noise sources, or instrumental glitches.

2.5 Stage E

After having rejected candidate data windows that fail any of the shape or similarity based criteria described above, we are left with a small number of windows, each of which taken singly would be an acceptable time window for measurement. As can be seen from Fig. 2(d) and the last panel of Fig. 4, the remaining windows may

overlap partially or totally with their neighbours. Such overlaps are problematic for automated measurement schemes, as they lead to multiple measurements of those features in the seismogram that lie within the overlapping portions. Resolving this overlap problem is the last step in the windowing process.

Overlap resolution can be seen as a set of choices leading to the determination of an optimal set of time windows. What do we mean by optimal? For our purposes, an optimal set of time windows contains only windows that have passed all previous tests, that do not overlap with other windows in the set, and that cover as much of the seismogram as possible. When choosing between candidate windows, we favour those within which the observed and synthetic seismograms are most similar (high values of CC). Furthermore, should we have the choice between two short windows and a longer, equally well-fitting one covering the same time-span, we may wish to favour the longer window as this poses a stronger constraint on the tomographic inversion.

The condition that optimal windows should have passed all previous tests removes the straightforward solution of merging overlapping windows. Indeed, given any two overlapping windows, we know that the window defined by their merger existed in the complete set of candidate windows obtained at the end of Stage B, and that its absence from the current set means it was rejected either because of the shape of its $E(t)$ time-series (Stage C), or because of an inadequate similarity between observed and synthetic waveforms (Stage D). It would therefore be meaningless to re-instate such a window at this stage. Any modification of current candidate windows would be disallowed by similar considerations. We must therefore choose between overlapping candidates.

We make this choice by constructing all possible non-overlapping subsets of candidate windows, and scoring each subset on three criteria: length of seismogram covered by the windows, average cross-correlation value for the windows, and total number of windows. These criteria often work against each other. For example, a long window may have a lower CC than two shorter ones, if the two short ones have different time lags $\Delta\tau$. Weighting of the three scores is

necessary, and is controlled by the three parameters w_{CC}, w_{len} and w_{nwin} in Table 1.

As can be seen in Fig. 6, the generation of subsets is facilitated by first grouping candidate windows such that no group overlaps with any other group. The selection of the optimal subsets can then be performed independently within each group. We score each non-overlapping subset of windows within a group using the following three metrics:

$$S_{CC} = \sum_{i}^{N_{set}} CC_i / N_{set}, \tag{20}$$

$$S_{len} = \left[\sum_{i}^{N_{set}} t_i^e - t_i^s \right] \Big/ \left[t_g^e - t_g^s \right], \tag{21}$$

$$S_{nwin} = 1 - N_{set}/N_{group}, \tag{22}$$

where CC_i is the cross-correlation value of the ith window in the subset, N_{set} is the number of windows in the subset, N_{group} is the number of windows in the group, and t_i^s, t_i^e, t_g^s and t_g^e are, respectively, the start and end times of the ith candidate window in the set, and of the group itself. The three scores are combined into one using the weighting parameters:

$$S = \frac{w_{CC}S_{CC} + w_{len}S_{len} + w_{nwin}S_{nwin}}{w_{CC} + w_{len} + w_{nwin}}. \tag{23}$$

The best subset of candidate windows within each group is the one with the highest combined score S. The final set of windows is given by concatenating the best subsets of candidate windows for each group. Fig. 7(a) shows an example of final windows selected on real data.

3 WINDOWING EXAMPLES

We present a set of examples showing the results of the FLEXWIN algorithm applied to real data. These examples illustrate the
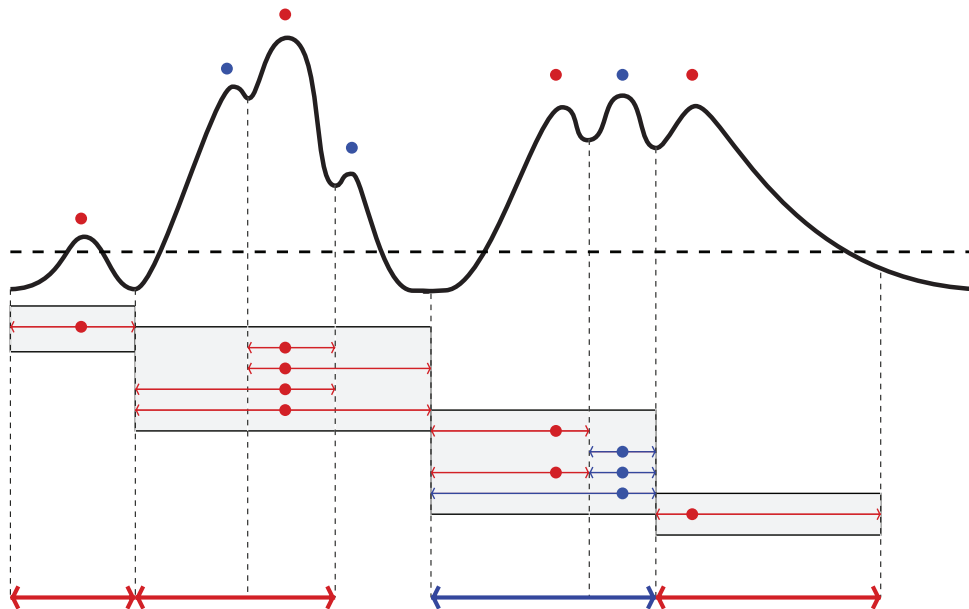

Figure 6. The selection of the best non-overlapping window combinations. Each grey box represents a distinct group of windows. Non-overlapping subsets of windows are shown on separate lines. Only one line from within each group will be chosen, the one corresponding to the highest score obtained in eq. (23). The resulting optimal set of data windows is shown by thick arrows.
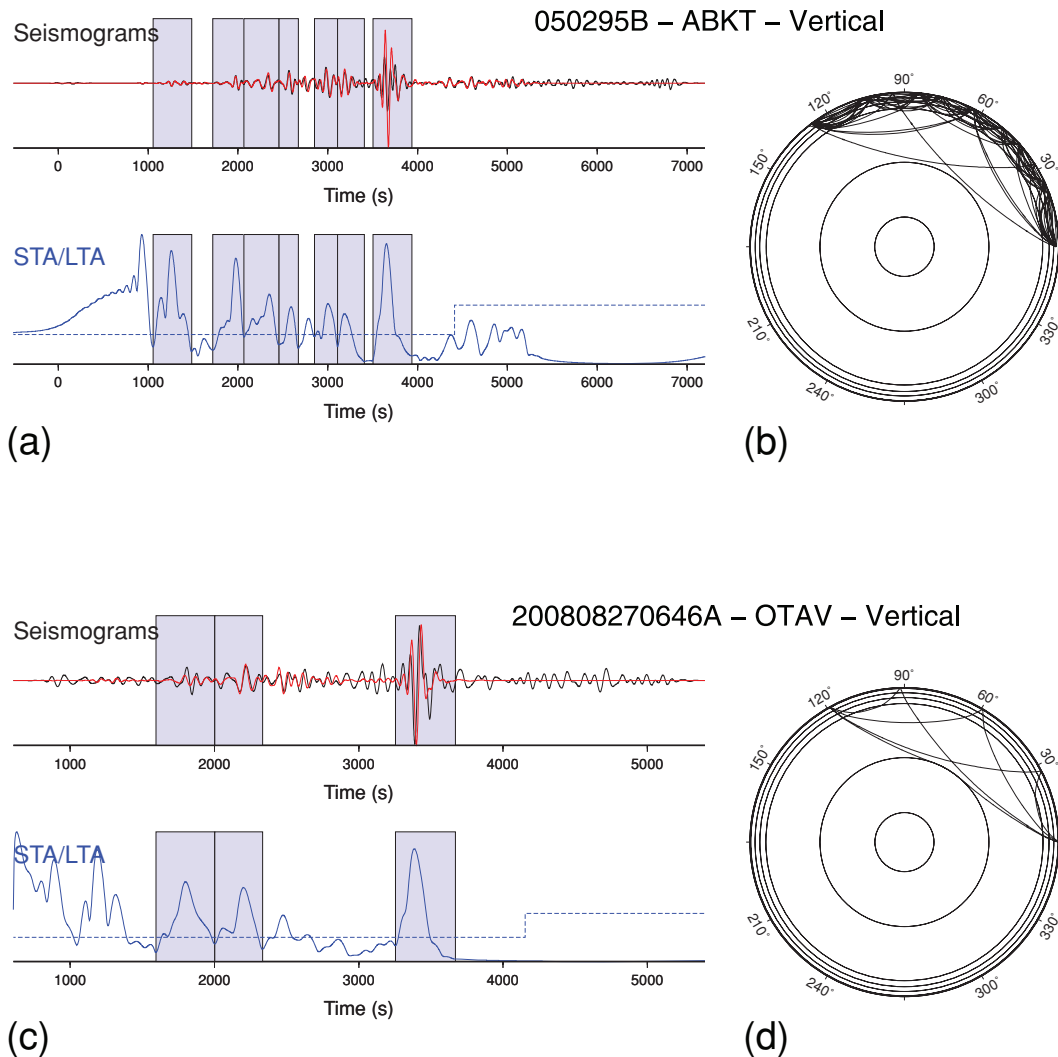
Figure 7. Window selection results for event 050295B from Table 2 recorded at ABKT (37.93°N, 58.11°E, $\Delta = 127°$, vertical component). (a) Top: observed and synthetic seismograms (black and red traces); bottom: STA:LTA timeseries $E(t)$. Windows chosen by the algorithm are shown using light blue shading. The phases contained within these windows are: (1) *PP*, (2) *PS + SP*, (3) *SS*, (4) *SSS*, (5) *S5*, (6) *S6*, (7) fundamental mode Rayleigh wave. (b) Ray paths corresponding to the body-wave phases present in the data windows in (a). (c) Window selection results for event 200808270646A from Table 2 recorded at OTAV (0.24°N, 78.45°W, $\Delta = 119°$, vertical component). Phases contained within selected windows: (1) S_{diff} and *PS + SP*, (2) *SS*, (3) fundamental mode Rayleigh wave. (d) Ray paths corresponding to the body-wave phases present in the data windows in (c).

robustness and flexibility of the algorithm. We have applied the algorithm to three tomographic scenarios, with very different geographical extents and distinct period ranges: long-period global tomography (50–150 s), regional tomography of the Japan subduction zone down to 700 km (6–120 s), and regional tomography of southern California down to 60 km (2–30 s). For each of these scenarios, we compare observed seismograms to spectral-element synthetics, using our algorithm to select time windows on the pairs of timeseries.

The windowing algorithm itself has little prior knowledge of seismology, other than in the most general terms: it considers a seismogram to be a succession of seismic phases indicated by changes in amplitude and frequency of the signal with time; it is based upon the idea that the short-term to long-term average ratio STA:LTA is a good indicator of the arrival of such phases; it has a notion of the characteristics of an optimal set of data windows. All other prior information—the frequency range to be considered, the portions of the seismogram to be excluded, the acceptable signal-to-noise

ratios, the tolerance of dissimilarity between the observed and synthetic seismogram—varies greatly between any two seismological studies. In order to ensure maximum flexibility of our windowing algorithm, all such scenario-dependent information is encapsulated in the tuning parameters of Table 1.

We tuned the windowing algorithm separately for each of the three scenarios we present here, and we present examples based on the events listed in Table 2. Tuning parameter values for each scenario can be found in Table 3, whereas the functional forms of the time-dependent parameters can be found in Appendix A1. Once tuned for a given scenario, the algorithm is applied to all its events without further modification.

3.1 Global tomography

Our first scenario is a global scale, long-period tomographic study. We calculate spectral-element synthetic seismograms through an Earth model for which the mantle is given by the S20RTS model

Table 2. Example events used in this study. The identifier refers to the CMT catalogue for global events and Japan events, and refers to the Southern California Earthquake Data Center catalogue for southern California events.

Identifier	Latitude	Longitude	Depth, km	Moment, N m	M_w	Location
			Global			
101895B	28.06	130.18	18.5	5.68e19	7.1	Ryukyu Islands
200808270646A	−10.49	41.44	24.0	4.68e17	5.7	Comoros Region
050295B	−3.77	−77.07	112.8	1.27e19	6.7	Northern Peru
060994A	−13.82	−67.25	647.1	2.63e21	8.2	Northern Bolivia
			Japan			
051502B	24.66	121.66	22.4	1.91e18	6.1	Taiwan
200511211536A	30.97	130.31	155.0	2.13e18	6.2	Kyuhu, Japan
091502B	44.77	130.04	589.4	4.24e18	6.4	Northeastern China
			Southern California			
9983429	35.01	−119.14	13.5	9.19e15	4.6	Wheeler Ridge, California
9818433	33.91	−117.78	9.4	3.89e15	4.3	Yorba Linda, California

Table 3. Values of standard and fine-tuning parameters for the three seismological scenarios discussed in this study.

	Global			Japan	S. California		
$T_{0,1}$	50, 150	24, 120	6, 30	6, 30	3, 30	2, 30
$r_{P,A}$	3.5, 3.0	3.5, 3.0	3.5, 3.0	3.0, 2.5	2.5, 3.5	2.5, 3.5
r_0	2.5	1.5	3.0	3.0	4.0	4.0
w_E	0.08	0.11	0.12	0.18	0.11	0.07
CC_0	0.85	0.70	0.73	0.71	0.80	0.85
$\Delta\tau_0$	15	12.0	3.0	8.0	4.0	3.0
$\Delta\ln A_0$	1.0	1.0	1.5	1.5	1.0	1.0
$\Delta\tau_{ref}$	0.0	0.0	0.0	4.0	2.0	1.0
$\Delta\ln A_{ref}$	0.0	0.0	0.0	0.0	0.0	0.0
c_0	0.7	0.7	0.7	0.7	1.3	1.0
c_1	4.0	3.0	3.0	2.0	4.0	5.0
c_2	0.3	0.0	0.6	0.0	0.0	0.0
$c_{3a,b}$	1.0, 2.0	1.0, 2.0	1.0, 2.0	3.0, 2.0	4.0, 2.5	4.0, 2.5
$c_{4a,b}$	3.0, 10.0	3.0, 25.0	3.0, 12.0	2.5, 12.0	2.0, 6.0	2.0, 6.0
$w_{CC}, w_{len}, w_{nwin}$	1, 1, 1	1, 1, 1	1, 1, 1	0.5,1.0,0.7	0.70,0.25,0.05	1,1,1

of Ritsema *et al.* (2004), and the crust by the CRUST2.0 model of Bassin *et al.* (2000). The degree-20 *S*-wave velocity model S20RTS defines isotropic perturbations to radially anisotropic PREM (Dziewonski & Anderson 1981); the SPECFEM3D implementation of S20RTS takes *P*-wave velocity anomalies from the degree-12 *P*-wave velocity model of Ritsema & van Heijst (2002). CRUST2.0 specifies a seven-layer crustal seismic velocity and density profile for each cell on a 2° grid. The S20RTS+CRUST2.0 combination produces synthetics that are a good match to observed seismograms for periods longer than 35–40 s. For our examples, we shall be working in the period range 50–150 s.

Here we discuss windowing results for shadow-zone seismograms of four earthquakes listed in Table 2: a shallow event in the Ryukyu Islands, Japan (101895B), a smaller magnitude shallow event in the Comoros region, between Mozambique and Madagascar (200808270646A), an intermediate depth event in northern Peru (050295B), and a strong deep event in northern Bolivia (060994A). We focus on shadow zone seismograms as these contain a large number of often poorly time-separated phases, and pose a greater windowing challenge than more commonly used teleseismic seismograms.

Windowing results for these seismograms (one example per earthquake) are shown in Figs 7(a), (c) and 8(a), (c). The first observation we make when looking at these examples is that the synthetics match the data well, indicating that the Earth model S20RTS+CRUST2.0 provides a good 3-D image of how the

Earth is seen by 50–150 s seismic waves. The fit is far from perfect, though, as is attested by the shape differences, time-lags and amplitude differences visible on many seismic phases; these indicate that there is room for improving the Earth model and possibly certain earthquake parameters even at these low frequencies.

The second observation we make is that our algorithm has placed time windows around most of the significant features that stand out in the STA:LTA timeseries $E(t)$ and in the seismograms themselves, and that the window limits also seem to be sensibly placed. These windows were selected according to the purely signal processing algorithm described in the previous section, which has no knowledge of Earth structure or of seismic phases and their traveltime curves. To demonstrate the ability of such an Earth-blind algorithm to set windows around actual seismic phases, we have identified the seismic arrivals contained within the chosen data windows, using standard PREM-based traveltime curves. We have found that most of the features within the windows in Figs 7(a), (c) and 8(a), (c) correspond to known seismic phases, which are listed in the corresponding figure captions. We have also traced the body-wave ray paths corresponding to these phases and show them in Figs 7(b), (d) and 8(b), (d); these ray path plots serve to illustrate the considerable amount of information contained in a single seismogram, even a long period seismogram, when all the usable seismic phases are considered. Fewer useable seismic phases are windowed for the smaller magnitude event in Fig. 7(c).
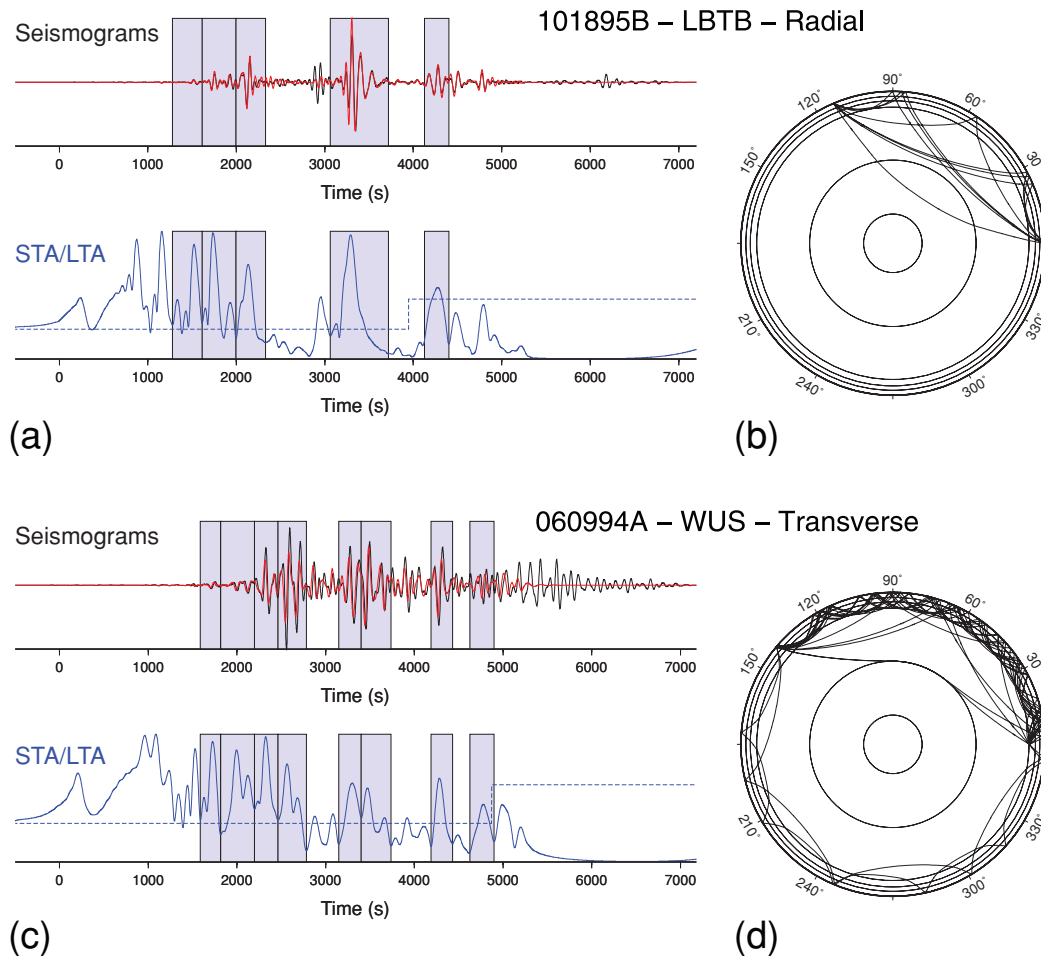
Figure 8. (a) Window selection results for event 101895B from Table 2 recorded at LBTB (25.01°S, 25.60°E, $\Delta = 113°$, radial component). Phases contained within selected windows: (1) *SKS*, (2) *PS* + *SP*, (3) *SS*, (4) fundamental mode Rayleigh wave (5) unidentified late phase. (b) Body-wave ray paths corresponding to data windows in (a). (c) Window selection results for event 060994A from Table 2 recorded at WUS (41.20°N, 79.22°E, $\Delta = 140°$, transverse component). Phases contained within selected windows: (1) $S_{\rm diff}$, (2) $sS_{\rm diff}$, (3) *SS*, (4) *sSS* followed by *SSS*, (5) *sS5* + *S6*, (6) *sS6* + *S7* followed by *sS7*, (7) major arc *sS4*, (8) major arc *sS6*. (d) Body-wave ray paths corresponding to data windows in (c).

Not all the features within a given seismogram are identifiable as seismic phases. For example, the second window in Fig. 8(b) seems to contain two features. When we look at periods shorter than 50 s, the first feature retains its character and is clearly identifiable as $sS_{\rm diff}$, whereas the second feature looses its character entirely and is more readily assimilated to a generic *S* wave coda than to a distinct seismic phase. This feature is present in both observed and synthetic seismograms, and undoubtedly contains information. The particularity of our windowing algorithm is to treat such features as information, without trying to identify their sources. A scheme that permits the computation of sensitivity kernels for such features (e.g. the adjoint scheme), would allow measurements made on them to be interpreted and inverted correctly. Other methods of determining measurement sensitivities may have more difficulty dealing with them. These considerations illustrate the strong ties that exist between the selection, measurement and interpretation stages of any study using seismological data.

We have described those seismic phases and other features in the seismograms that have been selected by our windowing algorithm. Equally important are the phases that have been rejected. Two such phases are $P_{\rm diff}$ and *S*4 on the vertical component seismogram in Fig. 7(a). We can identify the reasons for the rejection of these phases by comparing the selected time windows with the candidate

windows at each stage in the rejection process (Fig. 4). The $P_{\rm diff}$ phase, though small on the long period seismogram, gives rise to a strong maximum on the $E(t)$ timeseries and therefore to at least one candidate window. Candidate windows containing $P_{\rm diff}$ disappear from Fig. 4 at the SNR $_W$ based rejection stage, indicating that this phase was rejected for its low signal-to-noise ratio. The *S*4 phase also gives rise to a distinct maximum in $E(t)$, and to its own candidate window that is still present at the end of both window rejection stages (it corresponds to the fourth window from the right at the bottom of Fig. 4). As the *S*4 candidate and its neighbour the *S*5 candidate overlap, the algorithm has to choose between them using eq. (23). The fit to the shape of *S*4 is worse than that to *S*5, therefore the *S*4 window is discarded. It is helpful, when setting and tuning the values of the parameters in Table 1, to analyse the rejection and overlap resolution steps as we have done here for a number of representative seismograms, seeking to avoid the acceptance of unreasonable candidate windows, and to minimize the rejection of acceptable ones.

A further appreciation of the windowing results is given by event-based summaries such as those in Fig. 9, which show at a glance the geographical path distribution of records containing acceptable windows, the distribution of CC, $\Delta\tau$ and $\Delta\ln A$ values within the accepted time windows, and time-window record sections.
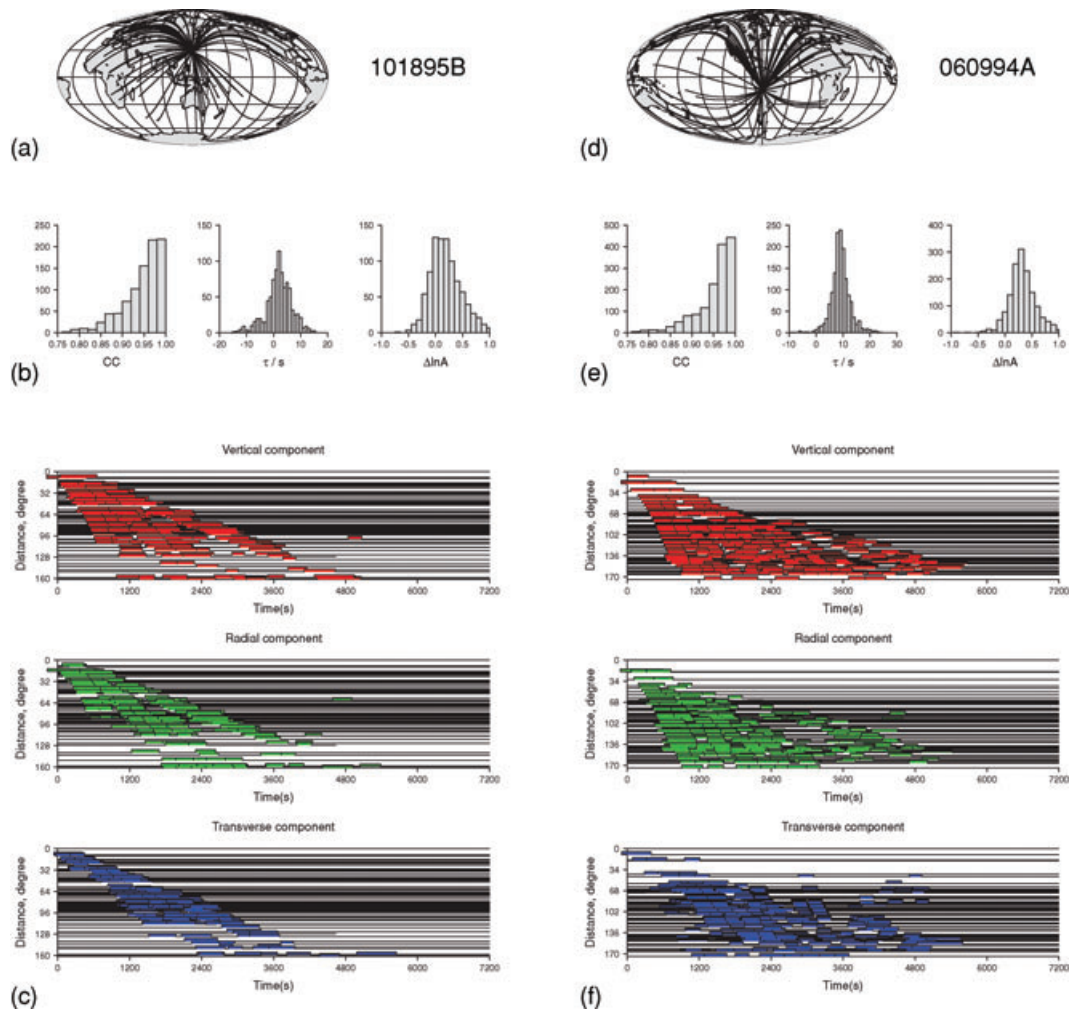
Figure 9. (a)–(c) Summary plots of windowing results for event 101895B in Table 2. (a) Global map showing great-circle paths to stations. (b) Histograms of number of windows as a function of normalised cross-correlation CC, time-lag τ and amplitude ratio $\Delta \ln A$; these give information about systematic trends in time-shift and amplitude scaling. (c) Record sections of selected windows for the vertical, radial and transverse components. The filled portions of the each record in the section indicate where windows have been selected by the algorithm. (d)–(f) Summary plots of windowing results for event 060994A in Table 2.

Comparison of the summary plots for the shallow Ryukyu Islands event and the deep Bolivia event (Fig. 9b and e, respectively) shows that both have similar one-sided distributions of CC values, strongly biased towards the higher degrees of similarity CC > 0.95. The two events also have similar two-sided $\Delta \ln A$ distributions that peak at $\Delta \ln A \simeq 0.25$, indicating that on average the synthetics underestimate the amplitude of the observed waveforms by 25 per cent. We cannot know at this stage if this anomaly is due to an underestimation of the seismic moments of the events, or to an overestimation of attenuation. The $\Delta \tau$ distributions for the two events are also two-sided. The shallow event $\Delta \tau$ values peak between 0 and 4 s, indicating that the synthetics are moderately faster than the observed records; the deep event $\Delta \tau$ distribution peaks at much higher time lags of 8–10 s. Possible explanations for these large average time lags include an origin time error, and/or an overestimation of the seismic velocity at the source location.

3.2 Regional tomography of the Japan subduction zone

Our second scenario is a regional-scale tomographic study of the Japan subduction zone, using a set of local events within the depth range 0–600 km. The lateral dimensions of the domain are

$44°(EW) \times 33°(NS)$ (108–152°E and 18–51°N). The initial model is constructed using the southeast Asia model of Lebedev & Nolet (2003) as the background model, with P-wave velocity anomalies added from a high-resolution Japan P-wave model (Zhao *et al.* 1994) and S-wave velocity anomalies scaled to P by a factor of 1.5 (Chen *et al.* 2007a). Two different crustal models are implemented in the spectral-element mesh: inside the region of the high-resolution model (32–45°N, 130–145°E and down to 500 km), the crustal model is derived from the arrival times of local shallow earthquakes (Zhao *et al.* 1992); outside this region, the crustal model is CRUST2.0 (Bassin *et al.* 2000).

We collected data for more than 200 events with M_w 4.5–8 that occurred between 2000 and 2006. The source locations and focal mechanisms are the centroid-moment tensor (CMT) solutions. We used a total of 818 stations from three different networks (GSN, F-net and Hi-net): the 119 stations of the GSN and F-net provide broad-band records, whereas the 699 Hi-net stations provide only high-frequency records. We use the one-chunk version of spectral-element code to calculate synthetic seismograms accurate at periods of ~6 s and longer (Chen *et al.* 2007a), and present results for two period ranges: 6–30 s, using all the records, and 24–120 s, using the broad-band records only.
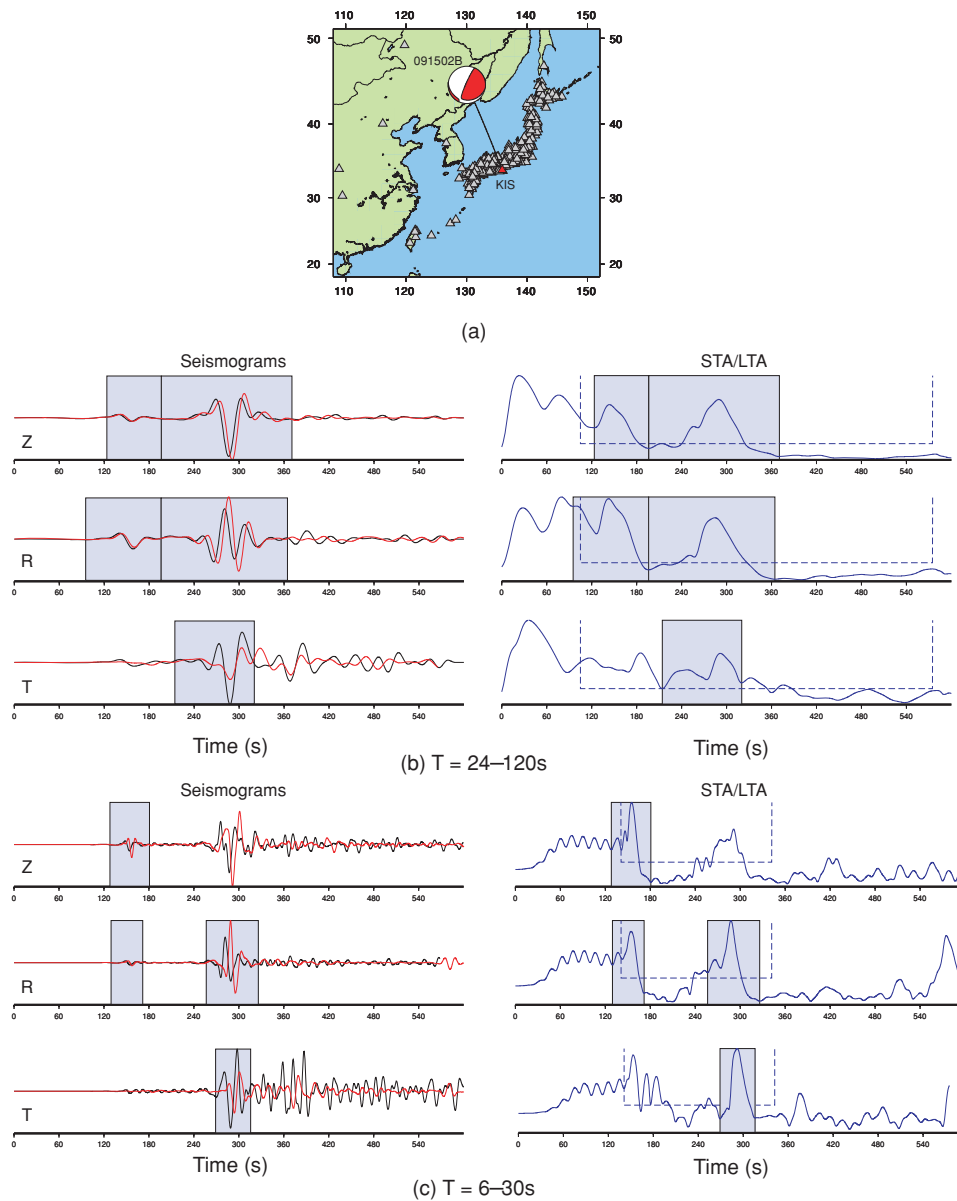
Figure 10. Window selection results for event 091502B from Table 2 recorded at station KIS ($\Delta = 11.79°$). (a) Map showing all stations with at least one measurement window for the period range 24–120 s for this event. Red triangle denotes station KIS. (b) Results for station KIS for the period range 24–120 s. Vertical (Z), radial (R), and transverse (T) records of data (black, left column) and synthetics (red, left column), as well as the STA:LTA records (right column) used to produce the window picks. (c) Results for station KIS for the period range 6–30 s.

Figs 10 and 11 show windowing results for two events at two different depths (Table 2): 091502B, 589.4 km deep, northeastern China; 200511211536A, 155 km deep, Kyushu, Japan. We have tuned the windowing algorithm using different sets of parameters for the two period ranges (see Table 3). In the period range 24–120 s, the water level is raised after the surface-wave arrivals to exclude the later arrivals that are not sensitive to upper-mantle structure. In the period range 6–30 s, the water level is raised after the *S*-wave arrivals to exclude the surface-waves, as the current crustal model is not detailed enough to predict the short-period surface-waves.

Fig. 10 shows an example of window picks for a deep event beneath northeastern China (091502B) recorded at station KIS. The seismograms from this event are relatively simple, containing only two major body-wave arrivals (*P* and *S*). The windowing algorithm's

similarity criterion comes into play here, causing it not to pick the short-period *S* arrival on the vertical component (Fig. 10c) as the distorted *S*-wave waveform of the data is quite different from the Gaussian shaped synthetics. The long-period *S*-wave arrival on the same component is selected due to higher data-synthetic waveform similarity (Fig. 10b).

The records of the intermediate-depth event (200511211536A) recorded by station SHR (Fig. 11) contain more seismic phases than the previous two examples. On the vertical component of the short-period seismogram (Fig. 11c), the *P*-wave arrives at ∼230 s, immediately followed by *pP* and *sPn*, and the *S*-wave arrives at ∼420 s, followed by *sS* and *PcP*. The windowing algorithm selects separate windows for the *P*, *sPn*, *S* and *sS* arrivals on the vertical component, and selects only the *P* and *sPn* arrivals on the radial component. In the period range 24–120 s (Fig. 11b), the *P* and *S*
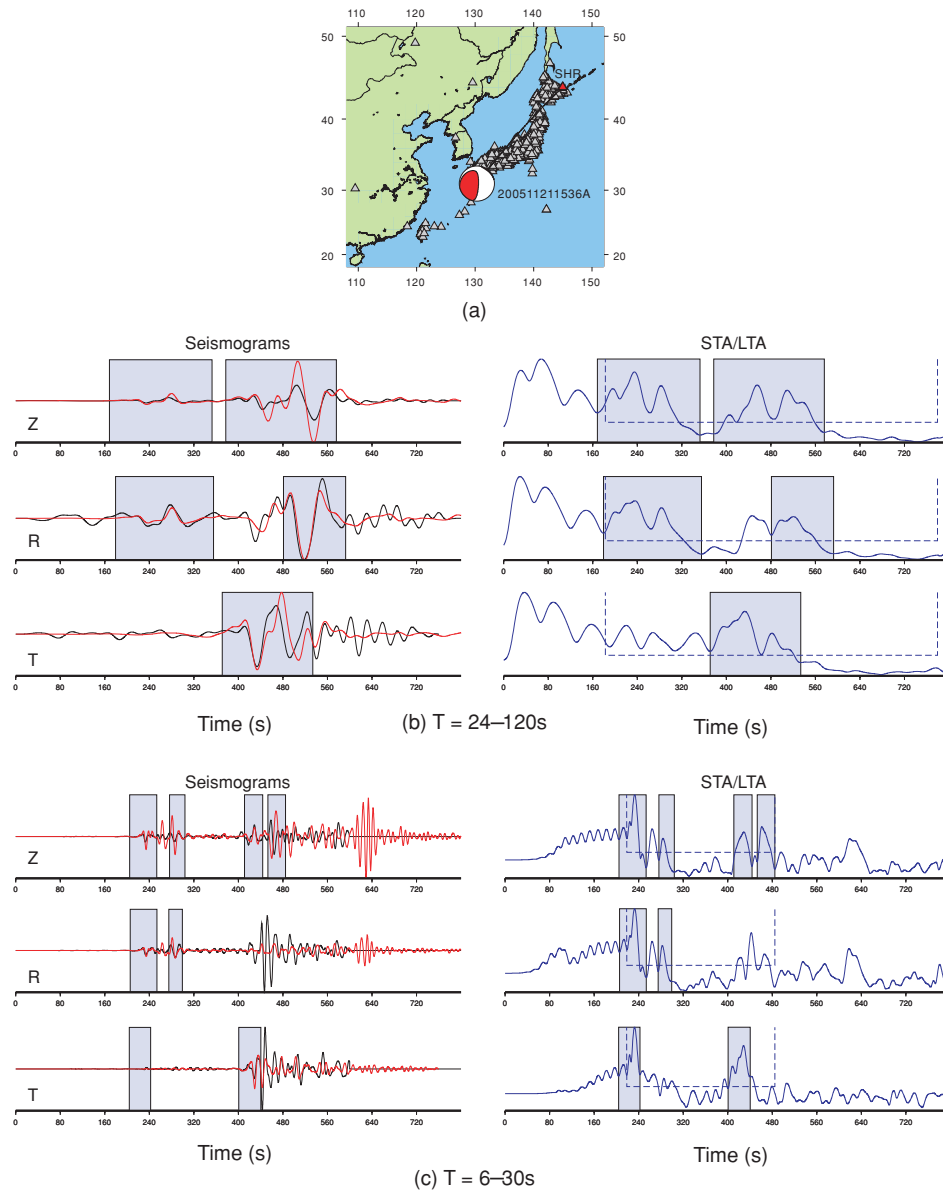
Figure 11. Window selection results for event 20051121536A from Table 2 recorded at station SHR ($\Delta = 17.47°$). (a) Map showing all stations with at least one measurement window for the period range 24–120 s for this event. Red triangle denotes station SHR. (b) Results for station SHR for the period range 24–120 s. Vertical (Z), radial (R), and transverse (T) records of data (black, left column) and synthetics (red, left column), as well as the STA:LTA records (right column) used to produce the window picks. (c) Results for station SHR for the period range 6–30 s.

waves merge with the arrivals that follow them, causing the windowing algorithm to select wave packets instead of single phases: $P + pP + sPn$ and $S + sS + PcP$ on the vertical and radial components, and $S + sS$ on the transverse component.

Summary plots of window picks for event 200511211536A in the two period ranges 6–30 s and 24–120 s are shown in Figs 12 and 13, respectively. On the short-period window record sections, the windows picked by the algorithm form two main branches that correspond to P and S arrivals. Some P arrivals are visible even on the transverse component (Fig. 12g). The number and width of windows for each trace varies with epicentral distance. On the vertical and radial components (Figs 12e and f), beyond a distance of 13° and after the P-arrival branch, there are two small branches corresponding to pP and sPn, whereas after the S-arrival branch

there is another branch corresponding to sS. The summary plot for the 24–120 s period range shows a single branch of windows on the vertical component, that splits up into separate P- and S-wave packets at distances greater than 15°. The same split is visible on the radial component, but occurs earlier (around 10°), whereas the transverse component windows form a single branch containing the merged $S + sS$ arrivals.

Comparison of the histograms in Figs 12 and 13 shows that windows selected on the 24–120 s seismograms tend to have higher degrees of waveform similarity than those selected on the 6–30 s records. $\Delta \tau$ values peak between −5 s and 0 s in both period ranges, indicating that the synthetics are slower than the observed records. The particularly large peak at −2 s in the $\Delta \tau$ distribution of Fig. 12(c) is probably due to the large number of Hi-net recordings
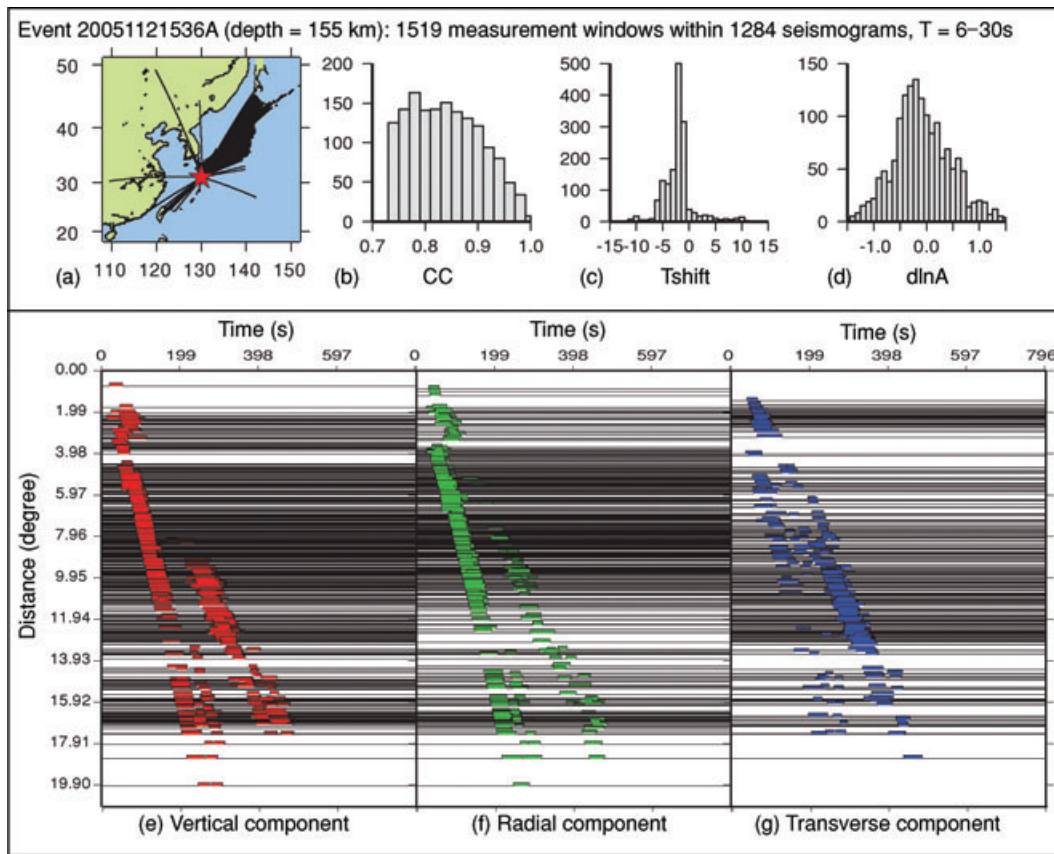
Figure 12. Summary plots of windowing results for event 200511211536A in Table 2, for the period range 6–30 s. (a) Map showing paths to each station with at least one measurement window. (b)-(d) Histograms of number of windows as a function of normalised cross-correlation CC, time-lag τ and amplitude ratio $\Delta \ln A$. (e)-(g) Record sections of selected windows for the vertical, radial and transverse components.

that make up the short-period range records. The $\Delta \ln A$ distribution peaks at $\Delta \ln A \simeq 0$ for 24–120 s (Fig. 13d), indicating the amplitude of the synthetics matches the amplitude of the data at long periods. The peak at $\Delta \ln A \simeq -0.2$ in Fig. 12(d) indicates that, on average, the synthetics overestimate the amplitude of the observed waveforms by 20 per cent at short-periods. We cannot know at this stage if this anomaly is due to an overestimation of short-period energy in the source spectra of the events, or to an underestimation of the seismic attenuation.

Fig. 14 shows summary plots of window picks for the shallow (051502B), intermediate (200511211536A) and deep events (091502B) for the period range 6–30 s. Note the very large numbers of measurement windows picked due to the over 600 Hi-net stations: 1361 windows for event 051502B, 1519 windows for event 200511211536A, and 2099 windows for event 091502B. Comparing the statistics for these three events, we see that the degree of similarity, CC, improves with increasing event depth, implying that the representation of mantle structure is better than that of crustal structure in the initial model. The $\Delta \ln A$ distributions of these three events have similar shapes, with peaks in the range of −0.5–0. However, the $\Delta \tau$ distributions have very different features: the shallow event (051502B) has a large peak at −9 s and another smaller peak at 8 s; the intermediate-depth event (200511211536A) has sharp peak at −2 s; the deep event(091502B) has a more distributed $\Delta \tau$ in the range −2 to −10 s. Possible explanations for these large average time lags include an origin time error, and/or an underestimation of the seismic velocity at the source location.

3.3 Local tomography in Southern California

Our last scenario is a local tomographic study of southern California. We apply the windowing algorithm to a set of 140 events within southern California, for which we have computed synthetic seismograms using the spectral-element method and a regional 3-D crustal and upper-mantle model (Komatitsch *et al.* 2004). This model contains three discontinuities: the surface topography (included in the mesh), the basement layer that separates the sedimentary basins from the bedrock, and the Moho, separating the lower crust from the upper mantle. The model includes several sedimentary basins, such as the Ventura basin, the Los Angeles basin, and the Salton trough (Komatitsch *et al.* 2004; Lovely *et al.* 2006). The smooth 3-D background velocity model used in Komatitsch *et al.* (2004) was determined by Hauksson (2000); we use an updated version provided by Lin *et al.* (2007b). The physical domain of the model is approximately 600 km by 500 km at the surface, and extends to a depth of 60 km. Our simulations of seismic waves are numerically accurate down to a period of 2 s.

The 140 events have M_w magnitudes between 3.5 and 5.5 and were recorded between 1999 and 2007. The locations and origin times are primarily from Lin *et al.* (2007a), and the focal mechanisms are from Clinton *et al.* (2006), Hardebeck & Shearer (2003), or Tan (2006).

We test the windowing code using three period ranges: 6–30 s, 3–30 s and 2–30 s. The parameters we use for the windowing code are listed in Table 3. Figs 15 and 16 show examples of the output
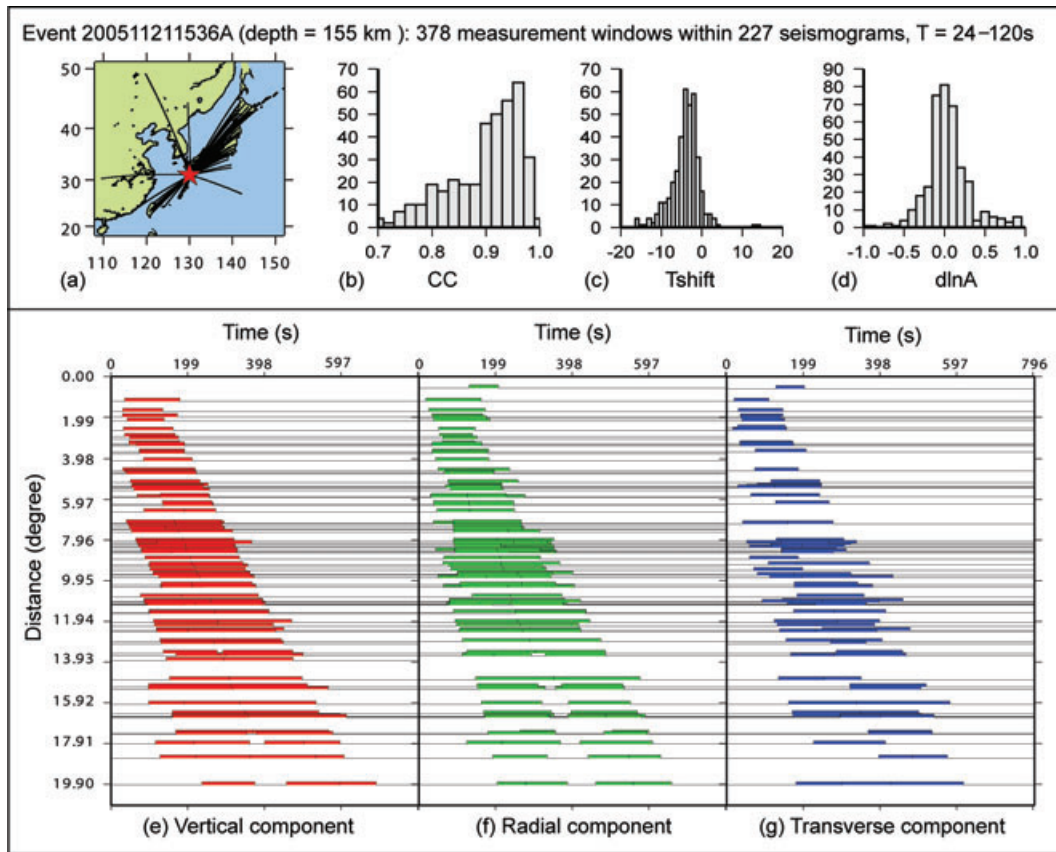
Figure 13. Summary plots of windowing results for event 200511211536A in Table 2, for the period range 24–120 s.

from the windowing algorithm for event 9818433 listed in Table 2 recorded at two different stations, whereas Fig. 17 shows a summary plot for event 9983429 in the 6–30 s period range.

The windowing algorithm tends to identify five windows on each set of three-component 6–30 s seismograms (Figs 15 and 17): on the vertical and radial components the first window corresponds to the body-wave arrival and the second to the Rayleigh wave, whereas windows on the transverse component capture the Love wave. The 2–30 s synthetic seismograms agree with the observed seismograms, only in the early part of the signal, leading to fewer picked windows. In Fig. 15c, only three windows are selected by the algorithm: the *P* arrival recorded on the radial component, the *S* arrival on the transverse component, and the Love-wave arrival on the transverse component. The *P* arrival (*PmP* or *Pn*) in fact appears on all three components on both data and synthetics. On the vertical component it is rejected because the cross-correlation value within the time window did not exceed the specified minimum value of 0.85 (Table 3). On the transverse component it does not have a large enough signal-to-noise ratio to be picked, but it is evident as a small peak at 36 s in the STA:LTA curve, and it is more conspicuous when zooming into the synthetics and data. The presence of the *P* arrival on the transverse component highlights the possibility of measuring subtle phases that may be present in 3-D synthetics.

Fig. 16 shows results for the same event as Fig. 15, but for a different station, FMP, situated 52 km from the event and within the Los Angeles basin. Comparison of the two figures highlights the characteristic resonance caused by the thick sediments within the basin. This resonance is beautifully captured by the transverse component synthetics (Fig. 16b, record T), thanks to the inclusion

of the basin in the model (Komatitsch *et al.* 2004). To pick such long time windows with substantial frequency-dependent measurement differences, we are forced to lower the minimum cross-correlation value CC_0 for the entire data set (0.71 in Table 3) and increase c_{4b} to capture the slow decay in the STA:LTA curves (Fig. 16b, record T). It is striking that although these arrivals look nothing like the energy packets typical for the global case, the windowing algorithm is still able to determine the proper start and end times for the windows. In Fig. 16c the windowing algorithm selects three short-period body-wave time windows with superb agreement between data and synthetics.

4 USING FLEXWIN FOR TOMOGRAPHY

The window selection algorithm we describe in this paper was designed to solve the problem of automatically picking windows for tomographic problems in which phase separation and identification are not necessary: 3D–3D numerical tomography, of which the adjoint tomography proposed by Tromp *et al.* (2005) and Tape *et al.* (2007) is an example. For these problems, our algorithm provides a window-selection solution that is midway between full-waveform selection—which carries the risk of including high-noise portions of the waveform that would contaminate the tomography—and the selection of known phases or phase-groups based on *a priori* arrival times—which carries the risk of missing the information contained in the non-traditional phases produced by fully 3-D structures.

FLEXWIN may also be used to select windows for tomographic problems in which separation of seismic arrivals is necessary and
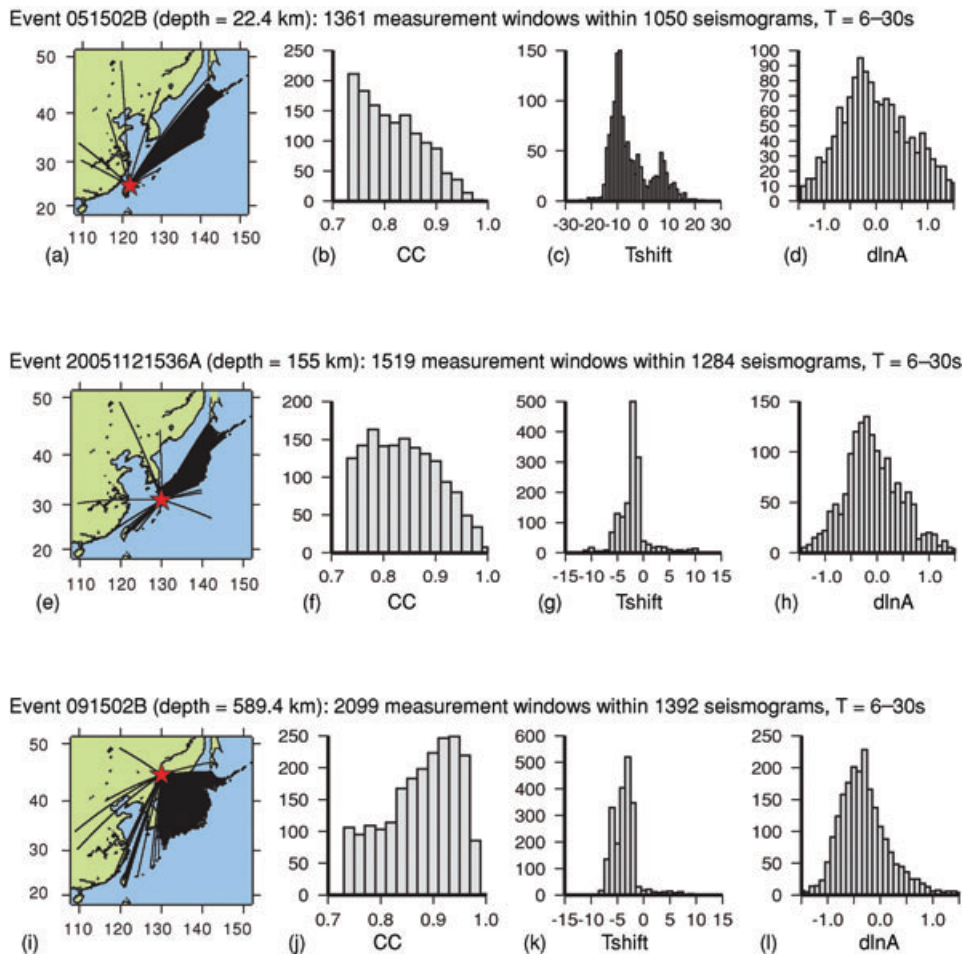
Figure 14. Summary statistics of windowing results for events 051502B, 200511211536A and 091502B in Table 2, for the period range 6–30 s.

occurs naturally (under certain frequency and epicentral distance conditions) by virtue of differences in traveltimes. It can straightforwardly be adapted to studies of distinct body-wave phases (e.g. Ritsema & van Heijst 2002) or to emulate the wavepacket selection of Panning & Romanowicz (2006) by modulating the $w_E(t)$ water-level using predicted phase arrival times, and selecting appropriate values for the signal-to-noise, cross-correlation and amplitude limits. The method can also be used to pre-select windows for studies of fundamental mode surface-waves (e.g. those based on the methods of Trampert & Woodhouse 1995; Ekström *et al.* 1997; Levshin & Ritzwoller 2001) by modulating $w_E(t)$ to exclude portions of the waveform that do not correspond to the correct group velocity window or epicentral distance range. Given the dispersed nature of surface-waves, synthetics produced by 1-D starting models often are considerably different in shape from the data, so the *CC* and ΔT conditions (but not the signal-to-noise or $\Delta \ln A$ conditions) should be relaxed in the window selection. These windows should then be passed on to specific algorithms used to extract the dispersion information. For this class of tomographic problems, the advantages of using FLEXWIN over manual or specifically designed automated windowing would be the encapsulation of the selection criteria entirely within the parameters of Table 1 (and their time-dependent modulation), leading to greater clarity and portability between studies using different inversion methods.

FLEXWIN is not intended for tomographic problems in which the extraction and separation of information from overlapping

portions of a single timeseries is required, for example studies of higher-mode surface-wave dispersion for which specific methods have been developed: mode branch stripping (Van Heijst & Woodhouse 1997), separation of secondary observables (Cara & Lévêque 1987; Debayle 1999), partitioned waveform and automated multimode inversion (Nolet 1990; Lebedev *et al.* 2005), and non-linear direct search (Yoshizawa & Kennett 2002; Visser *et al.* 2007).

4.1 Relevance to adjoint tomography

The full power of FLEXWIN can only be unleashed for problems—such as 3D–3D tomography—that do not require the separation (natural or otherwise) of seismic phases. The specificity of adjoint tomography, among the 3D–3D tomographic methods, is to calculate the sensitivity kernels by interaction between the wavefield used to generate the synthetic seismograms and an adjoint wavefield whose source term is derived from measurements of misfit between the synthetic and observed seismograms (Tromp *et al.* 2005; Liu & Tromp 2006, 2008). The manner in which the adjoint sources are constructed is specific to each type of measurement (e.g. waveform difference, cross-correlation time lag, multitaper phase and amplitude anomaly), but once formulated can be applied indifferently to any part of the seismogram. Adjoint methods have been used to calculate kernels of various body- and surface-wave phases with respect to isotropic elastic parameters and interface depths
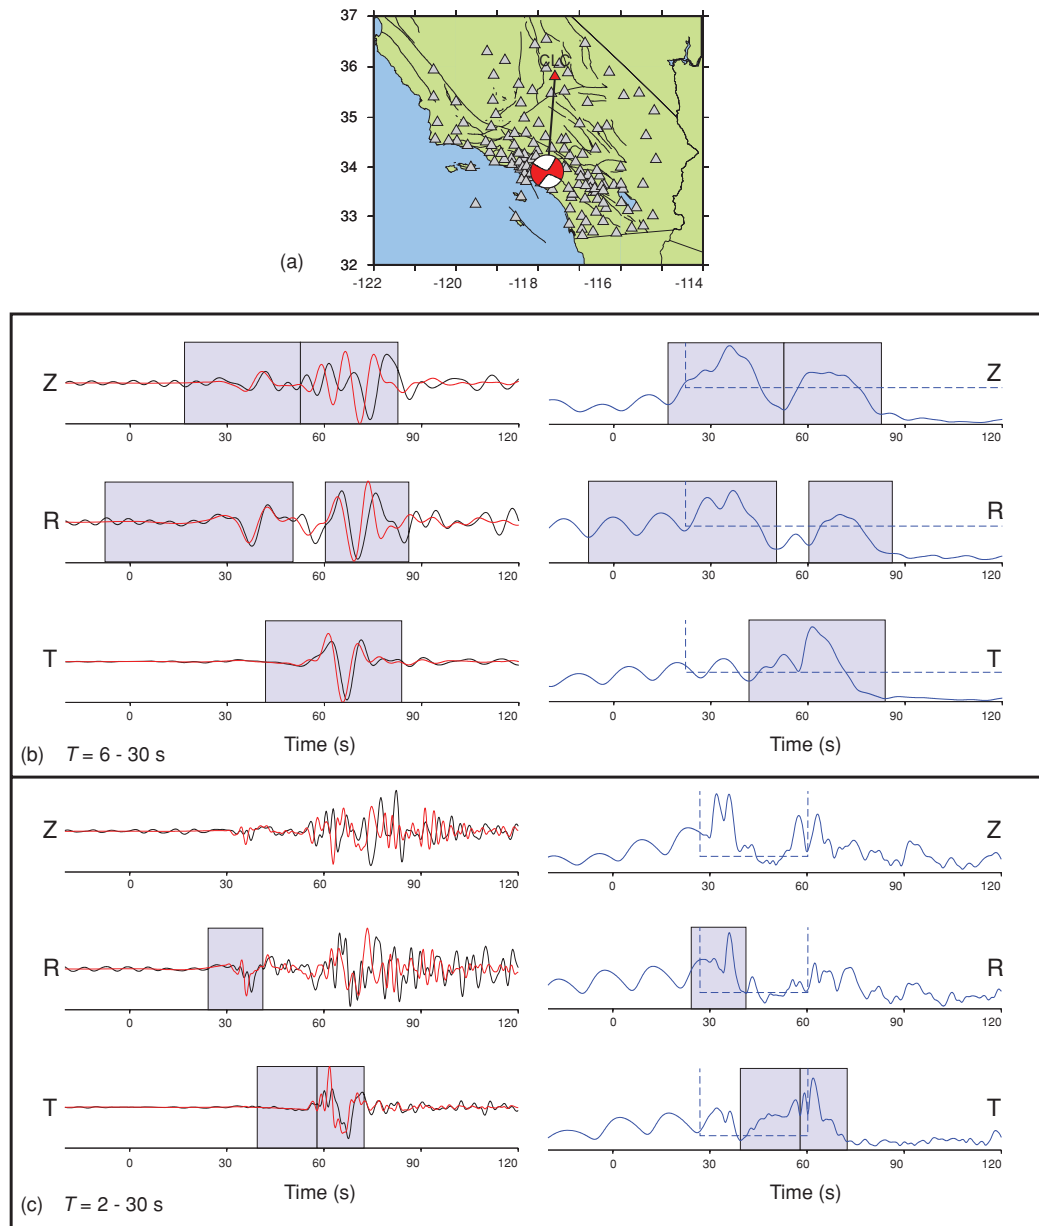
Figure 15. Window selection results for event 9818433 from Table 2 recorded at station CLC ($\Delta = 211.7$ km). (a) Map showing all stations with at least one measurement window for the period range 6–30 s for this event. Red triangle denotes station CLC. (b) Results for station CLC for the period range 6–30 s. Vertical (Z), radial (R), and transverse (T) records of data (black, left column) and synthetics (red, left column), as well as the STA:LTA records (right column) used to produce the window picks. (c) Results for station CLC for the period range 2–30 s.

(Liu & Tromp 2006,2008), and with respect to anisotropic elastic parameters (Sieminski *et al.* 2007a,b). Adjoint methods allow us to calculate kernels for each and every wiggle on a given seismic record, thereby giving access to virtually all the information contained within.

It is becoming clear, as more finite-frequency tomography models are published, that better kernels on their own are not the answer to the problem of improving the resolution of tomographic studies. Trampert & Spetzler (2006) and Boschi *et al.* (2007) investigate the factors limiting the quality of finite-frequency tomography images, and conclude that incomplete and inhomogeneous data coverage limit in practice the improvement in resolution that accurate finite-frequency kernels can provide. The current frustration with the data-induced limitations to the improvements in wave-propagation theory

is well summarized by Romanowicz (2008). The ability of adjoint methods to deal with all parts of the seismogram indifferently means we can incorporate more information from each seismogram into a tomographic problem, thereby improving data coverage.

The computational cost of constructing an adjoint kernel is independent of the number of time windows we choose to measure, and also of the number of records of a given event we choose to work with. It is therefore computationally advantageous to make measurements on as many records as possible for each event, while covering as much as possible of each record. There are, however, certain limits we must be aware of. As mentioned in the introduction, there is nothing in the adjoint method itself that prevents us from constructing a kernel from noise-dominated portions of the data. As the purpose of 3D–3D tomography is to improve the fine
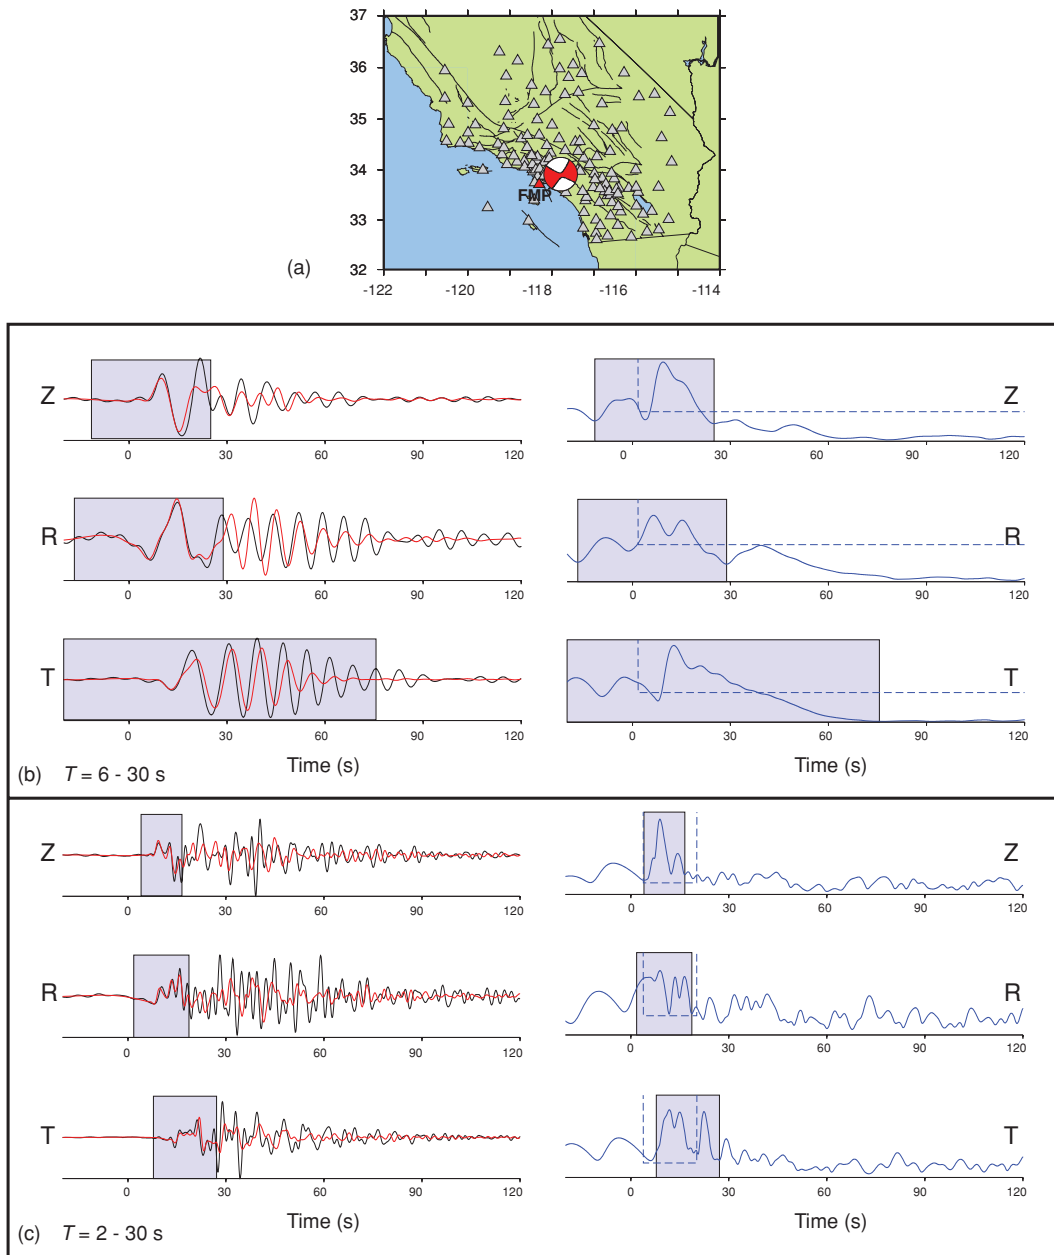
Figure 16. Window selection results for event 9818433 from Table 2 recorded at station FMP ($\Delta = 52.2$ km). Same caption as Fig. 15, only for a different station.

details of Earth models, it would be counterproductive to pollute the inversion process with such kernels. It is clear that the use of adjoint methods for tomography requires a strategy for selecting and windowing seismograms that avoids seismic noise while extracting as much information as possible from the signals.

The adjoint kernels are only strictly valid for the 3-D Earth model they were constructed in, and therefore need to be re-computed at each iteration of the tomographic inversion (Tape *et al.* 2007). At each iteration, the similarities between the synthetic and observed seismograms improve, such that for later iterations a greater proportion of the waveform is adequate for measurement. To take advantage of this extra information, the windowing method used to isolate the portions of the waveform to be measured needs to be automated. The method must also be adaptable to the features that exist in the seismograms themselves, because 3-D wave-

field simulations are able to synthesize phases that do not exist in 1-D simulations or traditional traveltime curves. All these considerations led us to favour a signal processing approach to the problem of data selection, an approach which in turn led to the development of the FLEXWIN algorithm we have presented here.

Finally, we note that the design of this algorithm is based on the desire *not* to use the entire timeseries of each event when making a measurement between data and synthetics. If one were to simply take the waveform difference between two timeseries, then there would be no need for selecting time windows of interest. However, this ideal approach (e.g. Gauthier *et al.* 1986) may only work in real applications if the statistical properties of the noise are well known, which is rare. Without an adequate description of the noise, it is more prudent to resort to the selection of time windows even
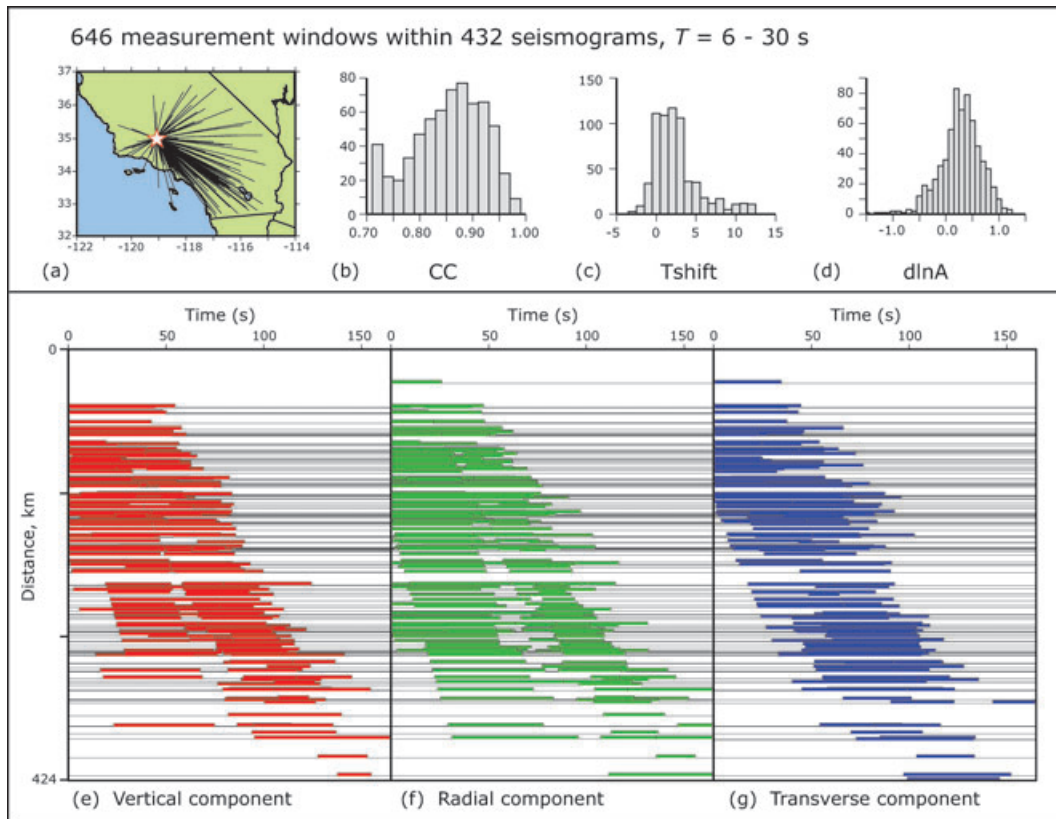
Figure 17. Summary plots of windowing results for event 9983429 in Table 2, for the period range 6–30 s. (a) Map showing paths to each station with at least one measurement window. (b)-(d) Histograms of number of windows as a function of normalised cross-correlation CC, time-lag τ and amplitude ratio $\Delta \ln A$. (e)-(g) Record sections of selected windows for the vertical, radial and transverse components. The two branches observed on the vertical and radial components correspond to the body-wave arrivals and the Rayleigh wave arrivals.

when tomographic inversion is performed on waveform difference measurements.

5 SUMMARY

The FLEXWIN algorithm was designed to automatically pick time windows for tomographic problems in which phase separation and identification are not necessary, however it can also be applied to problems in which phase separation is necessary and occurs naturally. It provides an automated window-selection solution that is midway between full-waveform selection and the selection of known phases or phase-groups based on *a priori* arrival times.

FLEXWIN has no *a priori* knowledge related to input model, geographic scale or frequency range. It is a configurable process that can be applied to different seismic scenarios by changing the handful of parameters in Table 1. The configuration process is data-driven: starting from the description of how each parameter influences the window selection (Section 2 and Appendix A), the user tunes the parameters using a representative subset of the full data set until the algorithm produces an adequate set of windows, then applies the tuned algorithm to the full data set. The choice of what makes an adequate set of windows remains subjective, as it depends strongly on the quality of the input model, the quality of the data, and the region of the Earth the tomographic inversion aims to constrain. We consider the algorithm to be correctly tuned when false positives (windows around undesirable features of the seismogram) are minimized, and true positives (window around desirable features) are maximized. For a given data set, the set of tuned parameters

and their user-defined time dependencies completely determine the window selection results, which are therefore entirely reproducible.

The desire to study regions with strong 3-D variations in Earth structure requires seismologists to deal with increasingly complex seismic records, and to use methods that take advantage of full wavefield simulations. Only by using all available information will tomographic inversions produce more accurate and higher resolution images of the Earth's interior. A window selection method such as FLEXWIN is necessary to fully unleash the potential of recent tomographic methods—and specifically of adjoint tomography—to exploit information from all parts of the waveform.

FLEXWIN is available as an open-source package through CIG (Computational Infrastructure for Geodynamics, http://www.geodynamics.org).

REFERENCES

Allen, R.V., 1982. Automatic phase pickers: their present use and future prospects, *B. Seismol. Soc. Am.,* **72,** S225–S242.

Aster, R. & Rowe, A., 2000. Automatic phase pick refinement and similar event association in large seismic data sets, in *Advances in seismic event location,* pp. 231–263. Kluwer Academic Publishers, Dordrecht.

Bai, C. & Kennett, B.L.N., 2000. Automatic phase-detection and identification by full use of a single three-component broadband seismogram, *B. Seismol. Soc. Am.,* **90,** 187–198.

Bai, C.Y. & Kennett, B.L.N., 2001. Phase identification and attribute analysis of broadband seismograms at far-regional distances, *J. Seismol.,* **5,** 217–231.

Bassin, C., Laske, G. & Masters, G., 2000. The current limits of resolution for surface wave tomography in North America, *EOS, Trans. Am. geophys. Un.,* **81,** 897.

Beucler, E., Stutzmann, E. & Montagner, J.-P., 2003. Surface wave higher-mode phase velocity measurements using a roller-coaster-type algorithm, *Geophys. J. Int.,* **155,** 289–307.

Boschi, L., Ampuero, J.P., Peter, D., Mai, P.M., Soldati, G. & Giardini, D., 2007. Petascale computing and resolution in global seismic tomography, *Phys. Earth planet. Inter.,* **163,** 245–250.

Capdeville, Y., 2005. An efficient Born normal mode method to compute sensitivity kernels and synthetic seismograms in the Earth, *Geophys. J. Int.,* **163,** 639–646.

Capdeville, Y., Chaljub, E., Vilotte, J.P. & Montagner, J.P., 2003. Coupling the spectral element method with a modal solution for elastic wave propagation in global earth models, *Geophys. J. Int.,* **152,** 34–67.

Cara, M. & Lévêque, J.J., 1987. Waveform inversion using secondary observables, *Geophy. Res. Lett.,* **14,** 1046–1049.

Chen, M., Tromp, J., Helmberger, D. & Kanamori, H., 2007a. Waveform modeling of the slab beneath Japan, *J. Geophys. Res.,* **112,** B02305, doi:10.1029/2006JB004394.

Chen, P., Zhao, L. & Jordan, T.H., 2007b. Full 3D tomography for the crustal structure of the Los Angeles region, *B. Seismol. Soc. Am.,* **97,** 1094–1120.

Clinton, J.F., Hauksson, E. & Solanki, K., 2006. An evaluation of the SCSN moment tensor solutions: Robustness of the M_w magnitude scale, style of faulting, and automation of the method, *B. Seismol. Soc. Am.,* **96,** 1689–1705.

Dahlen, F. & Baig, A., 2002. Fréchet kernels for body wave amplitudes, *Geophys. J. Int.,* **150,** 440–446.

Dahlen, F.A. & Zhou, Y., 2006. Surface-wave group-delay and attenuation kernels, *Geophys. J. Int.,* **165,** 545–554.

Dahlen, F.A., Hung, S.-H. & Nolet, G., 2000. Fréchet kernels for finite-frequency traveltimes – I. Theory, *Geophys. J. Int.,* **141,** 157–174.

Debayle, E., 1999. SV-wave azimuthal anisotropy in the Australian upper mantle: preliminary results from automated Rayleigh waveform inversion, *Geophys. J. Int.,* **137,** 747–754.

Dziewonski, M.A. & Anderson, D.L., 1981. Preliminary Reference Earth Model, *Phys. Earth planet. Inter.,* **25**(4), 297–356.

Earle, P.S. & Shearer, P.M., 1994. Characterization of global seismograms using an automatic picking algorithm, *B. Seismol. Soc. Am.,* **84,** 366–376.

Ekström, G., Tromp, J. & Larson, E.W.F., 1997. Measurements and global models of surface wave propagtaion, *J. geophys. Res.,* **102,** 8137–8157.

Gauthier, O., Virieux, J. & Tarantola, A., 1986. Two-dimensional nonlinear inversion of seismic waveforms: numerical results, *Geophysics,* **51,** 1387–1403.

Hardebeck, J.L. & Shearer, P.M., 2003. Using S/P amplitude ratios to constrain the focal mechanisms of small earthquakes, *B. Seismol. Soc. Am.,* **93**(6), 2434–2444.

Hauksson, E., 2000. Crustal structure and seismicity distribution adjacent to the Pacific and North America plate boundary in southern California, *J. geophys. Res.,* **105,** 13 875–13 903.

Houser, C., Masters, G. & Laske, G., 2008. Shear and compressional velocity models of the mantle from cluster analysis of long-period waveforms, *Geophys. J. Int.,* **174,** 195–212.

Kennett, B.L.N. & Engdahl, E.R., 1991. Traveltimes for global earthquake location and phase identification, *Geophys. J. Int.,* **105,** 429–465.

Komatitsch, D. & Vilotte, J.-P., 1998. The spectral element method: An efficient tool to simulate the seismic response of 2D and 3D geological structures, *B. Seismol. Soc. Am.,* **88,** 368–392.

Komatitsch, D., Ritsema, J. & Tromp, J., 2002. The spectral-element method, Beowulf computing, and global seismology, *Science,* **298,** 1737–1742.

Komatitsch, D., Liu, Q., Tromp, J., Süss, P., Stidham, C. & Shaw, J.H., 2004. Simulations of ground motion in the Los Angeles basin based upon the spectral-element method, *B. Seismol. Soc. Am.,* **94,** 187–206.

Laske, G. & Masters, G., 1996. Constraints on global phase velocity maps from long-period polarization data, *J. geophys. Res.,* **101,** 16 059–16 075.

Lawrence, J.F. & Shearer, P.M., 2008. Imaging mantle transition zone thickness with *sds-ss* finite-frequency sensitivity kernels, *Geophys. J. Int.* **174,** 143–158.

Lawrence, J.F., Shearer, P.M. & Masters, G., 2006. Mapping attenuation beneath north america using waveform cross-correlation and cluster analysis, *Geophy. Res. Lett.,* **33,** L07315, doi:10.1029/2006GL025813.

Lebedev, S. & Nolet, G., 2003. Upper mantle beneath southeast Asia from *S* velocity tomography, *J. geophys. Res.,* **108**(B1), doi:10.1029/2000JB000073.

Lebedev, S., Nolet, G., Meier, T. & van der Hilst, R.D., 2005. Automated multimode inversion of surface and S waveforms, *Geophys. J. Int.,* **162,** 951–964.

Levshin, A.L. & Ritzwoller, M.H., 2001. Automated detection, extraction, and measurement of regional surface waves, *Pure appl. Geophys.,* **158,** 1531–1545.

Li, X.D. & Tanimoto, T., 1993. Waveforms of long-period body waves in a slightly aspherical earth model, *Geophys. J. Int.,* **112,** 92–102.

Lin, G., Shearer, P.M. & Hauksson, E., 2007a. Applying a three-dimensional velocity model, waveform cross correlation, and cluster analysis to locate southern California seismicity from 1981 to 2005, *J. geophys. Res.,* **112,** B12309, doi:10.1029/2007JB004986.

Lin, G., Shearer, P.M., Hauksson, E. & Thurber, C.H., 2007b. A three-dimensional crustal seismic velocity model for southern California from a composite event method, *J. geophys. Res.,* **112,** B11306, doi:10.1029/2007JB004977.

Liu, Q. & Tromp, J., 2006. Finite-frequency kernels based on adjoint methods, *B. Seismol. Soc. Am.,* **96**(6), 2383–2397.

Liu, Q. & Tromp, J., 2008. Finite-frequency sensitivity kernels for global seismic wave propagation based upon adjoint methods, *Geophys. J. Int.,* **174,** 265–286.

Lovely, P., Shaw, J.H., Liu, Q. & Tromp, J., 2006. A structural V_P model of the Salton Trough, California, and its implications for seismic hazard, *B. Seismol. Soc. Am.,* **96,** 1882–1896.

Marquering, H., Dahlen, F.A. & Nolet, G., 1999. Three-dimensional sensitivity kernels for finite-frequency traveltimes: the banana-doghnut paradox, *Geophys. J. Int.,* **137,** 805–815.

Montelli, R., Nolet, G., Dahlen, F.A., Masters, G., Engdahl, E.R. & Hung, S-H., 2004. Finite-frequency tomography reveals a variety of plumes in the mantle, *Science,* **303,** 338–343.

Nolet, G., 1990. Partitioned waveform inversion ant two-dimensional structure under the Network of Autonomously Recording Seismographs, *J. geophys. Res.,* **95,** 8499–8512.

Panning, M. & Romanowicz, B., 2006. A three-dimensional radially anisotropic model of shear velocity in the whole mantle, *Geophys. J. Int.,* **167,** 361–379.

Ritsema, J. & van Heijst, H.J., 2002. Constraints on the correlation of *P*- and *S*-wave velocity heterogeneity in the mantle from *P*, *PP*, *PPP* and *PKPAB* traveltimes, *Geophys. J. Int.,* **149,** 482–489.

Ritsema, J., van Heijst, H.J. & Woodhouse, J.H., 2004. Global transition zone tomography, *J. geophys. Res.,* **106** (**B02302**), doi:10.1029/2003JB002610.

Romanowicz, B., 2008. Using seismic waves to image Earth's internal structure, *Nature,* **451,** 266–268.

Sieminski, A., Liu, Q., Trampert, J. & Tromp, J., 2007a. Finite-frequency sensitivity of surface waves to anisotropy based upon adjoint methods, *Geophys. J. Int.,* **168,** 1153–1174.

Sieminski, A., Liu, Q., Trampert, J. & Tromp, J., 2007b. Finite-frequency sensitivity of body waves to anisotropy based on adjoint methods, *Geophys. J. Int.,* **171,** 368–389.

Sigloch, K. & Nolet, G., 2006. Measuring finite-frequency body-wave amplitudes and traveltimes, *Geophys. J. Int.,* **167,** 271–287.

Sleeman, R. & van Eck, T., 2003. Single station real-time P and S phase pickers for seismic observatories, in *Methods and Applications of Signal Processing in Seismic Network Operations,* Vol. 98, pp. 173–194, eds Takami, T. & Kitagawa, G., Lecture Notes in Earth Sciences, Springer, New York.

Tan, Y., 2006. Broadband waveform modeling over a dense seismic network, *PhD thesis*, California Institute of Technology.

Tape, C., Liu, Q. & Tromp, J., 2007. Finite-frequency tomography using adjoint methods. Methodology and examples using membrane surface waves, *Geophys. J. Int.,* **168**(3), 1105–1129.

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

Trampert, J. & Spetzler, J., 2006. Surface wave tomography: finite-frequency effects lost in the null space, *Geophys. J. Int.,* **164,** 394–400.

Trampert, J. & Woodhouse, J.H., 1995. Global phase velocity of Love and Rayleigh waves between 40 and 150 seconds, *Geophys. J. Int.,* **22,** 675–690.

Tromp, J., Tape, C. & Liu, Q., 2005. Seismic tomography, adjoint methods, time reversal and banana-doughnut kernels, *Geophys. J. Int.,* **160,** 195–216.

Van Heijst, H.J. & Woodhouse, J.H., 1997. Measuring surface-wave overtone phase velocities using a mode-branch stripping technique, *Geophys. J. Int.,* **131,** 209–220.

vanDecar, J.C. & Crosson, R.S., 1990. Determination of teleseismic relative phase arrival times using multi-channel cross-correlation and least squares, *B. Seismol. Soc. Am.,* **80**(1), 150–169.

Visser, K., Lebedev, S., Trampert, J. & Kennett, B.L.N., 2007. Global love wave overtone measurements, *Geophy. Res. Lett.,* **34,** L03302, doi:10.1029/2006GL028671.

Wald, L.A., Hutton, L.K. & Given, D.D., 1995. The Southern California Network Bulletin: 1990–1993 summary, *Seismol. Res. Lett.,* **66**(1), 9–19.

Withers, M., Aster, R., Young, C., Beiriger, J., Harris, M., Moore, S. & Trujillo, J., 1998. A comparison of select trigger algorithms for automated global seismic phase and event detection, *B. Seismol. Soc. Am.,* **88**(1), 96–106.

Yoshizawa, K. & Kennett, B.L.N., 2002. Non-linear waveform inversion for surface waves with a neighbourhood algorithm—application to multimode dispersion measurements, *Geophys. J. Int.,* **149**(1), 118–133.

Zhao, D., Horiuchi, S. & Hasegawa, A., 1992. Seismic velocity structure of the crust beneath the Japan Islands, *Tectonophysics,* **212,** 289–301.

Zhao, D., Hasegawa, A. & Kanamori, H., 1994. Deep structure of Japan subduction zone as derived from local, regional, and teleseismic events, *J. geophys. Res.,* **99,** 22 313–22 329.

Zhao, L., Jordan, T.H. & Chapman, C.H., 2000. Three dimensional fréchet differential kernels for seismic delay times, *Geophys. J. Int.,* **141,** 558–576.

Zhao, L., Jordan, T.H., Olsen, K.B. & Chen, P., 2005. Fréchet kernels for imaging regional earth structure based on three-dimensional reference models, *B. Seismol. Soc. Am.,* **95,** 2066–2080.

Zhou, Y., Nolet, G., Dahlen, F.A. & Laske, G., 2006. Global upper-mantle structure from finite-frequency surface-wave tomography, *J. geophys. Res.,* **111,** B04304, doi:10.1029/2005JB003677.

APPENDIX A: TUNING CONSIDERATIONS

FLEXWIN is not a black-box application, and as such cannot be applied blindly to any given data set or tomographic scenario. The data windowing required by any given problem will differ depending on the inversion method, the scale of the problem (local, regional, global), the quality of the data set and that of the model and method used to calculate the synthetic seismograms. The user must configure and tune the algorithm for the given problem. In this appendix we shall discuss some general considerations the user should bear in mind during the tuning process. For more detailed information on tuning, and for further examples of tuning parameter sets, we refer the reader to the user manual that accompanies the source code.

The order in which the parameters in Table 1 are discussed in the main text of this paper follows the order in which they are used by the algorithm, but is not necessarily the best order in which to consider them for tuning purposes. We suggest the following as a practical starting sequence (the process may need to be repeated and refined several times before converging on the optimal set of parameters for a given problem and data-set).

$T_{0,1}$: In setting the corner periods of the bandpass filter, the user is deciding on the frequency content of the information to be used in the tomographic problem. Values of these corner periods should reflect the information content of the data, the quality of the Earth model and the accuracy of the simulation used to generate the synthetic seismogram. The frequency content in the data depends on the spectral characteristics of the source, on the instrument responses, and on the attenuation characteristics of the medium. As $T_{0,1}$ depend on the source and station characteristics, which may be heterogeneous in any given data-set, these filter periods can be modified dynamically by constructing an appropriate user function (e.g. *if station is in list of stations with instrument X then reset T_0 and T_1 to new values*).

$r_{P,A}$: In setting the signal-to-noise ratios for the entire seismogram the user is applying a simple quality control on the data. Note that these criteria are applied after filtering. No windows will be defined on data that fail this quality control.

$w_E(t)$: The short-term average/long-term average ratio $E(t)$ of a constant signal converges to a constant value when the length of the time-series is greater than the effective averaging length of the long-term average. This value would be 1 if $C_S = C_L$ in eqs (4) and (5) of the main text. For C_S and C_L given by equation (6) of the main text, $E(t)$ of a constant signal converges to a value close to 0.08, with only a weak dependence on T_0. We suggest the user start with a constant level for $w_E(t)$ equal to this convergence value. The time dependence of $w_E(t)$ should then be adjusted to exclude those portions of the waveform the user is not interested in, by raising $w_E(t)$ (e.g. to exclude the fundamental mode surface-wave: *if $t >$ fundamental mode surface-wave arrival time then set $w_E(t) = 1$*). We suggest finer adjustments to $w_E(t)$ be made after $r_0(t)$, $CC_0(t)$, $\Delta T_0(t)$ and $\Delta \ln A_0(t)$ have been configured.

$r_0(t)$, $CC_0(t)$, $\Delta\tau_{ref}$, $\Delta\tau_0(t)$, $\Delta\ln A_{ref}$ and $\Delta\ln A_0(t)$: These parameters—window signal-to-noise ratio, normalized cross-correlation value between observed and synthetic seismograms, cross-correlation time lag, and amplitude ratio—control the degree of well-behavedness of the data within accepted windows (Stage D). The user first sets constant values for these four parameters, then adds a time dependence if required. Considerations that should be taken into account include the quality of the Earth model used to calculate the synthetic seismograms, the frequency range, the dispersed nature of certain arrivals (e.g. *for t corresponding to the group velocities of surface-waves, reduce $CC_0(t)$*), and *a priori* preferences for picking certain small-amplitude seismic phases (e.g. *for t close to the expected arrival for P_{diff}, reduce $r_0(t)$*). $\Delta\tau_{ref}$ and $\Delta\ln A_{ref}$ should be set to zero at first, and only reset if the synthetics contain a systematic bias in traveltimes or amplitudes.

c_{0-4}: These parameters are active in Stage C of the algorithm, the stage in which the suite of all possible data windows is pared down using criteria on the shape of the $E(t)$ waveform alone. Detailed descriptions of the behaviour of each parameter are available in Section 2.3 and will not be repeated here. We suggest the user start by setting these values to those used in our global example (see Table 3). Subsequent minimal tuning should be performed by running the algorithm on a subset of the data and closely examining the lists of windows rejected at each stage to make sure the user agrees with the choices made by the algorithm.

w_{CC}, w_{len} and w_{nwin}: These parameters control the overlap resolution stage of the algorithm (Stage E), and are discussed in detail in Section 2.5. Values of $w_{CC} = w_{len} = w_{nwin} = 1$ should be reasonable for most applications.

The objective of the tuning process summarily described here should be to maximize the selection of windows around desirable features in the seismogram, while minimizing the selection of undesirable features, bearing in mind that the desirability or undesirability of a given feature is subjective, and depends on how the user subsequently intends to use the information contained within the data windows.

A1 Examples of user functions

As concrete examples of how the time dependence of the tuning parameters can be exploited, we present here the functional forms of the time dependencies used for the three example tomographic scenarios described in the text (Section 3). In each example we use predicted arrival times derived from 1-D Earth models to help modulate certain parameters. Note, however, that the actual selection of individual windows is based on the details of the waveforms, and not on information from 1-D Earth models.

A1.1 Global scenario

In the following, h indicates earthquake depth, t_Q indicates the approximate start of the Love wave predicted by a group wave speed of 4.2 km s^{-1}, and t_R indicates the approximate end of the Rayleigh wave predicted by a group wave speed of 3.2 km s^{-1}. To reduce the number of windows picked beyond R1, and to ensure that those selected beyond R1 are a very good match to the synthetic waveform, we raise the water level on the STA:LTA waveform and impose stricter criteria on the signal-to-noise ratio and the waveform similarity after the approximate end of the surface-wave arrivals. We allow greater flexibility in cross-correlation time lag $\Delta\tau$ for inter-

mediate depth and deep earthquakes. We lower the cross-correlation value criterion for surface-waves to retain windows with a slight mismatch in dispersion characteristics.

We therefore use the following time modulations:

$$w_E(t) = \begin{cases} w_E & t \leq t_R, \\ 2w_E & t > t_R, \end{cases} \tag{A1}$$

$$r_0(t) = \begin{cases} r_0 & t \leq t_R, \\ 10r_0 & t > t_R, \end{cases} \tag{A2}$$

$$CC_0(t) = \begin{cases} CC_0 & t \leq t_R, \\ 0.9CC_0 & t_Q < t \leq t_R, \\ 0.95 & t > t_R, \end{cases} \tag{A3}$$

$$\Delta\tau_0(t) = \begin{cases} \begin{cases} \tau_0 & t \leq t_R, \\ \tau_0/3 & t > t_R, \end{cases} & h \leq 70\,\text{km} \\ 1.4\tau_0 & 70\,\text{km} < h < 300\,\text{km}, \\ 1.7\tau_0 & h \geq 300\,\text{km}, \end{cases} \tag{A4}$$

$$\Delta\ln A_0(t) = \begin{cases} \Delta\ln A_0 & t \leq t_R, \\ \Delta\ln A_0/3 & t > t_R. \end{cases} \tag{A5}$$

A1.2 Japan scenario

In the following, t_P and t_S denote the start of the time windows for P- and S waves, as predicted by the 1-D IASPEI91 model (Kennett & Engdahl 1991), and t_{R1} indicates the end of the surface-wave time window. For the 24–120 s data, we consider the waveform between the start of the P wave to the end of the surface-wave. We therefore modulate $w_E(t)$ as follows:

$$w_E(t) = \begin{cases} 10w_E & t < t_P, \\ w_E & t_P \leq t \leq t_{R1}, \\ 10w_E & t > t_{R1}. \end{cases} \tag{A6}$$

For the 6–30 s data, the fit between the synthetic and observed surface-waves is expected to be poor, as the 3-D model used to calculate the synthetics cannot produce the required complexity. We therefore want to concentrate on body-wave arrivals only, and avoid surface-wave windows altogether by modulating $w_E(t)$ as follows:

$$w_E(t) = \begin{cases} 10w_E & t < t_P, \\ w_E & t_P \leq t \leq t_S, \\ 10w_E & t > t_S. \end{cases} \tag{A7}$$

We use constant values of $r_0(t) = r_0$, $CC_0(t) = CC_0$ and $\Delta\ln A_0(t) = \Delta\ln A_0$ for both period ranges. To allow greater flexibility in cross-correlation time lag $\Delta\tau$ for intermediate depth and deep earthquakes we use:

$$\Delta\tau_0(t) = \begin{cases} 0.08t_P & h \leq 70\,\text{km}, \\ \max(0.05t_P, 1.4\tau_0) & 70\,\text{km} < h < 300\,\text{km}, \\ \max(0.05t_P, 1.7\tau_0) & h \geq 300\,\text{km}. \end{cases} \tag{A8}$$

A1.3 Southern California scenario

In the following, t_P and t_S denote the start of the time windows for the crustal P wave and the crustal S wave, computed from a 1-D layered model appropriate to Southern California (Wald *et al.* 1995). The start and end times for the surface-wave time window, t_{R0} and t_{R1}, as well as the criteria for the time-shifts $\Delta\tau_0(t)$, are derived from formulae in Komatitsch *et al.* (2004).

For the 6–30 s and 3–30 s data, we use constant values of $r_0(t) = r_0$, $\mathrm{CC}_0(t) = \mathrm{CC}_0$, $\Delta\tau_0(t) = \Delta\tau_0$, and $\Delta\ln A_0(t) = \Delta\ln A_0$. We exclude any arrivals before the P wave and after the Rayleigh wave. This is achieved by the box-car function for $w_E(t)$

$$w_E(t) = \begin{cases} 10 w_E & t < t_P, \\ w_E & t_P \leq t \leq t_{R1}, \\ 10 w_E & t > t_{R1}, \end{cases} \tag{A9}$$

For the 2–30 s data, we avoid selecting surface-wave arrivals as the 3-D model used to calculate the synthetics cannot produce the required complexity. The water-level criteria then becomes:

$$w_E(t) = \begin{cases} 10 w_E & t < t_P, \\ w_E & t_P \leq t \leq t_S, \\ 10 w_E & t > t_S. \end{cases} \tag{A10}$$