

Automatic slowness vector measurements of seismic arrivals with uncertainty estimates using bootstrap sampling, array methods and unsupervised learning

J. Ward¹,^{*} M. Thorne,² A. Nowacki¹ and S. Rost¹

¹*School of Earth and Environment, University of Leeds, Woodhouse, Leeds LS2 9JT, United Kingdom. E-mail: eejwa@leeds.ac.uk*

²*Department of Geology and Geophysics, University of Utah, 115 S 1460 E, Salt Lake City, UT 84112, USA*

Accepted 2021 May 13. Received 2021 May 10; in original form 2020 November 20

SUMMARY

Horizontal slowness vector measurements using array techniques have been used to analyse many Earth phenomena from lower mantle heterogeneity to meteorological event location. While providing observations essential for studying much of the Earth, slowness vector analysis is limited by the necessary and subjective visual inspection of observations. Furthermore, it is challenging to determine the uncertainties caused by limitations of array processing such as array geometry, local structure, noise and their effect on slowness vector measurements. To address these issues, we present a method to automatically identify seismic arrivals and measure their slowness vector properties with uncertainty bounds. We do this by bootstrap sampling waveforms, therefore also creating random sub arrays, then use linear beamforming to measure the coherent power at a range of slowness vectors. For each bootstrap sample, we take the top N peaks from each power distribution as the slowness vectors of possible arrivals. The slowness vectors of all bootstrap samples are gathered and the clustering algorithm DBSCAN (Density-Based Spatial Clustering of Applications with Noise) is used to identify arrivals as clusters of slowness vectors. The mean of slowness vectors in each cluster gives the slowness vector measurement for that arrival and the distribution of slowness vectors in each cluster gives the uncertainty estimate. We tuned the parameters of DBSCAN using a data set of 2489 SKS and SKKS observations at a range of frequency bands from 0.1 to 1 Hz. We then present examples at higher frequencies (0.5–2.0 Hz) than the tuning data set, identifying PKP precursors, and lower frequency by identifying multipathing in surface waves (0.04–0.06 Hz). While we use a linear beamforming process, this method can be implemented with any beamforming process such as cross correlation beamforming or phase weighted stacking. This method allows for much larger data sets to be analysed without visual inspection of data. Phenomena such as multipathing, reflections or scattering can be identified automatically in body or surface waves and their properties analysed with uncertainties.

Key words: Structure of the Earth; Body waves; Surface waves and free oscillations.

1 INTRODUCTION

Seismic array techniques which measure the full horizontal slowness vector (backazimuth and inclination) of seismic arrivals have been used to investigate Earth structure for decades. These analyses have been applied to a wide variety of seismic arrivals and problems such as by using long period surface waves to identify upper mantle and surface heterogeneity (Ji *et al.* 2005; Maupin 2011; Xia *et al.* 2018), short period S waves to analyse lower mantle structure (Cottaar & Romanowicz 2012; Schumacher & Thomas 2016; Stockmann *et al.* 2019; Ward *et al.* 2020), high-frequency P waves to study scatterers in the mid and lower mantle

(Niu & Kawakatsu 1997; Thomas *et al.* 2002; Cao & Romanowicz 2007; Frost *et al.* 2013; Bentham & Rost 2014; Yang & He 2015; Ritsema *et al.* 2020), event detection and spatial location (Chevrot *et al.* 2007; Landès *et al.* 2010; Liu *et al.* 2016), ambient noise (Behr *et al.* 2013; Roux & Ben-Zion 2017), nuclear event detection (Bowers & Selby 2009; Gibbons & Ringdal 2011) and meteorological event spatial location (Gerstoft *et al.* 2006, 2008).

Past studies which analysed slowness vector properties using array methods (for a review see: Rost & Thomas 2002, 2009) were limited in terms of number of observations due to the usual requirement to visually inspect each observation to determine an arrivals slowness vector properties or if it is too noisy to use. In addition,

several studies have discussed the limitations of using beamforming or f - k methodology to identify phases and estimate their slowness vector properties (Berteussen 1976; Gibbons *et al.* 2008; Selby 2011) and methods have been developed to correct slowness vector measurements for Earth structure when locating events (Bondár *et al.* 1999; Koch & Kradolfer 1999; Gibbons *et al.* 2011; Schweitzer 2001). To clarify what limitations the uncertainty estimate is accounting for, we first discuss the assumptions and limitations of making one slowness vector measurement.

Each slowness vector in the beamforming grid search assumes the wave moves over the array with a constant horizontal slowness and arrives at the stations with a backazimuth equal to that along the great circle path from the relocated event location (for details see: Ward *et al.* 2020). The beamforming process does not account for limitations in heterogeneous station distribution, which can lead to heterogeneous sampling of the wavefield, and interference from noise may contribute to errors in slowness vector measurement. The waveforms of the arrivals are assumed to be coherent across over the array. Waveform incoherence of the signal across the array may result in deviations from the slowness vector prediction (Gibbons *et al.* 2008). Source complexity could lead to unusual waveforms recorded at the array, but should not affect the slowness vector measurement as source complexity should introduce consistent waveform complexity across the array. Local structure may deform the wavefield as it moves across the array such that the arrival times at the stations will deviate from the prediction (Gibbons *et al.* 2018). This may lead to slowness vector deviations depending on geometry of stations distribution and local velocity and topography structure.

The predicted backazimuth of the arrival is assumed along the great circle path between the event and the mean station location assuming a spherical Earth. The predicted horizontal slowness of the arrival is taken from ray tracing through a 1-D velocity model in a spherical Earth. Any structures local to the array or deeper with properties which differ from the 1-D velocity model may result in deviations from this prediction. It is difficult if not impossible to separate out these different contributions using just a single array measurement, let alone determine their relative contributions.

Automating the identification of arrivals and measuring their slowness vector properties would remove the time consuming and subjective process of visually inspecting each observation and could allow for larger data sets to be analysed. Estimating the uncertainty of these measurements allow for better interpretation of the observations, and the ability to rigorously accept or reject scientific hypotheses on Earth structure or its processes.

Previous efforts have been made in automating standard seismic processing techniques such as shear wave splitting (Teanby *et al.* 2004) and $H - \kappa$ stacking (Ogden *et al.* 2019). Methods also exist to estimate uncertainties in the beamforming methodology (Lin & Roecker 1996; Bear & Pavlis 1997; Ritsema *et al.* 2020) and to improve the detection of one or multiple arrivals (Gal *et al.* 2014, 2016; Schmidt 1986). The method we propose differs from these by automatically identifying the number of arrivals with their slowness vector properties and uncertainties. To our knowledge, no method has been proposed that does all of these at once. The method we present later uses a linear relative beamforming process; however, this method can be applied with other techniques such as phase weighted stacking (Schimmel & Paulssen 1997) or cross correlation beamforming (Ruigrok *et al.* 2017).

Machine learning methodologies are becoming more prevalent in the geosciences (for a review see: Bower *et al.* 2013) and seismology (for a review see: Kong *et al.* 2019) with methods used to

automate data selection (e.g. Valentine & Woodhouse 2010; Thorne *et al.* 2020) and extracting properties from data by mapping seismograms to lower dimensional space using autoencoders (Valentine & Trampert 2012) or sequence seismograms and identify features such as the presence of seismic scatterers (Kim *et al.* 2020). Here we use an unsupervised learning algorithm as part of our automation technique.

In the approach we present in this paper, we create subsets of waveforms using bootstrap sampling (Efron 1992). For each sample, beamforming (Rost & Thomas 2002) corrected for a curved wavefront (Ward *et al.* 2020) is used to search over a range of slowness vectors and recover the slowness vectors of potential seismic arrivals. The slowness vector measurements of all the individual bootstrap samples are collected and we use the DBSCAN (Density-Based Spatial Clustering of Applications with Noise) algorithm (Ester *et al.* 1996) to identify clusters of slowness vectors as seismic arrivals. DBSCAN is an unsupervised learning algorithm which uses the density of points to classify them as part of a cluster or as noise. For further details, see Section 2.

By bootstrapping the traces, and therefore creating random subsets of the stations in the array, the scatter of the measurements in each cluster can give an estimate of the combination of some of the previously mentioned uncertainties. The uncertainty estimate will account for the following phenomena which cause different subsets of stations to have different slowness vector measurements:

- (i) Signal aberration where the arrival time of the wave at stations deviates from the prediction due to local array structure.
- (ii) Incoherent or coherent noise.
- (iii) The horizontal slowness of the wave changing as it moves over the array, due to the size of the array, or unaccounted for velocity variations within the array.
- (iv) Heterogeneous distribution of the stations causing heterogeneous sampling of the wavefield.
- (v) Slowness resolution limitations of the array aperture.
- (vi) Wavelet shape changing over the array.

All of these can relate to local structure or effects within the array and the uncertainty estimate describes the combination of all these effects on the wavefield. If a measured slowness vector deviates from the 1-D earth model prediction and is not within the uncertainty estimate, then the cause of this deviation must be external to the array and local structure. Determining the cause of these deviations to structures such as a dipping Moho, or deeper structure requires additional information and might be resolvable through, for example forward modelling. We do not try to measure the uncertainties of that aspect, only those listed above.

We tune the parameters of the DBSCAN algorithm on a visually inspected data set where each observation is labeled as having either 0, 1 or 2 arrivals. More arrivals are possible, but in this data set the maximum number confidently observed is 2. In this data set, observations with more than one arrival are hypothesized to be caused by multipathing, one of many phenomena which can cause multiple arrivals. Multipathing occurs when the wavefront is incident of a sufficiently large velocity gradient causing different parts of the wavefield to move at different velocities, diffract and refract. Multipathing results in 2 arrivals arriving at the station at different times and different slowness vector properties. The predictions made by the method are compared to the labels given from visual inspection to find the best parameters for the DBSCAN algorithm. Following this, we show the effectiveness of this automated method on finding the slowness properties of short-period PKP scattering and long-period surface wave arrivals. Guidance on using the method is

given in Section 5. We find the parameters work well for our example applications with a minor change needed for the surface wave example. Tuning the algorithm can be done for specific applications.

2 METHOD OVERVIEW

This section outlines the method to automatically measure the slowness vector properties with uncertainty estimates. The process can be roughly broken down into the following steps with more detail given below.

- (i) Create a number of bootstrap subsamples (1000 here) through random sampling with replacement of a set of waveforms recorded at the seismic array in question.
- (ii) For each bootstrap sample, use beamforming (Rost & Thomas 2002) correcting for a curved wavefront (Ward *et al.* 2020) to search over a grid of slowness vectors and find how the power of coherent energy varies with backazimuth and horizontal slowness. Therefore, each bootstrap sample will have its own grid of power values.
- (iii) Calculate a noise estimate for the bootstrap sample by shifting each trace in the bootstrap sample with a randomly generated time. These scrambled traces are then stacked and the power of the beam is measured. This is repeated 1000 times and the mean power is taken as the noise estimate.
- (iv) Set all power values in the slowness grid below the noise estimate to zero.
- (v) From the resultant power distribution, take up to X peaks (in this study we take up to three peaks), which describe the slowness vectors of possible arrivals.
- (vi) Gather the locations for these peaks of all the bootstrap samples.
- (vii) Use DBSCAN, a density-based clustering algorithm, to identify the arrivals and measure their slowness properties with uncertainties.

2.1 Bootstrapping and peak recovery

One advantage of the bootstrap sampling process is that bootstrap samples of the stations in the array are used. Beamforming subsets of the array leads to different peak power in the beams which leads to variations in the recovered slowness vectors for each arrival. When all of the slowness vectors are taken into account, using all of the bootstrap sampled arrays, we obtain uncertainty estimates in the slowness vector. These uncertainty estimates will include the effect that array geometry and local structure has on the slowness vector measurements. For each bootstrap sample, we use a relative beamforming method where the traces are aligned on a target slowness before searching over the slowness vectors. After the beamforming, we calculate a noise estimate using the traces in the bootstrap sample with a similar method to Korenaga (2013). The traces are aligned using the slowness vector with the highest power. Then, they are randomly shifted in time, stacked and the power of the stack calculated. This is repeated 1000 times and the mean of all power estimates is used for the noise power estimate. All power values in the beamforming plot (Fig. 1) below three times this noise estimate are set to zero. Multiplying the estimate by three was determined by exploratory analysis and found to give the most satisfactory result. This can be changed depending on the application. To remove local power maxima, the power distribution is smoothed using a 2-D Gaussian filter. The 2-D Gaussian is formed by the product of two 1-D Gaussians. The standard deviation of the 1-D Gaussians is

equal to the grid spacing (0.05 s deg^{-1}), therefore will have a full width at half maximum of 0.12 s deg^{-1} .

The 2-D Gaussian acts as a point spread function and is convolved with the power plot to smooth it and remove local maxima. After this, the top X peaks are taken from the power distribution. The peaks are found with a maximum neighbourhood filter which identifies points with higher power values than those in the surrounding neighbourhood. Fig. 1 shows how the peaks are found for each bootstrap sample.

2.2 Identifying arrivals with cluster analysis

The peaks recovered for each bootstrap sample are then collected and the clustering algorithm DBSCAN (Ester *et al.* 1996) is used to find clusters. DBSCAN is an unsupervised learning algorithm which uses the density of points to identify clusters and noise. The algorithm takes a radius ϵ and a minimum number of points (MinPts) to define a minimum density for points to be a cluster. Here, we define MinPts as a fraction of the number of bootstrap samples. DBSCAN sorts the data into three categories as visualized in Fig. 2.

- (i) Core point: A point with at least MinPts points within its neighbourhood (i.e. within radius ϵ).
- (ii) Boundary point: A point within the neighbourhood of a core point, but without MinPts points in its own neighbourhood.
- (iii) Noise: Points that are not within ϵ of a core point and does not have MinPts points within its neighbourhood.

The DBSCAN algorithm begins at a random point and measures its density by the number of points within the radius ϵ (Fig. 2). If the density is lower than the threshold defined by ϵ and MinPts, the point is classified as noise (yellow points in Fig. 2) and the algorithm moves on to another random point. If the density is higher than the defined threshold, the point is classified as a core point and cluster formation begins (red points in Fig. 2). Points within ϵ of the core points then have the number of points in their neighbourhood measured. Those which do not have MinPts points within their neighbourhood are boundary points and are still part of the cluster (blue points in Fig. 2). The points which do have MinPts points in their neighbourhood are classified as core points and added to the cluster. The points within ϵ of these new core points are also searched and the cluster expands until it finds no new core points to add to the cluster. Once no new core points can be added, an unexamined point is chosen at random and the process begins again. This process continues until all points have been examined. In this manner, DBSCAN can separate high density clusters from low density noise. Fig. 3 shows the result of DBSCAN applied to the peaks recovered after the bootstrapping process.

DBSCAN has advantages over other clustering algorithms such as k -means (MacQueen *et al.* 1967) for this application such as:

- (i) It does not take the number of clusters as input so visual inspection before the clustering is not required.
- (ii) Not all points need to be part of a cluster allowing for noise.
- (iii) If clusters are not well separated or the data is noisy, clusters of non-hyperspherical shape can still be recovered unlike k -means (Ertöz *et al.* 2003; Celebi *et al.* 2013).

There are also disadvantages to DBSCAN:

- (i) If the range and data is not well understood, choosing the parameters can be challenging.




Figure 1. Example of recovery of peaks from a bootstrap sample of traces. The left figure shows a record section of data from the 05 April 1999 event recorded at the Kaapvaal array in Southern Africa (event metadata in the supplementary material). The traces are coloured by the number of times they have been sampled. The data had the instrument response removed and are filtered between 0.10 and 0.40 Hz before beamforming. The right figure shows the power distribution at each slowness vector with powers lower than the noise estimate set to zero and the 2-D Gaussian smoothing filter applied. Here each point on the grid represents a slowness vector described with their x (p_x) and y (p_y) components. In this example, two peaks have been recovered.




Figure 2. Cartoon illustrating what classifies as a core point, boundary point or noise. The neighbourhoods of the points are shown as a lighter colour of the point itself. The minimum number of points needed for a core point is 4 in this example. The red points all have at least 4 points in their neighbourhood, so are defined as core points. The blue points are within the neighbourhood of the core (red) points, but do not have 4 points in their own neighbourhood and are classified as boundary points. The yellow points are classified as noise because they are not in the neighbourhood of a core point and do not have 4 points within their own neighbourhood.

(ii) Clustering data with large variations in density is challenging because there may be no combination of ϵ and MinPts which will find all of the clusters.

(iii) Clusters separated by a distance smaller than ϵ will be combined into one cluster.

We tested other density-based clustering algorithms such as HDBSCAN (Campello *et al.* 2013, 2015) and OPTICS (Ankerst *et al.* 1999) but found that both techniques have issues for this

application. HDBSCAN (Hierarchical DBSCAN) searches over a range of ϵ values and measures over what length scales a cluster ‘persists’ while containing a minimum number of points to form a cluster. Using how long each cluster survives and how many points it contains at each ϵ , clusters are extracted with the excess of mass algorithm (EOM, McInnes & Healy 2017). HDBSCAN will preferentially return a large, single cluster because one large cluster will usually contain more ‘mass’ (for a detailed explanation, see McInnes & Healy 2017). To avoid one large cluster being returned when multiple clusters exist, HDBSCAN by default will not return a single cluster as an output. If this default is kept, instances with one arrival (cluster) will be misidentified. Changing the default and allowing HDBSCAN to return one cluster will mean phenomena causing multiple arrivals (such as multipathing) may not be identified as EOM will preferentially return a single cluster.

OPTICS (Ordering Points To Identify the Clustering Structure, Ankerst *et al.* 1999) is another density-based algorithm which specializes in identifying clusters of varying density. OPTICS orders the points to represent the clustering structure. From this, clusters can be extracted. When using OPTICS, we found the size of the clusters retrieved was too inconsistent to estimate the uncertainties of slowness vector properties. Because of these considerations, we decide to use DBSCAN instead of OPTICS or HDBSCAN.

2.3 Slowness vector uncertainty estimates

We estimate the uncertainty with the standard deviation of back-azimuths and horizontal slownesses in each cluster and also use the area of error ellipse of the clusters as a relative measure of uncertainty of each observation. The error ellipses are found by calculating the eigenvectors and eigenvalues of the covariance matrix for each cluster. These eigenvectors and eigenvalues give the directions and magnitudes of the maximum variances in the cluster which is used to determine the width, length and orientation of the ellipse. Fig. 4 shows clusters plotted with their error ellipses for 1, 2 and 3 standard deviations. We would like to highlight the importance of the slowness grid dimensions; if the slowness grid is too




Figure 3. Cluster retrieval from points recovered through bootstrap sampling the traces (Fig. 1). The left figure shows all the power peaks (blue dots) recovered using data from the 05 April 1999 event. The right image shows the clusters found by the DBSCAN algorithm (Ester *et al.* 1996) where MinPts is 0.25 and ϵ is 0.2 s deg⁻¹. The red and yellow points are classified as clusters 1 and 2 respectively and the black points are noise. The background power distribution is the mean of all the power distributions found from bootstrap sampling.




Figure 4. Example of error ellipses for 1, 2 and 3 standard deviations. The data are the same as used in Figs 1 and 3. The background power plot is the mean of the power plots searching over a range of slowness vectors from each bootstrap sample.

small, the arrivals may be truncated at the edge leading to a smaller cluster and underestimate the uncertainty.

3 PARAMETER TUNING

To find the best parameters to use with the DBSCAN algorithm (ϵ and MinPts), we compare the number of arrivals predicted by the algorithm to the number of arrivals identified from visual inspection. We use the same data set as Ward *et al.* (2020) which used SKS and SKKS data recorded at the Kaapvaal array in southern Africa. Ward *et al.* (2020) make observations at a range of frequency bands (Table 1) using the whole Kaapvaal array and several subarrays. The traces are first aligned on the predicted slowness of SKS or SKKS depending on the arrival of interest. The beamforming is conducted in a time window that is 20s before and 40s after the predicted arrival.

Table 1. The number of labels in each frequency band. Labels indicate the number of arrivals in that observation and 1–2 could be either 1 or 2. In total, there are 2628 labels with 2489 used in the tuning.

Frequency (Hz)	Number of arrivals			
	1	2	1–2	0
0.07–0.28	403	18	10	7
0.107–0.40	378	21	20	19
0.137–0.52	326	33	25	54
0.157–0.60	308	28	23	73
0.187–0.72	280	27	27	104
0.207–0.80	253	35	28	122
Total	1948	162	133	379

The data set provides a good test for the algorithm since it has clear single arrivals, multipathed arrivals (2 arrivals) and observations that are too noisy to identify any arrivals (0 arrivals). Each observation is labeled from visual inspection of the distribution and density of the points collected from all the bootstrap samples and the mean power distribution of all the bootstrap samples. If the algorithm predicts a higher number of arrivals than the human given labels, we assume here the algorithm has identified noise as arrivals. If the algorithm predicts a lower number of arrivals, the density threshold is too high for arrivals to be identified. Due to the subjective nature of the labelling this may not always be the case, but for the tuning process we assume the human labels are a ground truth. Observations where it was not clear whether there is one or two arrivals are labeled as ‘1–2 arrivals’ and excluded from this tuning process.

We searched over a range of ϵ and MinPts values and predict the number of arrivals in each observation. This is compared to the human labels in Table 1 and an accuracy score is calculated. The accuracy score is defined as the number of instances where the method correctly predicts the number of arrivals relative to the total number of instances ($\frac{\text{No. correct predictions}}{\text{Total instances}}$). Values of ϵ range from 0.05 to 1.0 s deg⁻¹ and MinPts is given as a fraction of the bootstrap samples (1000 here) and varies from 0.05 to 1.0. Fig. 5 shows how the accuracy varies in the parameter space. The grid search shows the sensitivity of our method to the DBSCAN parameters chosen.




Figure 5. Grid search of DBSCAN parameters ϵ and MinPts (given as a fraction of bootstrap samples). For each combination, the number of arrivals in each observation are predicted, compared to the true labels (Table 1) and the accuracy calculated. The location of the highest accuracy value is plotted as a red cross where $\epsilon = 0.20 \text{ s deg}^{-1}$ and MinPts = 0.25.

With some parameters, the accuracy can exceed 90 per cent while with others it can be less than 20 per cent. The method performs the worst with small ϵ and high MinPts meaning the minimum density criteria will be very high and very few arrivals will be found.

We test how well the algorithm generalizes using cross validation. Cross validation involves splitting the data set into N representative subsets (5 here). One of the subsets is removed and the grid search is conducted on the remaining $N - 1$ subsets and the best set of parameters recorded. The removed subset acts as a validation set. Then we take these best parameters and make predictions on the validation set. The accuracy of the predictions for the validation subset is measured and gives an indication of how well the algorithm generalizes. The process is repeated by sequentially removing one subset and tuning the parameters on the remaining $N - 1$ subsets. After the cross validation process, there are N estimates indicating how well the algorithm performs on unseen data. Here we split the data into five subsets because of the low number of multipathed (2 arrivals) and 0 arrivals samples. Cross validation and measuring the accuracy gave a mean accuracy of 0.939 with a standard deviation of 0.0090. In all the cross validation samples, the best parameters were $\epsilon = 0.20 \text{ s deg}^{-1}$ and MinPts = 0.25.

As there are many more instances of observations with one arrival, we also analyse each of the target labels (0, 1 or 2 arrivals) individually using the precision, recall and F_1 measures (defined below). These measures all depend on the number of true positive (TP), true negative (TN), false positive (FP) and false negative (FN) instances. These are best understood with an example. If the target label is '2', true positives are instances where the algorithm correctly identifies 2 arrivals in an observation. True negatives are instances correctly identified as not having 2 arrivals (1 or 0 arrivals). False positives are those incorrectly identified as having 2 arrivals. False negatives are instances where 2 arrivals have not been identified when they should have been.

From these measures, the precision is defined by $P = \frac{TP}{TP+FP}$. This is essentially the proportion of the target labels which have been correctly identified. The recall, $R = \frac{TP}{TP+FN}$, is a measure of how many of the target labels has been recovered by the algorithm. The F_1 score is the harmonic mean of the precision and recall and can be described as $F_1 = \frac{2}{\frac{1}{P} + \frac{1}{R}}$. The F_1 score is only large if both the recall

and precision are high. We only present the F_1 score as it shows which parameters have both high precision and recall. Fig. 6 shows how the F_1 score varies with different parameter combinations for each target label.

Figs 5 and 6 show that the method is capable of greater than 90 per cent agreement with the observations of a human. This is mainly from observations with one clear arrival, which makes up the majority of the observations. The algorithm also performs well with more complex observations of multipathing with a F_1 score of over 0.75. This method is quite insensitive to noise as it does not regularly incorrectly identify noisy observations as shown by a F_1 score of over 0.85 for observations with 0 arrivals. As with the accuracy, we use cross validation to see how well the parameters generalize with new data. Table 2 shows the mean F_1 scores for the individual labels. As in Fig. 5, there are DBSCAN parameters which perform very poorly showing the importance of the parameters used.

The cross validation analysis of all the labels and F_1 score on the individual labels show the parameters $\epsilon = 0.20 \text{ s deg}^{-1}$ and MinPts = 0.25 are consistently found to be the best. Inferring how well the parameters generalize from this analysis is limited because of the low number of cross validation samples (5 here). The low sample number was necessary because of the small number of observations with 2 and 0 arrivals. Despite this, the mean values obtained for the accuracy score and F_1 scores from the cross validation are very similar to that obtained by tuning with all the data (Figs 5 and 6). The standard deviations from the cross validation are low suggesting similar performance on similar data sets.

Due to the subjective nature of labelling each observation with the number of arrivals, some difference between the method's prediction and the human labels is acceptable. To analyse how reasonable the predictions are when the technique disagrees with the human labels, we create a confusion matrix using the predictions with parameters of $\epsilon = 0.20 \text{ s deg}^{-1}$ and MinPts = 0.25 (Fig. 7). In the confusion matrix, each row represents a true label (number of arrivals in this case) and each column the predicted arrivals. The values at each point in the matrix indicates how many times that true label is identified as the corresponding predicted labels. For example, for all instances with the true label of 1 arrival, the confusion matrix will show how many are correctly classified as having one arrival and how many are incorrectly identified with 0, 2 or 3 arrivals. We normalize the values along each row of the confusion matrix so for each true label, the columns show the proportion of the predictions given to that label. For example, for the instances with a true label of '0 arrivals', 80 per cent of the predictions are correctly identified as having 0 arrivals, 18 per cent are identified as having 1 arrival and so on.

The confusion matrix shows that when the method prediction differs from the human labels, the predictions it makes are not radically unreasonable. It is worth remembering the labeling process is quite subjective and just because the algorithm predicts a different number of arrivals to that labeled by a human, does not mean it is wrong. It is possible that some of the human labels with two arrivals only have one arrival or some have three arrivals. Equally, it is possible some instances labeled with no arrivals do have one arrival but a human could not confidently identify it above the noise. Fig. 7 shows the algorithm makes reasonable predictions in the vast majority of the cases for this data set using the parameters found from the tuning process and cross validation. Analysis of the uncertainty estimates show the slowness vector measurements have small variation with the mean standard deviation for backazimuth measurements of 1.2° and horizontal slowness of 0.14 s deg^{-1} . The mean area bounded by the 95 per cent confidence ellipse is $0.14 \text{ s}^2 \text{ deg}^{-2}$.




Figure 6. F_1 scores for combinations of DBSCAN parameters ϵ and MinPts where each plot represents a different target labels of 0 arrivals (left) one arrival (centre) and two arrivals (right). The location of the highest F_1 score is plotted as a red cross, which has parameters of $\epsilon = 0.20 \text{ s deg}^{-1}$ and MinPts = 0.25 for 1 and 2 arrivals and $\epsilon = 0.25 \text{ s deg}^{-1}$ and MinPts = 0.35.

Table 2. Table of the cross-validation result for each of the labels (0,1 or 2 arrivals) where the F_1 score is the measure of success. Note the standard deviation is an order of magnitude higher for labels 0 and 2, most likely because of the significantly fewer instances of those labels in the subsets created during cross-validation.

No. arrivals	Mean F_1 score	Standard deviation	Best parameters
0	0.86	0.030	$\epsilon = 0.35 \text{ s deg}^{-1}$ MinPts = 0.25
1	0.97	0.0063	$\epsilon = 0.20 \text{ s deg}^{-1}$ MinPts = 0.25
2	0.78	0.035	$\epsilon = 0.20 \text{ s deg}^{-1}$ MinPts = 0.25




Figure 7. Confusion matrix for predictions made with $\epsilon = 0.20 \text{ s deg}^{-1}$ and MinPts = 0.25. Each row represents a true label (number of arrivals) and each column the predicted arrivals. The values on the diagonal of the matrix show the percentage of correct predictions for the true label.

Analysis of the confusion matrix in addition to the findings from the cross validation process shows the parameters $\epsilon = 0.20 \text{ s deg}^{-1}$ and MinPts = 0.25 will give reasonable results that will generalize well. We use this parameters in other applications with a minor change for applications to surface waves (Section 4).

4 APPLICATIONS TO PKP SCATTERING AND RAYLEIGH WAVE MULTIPATHING

This section provides two example applications of this method to study Earth structure. First, we show an example identifying a PKP precursor in the high frequency teleseismic wavefield (0.5–2 Hz). Coherent precursors are indicative of scattering caused by small scale structures and our method can constrain uncertainties on their location. Then, we show an example of low frequency (0.04–0.06 Hz) Rayleigh wave multipathing. Using our method to identify Rayleigh wave multipathing, we can interpret possible causes of multipathing and provide uncertainties for phase velocity measurements. All measurements of backazimuth and horizontal slowness are shown with one standard deviation describing the uncertainties.

4.1 PKP precursors

Analysing the slowness vectors of PKP precursors is indicative of their location and whether they are caused by source or receiver side structure (Haddon & Cleary 1974). We use PKP data from Thomas *et al.* (1999) who observe several scatterers beneath Europe and Eastern Asia. Of the data used in Thomas *et al.* (1999), we focus on a single event occurring on 15 September, 1992 which shows clear PKP precursors. We only use data recorded at the Gräfenberg array and not the larger GRSN array to avoid spatial aliasing. In this example, the PKP precursors appear to be coherent from visual inspection of the seismograms (Fig. 8). Coherent precursors suggest they probably originate from localized scatterers such as an Ultra Low Velocity Zone (ULVZ, Ma & Thomas 2020).

Fig. 8 shows the traces used for this example and the clusters found by our algorithm. The data have the instrument response removed and are filtered between 0.5 and 2 Hz before the beam-forming process. We used a time window of 10 s before the predicted PKIKP arrival and the same DBSCAN parameters found from the tuning ($\epsilon = 0.20 \text{ s deg}^{-1}$ and MinPts = 0.25). The method identifies a single precursor arriving with a backazimuth of $58.6^\circ \pm 2.3^\circ$ and a horizontal slowness of $2.93 \pm 0.32 \text{ s deg}^{-1}$. This is similar to the slowness vector properties of the dominant arrival found by Thomas *et al.* (1999) arriving 6.5 s before PKIKP with a horizontal slowness of 2.8 s deg^{-1} and backazimuth of 53.6° . Unlike Thomas *et al.* (1999), we only identify one precursor rather than three. We believe this is because our time window encompasses all precursors meaning if one precursor has a significantly higher amplitude it may be the only one recovered. Furthermore, visual inspection of




Figure 8. Example application of the method on PKP precursors. This example uses data from the 15 September 1992 event recorded at the Gräfenberg array in Germany (GR) filtered between 0.5 and 2.0 Hz. The left subfigure shows the traces used in the example which are aligned on the predicted PKIKP arrival time and the time window for the analysis shown in red. On the right, the result of the algorithm with parameters of $\epsilon = 0.2 \text{ s deg}^{-1}$ and $\text{MinPts} = 0.25$.

waveforms suggests a single dominant precursor (Fig. 8). The range of possible horizontal slowness of this PKP precursor inferred from the uncertainty of the measurement ($2.93 \pm 0.32 \text{ s deg}^{-1}$) at a distance of approximately 140° means this precursor could originate from either source side or receiver side structure (Haddon & Cleary 1974).

4.2 Rayleigh wave multipathing

The second example shows the identification of multipathed Rayleigh waves. From this observation, the phase velocities and backazimuths of the multipathed arrivals can be measured and analysed with uncertainty bounds. Xia *et al.* (2018) identify multipathing in Rayleigh waves in the western US and suggest this is caused by the transition from continental to coastal to oceanic structure each with unique velocity profiles. We analyse Rayleigh waves from an event on 05 January 2013 recorded at the Southern California Seismic Array (CI) to identify multipathing and hypothesize some potential causes. The instrument response is removed and traces are filtered between 0.04 and 0.06 Hz. The time window used in the relative beamforming is 200 s before and after the predicted arrival time assuming a velocity of 3.5 km s^{-1} . In this example, the points in each cluster are distributed over a different slowness-space scale that is an order of magnitude lower than in the body wave examples. The difference is due to the Rayleigh wave velocity and the change in units. p_x/p_y for body waves will vary on the order of 10^0 , whereas for Rayleigh waves p_x/p_y vary on the order of 10^{-1} , an order of magnitude lower. Because of this, the ϵ parameter is also lowered by an order of magnitude from 0.20 s deg^{-1} found from tuning to 0.02 s km^{-1} .

Fig. 9 shows the result of the clustering method, which identifies three multipathed arrivals with backazimuths of $319 \pm 0.7^\circ$, $344 \pm 1.3^\circ$ and $299 \pm 1.4^\circ$ and velocities of 3.6 ± 0.025 , 3.5 ± 0.032 and $3.8 \pm 0.093 \text{ km s}^{-1}$, respectively. For each arrival, we mark the path from the mean station location along the mean backazimuth (dashed white line in Fig. 9) to determine a possible cause for the multipathing. Also shown are the paths showing the backazimuth uncertainty bounds (solid white lines in Fig. 9), which suggest it is reasonable to hypothesize possible causes of the measurements.

We investigate dispersion in the wave velocities by repeating the analysis in three frequency bands of 0.035–0.045, 0.045–0.055 and 0.055–0.065 Hz, finding differences in the number of arrivals and their backazimuths, but no absolute slowness variation between frequencies (Fig. S1). We argue this is a result of the different scale lengths of the structures which cause the observed multipathing, and not because of a property of the material the wave is travelling through.

The top and middle paths may come from interactions with the boundary between the continental and coastal regions, which agrees with the interpretation of Xia *et al.* (2018). The direction of the western most arrival suggests it could be caused by interacting with a coastal–ocean velocity transition or possibly due to more localized velocity variations. Further modelling is beyond the scope of this work, but our results demonstrate the potential of the method to investigate such phenomena in an efficient way.

The phase velocities of the arrivals may be indicative of azimuthal anisotropy beneath the array. The phase velocities of the central and eastern most arrival are the same within the uncertainties (3.6 ± 0.025 and $3.5 \pm 0.032 \text{ km s}^{-1}$, respectively). The western most arrival moves with a significantly higher phase velocity over the array ($3.8 \pm 0.093 \text{ km s}^{-1}$) along a backazimuth of $299^\circ \pm 1.4^\circ$. While we do not have enough measurements to fully explore the nature of this azimuthal anisotropy beneath the array, our observation of a faster arrival from 299° is in line with that found by Alvizuri & Tanimoto (2011) who report a fast direction of approximately 290° . Further analysis would be needed to recover the anisotropic properties, but this example shows how our technique can be used to identify statistically significant differences in phase velocity measurements.

5 CODE GUIDELINES

This section outlines some guidance to use this technique in terms of parameter selection and computation time. There are many potential aspects of a study that can influence the method's effectiveness such as frequency bands, array size and configuration or local receiver side structure. The tuning process (Section 3) shows we cover a range of frequency bands (Table 1) and array sizes (10–50 stations)




Figure 9. Example application of the method for identifying multipathing in surface waves. The left subfigure shows the ray paths (red lines) from the 05 January, 2013 event (white star) to the Southern California Seismic Array (CI) stations (green triangles). Before the beamforming, the data was filtered between 0.04 and 0.06 Hz. In this example, three arrivals have been identified by the algorithm (right subfigure). For each arrival, a path is marked from the mean station location along the mean backazimuth to a point with the same epicentral distance as the event (dashed white lines and circle). The solid white lines indicate the uncertainty bounds of the backazimuth for the measurement.

and the sub arrays have a wide range of configurations. For applications analysing body waves in similar frequency bands (0.1–1 Hz) with a similar array size (10–50), we recommend the parameters ($\text{MinPts} = 0.25$, $\epsilon = 0.2 \text{ s deg}^{-1}$) used here as a starting point and adjusted if necessary.

The number of peaks above the noise threshold should be equal to the maximum number of arrivals of interest or expect to be possible. The noise threshold was determined to be three times the noise estimate through exploratory analysis and found to give satisfactory results, but this can be changed depending on the application. DBSCAN parameters ϵ and MinPts of 0.20 and 0.25 s deg^{-1} , respectively, will work well for identifying single arrivals and is relatively intolerant to noise. If the study is searching for multipathing, changing MinPts to 0.15 and keeping ϵ as 0.20 s deg^{-1} increases the accuracy of the multipathed arrivals from 66 to 75 per cent but decreases the accuracy of the noisy arrivals from 80 to 44 per cent. These alternative parameters would require visual inspection of those identified as multipathing by the algorithm but would significantly reduce the amount of visual inspection as observations with one arrival need not be visually inspected.

For surface waves, the algorithm also works well after changing ϵ to 0.02 s km^{-1} . For applications with significantly different frequency bands or array size or searching for a very specific phenomenon, the DBSCAN parameters may need to be tuned to optimize performance (Section 3). The remaining parameters can be kept the same. Sensible beamforming practice such as avoiding spatial aliasing still applies when using this method.

The computationally intensive part of the method is the bootstrap sampling and the beamforming on each sample, which must be performed for each observation; the cluster analysis is comparatively quick. However, the code is trivially parallelizable over observations since each is independent of all the others. The code is written in Python, is easily editable and freely available (https://github.com/eejwa/Array_Seis_Circle). The code has been parallelized so the bootstrap sampling can be spread over several cores and uses Numba (Lam *et al.* 2015) to compile the functions into machine code before execution. Further improvements in efficiency could be made by rewriting the algorithm in more efficient

languages such as Julia, C++ or Fortran, and investigating further performance improvements possible with the existing code base. For an example array with 20 stations, a time window of 30 seconds, sampling rate of 0.05 s and searching over a grid of slowness vector properties with 14 641 vectors (a grid where each axis covers 6 s deg^{-1} in increments of 0.05 s deg^{-1}), each bootstrap sample takes approximately 1.6 s to process. This makes tens of observations viable on a handful of cores such as on a desktop machine. Larger data sets (thousands of observations) can be processed on the order of hours using hundreds of cores.

6 CONCLUSIONS

Slowness vector measurements have been used to understand a variety of Earth structures and phenomena. They are typically used to identify wavefield perturbations, scattering and event/noise source localization. While this analysis is a common tool used by seismologists, studies are limited because of the necessary and subjective visual inspection of observations. Interpretation of the measurements is limited by uncertainties such as the contribution of array geometry, noise and local structure. These may result in different slowness vector measurements depending on which stations are used in the analysis.

In this study, we described a method to automate slowness vector measurement, estimate the uncertainties and identify the number of possible arrivals. To do this, we bootstrap sample the waveforms and in each sample use a relative beamforming process to measure the coherent power and recover slowness vector properties of potential arrivals. These slowness vector properties are collected and the clustering algorithm DBSCAN is used to identify arrivals. The mean of the clusters gives the backazimuth and horizontal slowness and the spread of the cluster gives uncertainty estimates of phenomena which may vary the slowness vector measurement depending on which subset of stations are used. We use a linear beamforming approach but other beamforming methods such as phase weighted stacking (Schimmel & Paulssen 1997) and cross correlation beamforming (Ruigrok *et al.* 2017) can be used.

We tuned the DBSCAN parameters on a data set with 0, 1 and 2 arrivals and achieved >90 per cent accuracy in recovering these arrivals. We present examples of analysis of scattered *P*-wave energy and Rayleigh wave multipathing. The advantage this method brings to these applications is the ability to automatically identify the arrivals and measure the slowness vectors with uncertainty estimates. The difference in spatial scale and wavelengths used in these examples shows that our approach is applicable to studying Earth properties at a wide variety of spatial scales. Using this method, it may be possible to analyse slowness vector properties on larger data sets with reduced need for subjective visual inspection. In addition, uncertainties can also be quantified and used alongside the measurements. This technique makes 1000s of observations feasible in a matter of hours and allows for global-scale slowness vector observations to be made.

ACKNOWLEDGEMENTS

We thank reviewers Steve Gibbons and Yu Gu for their constructive reviews which has improved this manuscript. ObspyDMT (Hosseini & Sigloch 2017) was used to download data. GMT (Wessel et al. 2013) was used to make some of the figures. Predictions from 1-D velocity models were made using the TauP toolkit Crotwell et al. (1999). We thank the Department of Geology and Geophysics, University of Utah for hosting a collaborative visit and the Centre for High Performance Computing (CHPC) for computer resources and support. JW is supported by the Leeds-York Natural Environment Research Council Doctoral Training Partnership NERC DTP Spheres grant NE/L002574/1, AN was funded by NERC standard grant NE/R001154/1 (REMIS: Reliable Earthquake Magnitudes for Induced Seismicity) and MT was partially supported by NSF grant EAR-1723081.

DATA AVAILABILITY

Code to perform the analysis is available at: <https://github.com/ejwa/Array-Seis-Circle>. Data used for tuning and the examples is available to download from: <https://figshare.com/s/fbcb167ad15d581cfd4e>. Seismic arrays used were the Kaapvaal array (James et al. 2001), the Gräfenberg array (BGR SZO Data Center 1976) [<https://doi.org/10.25928/mbx6-hr74>] and the Southern California Seismic Network [<https://doi.org/10.7914/SN/CI>] (California Institute of Technology and United States Geological Survey Pasadena 1926).

REFERENCES

- Alvizuri, C. & Tanimoto, T., 2011. Azimuthal anisotropy from array analysis of Rayleigh waves in Southern California, *Geophys. J. Int.*, **186**(3), 1135–1151.
- Ankerst, M., Breunig, M.M., Kriegel, H.-P. & Sander, J., 1999. OPTICS: ordering points to identify the clustering structure, *ACM. Record*, **28**(2), 49–60.
- Bear, L.K. & Pavlis, G.L., 1997. Estimation of slowness vectors and their uncertainties using multi-wavelet seismic array processing, *Bull. seism. Soc. Am.*, **87**(3), 755–769.
- Behr, Y., Townend, J., Bowen, M., Carter, L., Gorman, R., Brooks, L. & Bannister, S., 2013. Source directionality of ambient seismic noise inferred from three-component beamforming, *J. geophys. Res.*, **118**(1), 240–248.
- Bentham, H. & Rost, S., 2014. Scattering beneath Western Pacific subduction zones: evidence for oceanic crust in the mid-mantle, *Geophys. J. Int.*, **197**(3), 1627–1641.
- Berteussen, K., 1976. The origin of slowness and azimuth anomalies at large arrays, *Bull. seism. Soc. Am.*, **66**(3), 719–741.
- BGR SZO Data Center. German Regional Seismic Network (GRSN), *Federal Institute for Geosciences and Natural Resources (BGR)*; 1976. <https://doi.org/10.25928/MBX6-HR74>.
- Bondár, I., North, R.G. & Beall, G., 1999. Teleseismic slowness-azimuth station corrections for the International Monitoring System seismic network, *Bull. seism. Soc. Am.*, **89**(4), 989–1003.
- Bower, D.J., Gurnis, M. & Seton, M., 2013. Lower mantle structure from paleogeographically constrained dynamic Earth models, *Geochem. Geophys. Geosyst.*, **14**(1), 44–63.
- Bowers, D. & Selby, N.D., 2009. Forensic seismology and the Comprehensive Nuclear-Test-Ban Treaty, *Ann. Rev. Earth planet. Sci.*, **37**, 209–236.
- California Institute of Technology and United States Geological Survey Pasadena, 1926. Southern California seismic network, *International Federation of Digital Seismograph Networks, Dataset/Seismic Network*.
- Campello, R.J., Moulavi, D. & Sander, J., 2013. Density-based clustering based on hierarchical density estimates, in *Pacific-Asia conference on knowledge discovery and data mining*, pp. 160–172, Springer.
- Campello, R.J., Moulavi, D., Zimek, A. & Sander, J., 2015. Hierarchical density estimates for data clustering, visualization, and outlier detection, *ACM Trans. Knowled. Discov. Data (TKDD)*, **10**(1), 1–51.
- Cao, A. & Romanowicz, B., 2007. Locating scatterers in the mantle using array analysis of PKP precursors from an earthquake doublet, *Earth planet. Sci. Lett.*, **255**(1–2), 22–31.
- Celebi, M.E., Kingravi, H.A. & Vela, P.A., 2013. A comparative study of efficient initialization methods for the k-means clustering algorithm, *Expert Syst. Appl.*, **40**(1), 200–210.
- Chevrot, S., Sylvander, M., Benahmed, S., Ponsolles, C., Lefevre, J. & Paradis, D., 2007. Source locations of secondary microseisms in western Europe: Evidence for both coastal and pelagic sources, *J. geophys. Res.*, **112**(B11), doi:10.1029/2007JB005059.
- Cottaar, S. & Romanowicz, B., 2012. An unusually large ULVZ at the base of the mantle near Hawaii, *Earth planet. Sci. Lett.*, **355**, 213–222.
- Crotwell, H.P., Owens, T.J. & Ritsema, J., 1999. The TauP Toolkit: flexible seismic travel-time and ray-path utilities, *Seismol. Res. Lett.*, **70**(2), 154–160.
- Efron, B., 1992. Bootstrap methods: another look at the jackknife, in *Breakthroughs in Statistics*, pp. 569–593, Springer.
- Ertöz, L., Steinbach, M. & Kumar, V., 2003. Finding clusters of different sizes, shapes, and densities in noisy, high dimensional data, in *Proceedings of the 2003 SIAM International Conference on Data Mining*, pp. 47–58, SIAM.
- Ester, M., Kriegel, H.-P., Sander, J., Xu, X., et al., 1996. A density-based algorithm for discovering clusters in large spatial databases with noise, in *KDD '96: Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, August 1996, pp. 226–231.
- Frost, D.A., Rost, S., Selby, N.D. & Stuart, G.W., 2013. Detection of a tall ridge at the core–mantle boundary from scattered PKP energy, *Geophys. J. Int.*, **195**(1), 558–574.
- Gal, M., Reading, A., Ellingsen, S., Koper, K., Gibbons, S. & Näsholm, S., 2014. Improved implementation of the fk and Capon methods for array analysis of seismic noise, *Geophys. J. Int.*, **198**(2), 1045–1054.
- Gal, M., Reading, A., Ellingsen, S., Koper, K., Burlacu, R. & Gibbons, S., 2016. Deconvolution enhanced direction of arrival estimation using one-and three-component seismic arrays applied to ocean induced microseisms, *Geophys. J. Int.*, **206**(1), 345–359.
- Gerstoft, P., Fehler, M.C. & Sabra, K.G., 2006. When Katrina hit California, *Geophys. Res. Lett.*, **33**(17), doi:10.1029/2006GL027270.
- Gerstoft, P., Shearer, P.M., Harmon, N. & Zhang, J., 2008. Global P, PP, and PKP wave microseisms observed from distant storms, *Geophys. Res. Lett.*, **35**(23), doi:10.1029/2008GL036111.
- Gibbons, S.J. & Ringdal, F., 2011. Seismic monitoring of the North Korea nuclear test site using a multichannel correlation detector, *IEEE Trans. Geosci. Remote Sens.*, **50**(5), 1897–1909.
- Gibbons, S.J., Ringdal, F. & Kverna, T., 2008. Detection and characterization of seismic phases using continuous spectral estimation on incoherent and partially coherent arrays, *Geophys. J. Int.*, **172**(1), 405–421.

- Gibbons, S.J., Schweitzer, J., Ringdal, F., Kväerna, T., Mykkeltveit, S. & Paulsen, B., 2011. Improvements to seismic monitoring of the European Arctic using three-component array processing at SPITS, *Bull. seism. Soc. Am.*, **101**(6), 2737–2754.
- Gibbons, S.J., Näsholm, S., Ruigrok, E. & Kväerna, T., 2018. Improving slowness estimate stability and visualization using limited sensor pair correlation on seismic arrays, *Geophys. J. Int.*, **213**(1), 447–460.
- Haddon, R. & Cleary, J., 1974. Evidence for scattering of seismic PKP waves near the mantle-core boundary, *Phys. Earth planet. Inter.*, **8**(3), 211–234.
- Hosseini, K. & Sigloch, K., 2017. obspyDMT: a Python toolbox for retrieving and processing of large seismological datasets, *Solid Earth*, **8**, 1047–1070.
- James, D., Fouch, M., VanDecar, J., Van Der Lee, S. & Group, K.S., 2001. Tectospheric structure beneath southern Africa, *Geophys. Res. Lett.*, **28**(13), 2485–2488.
- Ji, C., Tsuboi, S., Komatschik, D. & Tromp, J., 2005. Rayleigh-wave multipathing along the west coast of North America, *Bull. seism. Soc. Am.*, **95**(6), 2115–2124.
- Kim, D., Lekić, V., Ménard, B., Baron, D. & Taghizadeh-Popp, M., 2020. Sequencing seismograms: a panoptic view of scattering in the core-mantle boundary region, *Science*, **368**(6496), 1223–1228.
- Koch, K. & Kradolfer, U., 1999. Determination of mislocation vectors to evaluate bias at GSETT-3 primary stations, *J. Seismol.*, **3**(2), 139–151.
- Kong, Q., Trugman, D.T., Ross, Z.E., Bianco, M.J., Meade, B.J. & Gerstoft, P., 2019. Machine learning in seismology: turning data into insights, *Seismol. Res. Lett.*, **90**(1), 3–14.
- Korenaga, J., 2013. Stacking with dual bootstrap resampling, *Geophys. J. Int.*, **195**(3), 2023–2036.
- Lam, S.K., Pitrou, A. & Seibert, S., 2015. Numba: a LLVM-based Python JIT compiler, in *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, pp. 1–6.
- Landès, M., Hubans, F., Shapiro, N.M., Paul, A. & Campillo, M., 2010. Origin of deep ocean microseisms by using teleseismic body waves, *J. geophys. Res.*, **115**(B5), doi:10.1029/2009JB006918.
- Lin, C.-H. & Roecker, S., 1996. P-wave backazimuth anomalies observed by a small-aperture seismic array at Pinyon Flat, southern California: implications for structure and source location, *Bull. seism. Soc. Am.*, **86**(2), 470–476.
- Liu, Q., *et al.*, 2016. Source locations of teleseismic P, SV, and SH waves observed in microseisms recorded by a large aperture seismic array in China, *Earth planet. Sci. Lett.*, **449**, 39–47.
- Ma, X. & Thomas, C., 2020. Small-scale scattering heterogeneities in the lowermost mantle from a global analysis of PKP precursors, *J. geophys. Res.*, **125**(3), e2019JB018736.
- MacQueen, J. *et al.*, 1967. Some methods for classification and analysis of multivariate observations, in *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, Vol. 1, pp. 281–297, Oakland, CA, USA.
- Maupin, V., 2011. Upper-mantle structure in southern Norway from beam-forming of Rayleigh wave data presenting multipathing, *Geophys. J. Int.*, **185**(2), 985–1002.
- McInnes, L. & Healy, J., 2017. Accelerated hierarchical density based clustering, in *Proceedings of the 2017 IEEE International Conference on Data Mining Workshops (ICDMW)*, pp. 33–42, IEEE.
- Niu, F. & Kawakatsu, H., 1997. Depth variation of the mid-mantle seismic discontinuity, *Geophys. Res. Lett.*, **24**(4), 429–432.
- Ogden, C., Bastow, I., Gilligan, A. & Rondenay, S., 2019. A reappraisal of the H- κ stacking technique: implications for global crustal structure, *Geophys. J. Int.*, **219**(3), 1491–1513.
- Ritsema, J., Kaneshima, S. & Haugland, S.M., 2020. The dimensions of scatterers in the lower mantle using USArray recordings of S-wave to P-wave conversions, *Phys. Earth planet. Inter.*, **306**, 106541.
- Rost, S. & Thomas, C., 2002. Array seismology: methods and applications, *Rev. Geophys.*, **40**(3), 2-1–2-27.
- Rost, S. & Thomas, C., 2009. Improving seismic resolution through array processing techniques, *Surv. Geophys.*, **30**(4–5), 271–299.
- Roux, P. & Ben-Zion, Y., 2017. Rayleigh phase velocities in Southern California from beamforming short-duration ambient noise, *Geophys. J. Int.*, **211**(1), 450–454.
- Ruigrok, E., Gibbons, S. & Wapenaar, K., 2017. Cross-correlation beam-forming, *J. Seismol.*, **21**(3), 495–508.
- Schimmel, M. & Paulssen, H., 1997. Noise reduction and detection of weak, coherent signals through phase-weighted stacks, *Geophys. J. Int.*, **130**(2), 497–505.
- Schmidt, R., 1986. Multiple emitter location and signal parameter estimation, *IEEE Trans. Antenn. Propagat.*, **34**(3), 276–280.
- Schumacher, L. & Thomas, C., 2016. Detecting lower-mantle slabs beneath Asia and the Aleutians, *Geophys. J. Int.*, **205**(3), 1512–1524.
- Schweitzer, J., 2001. Slowness corrections—one way to improve IDC products, *Pure appl. Geophys.*, **158**(1), 375–396.
- Selby, N.D., 2011. Improved teleseismic signal detection at small-aperture arrays, *Bull. seism. Soc. Am.*, **101**(4), 1563–1575.
- Stockmann, F., Cobden, L., Deschamps, F., Fichtner, A. & Thomas, C., 2019. Investigating the seismic structure and visibility of dynamic plume models with seismic array methods, *Geophys. J. Int.*, **219**(Suppl. 1), S167–S194.
- Teanby, N., Kendall, J.-M. & Van der Baan, M., 2004. Automation of shear-wave splitting measurements using cluster analysis, *Bull. seism. Soc. Am.*, **94**(2), 453–463.
- Thomas, C., Weber, M., Wicks, C. & Scherbaum, F., 1999. Small scatterers in the lower mantle observed at German broadband arrays, *J. geophys. Res.*, **104**(B7), 15073–15088.
- Thomas, C., Kendall, J.-M. & Weber, M., 2002. The lowermost mantle beneath northern Asia—I. Multi-azimuth studies of a D heterogeneity, *Geophys. J. Int.*, **151**(1), 279–295.
- Thorne, M.S., Pachhai, S., Leng, K., Wicks, J.K. & Nissen-Meyer, T., 2020. New Candidate ultralow-velocity zone locations from highly anomalous SPdKS waveforms, *Minerals*, **10**(3), 211.
- Valentine, A.P. & Trampert, J., 2012. Data space reduction, quality assessment and searching of seismograms: autoencoder networks for waveform data, *Geophys. J. Int.*, **189**(2), 1183–1202.
- Valentine, A.P. & Woodhouse, J.H., 2010. Approaches to automated data selection for global seismic tomography, *Geophys. J. Int.*, **182**(2), 1001–1012.
- Ward, J., Nowacki, A. & Rost, S., 2020. Lateral velocity gradients in the African lower mantle inferred from slowness space observations of multipathing, *Geochem. Geophys. Geosyst.*, **21**(8), e2020GC009025.
- Wessel, P., Smith, W.H., Scharroo, R., Luis, J. & Wobbe, F., 2013. Generic mapping tools: improved version released, *EOS, Trans. Am. Geophys. Un.*, **94**(45), 409–410.
- Xia, Y., Ni, S. & Tape, C., 2018. Multipathing Rayleigh waves from long-distance noise cross correlation along an ocean-continent boundary (Alaska to California), *Geophys. Res. Lett.*, **45**(12), 6051–6060.
- Yang, Z. & He, X., 2015. Oceanic crust in the mid-mantle beneath west-central Pacific subduction zones: evidence from S to P converted waveforms, *Geophys. J. Int.*, **203**(1), 541–547.

SUPPORTING INFORMATION

Supplementary data are available at [GJI](https://doi.org/10.1093/gji/ggk117) online.

event_metadata.txt

Figure S1. Three results when using the automated method on Rayleigh wave data in three different frequency bands do investigate the effect of dispersion. Data is from the 05 January 2013 event recorded at the Southern California Seismic Array (CI).

Please note: Oxford University Press is not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the paper.