



HAL
open science

Mapping structures on the core–mantle boundary using Sdiff postcursors: Part I. Method and Validation

Carl Martin, Thomas Bodin, Sanne Cottaar

► To cite this version:

Carl Martin, Thomas Bodin, Sanne Cottaar. Mapping structures on the core–mantle boundary using Sdiff postcursors: Part I. Method and Validation. *Geophysical Journal International*, 2023, 235 (3), pp.2385 - 2398. 10.1093/gji/ggad340 . hal-04496308

HAL Id: hal-04496308

<https://hal.science/hal-04496308v1>

Submitted on 8 Mar 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Mapping structures on the core–mantle boundary using Sdiff postcursors: Part I. Method and Validation

Carl Martin ¹, Thomas Bodin ² and Sanne Cottaar ¹

¹*Bullard Laboratories, Department of Earth Sciences, University of Cambridge, CB3 0EZ 69622, UK. E-mail: cm828@cam.ac.uk*

²*Laboratoire de Géologie de Lyon, UMR 5276, Université de Lyon, Villeurbanne, France*

Accepted 2023 August 21. Received 2023 August 10; in original form 2022 November 10

SUMMARY

Ultra-low velocity zones (ULVZs) are patches of extremely slow seismic velocities on the core–mantle boundary (CMB). Here, we target them using the postcursors to S core-diffracted phases (Sdiff) caused by ULVZs. We use traveltimes of these postcursors to make probabilistic maps using a reversible-jump Markov chain Monte Carlo inversion setup. For the forward model, we extend 2-D wave front tracking (2DWT) software, previously developed for surface wave multipathing studies, to the CMB. The 2DWT is able to model the full multipathing behaviour of Sdiff postcursors and compute arrival times for a given ULVZ input velocity structure on the order of a few CPU seconds, as opposed to 100s of CPU hours required for 3-D full waveform synthetics. We validate the method using synthetic data sets produced by the 2DWT, as well as 3-D full waveform synthetics, using a parametrisation formed from a collection of ellipses. We also test idealistic data coverage versus a case of more realistic coverage. We show ULVZ size and velocity reduction can typically be well recovered, and our maps show the inherent trade-off between these parameters around the edge of the ULVZ. Our method cannot directly constrain ULVZ height; tests show that it underestimates ULVZ velocity reductions and overestimates ellipticity for thinner ULVZs due to neglecting mantle effects.

Key words: Composition and structure of the mantle; Body waves; Computational seismology; Wave scattering and diffraction.

1 INTRODUCTION

Broad-scale low velocity features known as large low shear velocity provinces (LLSVPs) dominate the degree-two structure of the lower mantle, beneath Africa and the Pacific, and have been the subject of numerous studies in recent years (e.g. Cottaar & Lekić 2016; Garnero *et al.* 2016). On the smaller scale, extreme features on the core–mantle boundary (CMB), known as ultra-low velocity zones (ULVZs), have been identified which show a significant reduction in V_s of 5–50 % and V_p 5–25 % (Yu & Garnero 2018). Understanding ULVZ velocity reductions and morphology is important to understand their composition, origin, and role in mantle dynamics.

High-resolution studies using various seismic phases (e.g. ScS, Jenkins *et al.* 2021; ScP, Pachhai *et al.* 2014; Sdiff, Cottaar & Romanowicz 2012; SPdKS, Thorne *et al.* 2013) have identified ULVZs between 5–40 km thick and 50–900 km wide. The largest of these have been identified near Samoa (Thorne *et al.* 2013), Hawaii (Cottaar & Romanowicz 2012), Iceland (Yuan & Romanowicz 2017) and the Galápagos Islands (Cottaar *et al.* 2022). Identifying and characterising these anomalies is difficult owing to poor data coverage (e.g. a lack of dense arrays of seismometers in suitable geome-

tries, and poor spatial distribution of earthquakes), as well as the computational expense of modelling the full wavefield up to the required frequencies to observe postcursors (~ 10 s period, Cottaar & Romanowicz 2012).

Whilst a number of phases can be used to identify ULVZs (Yu & Garnero 2018), this study uses S core-diffracted waves (Sdiff) which travel a significant distance around the CMB and have been observed to show extreme multipathing behaviour where a postcursor phase (Sdiff+) is observed ≥ 30 s later (Cottaar & Romanowicz 2012; Yuan & Romanowicz 2017). As Sdiff+ traveltimes depend nonlinearly on the ULVZ properties, conventional linearised tomography cannot be applied. Typically, ULVZs have been constrained using Sdiff+ through full waveform modelling (Cottaar & Romanowicz 2012; Yuan & Romanowicz 2017; J. Li *et al.* 2022, Z. Li *et al.* 2022; Cottaar *et al.* 2022), but the computational expense of this approach limits the ULVZ parameters that can be explored.

Whilst first arrivals are relatively easy to compute, later arrivals are more difficult to track due to triplication of the wave front or ‘swallow-tailing’ (Fig. 1). Hauser *et al.* (2008) developed software using a reduced phase-space Lagrangian formulation of the problem which allows later arrivals to be tracked in a computationally




Figure 1. Wave fronts (grey lines) propagating from an earthquake source (star) in Papua New Guinea to receivers (triangles) in North America and Central America. Triplication, referred to as a ‘swallow-tail’, occurs as the wave front passes through a circular slow velocity region near Hawaii. The first (green) and second (orange) ray path arrivals pass around the anomaly, with the faster arrival taking the more direct route; the third (blue) ray path arrival passes through the slow velocity region, and is delayed as a result. This third arrival is what we refer to as the Sdiff+. The relative amplitudes of the arrivals will depend on the velocity structure and geometry of source and receivers. Modelled after Hauser (2007).




Figure 2. A ray diagram showing the S (blue) and Sdiff (black) from an event at 500 km depth to stations at 80°, 100° and 120° epicentral distance. The red lines show an estimate of the region of CMB Sdiff sensitivity. The traveltimes, τ , is calculated by decomposing into and integrating across the horizontal, p , and vertical, η , components of the slowness. By taking measurements with respect to a reference model, vertical traveltimes effects can be removed to first order. As a result, we effectively ‘project’ the sources and receivers (grey) onto the CMB (white) and model the physics with a 2DWT. We then make the assumption that variations in the traveltimes are then caused by a region of the velocity structure where the Sdiff ray paths are sensitive to the CMB (dashed red).




Figure 3. A summary of allowed proposal changes for a starting model with one node. Properties of ellipses that can be changed are: rotation, size and velocity, ellipticity and position; or, ellipses can be added (‘birth’) or deleted (‘death’). Each of these proposed changes occur one at a time, except for size and velocity due to the intrinsic trade-off between these properties (Fig. S1, Supporting Information). Due to the nature of strong nonlinear effects of multipathing, multiple ellipses are only used for sensitivity testing; that is, determining regions where Sdiff+ ray paths have poor sensitivity to velocity variations.




Figure 4. Wave fronts are defined as a string of nodes (circles) and interpolated between to form a smooth wave front (solid line). As the wave front progresses (dashed lines), the distance between neighbouring nodes can increase or decrease. This is the gsf, that is, the node density of the wave front with respect to the initial density. If nodes get too close together (red circle), they are removed in the next wave front; if they are too far apart, a node is added (green circle). When two successive wave fronts, t and $t + \Delta t$, pass over a receiver (triangle), the arrival time and other properties are calculated by interpolating between the enclosing nodes on each wave front.

efficient manner. Using the core-diffracted waves (Sdiff and Sdiff+) we are able to project sources and their respective receiver stations to the CMB to track wave fronts as they propagate around a spherical shell for a given velocity structure (Fig. 2). These calculated arrival times can be adjusted for the up- and downgoing legs by comparison to 1-D PREM (Preliminary Reference Earth Model) arrival times (Dziewonski & Anderson 1981). Other effects such as mantle structure and ellipticity corrections are more difficult to adjust for since the ray path required for such correction calculations are dependent on the ULVZ 3-D morphology; these effects are on the order of a few seconds (Russell *et al.* 2022), compared to the 10s of seconds observed in Sdiff+, so are not considered further.

Such a setup could allow us to invert for ULVZ morphology since the forward model can be computed on the order of a few CPU seconds instead of 100–300 CPU hours required for full waveform synthetics up to frequencies required for Sdiff+ (0.1 Hz). However,




Figure 5. A simplified schematic of the rj-McMC Bayesian inversion. In one loop, the forward model is run for a model and the Metropolis–Hastings algorithm for model acceptance compares the acceptance probability α (a function of the movement between an initial model m and proposed model m' , where one parameter has been changed from m) and p (a random number between 0 and 1). The next iteration is then updated according to this value. A sparsely sampled ensemble of models is then taken to be representative of the underlying posterior distribution, from which we can make statistical inferences.

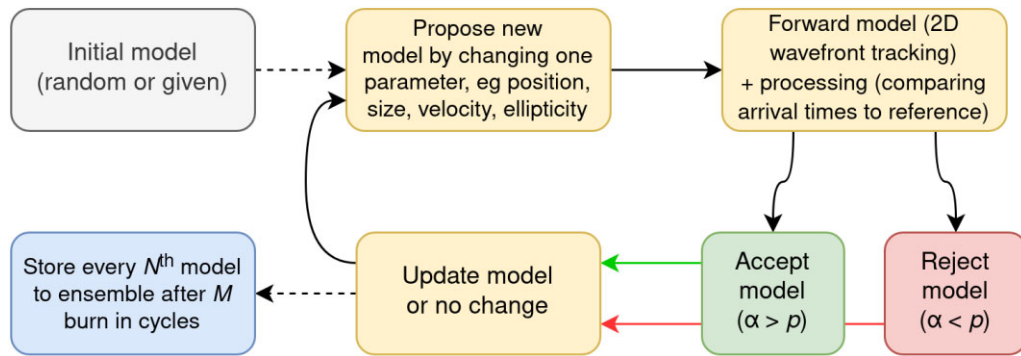


Figure 6. Left: gsf of wave fronts for the ULVZ model in Fig. 7, normalised by the geometric spreading of spherical curvature for the uniform background velocity model. The ULVZ results in lens focussing (red, lower gsf) in the wave fronts immediately after the anomaly; and then the wave front triplications spread out with time (green, higher gsf). The wave front at the back of the triplication has a lower gsf than that of the intermediate wave fronts; this is analogous to having a higher energy density, and therefore a higher amplitude. Right: normalised gsf calculated for the arrivals at each of the receivers. The main Sdiff arrival is always taken to be the first arrival at each station, regardless of gsf. Only one postcursor is present in the data at the 10–20 s period range, so the postcursor is chosen as the arrival with the lowest gsf after the main arrival. This is (usually) a continuous postcursor, though an allowance in gsf value is factored in to account for numerical approximation and discretisation effects.

Table 1. Earthquakes used in this study. All earthquake parameters are from the Global Centroid-Moment-Tensor project (Ekström *et al.* 2012). # is number of receivers with a positive postcursor detection. Source–receiver pairs have been extracted for positive postcursor observations, and then synthetically modelled. Ray paths of the postcursor coverage for these events are plotted in Fig. 7.

Event	Date	Lon	Lat	Depth	Mag	#	Region
1	2005/08/08	140.23	−3.73	17.0	5.9	96	Irian Jaya, Indonesia
2	2006/03/14	127.31	−3.35	13.0	6.7	100	Seram, Indonesia
3	2008/08/30	147.49	−6.33	88.25	6.4	477	Eastern New Guinea Region, PNG
4	2015/01/23	168.36	−17.06	231.0	6.8	203	Vanuatu Islands
5	2016/09/08	158.47	−54.51	13.7	6.1	206	Macquarie Island Region




Figure 7. An equal area projection map centred on the preferred ULVZ model proposed by Z. Li et al. (2022). Five events and their Sdiff ray paths close to the CMB are shown in green (Event 1), purple (Event 2), red (Event 3), blue (Event 4), and orange (Event 5), representing the widest range of source–ULVZ–receiver geometries presently in the literature (Table 1). In the background, an idealised array of sources and receivers is generated for five events 50° from the ULVZ centre (each with 216 stations each, 50° – 70° from the ULVZ with 20° azimuthal spread, corresponding to approximately the same number of source–receiver pairs as the real coverage; i.e. 1080 idealised vs 1082 real). The orange dots are surface hotspots, with annotations for guidance: H = Hawaii; C = Caroline; S = Samoa; E = Easter and G = Galapagos. Each concentric ring represents 10° epicentral distance from the centre of the ULVZ model.

the simplified 2DWT approach means we cannot account for finite-frequency effects, particularly those caused by the finite height of the ULVZ or wave front healing.

We build on previous studies which have performed Bayesian inversions for vertical ULVZ structure (e.g. Thorne *et al.* 2020; Pachhai *et al.* 2022), and present a framework to determine the lateral velocity structure of a ULVZ. Here, we use the 2DWT software as a forward modelling approach within a Bayesian inversion framework, where an ellipse parametrisation is used. We demonstrate the benefits of our inversion scheme on two types of synthetic tests: (i) observed Sdiff+ traveltimes generated with the 2DWT; and (ii) observed traveltimes extracted from several 3-D full waveform synthetics. This approach allows us to characterise the trade-offs between size and shear velocity reduction, and therefore the resolvability of ULVZ morphology using Sdiff+ postcursors.

An application of this methodology to a real data set sampling the Hawaiian ULVZ is shown in the companion paper (Martin *et al.* 2023).

2 METHOD

2.1 Forward model and Bayesian inversion framework

2.1.1 Parametrisation

Beamforming of Sdiff+ arrivals from data from previous studies suggest that the mega-ULVZs are well-approximated by axisymmetric anomalies placed on the CMB (e.g. Cottaar & Romanowicz 2012; Yuan & Romanowicz 2017; Z. Li et al. (2022); Cottaar *et al.*

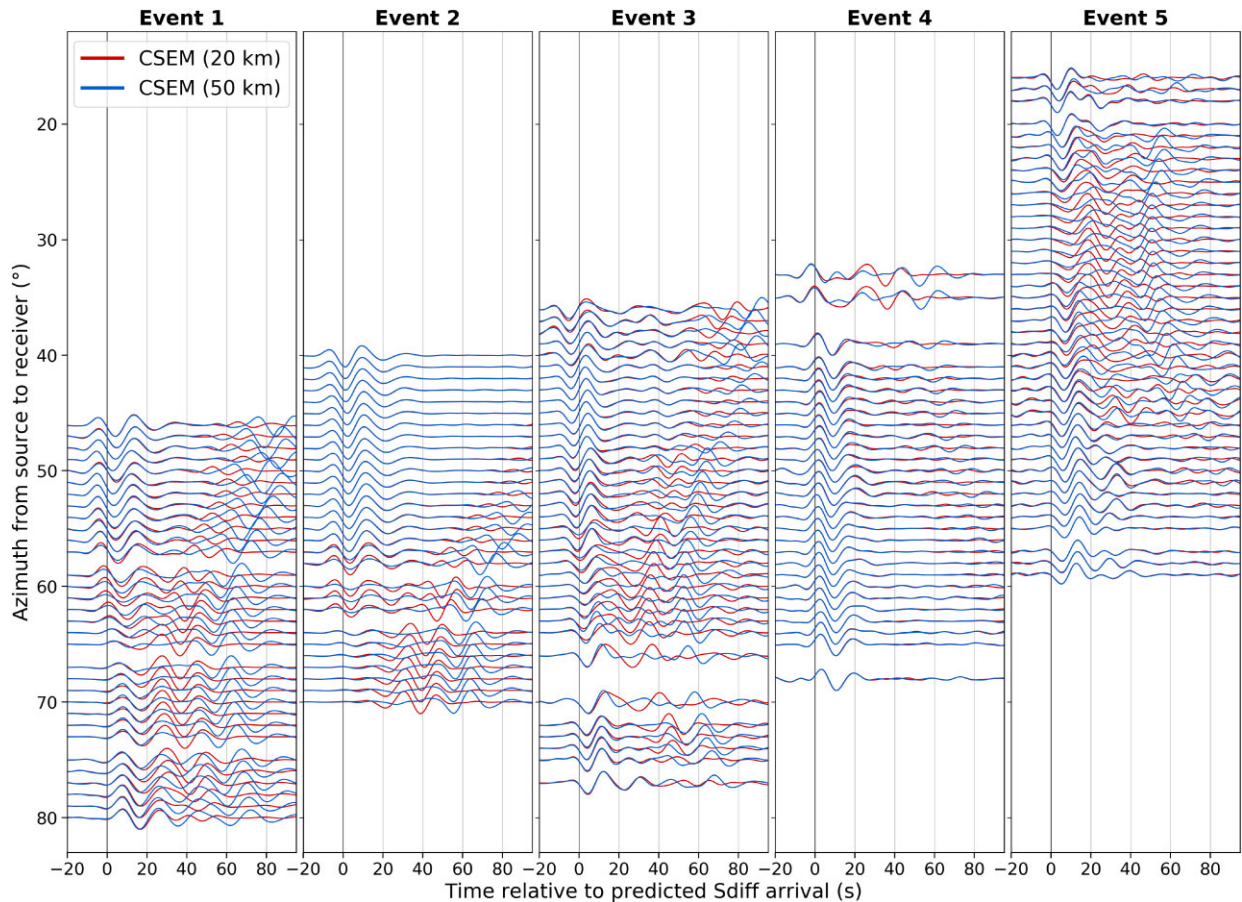


Figure 8. 3-D full waveform synthetics for the five events with 20 km (red) and 50 km (blue) thick ULVZs. Waveforms are sorted according to azimuth between source and receiver. Synthetics have been calculated using CSEM (Capdeville *et al.* 2002, 2003), and aligned with respect to predicted Sdiff arrival times from PREM (Dziewonski & Anderson 1981) calculated using the TauP toolkit (Crotwell *et al.* 1999). While the postcursor arrivals are significantly weaker in amplitude than the main arrivals in Event 4, they can be distinguished due to the SNR.

2022). However, studies in the Hawaiian ULVZ region using different phases find evidence for height-varying ULVZ landscape (e.g. Jenkins *et al.* 2021) or more complex lateral structures (e.g. J. Li *et al.* 2022). The mega-ULVZ beneath Samoa appears a lot more elongated (e.g. Thorne *et al.* 2020). Several studies have found that in 2-D and 3-D synthetic waveforms, data are not able to distinguish between boxcar or dome-like models (e.g. Vanacore *et al.* 2016; J. Li *et al.* 2022). Considering also the Fresnel half-zone width at 10 s period is 500 km (Z. Li *et al.* (2022)), the diffracted phases are only sensitive to the dominant broad-scale ULVZ structure.

Most parametrisation schemes for linearised tomography use Voronoi cells (e.g. Bodin & Sambridge 2009; Galetti *et al.* 2017) or its analogue Delaunay triangulation (e.g. Sambridge *et al.* 1995); a series of nuclei defined by a location and a velocity such that the velocity at each gridpoint is given by the nearest nucleus (Voronoi cell) or natural neighbours (Delaunay triangulation). Postcursors are nonlinearly generated by the velocity field, and vertices with strong velocity gradients in such cell-based parametrisation schemes result in spurious additional postcursor arrivals which are ‘unrealistic’ (i.e. not observed in the real data—see Martin *et al.* 2023). Furthermore, initial tests indicated that complex velocity models with large numbers of parameters resulted in strong non-uniqueness and/or non-convergence of the inversion.

We therefore seek a parametrisation which is (i) simple with few parameters, (ii) produces realistic Sdiff+ wave front arrivals and (iii) is able to reproduce the broad-scale structures of mega-ULVZs which have been previously reported. We settle on a level-set parametrisation where the velocity model is described as an ellipse of a given shear velocity, size and ellipticity. This ensures that every trial model is realistic, and similar to the previously modelled ULVZ morphologies. Fig. 3 demonstrates the ways in which parameters can be varied throughout the inversion cycle: changing the position, size and velocity, rotation, and ellipticity.

The prior distribution for most parameters is uniform within an allowed range: location is within 15° of the centre point (190° , 15°), velocity is 0.5–1.0 as a fraction of background velocity, semi-minor axis length is 50–800 km, and the rotation can be any value up to 2π . The eccentricity is chosen as the absolute value of a Gaussian distribution ($\mu = 0$ and $\sigma = 0.3$) where $e^2 < 0.75$. All subsequent proposals are then altered using a Gaussian distribution with a variance that is defined for each parameter change (position, velocity, radius, azimuth, eccentricity, and data noise). These Gaussian distribution parameter variances are originally input values, but change at intervals throughout the inversion to push the chain’s acceptance ratio towards 30–50 %.




Figure 9. Median (upper) and standard deviation (lower) velocity maps of the ensemble of models of single-ellipse inversions run for the idealistic (left) and realistic (right) receiver arrays defined in Fig. 7. The inner black dashed line represents the input ULVZ models, and the blue dashed line the defined area of sensitivity on the velocity map.




Figure 10. Median (upper) and standard deviation (lower) velocity maps of the ensemble of models of multiple-ellipse inversions run for the idealistic (left) and realistic (right) receiver arrays defined in Fig. 7; that is, with transdimensionality allowed. The inner black dashed line represents the input ULVZ models, and the blue dashed line the defined area of sensitivity on the velocity map.




Figure 11. Maps of the median velocity of the ensembles for inversions varying ULVZ radius and dVs. The realistic array from Fig. 7 is used. To the first order, the inversion produces results similar to the input models (dashed black) in terms of radius and dVs. The velocity probability density function for the midpoints of the ULVZ is shown in Fig. 13. The dVs cross-section of the dashed line AB is shown in Fig. 14.

2.1.2 The forward model

A wave front in a 2-D medium can be represented in a Lagrangian framework by a connected series of nodes, known as a bicharacteristic strip (Hauser 2007, Fig. 4). The nodes, which are then interpolated to reproduce the wave front, can then be advanced using local ray tracing from the eikonal equations in phase space (e.g. Cerveny 2001; Engquist *et al.* 2002):

$$\begin{aligned} \frac{d\mathbf{x}}{dt} &= c^2 \mathbf{p} \\ \frac{d\mathbf{p}}{dt} &= c \nabla \frac{1}{c} = -\frac{1}{c} \nabla c \end{aligned} \quad (1)$$

where $\mathbf{p} = \nabla c$ is the slowness vector, \mathbf{x} is the position vector in real space, c is the scalar velocity, and t is time.

It is also possible to represent these equations in a reduced phase space, where the bicharacteristic strip is defined by three components in 2-D space (Osher *et al.* 2002; Hauser 2007). This is done by combining the components of the slowness vector into the local

direction of the wave front, α , using

$$\tan \alpha = c_y / c_x \quad (2)$$

In a Cartesian plane, the equations of motion can then be written (e.g. Cerveny 2001; Osher *et al.* 2002; Hauser 2007):

$$\frac{d}{dt}(x, y, \alpha) = (c \cos \alpha, c \sin \alpha, c'_x \sin \alpha - c'_y \cos \alpha) \quad (3)$$

and can be readily extended to motion on a 2-D spherical shell by a coordinate system mapping:

$$\mathbf{p} = (p_x, p_y) = \frac{1}{c} (\cos \alpha, \sin \alpha) \rightarrow \mathbf{p} = (p_\alpha, p_\lambda) = \frac{1}{cR} \left(\frac{\cos \alpha}{\sin \lambda'}, \sin \alpha \right) \quad (4)$$

which gives the final equations of motion on a 2-D spherical shell (e.g. Julian 1970; Bukchin *et al.* 2006):

$$\frac{d}{dt}(x, y, \alpha) = \frac{1}{R} \left(\frac{c \cos \alpha}{\sin \lambda'}, c \sin \alpha, \left[\frac{c'_x \sin \alpha}{\sin \lambda'} - c'_y \cos \alpha + \frac{c \cos \alpha}{\tan \lambda'} \right] \right) \quad (5)$$

where $c(x, y)$ defines the wave speed, c'_x and c'_y are the derivatives in x and y , respectively, $-\pi \leq \alpha \leq \pi$ is the inclination angle of the




Figure 12. Maps of the standard deviation of velocity of the ensembles for inversions varying ULZ radius and dVs. The realistic array from Fig. 7 is used. The dVs cross-section of the dashed line AB is shown in Fig. 14.

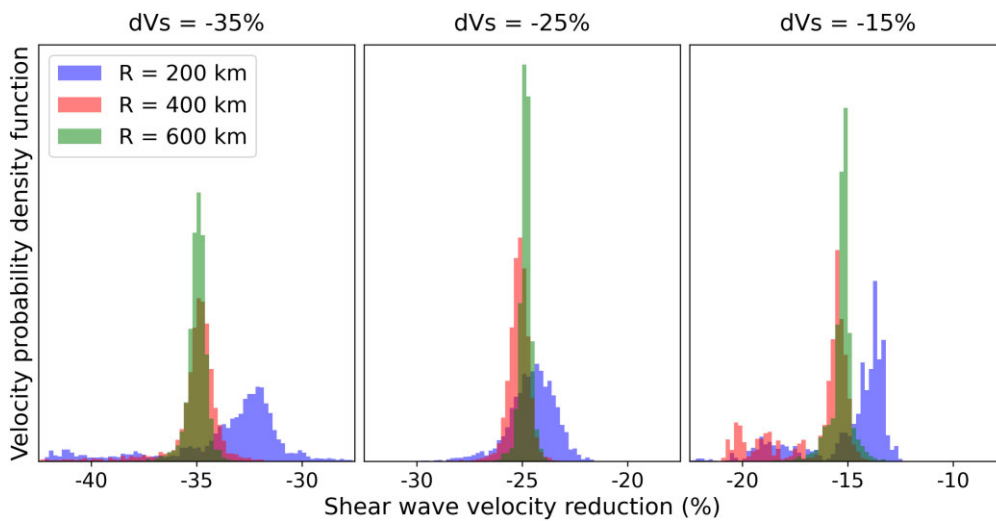


Figure 13. The probability density function of the dVs of the midpoints of the inversion ensembles from Fig. 11. Shown are dVs -35%, -25% and -15%; with radii of 200 (blue), 400 (red) and 600 km (green).




Figure 14. The probability density function of dVs of the cross-section of the ensemble of the -25% dVs , 400 km radius model inversion in Figs 11 and 12. The input ULVZ is in black (solid); and smoothed (dashed) according to the parametrization procedure. The mean (red) and median (blue) are overlaid.




Figure 15. The reference arrival times for 3-D full waveform CSEM synthetics (colours) and the 2DWT (greys). Arrivals have been extracted using the method outlined in Fig. S2, Supporting Information, for synthetics with ULVZ thicknesses: 20 km (reds), 30 km (oranges), 40 km (greens) and 50 km (blues). The arrival times are plotted against the effective azimuth, which is the angle between the azimuth of source to ULVZ midpoint, and the azimuth of ULVZ midpoint to receiver.

ray (x, y) , $\lambda'(y)$ is the colatitude, and R is the radius of the spherical shell.

By using the Lagrangian metric defined in reduced phase space, it is possible to track multipathing arrivals in a computationally efficient and accurate manner (Hauser 2007; Hauser *et al.* 2008). Whilst the scheme has been extended to account for reflections, refractions and triplications in layered media, for simplicity and reduced computational expense we consider only smoothly varying

media (i.e. no sharp discontinuities) which allows refraction and triplications.

We initiate a wave front as a point source at the location of the earthquake. The wave front is then propagated in time steps of Δt in a two-stage process (summarised in Fig. 4). First, all of the nodes are evolved using a fourth-order Runge–Kutta scheme using the kinematic ray tracing equations on a 2-D spherical shell (eq. 5). Second, nodes are added or removed to ensure uniform sampling in reduced phase space to prevent undersampling (insufficient nodes to capture triplication) and oversampling (unnecessary additional computation for the required accuracy, Hauser 2007). This sequence is then repeated until completed.

Application of this forward model to the CMB then takes the form of setting the spherical shell radius to that of the CMB ($R = 3481$ km), defining a scalar horizontal velocity field (taken with respect to the PREM horizontal velocity at the CMB, 7.2996 km s^{-1} ; Dziewonski & Anderson 1981), and projecting the sources and receivers down to the CMB. This process is outlined in Fig. 2. This forward model is referred to as the 2DWT.

2.1.3 Arrival calculations

Each wave front is stored as a series of nodes, with associated metadata, which represent each time step Δt . For each time step after that of the background velocity great circle path arrival time ($\Delta_{\text{path}}/v_{\text{PREM}}$), we check to see if any nodes on the wave front are within a critical distance of a receiver. If it is, the algorithm processes each of the nodes along the wave front to find two nodes forming a polygon with two other nodes on an adjacent wave front which minimally enclose the receiver location (Fig. 4).

2.1.4 Geometric spreading factor

Although various methods exist for calculating estimates of the amplitude of wave fronts (Osher *et al.* 2002; Hauser 2007), we choose to use the geometric spreading factor (gsf), which is a measure of closeness of neighbouring nodes on a wave front.




Figure 16. Median (upper) and standard deviation (lower) velocity maps of the ensemble of models of single-ellipse inversions run for the realistic receiver array defined in Fig. 7 with different input ULVZ thicknesses in the 3-D full waveform synthetics. The inner black dashed line represents the input ULVZ models (dVs -25% , radius 455 km), and the blue dashed line the defined area of sensitivity on the velocity map.

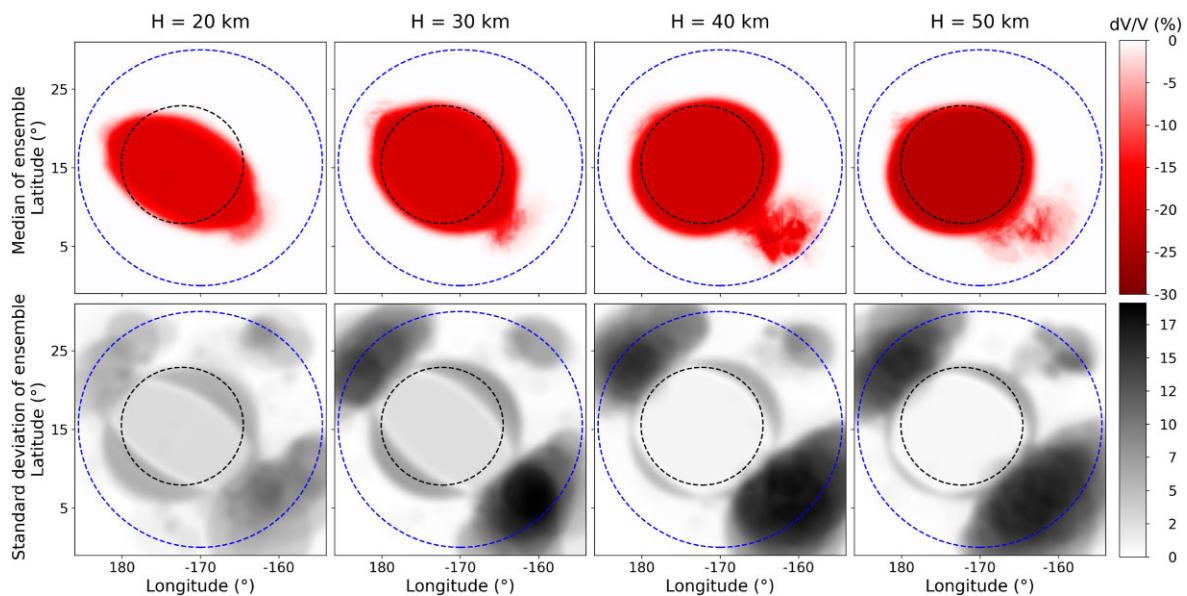


Figure 17. Median (upper) and standard deviation (lower) velocity maps of the ensemble of models of multiple-ellipse inversions run for the realistic receiver array defined in Fig. 7 with different input ULVZ thicknesses in the 3-D full waveform synthetics; that is, with transdimensionality allowed. The inner black dashed line represents the input ULVZ models (dVs -25% , radius 455 km), and the blue dashed line the defined area of sensitivity on the velocity map.

For each time step, the *gsf* is updated during wave front progression by taking into account the change in distance between neighbouring nodes, and also takes into account conditions when nodes are created and destroyed according to the wave front tracking formulation (Hauser 2007, Fig. 5). The spreading factor at the receiver is a weighted average of the values for the four nodes minimally enclosing the receiver location.

Since the *gsf* represents the cumulative curvature of the wave front at a given point, we expect a node with smaller total curvature

to correspond to a higher energy density and therefore a larger wave amplitude. As such, we expect to be able to use the *gsf* to aid in the identification of the appropriate postcursor arrival in the simulation. For example, in Fig. 1 we identify the third arrival as the postcursor phase, *Sdiff+*.

Traveltimes for models run in the 2DWT are extracted based on the lowest *gsf*. To take into account potential numerical error, we include an allowance of 20% from the lowest *gsf*. The minimum misfit contribution of this subset of traveltimes is then used in the misfit calculation.




Figure 18. The probability density function of the dVs of the midpoints of the inversion ensembles from Fig. 16, for CSEM ULVZ thicknesses: 20 km (red), 30 km (orange), 40 km (green) and 50 km (blue).

It is noted that this forward model does not capture the physics of wave front healing, so the Sdiff arrivals extracted from (real or 3-D full waveform synthetic) data cannot be compared with those calculated using the 2DWT. Therefore, only the Sdiff+ arrivals are inverted for in this scheme.

2.1.5 Bayesian inference

The probability of moving between states is given by Bayes' theorem, which states that the posterior distribution function is

$$\text{posterior} = \frac{\text{likelihood} \times \text{prior}}{\text{evidence}}, \quad p(\mathbf{m}|\mathbf{d}_{\text{obs}}) = \frac{p(\mathbf{d}_{\text{obs}}|\mathbf{m})p(\mathbf{m})}{p(\mathbf{d}_{\text{obs}})} \quad (6)$$

where $p(\mathbf{d}_{\text{obs}}|\mathbf{m})$ is the probability of \mathbf{d}_{obs} , the observed data, being observed given \mathbf{m} , a particular model. $p(\mathbf{m})$ is the 'a priori', which we treat as a constant within bounds, and $p(\mathbf{d}_{\text{obs}})$ is the 'evidence' (Bodin & Sambridge 2009; Bodin *et al.* 2009; Galetti *et al.* 2017).

The misfit value

$$\Phi(\mathbf{m}) = \left\| \frac{g(\mathbf{m}) - \mathbf{d}_{\text{obs}}}{\sigma_d} \right\|^2 \quad (7)$$

gives the squared differences between values predicted by a given model, $g(\mathbf{m})$, and the observed data, normalised with respect to an estimate of the standard deviation of the data noise, σ_d . The likelihood function, a measure of goodness of fit of the model to the data, is then given by

$$p(\mathbf{d}_{\text{obs}}|\mathbf{m}) \propto \exp(-\Phi(\mathbf{m})/2) \quad (8)$$

The form of the prior function, $p(\mathbf{m})$, depends on what changes are being proposed and the choice of parametrisation.

Further, we use a hierarchical Bayes method to the estimate of the standard deviation of the data noise, σ_d , which is also treated as an unknown in the inversion (Malinverno & Briggs 2004). In order to promote chain convergence, we perform simulated annealing on σ_d (Wéber 2000).

2.1.6 Markov chain Monte Carlo sampling

Where Bayesian inference is used to formulate the inverse problem, Markov chain Monte Carlo (MCMC) sampling is a tool amongst many to approximate the posterior distribution.

To extract a probabilistic velocity map of the CMB, we perform a Bayesian inversion using a reversible-jump MCMC (rj-MCMC)

algorithm (Fig. 6). Our implementation is heavily based on the work of Bodin & Sambridge (2009), Bodin *et al.* (2009), and a similar procedure carried out using surface waves by Galetti *et al.* (2017).

Starting from an initial state in a MCMC system in which an infinite number of states exist, we randomly move to another based on the probabilities associated with moving between these states. By the law of large numbers, once a sufficient number of moves has occurred such that the transient behaviour of the choice of initialisation is lost, the frequency of occupation of states is proportional to the posterior distribution of the system we are attempting to sample. That is, the properties of the system can be determined by the distribution of occupation of states over a large number of iterations. In our case, the system is the velocity structure of the CMB as sensed by Sdiff+ waves, and the states are velocity models proposed as a set of ellipses.

If we consider an updated proposal, \mathbf{m}' , to the model, \mathbf{m} , the acceptance probability is then given by

$$\alpha(\mathbf{m}'|\mathbf{m}) = \min \left[1, \frac{p(\mathbf{m}'|\mathbf{d}_{\text{obs}})}{p(\mathbf{m}|\mathbf{d}_{\text{obs}})} \frac{q(\mathbf{m}|\mathbf{m}')}{q(\mathbf{m}'|\mathbf{m})} \right] \quad (9)$$

where $q(\mathbf{m}'|\mathbf{m})$ is the probability that the proposal model is \mathbf{m}' given that the current model is \mathbf{m} . If a velocity value is being updated by picking a symmetrically distributed value centred on the previous position (e.g. uniform or Gaussian), the probability of $q(v'|v)$ is the same as its inverse $q(v|v')$, and the acceptance probability reduces to

$$\alpha(\mathbf{m}'|\mathbf{m}) = \min \left[1, \exp([\Phi(\mathbf{m}') - \Phi(\mathbf{m})]/2) \right] \quad (10)$$

We note that we have also dropped the prior ratio, so this is only true in the case where $p(\mathbf{m}) = p(\mathbf{m}')$.

After running the inversion for a long time, a sparsely sampled frequency distribution of states of the inversion is taken to be the posterior distribution for the system. From this, we extract a median velocity map of the inversion ensemble, and an estimate of the error by taking the standard deviation of the velocity at each gridpoint (Bodin *et al.* 2009).

Due to the range of possible complexities, we make simplifying assumptions about the nature of the interconnectedness of parametrisation. One obvious connection is the trade-off between velocity and size of the anomaly; that is, as the velocity becomes more negative, the size of anomaly required to fit the arrival times is smaller, and vice versa. The trade-off between the radius of a circular anomaly and the velocity reduction used in this study is demonstrated in Fig. S1a, (Supporting Information); we note that these are determined based on explicit forward modelling tests. By coupling these parameters, we expect to increase the exploration rate of the parameter space by moving parallel and tangentially to the preferred axis. We also test the coupling between longitude and latitudinal position in Fig. S1b, (Supporting Information) to adapt the form of the proposal to the form of the posterior. However, the biggest effect is from the source–receiver geometry rather than any underlying connection. As such, we do a random walk when varying the ellipse position; that is, there is no preferred change.

2.1.7 Transdimensionality

Additionally, transdimensionality can be introduced; that is, allowing the number of parameters to itself be a parameter determined by the inversion and the data (Bodin & Sambridge 2009) by adding ('birth') or deleting ('death') ellipses to the velocity map. In the

case where two ellipses overlap, the velocity at that point is taken to be the lowest of the two ellipses.

Where a model produces a better fit with fewer ellipses, the dimension of the model can be decreased; conversely, if a model improves the fit by increasing the number of ellipses, a threshold must be reached to prevent overdetermination. This preference for simple models is referred to as the ‘natural parsimony’ of the Bayesian formulation (e.g. Bodin *et al.* 2012).

To do this the acceptance probability is modified:

$$\begin{aligned} \alpha(\mathbf{m}'|\mathbf{m}) &= \min\left[1, \text{priorratio} \times \text{likelihoodratio} \times \text{proposalratio} \times |\mathbf{J}|\right] \\ &= \min\left[1, \frac{p(\mathbf{m}')}{p(\mathbf{m})} \frac{p(\mathbf{m}'|\mathbf{d}_{\text{obs}})}{p(\mathbf{m}|\mathbf{d}_{\text{obs}})} \frac{q(\mathbf{m}|\mathbf{m}')}{q(\mathbf{m}'|\mathbf{m})} |\mathbf{J}|\right] \end{aligned} \quad (11)$$

where the matrix \mathbf{J} is the Jacobian to account for the change in dimensionality between \mathbf{m}' and \mathbf{m} (Bodin & Sambridge 2009). The specific forms of the terms then depend on the choice of parametrisation and the type of change being attempted.

To take this into account, we use an algorithm where the prior is equal to the proposal, so the prior ratio and proposal ratio terms cancel (Mosegaard & Tarantola 1995). We therefore define a prior distribution on the number of ellipses that is $(1/k)^5$, where k is the number of ellipses.

Due to the complexities of arrivals in the case of multiple ellipses, synthetic tests show that more information is required in the observed data than one postcursor to resolve multiple ellipses. As such, since we only have (or can pick) one postcursor present in our data set, we use the transdimensional inversion only as a sensitivity test, to show the resolving limit of our data set, as well as the level of uncertainties in the recovered velocity maps. The sensitivity maps are then taken to be the equilibrium ensemble of the multiple-ellipse inversion chains initiated from the endpoints of the converged single-ellipse inversion chains.

2.2 Validation tests

We run a number of tests to validate the methodology. Initially, we construct a set of traveltimes for synthetic models using the 2DWT, and add 1.5 s Gaussian noise to approximate a realistic noise distribution (consistent with the ~ 1.7 s noise observed in real data, Martin *et al.* 2023). In order to test resolvability, we create ‘idealised’ sources and receivers; five equidistant events from the centre of the ULVZ input model, with a total of 1080 receivers organised in five 12×18 grids, respectively. This can be compared to a ‘realistic’ array, which represents five events in our observed data set for the Hawaiian ULVZ (Table 1; a full data set is available in Martin *et al.* 2023). Across the five events there are 1082 positive detections of the postcursor. Both the ‘idealistic’ and the ‘realistic’ arrays are shown in Fig. 7. Our reference ULVZ has a radius of 455 km and velocity reduction of 25 %, centred at $(187.7^\circ, 15.4^\circ)$ (Cottaar & Romanowicz 2012; Z. Li *et al.* 2022). We also explore synthetic models which range in size and velocity reduction.

Next, we use the realistic source–receiver array in 3-D full waveform synthetic modelling. The synthetics are calculated using the ‘sandwiched’ coupled spectral element method (CSEM), a computationally efficient method by coupling a spectral element solution in the lowermost 370 km of the mantle, to a normal mode solutions in the rest of the mantle and core (Capdeville *et al.* 2002, 2003).

Using the 3-D full waveform synthetics, we illustrate the limitations of our methods not accounting for full waveform, finite-frequency effects, and mantle effects. CSEM synthetics are run for

our reference ULVZ with heights of 20, 30, 40 and 50 km; waveforms for the 20 and 50 km models are shown in Fig. 8 for the five events. The 3-D full waveform synthetics are calculated with SEMUCB-WM1 in the model for the lowermost 370 km, which emulates mantle effects that are not captured by the 1-D PREM reference.

To pick the traveltimes in the full waveform synthetics, we use the same procedure as we apply to our observed data (see Martin *et al.* 2023). 1-D synthetics are calculated using PREM in Instaseis (van Driel *et al.* 2015). Data and 1-D synthetics are filtered between 10–20 s period. We correlate a snippet of the 1-D synthetics waveform around the predicted PREM Sdiff arrival with the data. By inspecting the positive sign correlations and plotting with respect to source–receiver azimuth, we are then able to identify the effective wave front arrival time of Sdiff+ in waveform data. This process is outlined in Fig. S2, Supporting Information.

3 RESULTS AND DISCUSSION

3.1 Tests using 2DWT synthetics

For the idealistic and realistic source–receiver arrays (Fig. 7), we run single-ellipse inversions for a reference circular ULVZ input model [Fig. 9; -25 % dVs, 455 km radius, at $(187.7^\circ, 15.4^\circ)$]. For both arrays there is excellent resolution of the input model, recovering the shape, size, and dVs (-25.0 ± 0.3 % and -24.5 ± 0.4 % for idealistic and realistic arrays, respectively). There is a slight asymmetry and smearing of the ULVZ in the SW–NE direction for the realistic array as a result of the source–receiver distribution, but the overall result is largely the same. In the standard deviation map of the ensemble, there is a clear ‘loop’ of uncertainty which is interpreted as a discontinuity (Galetti *et al.* 2015), that is, the uncertainty in the boundary of the ULVZ.

To estimate the sensitivity of the ray path coverage, we extend the inversion to multiple ellipses (Fig. 10). In this case, the inversion is started from the final models from the single-ellipse inversions and allowed to reach a new equilibrium before the ensemble is recompiled. For the idealistic array, there is excellent resolution of the input model in the median velocity map; and there are several small ‘shadow zones’ in the standard deviation map of the ensemble, which correspond to regions of low sensitivity due to the geometry of the source–receiver coverage. The shadow zones are asymmetric due to the asymmetry of the input ULVZ model with respect to the defined area of sensitivity (note that the idealistic array is symmetric around the ULVZ model). For the realistic array, we see two large shadow zones to the NW and SE of the ULVZ, and a smaller zone to the NE. This again corresponds to the geometry of sources and receivers, where there is a dominant SW–NE directionality. There is large uncertainty in these two shadow zones, indicating low sensitivity of the postcursor data to these regions.

To check that we are using an appropriate parametrisation and inversion scheme, we perform inversions for ULVZs for combinations of 200, 400, and 600 km radii, with velocity reductions of -15 %, -25 %, and -35 % (Figs 11 and 12). For the median velocity maps of the ensembles, we recover the location and size to a reasonable degree of accuracy. In all velocity standard deviation maps of the inversion ensembles we see a characteristic ring of uncertainty, indicating where the uncertainty in the ULVZ boundary lies. For the smallest velocity reduction of -15 %, there appears to

be some bimodality of the circular input model with an elliptical model tangential to the SW–NE direction of the source–receiver geometry. As the dVs becomes more negative, the thickness of this ring and the magnitude of standard deviation increases, indicating there is a stronger trade-off between dVs and size of the ULVZ.

When looking at the velocity probability density distributions of these inversion ensembles (Fig. 13) the 400 and 600 km inversions match the input velocities within error, and produce well-defined Gaussian distributions for the -35% and -25% ensembles. The -15% ensemble is less well-defined, with an asymmetric distribution tail towards more negative dVs . The 200 km ensembles are also not well approximated by Gaussian distributions. Taken together, this suggests $dVs -15\%$ and ~ 200 km radius represent the limit of ULVZ resolution with the 2DWT inversion scheme. We also expect that the *Sdiff+* for such small ULVZs will be difficult to observe and pick in real data.

Inspecting the posterior velocity profile in a cross-section shows there is good agreement in the dVs and ULVZ morphology (Fig. 14). The broader section of the uncertainty ring on the NE side of the ULVZ indicates the increased uncertainties due to the source–receiver geometry. The uncertainty in ULVZ dimensions are on the order of ~ 50 km. The trade-off between dVs and radius can also be seen here: the more negative dVs values have a smaller ULVZ width.

Further synthetic tests are shown in the Supporting Information, including the effect of smoothing on the velocity map (Fig. S5), the effect of inversion noise and simulated annealing (Fig. S6 and S7) and the effect of a linear shift on traveltimes on the inversion (Fig. S8 and S9).

3.2 Tests using 3-D full waveform synthetics

In the 3-D full waveform synthetics, ULVZ models with larger thicknesses produce more delayed *Sdiff+* but the *Sdiff* arrival is virtually unaffected by the ULVZ thickness (Fig. 8). The *Sdiff+* arrival times are plotted against the effective azimuth, which is the angle between the azimuth of source to ULVZ midpoint, and the azimuth of ULVZ midpoint to receiver (Fig. 15). If a ULVZ is axisymmetric, the postcursor arrivals collapse onto a single parabola as they do for the picked arrivals, with some scatter due to mantle effects.

Figs 16 and 17 show the inversion results for the single- and multiple-ellipse inversions for the CSEM reference data, respectively. In the single-ellipse inversion, the input model is relatively well recovered in all cases; however for the smaller thicknesses, some ellipticity is introduced which is less pronounced at higher thicknesses. This is likely as a result of the comparatively larger ratio of geometry-dependent 3-D mantle effects to the *Sdiff+* traveltimes. Notably, when the ULVZ is 10% larger than the input, the velocity reduction is underestimated by 40%. Overall, we observe that we decreasingly recover the velocity for thinner ULVZs (Fig. 18). This is due to the finite-frequency waves having increased sensitivity to the structure above the ULVZ. Our maps therefore represent the velocities observed for *Sdiff* waves within a certain period band (10–20 s).

The multiple-ellipse sensitivity inversions show a very similar characteristic shadow zone pattern to the 2DWT synthetic inversion results in Fig. 11; two large regions of high uncertainty to the NW and SE, and a smaller patch hidden behind the NE side of the ULVZ. This, again, is a characteristic resulting from the source–receiver

geometries. Note that the new patches in the median velocity map are associated with the much larger uncertainties of the shadow zones, and can therefore be ignored.

4 SUMMARY

We present a new method to map ULVZs on the CMB using *Sdiff+* phases. By projecting earthquake sources and receivers to the CMB, we are able to efficiently and rapidly model 2-D wave front physics. The methodology we present here can be readily applied to an *rj*-McMC Bayesian inversion, which allows us to model several toy examples, inverting for velocity structure: built with a 2DWT reference; and built with the CSEM.

Despite the significant differences between 2-D and 3-D synthetic modelling, there is good agreement between the inversion results from the 3-D full waveform synthetics to the 2DWT approximation. As such, we have demonstrated that this scheme could be readily extended to invert for real traveltime data.

SUPPORTING INFORMATION

Supplementary data are available at *GJI* online.

suppl_data

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.

ACKNOWLEDGMENTS

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 804071—ZoomDeep). This work was performed using resources provided by the Cambridge Service for Data Driven Discovery (CSD3) operated by the University of Cambridge Research Computing Service (www.csd3.cam.ac.uk), provided by Dell EMC and Intel using Tier-2 funding from the Engineering and Physical Sciences Research Council (capital grant EP/T022159/1), and DIRAC funding from the Science and Technology Facilities Council (www.dirac.ac.uk). CM thanks Nick Rawlinson for sharing the 2DWT software with us, and to Florian Millet and Ved Lekic for helpful discussions. We would like to thank the Editor, Ebru Bozdog, and our reviewers (Michael Thorne and two anonymous reviewers) for their comments which improved the quality of this manuscript. The authors declare that they have no competing interests.

DATA AND CODE AVAILABILITY

The facilities of IRIS Data Services (www.iris.edu), and specifically the IRIS Data Management Center, were used for access to waveforms and related metadata. IRIS Data Services are funded through the Seismological Facilities for the Advancement of Geoscience (SAGE) Award of the National Science Foundation under cooperative support agreement EAR-1851048. Data from the TA network were made freely available as part of the EarthScope USArray facility, operated by Incorporated Research Institutions for Seismology (IRIS) and supported by the National Science Foundation, under cooperative agreements EAR-1261681. Earthquake parameters were

used from the Global CMT Project (www.globalcmt.org). Codes used will be made available by CM or SC upon request.

REFERENCES

- Bodin, T. & Sambridge, M., 2009. Seismic tomography with the reversible jump algorithm, *Geophys. J. Int.*, **178**, 1411–1436.
- Bodin, T., Sambridge, M. & Gallagher, K., 2009. A self-parametrizing partition model approach to tomographic inverse problems, *Inverse Probl.*, **15**(5), doi:10.1088/0266-5611/25/5/055009.
- Bodin, T., Sambridge, M., Tkalčić, H., Arroucau, P., Gallagher, K. & Rawlinson, N., 2012. Transdimensional inversion of receiver functions and surface wave dispersion, *J. geophys. Res.: Solid Earth*, **117**(B2), doi:10.1029/2011JB008560.
- Bukchin, B., Yanovskaya, T., Montagner, J.-P., Mostinskiy, A. & Beucler, E., 2006. Surface wave focusing effects: numerical modeling and statistical observations, *Phys. Earth planet. Inter.*, **155**(3–4), 191–200.
- Capdeville, Y., Larmat, C., Vilotte, J. & Montagner, J., 2002. A new coupled spectral element and modal solution method for global seismology: a first application to the scattering induced by a plume-like anomaly, *Geophys. Res. Lett.*, **29**, 32, doi:10.1029/2001GL013747.
- Capdeville, Y., To, A. & Romanowicz, B., 2003. Coupling spectral elements and modes in a spherical Earth: an extension to the ‘sandwich’ case, *Geophys. J. Int.*, **154**(1), 44–57.
- Cerveny, V., 2001. *Seismic Ray Theory*, Cambridge University Press.
- Cottaar, S. & Lekić, V., 2016. Morphology of seismically slow lower-mantle structures, *Geophys. Suppl. Mon. Notices R. Astron. Soc.*, **207**(2), 1122–1136.
- Cottaar, S. & Romanowicz, B., 2012. An unusually large ulvz at the base of the mantle near hawaii, *Earth planet. Sci. Lett.*, **355–356**, 213–222.
- Cottaar, S., Martin, C., Li, Z. & Parai, R., 2022. The root to the Galápagos mantle plume on the core-mantle boundary, *Geophys. J. Int.*, **195**(2), 1184–1195.
- Crotwell, H., Owens, T. & Ritsema, J., 1999. The taup toolkit: Flexible seismic travel-time and ray-path utilities, *Seismol. Res. Lett.*, **70**, 154–160.
- Dziewonski, A. & Anderson, D., 1981. Preliminary reference earth model, *Phys. Earth planet. Inter.*, **25**, 297–356.
- Ekström, G., Nettles, M. & Dziewoński, A., 2012. The global cmt project 2004–2010: centroid-moment tensors for 13,017 earthquakes, *Phys. Earth planet. Inter.*, **200**, 1–9.
- Engquist, B., Runborg, O. & Tornberg, A.-K., 2002. High-frequency wave propagation by the segment projection method, *J. Comput. Phys.*, **178**(2), 373–390.
- Galetti, E., Curtis, A., Meles, G. A. & Baptie, B., 2015. Uncertainty loops in travel-time tomography from nonlinear wave physics, *Phys. Rev. Lett.*, **114**(14), 148501. doi: 10.1103/PhysRevLett.114.148501.
- Galetti, E., Curtis, A., Baptie, B., Jenkins, D. & Nicolson, H., 2017. Trans-dimensional love-wave tomography of the british isles and shear-velocity structure of the east irish sea basin from ambient-noise interferometry, *Geophys. J. Int.*, **208**(1), 36–58.
- Garnero, E., McNamara, A. & Shim, S.-H., 2016. Continent-sized anomalous zones with low seismic velocity at the base of earth’s mantle, *Nat. Geosci.*, **9**(7), 481–489.
- Hauser, J., 2007. *Multi arrival wavefront tracking and seismic imaging*, PhD thesis, The Australian National University.
- Hauser, J., Sambridge, M. & Rawlinson, N., 2008. Multiarrival wavefront tracking and its applications, *Geochem. Geophys. Geosyst.*, **9**, doi:10.1029/2008GC002069.
- Jenkins, J., Mousavi, S., Li, Z. & Cottaar, S., 2021. A high-resolution map of hawaiian ulvz morphology from scs phases, *Earth planet. Sci. Lett.*, **563**, 116885. doi: 10.1016/j.epsl.2021.116885.
- Julian, B. R., 1970. *Ray tracing in arbitrarily heterogeneous media*, Tech. rep., Massachusetts Inst of Tech Lexington Lincoln Lab.
- Li, J., Sun, D. & Bower, D. J., 2022. Slab control on the mega-sized north pacific ultra-low velocity zone, *Nat. Commun.*, **13**(1), 1–11.
- Li, Z., Leng, K., Jenkins, J. & Cottaar, S., 2022. Kilometer-scale structure on the core–mantle boundary near hawaii, *Nat. Commun.*, **13**(1), 1–8.
- Malinverno, A. & Briggs, V. A., 2004. Expanded uncertainty quantification in inverse problems: Hierarchical bayes and empirical bayes, *Geophysics*, **69**(4), 1005–1016.
- Martin, C., Bodin, T. & Cottaar, S., 2023. Mapping structures on the core-mantle boundary using Sdiff postcursors: Part II. Application to the Hawaiian ULVZ, *Geophys. J. Int.*, **X**(Y), Zgji.
- Mosegaard, K. & Tarantola, A., 1995. Monte carlo sampling of solutions to inverse problems, *J. geophys. Res.: Solid Earth*, **100**(B7), 12431–12447.
- Osher, S., Cheng, L.-T., Kang, M., Shim, H. & Tsai, Y.-H., 2002. Geometric optics in a phase-space-based level set and eulerian framework, *J. Comput. Phys.*, **179**(2), 622–648.
- Pachhai, S., Tkalčić, H. & Dettmer, J., 2014. Bayesian inference for ultralow velocity zones in the earth’s lowermost mantle: complex ulvz beneath the east of the philippines, *J. geophys. Res.: Solid Earth*, **119**(11), 8346–8365.
- Pachhai, S., Li, M., Thorne, M.S., Dettmer, J. & Tkalčić, H., 2022. Internal structure of ultralow-velocity zones consistent with origin from a basal magma ocean, *Nature Geoscience*, **15**(1), 79–84.
- Russell, S., Rudge, J.F., Irving, J.C. & Cottaar, S., 2022. *A re-examination of ellipticity corrections for seismic phases*, preprint (arXiv:2205.08229).
- Sambridge, M., Braun, J. & McQueen, H., 1995. Geophysical parametrization and interpolation of irregular data using natural neighbours, *Geophys. J. Int.*, **122**(3), 837–857.
- Thorne, M., Garnero, E., Jahnke, G., Igel, H. & McNamara, A., 2013. Mega ultra low velocity zone and mantle flow, *Earth planet. Sci. Lett.*, **364**, 59–67.
- 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. doi:10.3390/min10030211.
- van Driel, M., Krischer, L., Stahler, S., Hosseini, K. & Nissen-Meyer, T., 2015. Instaseis: instant global seismograms based on a broadband waveform database, *Solid Earth*, **6**, 701–717.
- Vanacore, E., Rost, S. & Thorne, M., 2016. Ultralow-velocity zone geometries resolved by multidimensional waveform modelling, *Geophys. J. Int.*, **206**(1), 659–674.
- Wéber, Z., 2000. Seismic traveltimes tomography: a simulated annealing approach, *Phys. Earth planet. Inter.*, **119**(1), 149–159.
- Yu, S. & Garnero, E., 2018. Ultralow velocity zone locations: a global assessment, *Geochem. Geophys. Geosyst.*, **19**, doi:10.1002/2017GC007281.
- Yuan, K. & Romanowicz, B., 2017. Seismic evidence for partial melting at the root of major hot spot plumes, *Science*, **357**, 393–397.